Cover Sheet for Bids ***(All sections must be completed)***	

Name of JISC Initiative: JISC UMF Cloud Computing Programme

Name of Lead Institution: University of Oxford

Name of Proposed Project: Virtual Infrastructure with Database as a Service (VIDaaS)

Name(s) of Project Partners(s) n/a

This project involves one or more commercial sector partners **~~YES~~ / NO**	**Name(s) of any commercial partner company(ies)**

Full Contact Details for Primary Contact:

Name: Professor Paul Jeffreys **Position:** Director of IT
Email: paul.jeffreys@odit.ox.ac.uk **Tel:** 01865 273229
Address: University Computing Services, 13 Banbury Road Oxford, OX2 6NN

Length of Project: 12 months

Project Start Date: 01/04/2011 **Project End Date:** 31/03/2012

Total Funding Requested from JISC:	£ 1,201,469
Funding requested from JISC broken down across Financial Years (Apr-March)	
1 April 2011 – 31 March 2012	£ 1,201,469
Total Institutional Contributions:	£161,391

Outline Project Description

The Virtual Infrastructure with Database as a Service (VIDaaS) Project aims to deliver cost savings and improve the efficiency of UK HE research through the creation of a production-ready hybrid virtual infrastructure (VI) hosting a 'Database as a Service' (DaaS) system. 'Hybrid' is used in the sense of having local and central VI components deploying common hypervisor technology. The hybrid 'Infrastructure as a Service' (IaaS) component provides a virtual infrastructure upon which the DaaS and other 'Software as a Service' components can be hosted. This involves enabling seamless capacity management and the sharing of infrastructure and services between universities and/or other academic institutions, with a HE VI based initially at Eduserv. The DaaS component will be developed from the prototype created for the Supporting Data Management Infrastructure for the Humanities (Sudamih) Project to a full production service, the prototype is currently in use by researchers in Oxford. Functionality will be extended to meet common requirements of researchers beyond the humanities and to cater for different database types besides the traditional relational database. From the user perspective, it will offer an intuitive means of creating, editing, and querying databases through a simple Web interface. The DaaS component will also be extended from the prototype to enable it to run in a hybrid virtualized environment. VIDaaS will create a full production-level service, implementing a standards-based federated identity and access management (IAM) solution and providing processes and solutions for billing and accounting. Researcher training will be a major focus of the project, with user support materials, courses, and documentation developed in close collaboration with the DCC. Oxford will also collaborate with Eduserv to develop the hybrid VI, and with JANET to investigate hosting DaaS on private VIs. A full ROI will be evaluated (an initial one is included in this bid) and a Business Plan created in collaboration with the DCC for development, and cost recovery operation beyond the life of the project. Oxford will offer 'warranty' until a full production service is offered for the HE community.

I have looked at the example FOI form at Appendix A and included an FOI form in this bid	**YES**
I have read the Funding Call and associated Terms and Conditions of Grant at Appendix B	**YES**
<u>For FE institutions only:</u> Please tick this box if you are an FE institution in England, please tick this box to confirm that you meet the eligibility requirement of teaching HE to more than 400 FTE	☐

Virtual Infrastructure with Database as a Service

1. Appropriateness and Fit to Programme Objectives and Overall Value to JISC Community

The VIDaaS proposal comprises two fundamental elements: the deployment of a 'Database as a Service' (DaaS) production-quality software service; and the implementation of a hybrid virtual infrastructure (VI) upon which Software as a Service applications (SaaS), including DaaS, can be developed, refined, run and managed.

The hybrid VI developed by VIDaaS will be based on VMware, the market leader in virtualisation technology. Oxford has extensive expertise having adopted the ESX platform in 2005, and with VMware Certified Professionals (VCP) who have presented at international Conferences. 'Hybrid' is used in the sense of having local and central VI components deployed with common hypervisor technology (although other hypervisors may be deployed in due course). The resulting Infrastructure as a Service (IaaS) will be secure, resilient, reliable, and cost-effective and necessarily include access management and accounting tools. A key aspect of the VIDaaS project will be to understand how to deliver the necessary management capabilities (security, identity and access management, usage tracking, billing, quality of service). The project will initially implement these within a local VI, and then extend the tools to a HE VI and beyond. The local VI will serve as a staging post for developing SaaS to operate on a HE VI based at Eduserv and will also have the potential to become a node on any distributed HE VI. The Eduserv VI is likely to serve in turn as a staging post for subsequent extension to the commercial cloud (Logicalis, Amazon, Salesforce, …), and also to HEI institutions that wish to make resources available from their VIs, all of which will be arranged through the JANET brokering service. The over-arching aim of VIDaaS, therefore, is to build on and transfer knowledge gained within Oxford University through both the Supporting Data Management Infrastructure for the Humanities (Sudamih) Project and the operation of enterprise-level virtual infrastructure so that the resulting IaaS and SaaS 'stack' has the potential for bi-directional operation between Oxford, Eduserv and other IaaS brokered through JANET.

