



eReSS Study

Initial phase report

Standards and technologies in the VRE domain

September 2006

A. Introduction

The eResearch Standards and Specifications (eReSS) study has been funded by the JISC to undertake a review and assessment of the use of standards, specifications and relevant technologies within the eResearch community, and particular in relation to their applicability to the development of Virtual Research Environments (VREs). The study is taking place over the period 2006-9, allowing a long-term view of how standards and technologies can be and are being used in the VRE domain. It will run in parallel with the JISC-funded VRE Phase 2 Programme as well as continuing VRE Phase 1 projects and will provide support and information to the projects in coordination with the JISC Programme Manager.

eReSS is being undertaken by a consortium coordinated through the University of Hull, involving national as well as international experts in a variety of technical and standards fields. The project commenced in March 2006 with an initial six-month phase to capture current usage of standards, specifications, tools and technologies within the current VRE Programme and issues for future consideration. This work is forming the baseline upon which subsequent activity within the study will be based, and is summarised in this document.

Why eReSS?

eResearch, in both the sciences and humanities, has attracted much attention in recent years, and has also, through programmes such as the e-Science Core Programme, received funding to investigate the capabilities that eResearch can provide. Within such developments there is much ongoing software development. Unfortunately, software developers often fail to use standards when developing this software, leading to difficulties when bringing systems together to facilitate the interaction between research groups that eResearch promotes: similarly, software development also risks duplication if research groups are working in the same area. The costs of “patching” interoperability to make systems work together and the costs of duplicating research are high. This is not by intent, though. When software developers are aware of applicable standards, they adopt them. When researchers duplicate effort—often identified after the work is completed and published—they wished they could have moved much more quickly and produced more relevant results.

Three factors are involved in this loss of productivity.

- Having a source of information that can provide the link between a project and appropriate standards/specifications or work that others have done elsewhere. The Internet and sophisticated search tools are helping, but the exponential growth of information and data makes the task more difficult. An intermediary is required to assist developers and researchers find the resources they need.
- The pressure for speed. Knowing of the individuals and organisations that are contributing to different fields can be one of the most productive sources of immediate and relevant information to ensure eResearch can take place quickly.
- Globalization. A group in another country is often carrying out the work of interest, and communication may falter, preventing awareness. Conversely, some open source software projects, such as uPortal, intentionally create links among developer groups around the globe, enabling, for example, work in France to benefit users in Australia and implementations in the U.K. to serve as references for Japanese universities.

The eReSS study is focusing on addressing these factors by providing a source of information that can be relied upon to guide software developers and researchers in the right directions, both to facilitate adoption of standards, and especially open standards, and ensure complementary research. The study is seeking to provide an authoritative and sustainable living reference point that, through experience and development, can serve the VRE and eResearch community in the future.

Aims and objectives

The main aims and objectives of the study overall are to:

- Review on an ongoing basis interoperability standards, frameworks and technology platforms and identify their applicability to VRE development activities. It is recognised that these might have their origin within the eResearch arena or have been developed by other sectors and have potential applicability to eResearch.
- Review current VRE development work and identify requirements for interoperability. This task will be imperative to provide a basis upon which the pragmatic implementation of appropriate interoperability recommendations can be made.
- Propose suitable interoperability solutions and provide guidance to current and future VRE projects in their implementation.
- Propose recommendations to the JISC on the development of the e-Framework for Education and Research. It is clear that experience within the e-learning and information environment sectors can contribute to eResearch interoperability, and it will be essential to reciprocate this where applicable to ensure the fullest view of education and research the e-Framework can provide.

Work has initially been divided into an initial phase, which this report covers, and a secondary phase, which is being scoped. Subsequent reports will provide snapshots of information gathered as well as highlighting key issues of relevance to VRE development.

Gathering information

The eReSS consortium has proposed the use of a wiki as the means through which information on the use of standards and specifications will be captured, organised and presented. The wiki is being structured into three sections, each presenting information on particular specifications and standards in different ways. The three sections are:

- An overview and synthesis of information about standards and technologies
- Details of each standard (facts and figures)
- Contextual use of the standard

Contributions to the wiki are being gathered and sought from a number of sources.

- Consortium members are contributing knowledge in their areas of expertise
- VRE 1 project staff have been invited to contribute their knowledge of using standards and technologies within their projects, bringing real world VRE experience to the wiki
- Where gaps are identified those with relevant knowledge are being invited to contribute: voluntary contributions are also welcome.
- VRE 2 project staff will be invited to make their own contributions as the work of their projects develops.

There is recognition that information about standards and how they are being used is continuously developing. It will be the ongoing nature of contributions from all parties that will lead to the development of the living reference point that eReSS is building, and ensure its sustainability beyond the current lifetime of the study.

The study has a coordinating editor in place, who will focus their attention on compiling the synthesis section for each standard. Such syntheses will develop over time as information becomes available and will provide an overview of the development of standards and technologies within their lifecycle.

Next steps

Development of the wiki has been internal to the eReSS consortium to date. It is planned to open up the wiki as part of the preparation for the VRE 2 Programme, to help inform project planning and widen awareness of the study. The eReSS wiki aims to be a reflection of the knowledge and expertise that is being built up through the JISC VRE Programmes. Contribution from those closest to VRE development is vital and the wiki offers a route for disseminating experience for others to benefit from.

In addition to capturing direct VRE and eResearch knowledge of relevant standards and technologies, it will also be essential to coordinate this with standards activities in other sectors. To this end, communication with representatives of UKOLN, CETIS and the JISC has taken place: a first meeting to facilitate the coherent provision of standards information on the web, including the eReSS wiki is scheduled for 17th October 2006.

eReSS will be a developing resource over its lifetime as VRE and eResearch developments maintain their momentum. If one thing is certain, by the end of the study's lifetime in 2009 we shall know far more than we do today about how VREs and eResearch can facilitate future research & development: and much of this knowledge will be captured within the eReSS wiki.

B. Standards and technologies

The initial phase of the eReSS study has been required to take account of the high level of existing activity already. The VRE 1 Programme has been active since Autumn 2004 and eResearch activity for some time prior to this. The use of standards and technologies over this time has been widespread but often uncoordinated. eReSS is capturing this background to act as a baseline to inform subsequent activity and help provide a level of coordination.

Standards

Standards can take many forms and have many origins. Many are 'de facto' standards, technical approaches that have become widely used through practice and word of mouth. Many others have been developed to meet a set of requirements and their development has been overseen by a ratifying organisation. Many standards are closed, managed by their creators but not available for the community to take forward themselves (e.g., Adobe's PDF standard): many others are open standards, developed in an open way such that anyone can contribute and ensure that the final outcome meets the needs of those who will be using the standard in practice.

eReSS will largely focus on the use of open standards in its information gathering, as these open standards permit wide adoption and exchange of experiences within the VRE community, as well as facilitating interoperability. The study will not, though, neglect other standards developments and will keep track of these where there is potential applicability within VREs.

Identifying relevant standards for eResearch, and related sectors, is not a new activity. eReSS has sought to build on this experience rather than repeat existing work. In particular, work carried out through the JISC e-Framework Working Group has been of value. A glossary of standards is available at <http://www.grids.ac.uk/ReDRESS/glossary/glossary.html>: eReSS has grouped and expanded this glossary for presentation through the wiki.

Analysis of available standards has led to the following categories being proposed.

Standards category	Description	Examples
Accessibility	Standards concerned with the accessible presentation of information to all end-users, taking into account different accessibility needs.	CSS, WAI-WCAG, XHTML, XSL(T)
Business processing	Standards that are designed to structure interactions between different systems	WSDL, UDDI, UML, BPEL, Quali
Communications	Standards that are involved in communicating and passing information between systems	HTTP(S), JMS, RSS, OAI-PMH, SOAP, SRW/U, POP
Data	Standards that determine the organisation of information, including metadata	Dublin Core, MPEG-7, SVG, XML
DRM-License	Standards for structuring and stating the rights of access to information or software	BSD, MPEG-21, ODRL
Frameworks	Standards that provide the scaffolding within which interactions can be built	OGCE, WSRF, Globus Toolkit
Information	Standards for structuring information	RDF, OWL
Presentation	Standards concerned with the presentation of information	HTML, PDF, ECMAScript
Security	Standards concerned with controlling and managing access	Active Directory, LDAP, Shibboleth, X.509