The SaaS component of the project will deliver a production-quality DaaS service initially on the local VI platform, then replicated on the Eduserv platform and by Southampton for use by their researchers. In addition, through the IaaS approach outlined above, Oxford will investigate the capability of a local instance of DaaS seamlessly to make use of additional capacity via the Eduserv platform. Oxford will collaborate closely with Eduserv to develop the hybrid VI, and with JANET to investigate hosting DaaS on a commercial VI. Oxford will also collaborate with Southampton to pilot their DataDrive SaaS (offered via a combination of the local and Eduserv VI). The implementation of a hybrid VI is essential to enable institutional IT providers to understand the opportunities and challenges of offering production IT services in this way. VIDaaS can be seen as delivering a *platform* between researchers and the low level infrastructure services found on the commercial cloud, which will avoid users being locked in to any one commercial cloud provider.

The Sudamih Project originally intended DaaS to be hosted on conventional servers for use in Oxford, and a fully operational prototype has been delivered, and is currently being used by researchers. DaaS offers a simple and intuitive means to create or import databases of research data. It is intended to be usable without specialist knowledge or expertise, and to be sufficiently flexible to meet a wide variety of researchers' needs. It is targeted at small to medium database research requirements, and not, for example, large petabyte-scale activities (for which existing High Performance Computing (HPC) solutions are more appropriate). VIDaaS will deliver DaaS as a full production service with extended functionality to cater for a broader range of academic disciplines, operating on the hybrid VI described above. As one of the first HE production-quality services running on a shared hybrid VI much will be learnt in through the process of delivering it. An essential part of VIDaaS will be the development of training and support materials undertaken with the DCC to ensure sustainability and effective use of DaaS. The hybrid VI on which DaaS is provided will offer savings through economies of scale, whilst research efficiency improvements in the creation and use of research databases will be generated. A full ROI will be evaluated and a business plan created with the DCC for cost recovery operation beyond the life of the project.

1.1 Value for Money

Initial findings suggest that demand for the DaaS amongst humanities researchers conducting data-based research is high. A survey of researchers who attended the second Sudamih workshop indicated that 56% of attendees would definitely consider using the DaaS, whilst a further 44% gave a more cautious 'possibly'. None of the respondents said that they would not consider using the

DaaS at all. Ten respondents mentioned specific existing or planned databases which they would consider using the DaaS to support. We have received expressions of interest in the DaaS from researchers at Bristol, Leicester, Southampton and a number of other universities, and DaaS is already in use by Oxford researchers, and was demonstrated at the JISC Conference in March.

The Sudamih Project has been working closely with the Oxford Roman Economy Project (OXREP) as a DaaS testbed. The Assistant Director of the OXREP estimated that by using DaaS for work undertaken in 2010 rather than an Access database, they would have saved approximately 21% on staff time, primarily due to simplifications to database structuring and through controlling and standardising data contributions. He indicated that for other projects the DaaS would also be likely to save money by reducing specialist staff requirements.

It is estimated that the savings made by moving the OXREP to a centrally-hosted Virtual Machine will be even greater. The IT Officer of the Classics Faculty, which currently hosts the OXREP database and Web front-end on a departmental VM, is planning to move it to the local VI (as an initial step). The cost savings of so doing, when the staff time needed to look after the VI is taken into account, amounts to approximately 37%. Given the economies of scale offered by a centrally-supported VI, this saving may be expected to increase further as the infrastructure is enlarged.

Figures obtained from HESA indicate that in 2009/10 the UK research funding councils along with the Royal Society and British Academy invested a total of £48.4m into UK humanities research (including archaeology and modern languages). It is estimated that about a third of AHRC-funded research involves some sort of database component. If it is assumed that a similar proportion of projects funded by the other funders include a database component, and that switching to the DaaS hosted on a VI would save 20% on current costs (a much more conservative figure than that provided by our case study), this would save over £3 million per year in the humanities alone. Expanded to other academic disciplines, which attract far more funding than the Humanities, the savings could be very significantly higher.