Standards are not exclusive to any one group and many will fit into more than one according to their role. For example, the OpenURL standard is a means of organising information and also a means for communicating this information between systems. There is also a clear overlap between Accessibility and Presentation, though it is considered that the former requires special attention to ensure it is properly addressed when developing software and systems.

Practical experience¹

Practical experience of standards within VRE 1 projects has been widespread. Notwithstanding the intended investigation of particular standards, many projects apply standards by default, for example through the use of HTML or HTTP(S). These are employed because of the interoperable basis they then provide upon which projects such as those within the VRE Programme can take place. It is to be hoped that investigation of specific standards within the VRE projects can assist in moving these to a stage where they are employed by default in the future.

Accessibility – The use of accessibility standards is an evolving one, guided by the latest guidelines available. WAI-WCAG is widely followed, though is clearly stated as current best practice only, not definitive. The adoption of accessibility standards is driven by external factors, most often legal factors, though there remains ongoing debate about the best ways to achieve the accessibility required. Nevertheless, the use of accessibility standards is moving towards a default (as described above) as a result of these external factors. As such all VRE 1 projects have sought to be accessible as required.

Business processing – The standards within this group are largely Web Services based, addressing the mechanisms through which business processes can be enabled technically. They have not been widely used within VRE 1, though for certain projects (GROWL, IUGO,) Web Services have been at the centre of the intended outputs. A number of others have explored the

¹ Project names and acronyms are as used in the eReSS wiki

use of WSDL and related Web Service standards as part of their work. Web Service standards are still very much in development and there is a potential issue about adopting a moving target within a project. Nevertheless, wide adoption, required because of the binary nature of Web Service interactions, will assist in validating their role and benefit. UML, a non-technical planning and design standard, has been used within related JISC Programmes (Digital Repositories, e-Learning) but was not formally adopted within VRE 1.

Communications – The Communications group contains standards that facilitate many different levels of communication: email communication (IMAP, POP, SMTP), videoconferencing (H.323, RTP, RTSP), internal software communication (JMS, SOAP, WS-*), web communication (HTTP(S)), communication of search parameters (Z39.50, SRW/U), communication of metadata (OAI-PMH, RSS), standards that facilitate communication through identification (URI), and file transfer (WebDAV). These have been used within VRE 1 projects extensively to a greater or lesser degree. Memetic, SAGE and PoIDis have investigated videoconferencing via the Access Grid. EVIE has examined a number of communication standards in the context of a portal framework and the services that can be provided through this: a key service, searching via SRW/U or Z39.50, has also been explored by Cheshire-Kepler-Sakai and BVREH. Sakai EduResearch, CORE, and IUGO have used OAI-PMH and/or RSS to communicate metadata, and a number of projects (including Cheshire-Kepler-Sakai, CORE, and IBVRE) have looked to work with Web Services.

Data – Many of the standards used under the banner of communications also apply in this grouping. In addition, Dublin Core metadata has been frequently used as part of using other standards, for example as a standard response format in SRW and OAI-PMH (Sakai demonstrator and EVIE); MPEG-7 has been used to describe multimedia content within an Access Grid context (Memetic), whilst XML and, in one case, SVG, has been used to organise data (Cheshire-Kepler-Sakai, ELVI, EVIE, Memetic and OGHAM)

DRM-License – No practical implementation or testing of licensing standards (either as licences or standards for describing licence terms) has specifically taken place in VRE 1 projects, though there remains the possibility that software outputs will emerge that require appropriate licensing. This can be guided by JISC's Open Source policy and guidance from OSS-Watch. It is notable that the choice of technologies to use within VRE projects has, though, been influenced by open source availability in many cases and it will be useful to assess the relative impact of open source vs. standards compliance in deciding on VRE platforms.

Frameworks – Because of the developing nature of many framework standards, many projects stated an intention to comply where feasible, or an intention to track a standard for potential use. JSR 168 has been used for portlet development (CORE, IBVRE, ISME, IUGO, Sakai demonstrator): JMF has been used by SAGE to assist general programming. Framework standards have also been used to guide development, rather than be used specifically within it. For example, WS-I and WSRF have been used as a guide to support interoperability (GROWL): in particular, WSRF can be used to help extend other Web Service standards and in turn can itself be extended.

Information – The use of RDF and OWL has been largely an prospective one within VRE 1 projects, though they have been actively used within Sakai EduResearch and Memetic.

Presentation – Many presentation standards fulfil the criteria indicated above of default standards that need to, or should, be used for delivery via the Web (HTML, XHTML in some cases). Some projects have looked beyond these to provide alternative output formats (e.g., ECMAScript in the OGHAM project). Elsewhere, standards have been used to establish presentation in novel ways: content has been displayed in 3D in PDF using U3D (ISME) and SVG used within the OGHAM project. Presentation is also key to accessibility and much presentation work has been based on a need to address accessibility issues.

Security – Although not an area widely investigated, a few VRE 1 projects have been conscious of the need to apply security within VREs, looking to make use of one or more of Active Directory, Shibboleth, LDAP and Kerberos. The Globus Security Infrastructure (GSI) has been used within the GROWL project and a number have tracked Shibboleth development, particularly the Sakai demonstrator project. Other potential security standard components in a security infrastructure include X.509 and SSL.

Technologies

In implementing a VRE it is unlikely that individual components using a single standard will be used (though they may be incorporated within the overall environment). Applications and frameworks will be combined to provide the overall virtual research environment. Technologies are defined by eReSS as compound applications or systems that use a range of standards within them. eReSS will identify relevant technologies that are applicable to VREs and eResearch and highlight both the use of standards within them and where standards are not being used. Quite often a technology is implemented because of stated standards compliance only to discover that this only goes so far. Suggestions for appropriate technologies to investigate are welcome, especially in areas of new standards development.

An initial set of technologies has been identified and information on these is being gathered.

Technology	Description	Examples of standards used
Access Grid	This videoconferencing system is being used by a growing number of universities, facilitating eResearch interaction	RTP, RTSP, MPEG-7 and related video and streaming standards
DSpace	An open source repository developed by MIT	OAI-PMH, Dublin Core and other metadata standards, SRW/U
Fedora	An open source repository framework developed by Cornell and Virginia Universities	OAI-PMH, Dublin Core and other metadata standards, RDF
GridSphere	A Java-based open source web portal developed for specific use with Grid tools	JSR 168, WSRP, JSF
Sakai	A Java-based open source collaboration & learning environment that has been used within a number of VRE 1 projects	SCORM, XPath, JSR 168, IMS CP, IMS QTI, IMS TI
uPortal	A Java-based open source web portal developed for use in Higher Education and used within a number of VRE 1 projects	JSR 168

It will be immediately apparent that there are likely to be many technologies that could be potentially used within VREs, and which will have issues with the standards they are using (where they are using them – new research will frequently be breaking ground where there are no standards yet). eReSS will endeavour to address technologies where relevant and most applicable to VREs, accepting the broad sweep of research that a VRE can potentially cover.

Practical experience²

Access Grid – this has been used within the Memetic, SAGE, PoDis, and ISME projects. Memetic investigated the importance of meetings in fostering collaboration and how technology can assist, aiming to make online meetings as productive as face-to-face. It is building on tools developed within the e-Science Programme’s CoAKTinG project (Collaborative Advanced Knowledge Technologies in the Grid) to provide a structured annotation layer over video: MPEG-

² Project names and acronyms are as used in the eReSS wiki

7 is used to describe this multimedia content. A separate tool, Arena, has been developed for recording Access Grid sessions. The RTSP/RTP common standards are being used as part of Access Grid and Arena and are important for extensibility of the Arena tool into H.323 videoconferencing in general. Elsewhere, SAGE has developed the stereoscopic use of Access Grid through an extension, whilst PoIDis has focused on the use of Access Grid for teaching a postgraduate MA: they are looking to foster collaboration using a combination of Access Grid and a wiki within Sakai. ISME used Access Grid to facilitate operation of the project itself across a distributed number of partners: this was affected by restrictions on Access Grid use in France where there is much more controlled access to this technology.