An initial ROI has been calculated for the VIDaaS service based on the delivery of the DaaS alone (excluding efficiencies gained from effective use by researchers). This will be refined through the lifetime of the project, and form an important output from VIDaaS. A target DaaS systems has been defined (2-8 GB RAM, share of quad core CPU, 5GB disk, 50 million database transactions, 5.1GB/month in, 10GB/month out), and costed to include OS and platform maintenance. The costings, which are necessarily ‘best efforts’ at this stage, are given per year:

- Single physical server running 30 2GB database instances: £125
- Oxford virtual machine running on local VI with 100 2GB instances: £ 79
- Oxford virtual machine running on local VI with 100 8GB instances: £109
- Eduserv virtual machine running on VI with 500 8GB instances: £76-98
- Amazon virtual machine with 8GB instances: £660-744

Humanities projects with substantial database elements recently funded within Oxford range from £4.5k to £73k in allocation. This gives the scale of investment currently being made for each database which is likely to be in the range £4–10k.

A ROI can be calculated in a number of different ways; here it is estimated by considering the creation of a database and operating it over the duration of a year [full details of the analysis are available on request]. Considering within Oxford first, the cost of running DaaS on a VI is estimated to be 63% of running on a server (with 2GB RAM). This reduction is sufficient to make this an attractive proposition. In addition, Eduserv estimates that operating on their VI will be between 10-30% less than the equivalent Oxford VI figures (using 8GB RAM). Using the Amazon VI would be very much more expensive (£744 for a 2GB instance and £660 if 4 databases were run on a 8GB instance) as their provision is designed for short term use and not providing dedicated long term resources. It is quite possible the Amazon charging model will change, especially if JANET brokers a deal where the JANET backbone network is used, but for now Amazon is not competitive for use over the period of a year. It would be good if the HE community were able to benefit from the full economies of scale available through Amazon or one of the large commercial cloud providers in due course. All the costings are very provisional at this stage, but assuming delivering DaaS on an Eduserv VI is very approximately 50% of the cost of running on a server, the saving is estimated to be circa £60/database/year. These savings can be scaled by the number of potential users estimated above. A more detailed analysis of the costs and benefits of the DaaS will be used to inform future service provision (see section 2.8 for exit and sustainability plans).