DSpace/Fedora – DSpace has been used by the Sakai EduResearch project: Fedora has not been used within VRE 1, but has been used in an eResearch context (e.g., the US National Science Digital Library (NSDL))

The Sakai EduResearch project has used DSpace in tandem with collaboration tools (Sakai) to promote data sharing and report dissemination: its OAI-PMH interface has been used to help build theme-based websites using RDF metadata based on the content captured within the repository. The use of OAI-PMH has resulted in the use of Dublin core metadata, mandated as the basic requirement within OAI-PMH. This is a case of one standard influencing the use of another (albeit that other metadata standards can be used with OAI-PMH). The project also found that this was a wider issue, with potential constraints on the use of standards being imposed by applications. Fedora provides another repository platform that VREs could use. It is currently used mostly to support large and often complex digital collections, but these have the potential to be used within a VRE, e.g., the NSDL usage for an online library of scientific-related resources. Fedora uses RDF and Dublin Core metadata internally and presents an OAI-PMH interface for harvesting.

It is recognised that the repository field is rapidly evolving. As such, other standards will be adopted over time as they mature, particularly in the areas of preservation and search, and as digital repository software makes them available.

Sakai – this has been used within the Sakai EduResearch, Sakai demonstrator, PoIDis, and IUGO projects.

Sakai conforms to a number of standards, mainly within three areas: W3C, learning content (IMS), and presentation (information on these has been provided by Anthony Whyte from the Sakai Foundation staff and is available via the wiki). The Sakai EduResearch project has worked with all of these, noting the dependency on external software development factors, but with the aim of using the standards where they are available. The Sakai demonstrator project has focussed its investigations around the use of portlet standards for presentation of services. IUGO has approached the issue from the other way round, investigating Sakai as the potential recipient of services that conform to presentation standards such as WSRP and JSR 168. In the PoIDis project's use of Sakai, the potential of using the standards it offers are the facilitation of re-use in different circumstances, reducing duplication of effort.

GridSphere/DotNetNuke – these have been evaluated by the ISME project, which looked at a number of platforms to establish a shared workspace. A key factor in this evaluation was the ability of technologies to adhere to other related standards (e.g., portlet standards and WSRF).

uPortal – this has been used within the ELVI, EVIE and IBVRE projects.

The ELVI project investigated the requirements for a large-scale VRE in an institution that was already running an institutional portal based on a commercially supported version of uPortal, Sungard's Luminis. From the start, therefore, there was recognition that this product may constrain full use of relevant standards, at least at this time. Evaluation of VRE requirements needed to be based on a balance between what was desired and what was practical, an issue that also faces developments where the standards or technologies available are not as mature as they will be. The EVIE project was based at an institution that adopted Luminis halfway through the project, and has had to adopt its work to meet local factors. The project has recognised that the use of a portal to underpin a VRE requires compliance with a wide range of standards,

including those for email, middleware, distributed searching and retrieval, security, and accessibility. The IBVRE project followed the same path as IUGO, initially focusing on underlying presentation standards and later adopting uPortal to help put the outputs from these investigations into practice.

Architectures

Alongside individual standards and specifications, and related technologies that make use of these, there have been a number of developments taking a broader view of the technical environment that have applicability to eResearch and VREs.

- Open Grid Services Architecture (OGSA)
- Service Oriented Architecture (SOA)

OGSA can be regarded as the blueprint for the Grid's architecture. It defines a set of core services and looks to standardise these, thus achieving interoperability and extensibility of use. SOA is an approach to developing services that allows individual components to interact as required, and adapting to change flexibly over time.

These architectures are intended to be high level, identifying requirements that can be met through adoption of the standards and specifications indicated above. As experience with implementing these architectures is gathered it is likely that they will evolve over time. However, they provide both a structure and guiding framework within which adoption of standards can take place.

C. Issues arising

The analysis of standards and technologies, and evidence of their usage, has led to the identification of a number of issues that need to be considered in the development of VREs, as well as proposing routes for addressing these and allowing further development to take place.

1. *User Facing Integration, Customisation, and Personalisation*
 - a. It is notable that there has been a high level of usage of both the Sakai Collaboration Environment (as specified in a number of project plans) and the uPortal web portal framework (as adopted subsequently by a number of projects) within VRE 1 projects. Both systems seek to offer an integrated and potentially customisable environment that can be adapted to specific needs.
 - b. Whilst projects in VRE 1 have undertaken valuable work with both the Sakai Collaboration Environment and uPortal web portal framework, the rich integration of portal framework and collaboration environment to provide a framework for fully personalised and customisable research space has not been addressed in detail. Such a combination provides a potentially powerful, rich and flexible toolset to address the needs of a eResearch, being carried out by a Virtual Research Community (VRC), formal or informal research groups, or individual researchers.
 - c. The decomposition of services provided by a Collaboration Environment such as Sakai and the recomposition of those services surfaced within a portal framework, or by desktop applications, may have significant impact on embedding VRE functionality within an institutional context, particularly given the increasing number of institutional portal deployments within the UK and internationally.
 - d. Strands of work/investigation to support these approaches could include;
 - i. Technical: assess framework requirements and adaptations of frameworks where appropriate. (For example WSRP/JSR 168 > WSRP

- 2.0/JSR 286 – whilst noting that these specifications at present represent something of a moving target).
- ii. Human Factors: needs/requirements/desires analysis from a variety of users in a variety of contexts
 - iii. Modelling how this approach could begin to more seamlessly integrate
 - 1. Research/Collaboration tools
 - 2. Administrative information relevant to research process
 - 3. Resources available from IE, Libraries, remote datasets
 - 4. Institutional repository initiatives – the scholarly communications process and data storage/access
 - iv. Web/desktop interface: the use of Sakai and uPortal has suggested the development of VREs in web space predominantly. Many researchers work with tools on their desktop: establishing how these can integrate with the web environments will allow researchers to blend their own activities with those of the wider VRE community.
 - e. Web 2.0 – Since the VRE 1 Programme commenced a variety of tools and technologies have emerged and coalesced around the concept of Web 2.0. This concept represents a trend toward promoting the flexible re-use and re-purposing of information. Although not directly connected with conventional research space it is likely that a Web 2.0 approach can benefit VRE development and support VRCs: this is especially likely to be the case over the lifetime of the VRE 2 Programme as the use of Web 2.0 technologies matures. There is scope in evaluating the Web 2.0 approach for its relevance and value to VRE space, and potentially the testing of integrative tools that adhere to this concept.

2. *Extensibility*

The VRE 1 projects have demonstrated a wide range of different tools to support the research process. The majority of these have been designed for use in specific circumstances and in specific domains. In order to address the interoperability of such tools there is value in investigating the extensibility of these tools for use in other research domains and circumstances. The following areas of work suggest themselves to achieve this goal:

- a. Further work on the development of portlets as pluggable spaces that enable user-facing integration. The development of more mature portlet standards (WSRP 2.0 and JSR 286) has been noted already and these might be the source of extensibility. It should be noted within this that the use of commercially supported open source systems, such as the Luminis uPortal-based portal from Sungard, can potentially compromise such extensibility and interoperability across all uPortal users by virtue of the different software versions and extensions being employed (the version of uPortal in Luminis is typically a couple of releases behind the Open source product). This variation on versions requires close attention to ensure interoperability applies both within and between versions wherever possible.
- b. Additional investigation of the variety of ways in which Sakai can embed information/services and have these embedded within itself, in order to allow VRE tools to be extended into the collaboration space and from the collaboration space into other areas.
- c. The desire for extensibility suggests the need to be as flexible as possible in promoting the re-use of different tools and services (and indeed promoting an eResearch variation of Web 2.0). A move towards a service-oriented architecture will allow such extensibility and warrants greater attention across VRE 2 projects to facilitate interoperability between them.