2. Quality of Proposal and Robustness of Workplan

2.1 Technical Architecture

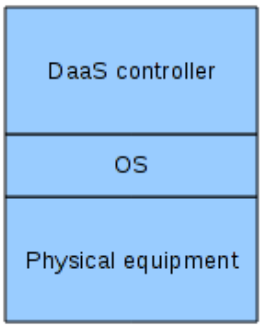

Figure 1. Sudamih DaaS architecture

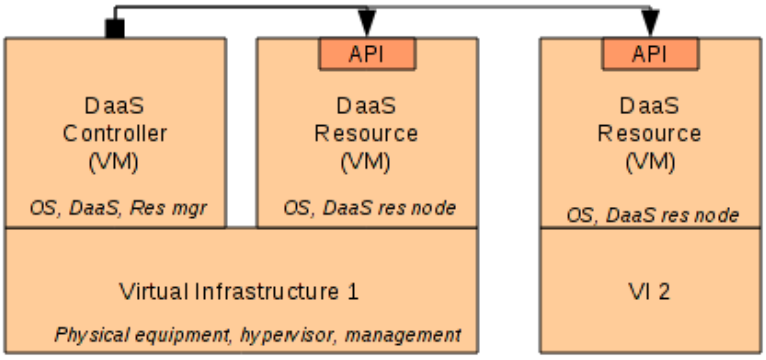

Figure 2. Proposed DaaS architecture

The outline architecture envisaged for the original implementation of the DaaS in the JISC funded Sudamih Project is shown in figure 1. In this model, the DaaS controller coordinates the use of dedicated physical equipment, and performs administrative operations such as database and Web interface creation, schema upload, and initial data conversion. Users interact with the DaaS controller through a Web interface, and the DaaS controller communicates with underlying database and web providers using existing application management interfaces on the physical servers held at the University of Oxford. Figure 2 illustrates the enhanced architecture that VIDaaS proposes to develop. In this model the system comprises a cluster of virtual machines instantiated from two ready-to-run virtual machine images. The DaaS controller is extended to include resource management capabilities necessary to coordinate the dynamic acquisition, management, and release of computation and storage resources. Individual resource nodes can be run on any number of virtual infrastructures, which may be sourced from different suppliers, offering a standardised API (integration layer) to the DaaS controller. In the case of traditional virtual infrastructure being used (virtual machine running on a hypervisor) the API can be presented directly by the virtual machine(s). Integration with cloud services (e.g. Amazon EC2, Microsoft Azure, or private Eucalyptus-based clouds) is also possible through the use of a lightweight adapter which provides linkage between the requisite cloud computing interface and the resource provider API. This enables immediate support for understood and mature virtual infrastructure, but also provides a high degree of flexibility to adapt smoothly as cloud computing matures.

The hypervisor will use VMware ESX4.1 in the initial phase of development. Later phases will use additional tools above the hypervisor (vSphere, vCloud director, Eucalyptus) to ensure compatibility across institutions and with the Higher Education virtual infrastructure. The overall approach in VIDaaS therefore is to deliver a robust production quality software product DaaS, deployed on an IaaS, a VMware platform, initially to serve researchers in the humanities, but followed rapidly by a more general version, and with DaaS capable of expanding its operation from a local VMware virtual infrastructure into a higher education VI at Eduserv and ultimately into the commercial cloud.

2.2 Project Plan

Work Package	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	FEB	MAR
WP1 – Project Initiation												
WP2 – Project Management												
WP3 – Virtual Infrastructure												
WP4 – Shared Services												
WP5 – Requirements Analysis												
WP6 – DaaS Productionization												
WP7 – DaaS Expansion												
WP8 – DaaS Service Design												
WP9 – DaaS to External VI												
WP10 – Documentation												
WP11 – Training												
WP12 – Costs & Benefits												
WP13 – Communications												
WP14 – Acceptance Testing												

WP1 – Project Initiation

Set up the project and the tools to assist management and dissemination. Includes preparation of a detailed project plan, website and blog. *[April – May 2011]*

- Establish project working environment, including management tools *(April 2011)*
- Build and publish project website & blog *(April 2011)*
- Undertake comprehensive stakeholder and risk analysis *(April 2011)*
- Produce detailed project plan *(April 2011)*
- Complete secondment arrangements for staff not already assigned at beginning of project *(May 2011)*

Manager (0.6 FTE); Analyst (0.5 FTE); Lead Architect (0.1 FTE)

WP 2 – Project Management

Ensure timely and efficient delivery of project deliverables, maintaining project schedule and resource allocation matrix; updating plans as required and writing regular progress reports both internally and for JISC. Manage the finances and resource allocation. Liaise between project collaborators at DCC, Eduserv, Southampton and JANET. *[April 2011 – March 2012]*

- Write monthly OUCS Project Progress Reports
- Hold monthly meetings of the Project Working Group
- Write termly Oxford PICT Reports
- Write two JISC Progress Reports
- Stage two meetings of Project Steering Group
- Produce JISC Final Report in draft *(February 2012)*
- Produce JISC Final Report *(March 2012)*

Project Manager (0.4 FTE)

WP3 – Development of Production VI

WP3a – Basic Virtual Infrastructure (VI)

Develop a small local VI 'Infrastructure as a Service' (IaaS) upon which production-quality Software as a Service (such as the DaaS) can be hosted. Specific tasks include: developing and deploying a core virtual infrastructure built on VMware vSphere

WP3b – Extension to Hybrid VI

Investigate and develop a hybrid VI, comprising the local VMware-based VI linking to other suitably configured VMware-based VIs, enabling seamless capacity management (including storage) and enabling IaaS for other private (institutional) or semi-public VI providers within the academic sector, such as Eduserv. More specific tasks include: design of a resource management and monitoring layer independent of hosting infrastructure; implementation of 'self-service' components for authorised users of the VI and investigating the automatic transfer of workloads from local to remote VI systems. Work with JANET to try extending the VI to a commercial VI, such as Logicalis. *[April 2011 – February 2012]*