3. *Workflow*

Issues surrounding workflow have been raised by a significant number of projects in VRE 1, and it is the focus of considerable current activity inside and outwith the eResearch community. This

is unsurprising, given the programme focus of supporting research lifecycle/s, and the potential role for lifecycle activity to be underpinned by systems supporting workflow. It should be noted, however, that the majority of systems embedding or instantiating workflow, and the specifications and standards supporting them, might best be described as “emergent” and relatively immature. They tend to be either designed around workflow for fairly generic business processes (“top-down” examples of this are the BPEL family of specifications emerging from OASIS, a “bottom up” example being Kualo workflow and notifications emerging from work undertaken by Indiana and Cornell in the Kualo Project), or designed around domain specific needs. Examples of the latter include Kepler and Taverna from the eResearch community, and arguably Learning Design and admixtures of SCORM and Simple Sequencing in “learning flow”.

The term “workflow” may not always be used consistently, then, and the somewhat nuanced interpretations across domain boundaries are amongst the issues for eReSS to elaborate over the remainder of the study. The following broad questions have risen sharply in the first six months of eReSS, and suggest themselves as the focus of future investigation or potential development effort

- a. How far do emerging generic workflow specifications such as BPEL or Kualo workflow enable the modelling of research processes, especially research administrative processes? Given the labour intensive nature of much workflow design, are additional tools required, particularly for visualisation and ease of adoption, adaptation and deployment?
- b. How far can generic workflow specifications such as BPEL or Kualo workflow be loosely or tightly integrated with existing frameworks supporting eResearch (e.g., Sakai, uPortal, institutional repositories). What integration points exist or need to be added?
- c. What are the meeting or touchpoints between generic workflow specifications such as BPEL or Kualo workflow and specialist scientific, instrumentation, or other research workflows such as Taverna or Kepler? Is a high level of workflow choreography necessary or desirable in the context of a VRE, or research process? How might this be triangulated with 3b above? Is the development of an Enterprise Service Bus (ESB) approach relevant in this space, either from the perspective of eResearch systems and resources integration, or their integration with institutional and other systems?
- d. How might any of the above interface (in technical or other terms) with the research publications process explored in JISC’s Digital Repositories Programme? What other overlaps exist between workflow in the space mapped out as “VRE” and that covered by IR’s?

4. *Validation*

The VRE 1 Programme has potentially eliminated a significant amount of wheel-re-invention by the evaluation of collaboration environments and portal frameworks in a limited range of research contexts. It should be noted that the primary technologies used and validated within VRE 1 have partial support for common specifications and standards, and that the environments in question are undergoing rapid evolution. This is to be expected as the communities developing these technologies continue to balance innovation and interoperability. VRE 2 should continue to undertake.

- a. Further validation of collaboration environments in a variety of research collaboration contexts
- b. Further validation of portal frameworks in a variety of research collaboration contexts

5. *Documentation*

Whilst not specific to this programme, or to the open source software (OSS) “products” used by it, issues around adequate (or, more truthfully, inadequate) documentation have been raised by

several projects. In part this may reflect requirements of cultural accommodation to OSS development, where documentation is typically more distributed and less conventional. Whichever is the case, in whatever proportion, this remains an issue of concern, as extensibility and interoperability of tools and resources in terms of technology is significantly reduced in practice by inadequate documentation. A number of actions suggest themselves as worthy of further consideration.

- a. It may be possible to design “bridging work” to culturally acclimatise projects in VRE 2 to work in OSS communities (perhaps in conjunction with OSS Watch).
- b. Promote good documentation practices in development activity, ranging from code comments and appropriate engagement in discussion lists through to more formal documentation.
- c. Fund study/ies of successful OSS documentation work, and attempt to draw/generalise lessons from these.

6. *Security*

An overview of the standards mentioned by VRE projects in their proposals and project plans reveals that whereas some security standards are mentioned by a few projects the majority make no mention of this aspect of developing VREs. Whilst this may have been appropriate in a first cut of projects investigating VRE space, subsequent VRE activity will need to address this area more proactively to ensure that researchers are not compromised by working in a virtual environment. Two factors will influence these investigations.

- a. It will be important for future VRE developments to ensure that consideration of security does not take place in isolation of the VRE’s wider environment. This will build on the awareness of this issue that VRE 1 projects have developed. Whilst for certain applications individual security requirements may be required, the advent of widely adopted security infrastructures, for example Shibboleth for resource and service access and the GSI for grid applications, mean that VREs can tap into and take advantage of available security as required without re-inventing the wheel. This will assist with the embedding of VREs within and across institutions.
- b. In the light of use of the uPortal and Sakai platforms within the VRE 1 projects, how these and other intermediary frameworks manage secure access to remote resources will also require further investigation. Work has already taken place, notably through the SPIE project within JISC’s AAA Programme and the Sakai demonstrator project in VRE 1, and it is vital that this is built on to ensure these and related platforms are able to provide the appropriate levels of security for access to the services and resources they present.

7. *Grid standards*

In the development of the grid to support eResearch there are many different technologies involved. Due to the need to ensure that the grid can be used nationally and globally there are also a large number of standards activities to support interoperability between different parts and facilitate eResearch across distributed research groups. These cover a wide spectrum of functionality and also a wide number of standards bodies, including W3C, Internet2, OASIS, GGF, DTMF, Liberty Alliance and WS-I.

- a. As standards emerge from these organisations they will need to be assessed against existing standards and how they fit in to the wider picture. Due to commercial pressures there are often competing standards, particularly in the Web Services area, and it can be uncertain which standard to adopt.
- b. The development of new standards for the grid has, in many cases, followed the model of creating the standard up front to help inform practice in this rapidly developing area, as opposed to a standard emerging from agreed practice over time. It is important that standards, where they fulfil appropriate criteria, be

tested in real world situations to see whether the theoretical nature of their origins stands up to practice.

- c. Alongside testing and monitoring, the development of standards can be influenced through participation in the standards bodies listed above. This requires commitment, though can ensure that a standard is going to achieve and deliver what is required.

8. *Digital Rights Management*

DRM is a widespread issue that affects not just VREs but all digital environments, especially where externally sourced applications and/or resources are being used. Whilst there is ongoing activity to examine the issues of DRM broadly, it is important to establish the different contexts and perspectives that will affect how DRM solutions are implemented. A number of the JISC Digital Repositories Programme projects have incorporated a DRM workpackage and it is appropriate for VRE 2 projects to do likewise. This may fit best alongside considerations of security (see above) because of the close relationship between the two.

D. Standards in VRE development

Information gathered so far from projects has suggested a number of possible slices that can be cut across the wiki, highlighting trends that are emerging from often disparate projects addressing different areas. Barriers to the implementation of standards, for example, are varied, but it is hoped can be identified to prevent duplication of bad experiences. There is also the concept of a standards lifecycle, tracing the life of a standard from conception through to ratification and obsolescence, and usage throughout this period. Both of these will influence decisions about when and where to adopt a particular standard.

In all VRE development it is urged that where technologies are to be used that their use of standards be carefully assessed to prevent problems occurring further down the line. The relationship and link between areas of functionality and potential standards to support that functionality should also be assessed to identify open and standard means of implementation that will help achieve interoperability. These initial assessments will be vital in facilitating interoperable eResearch not just within the domain of an individual VRE development but across different, distributed research groups as well.

Conclusions

The use of standards by VRE 1 projects has been widespread if limited within the domains that each project is working on. Ideally, a goal of 'infrastructure independence', a term used within the Cheshire-Sakai-Kepler VRE 1 project, is the long-term goal, allowing individual components of a VRE to be swapped out and in as required, simply through the use of relevant specifications and standards and the interfaces these provide. We may be some way from reaching this goal as yet, but the VRE 1 projects have moved the eResearch community forward towards this goal and there is little doubt that the VRE 2 projects will be able to do likewise.

The eReSS Consortium
September 2006