- Produce VI design documentation *(April 2011)*
- Produce User Acceptance Tests (UAT) for Basic VI *(May 2011)*
- Procure hardware and licenses (*May 2011*)
- Implement basic VMware vSphere VI per VI design documents. *(June 2011)*
- Perform UAT *(July 2011)*
- Produce design documents for self-service components and system integration with existing authentication mechanisms *(July 2011)*
- Produce UAT for self service, and management and monitoring components *(August 2011)*
- Implement resource monitoring and management solution *(September 2011)*
- Implement self-service components *(Sept 2011)*
- Undertake UAT for resource management, monitoring and self-service elements *(October 2011)*
- Investigate automatic transfer of workloads between virtual infrastructures *(December 2011)*
- Explore with JANET extending VI to include private cloud suppliers *(February 2012)*

Lead Architect (0.3 FTE); VI specialist (2 FTE); Shared Service Developer (0.5 FTE)

WP4 – Shared Services

Implement the frameworks around the core system to allow for successful delivery of a production service.

WP4a – Accounting and Billing

Implement technology to enable, control, and track usage, and associated processes and solutions to support billing and accounting. This work package will evaluate the National Grid Service (NGS) accounting solution, and VMware VI chargeback mechanisms and ensure local solutions interoperate with appropriate solutions elsewhere in the UMF programme. The solution will deal with resource management at both the VM and application level.

WP4b : Identity and Access Management (IAM)

Implement technology to integrate with appropriate IAM solutions. This work package will examine existing IAM solutions, including federated access management, and where appropriate integrate with these to allow for ease of use of the system and its resources. *[April 2011 – January 2012]*

- Produce evaluation plan for existing IAM, accounting and billing technologies *(May 2011)*
- Evaluate existing technologies against the evaluation plan. *(June 2011)*
- Produce design documents for IAM and billing and accounting systems *(July 2011)*
- Produce UAT plans for IAM and billing and accounting *(July 2011)*
- Implement IAM *(November 2011)*
- Implement billing and accounting systems *(November 2011)*
- Test against UAT plans *(December 2011)*
- Investigate possibilities for extending accounting and billing into private cloud *(January 2012)*

Lead architect (0.2 FTE); Shared service developer (1.0 FTE)

WP 5 – Requirements and Analysis

Improve understanding of user requirements both

- Conduct literature review focussing on academic researchers' database behaviour and use *(April 2011)*

within the humanities and extending to other academic disciplines. Includes desk-based research, interviews with potential users, and the extension of the DaaS trial users group to include representation from medical sciences, physical sciences, and the social sciences. Also includes coordination of trialists who will offer feedback on DaaS development, documentation (WP 11), and training (WP 12). Whilst the test-bed user-base will be primarily located in Oxford, user engagement activities will also include engagement with potential users within other institutions participating in this programme as well as more generally (e.g. via a UCISA community survey). *[April – July 2011]*

- Conduct interviews with at least eight researchers involved in data-driven research from the three academic divisions besides the Humanities at Oxford *(May 2011)*
- Identify and invite new DaaS trialists from academic divisions outside the humanities *(May 2011)*
- Design and distribute national requirements and DaaS uptake survey *(June 2011)*
- Analyse results of national survey (feeds in to WPs 7, 11, 12, 13, and 14) *(July 2011)*
- Prioritize additional DaaS functionality *(July 2011)*

Analyst (0.5 FTE); Project Manager (0.2 FTE); DaaS Develop (0.1 FTE).
The trialist group will form the base group for UAT at various points throughout the implementation phase

WP 6 – DaaS Productionization

Extend existing DaaS prototype to run in a production-quality virtualized environment, to meet scalability and reliability requirements, and to support repeatable deployment to execution environments. Specific extensions will include: mechanisms to engage, manage, and release infrastructure resources; data encryption strategies to protect data, i.e. stored data and data in transition; provision for load balancing, provision for fault tolerance and disaster recovery, provision for automatic and on-demand backup. *[April – December 2011]*

- Extend DaaS hosting environment for compatibility with VI (WP3) *(May 2011)*
- Develop strategies to protect user data *(July 2011)*
- Develop support for backup of user data *(Sept 2011)*
- Develop modules for fault tolerance and disaster recovery *(October 2011)*
- Develop installation mechanism and packaging of DaaS *(November 2011)*
- Develop load-balancing mechanism across virtual host and databases *(December 2011)*

DaaS Developer (1 FTE)

WP 7 – DaaS Expansion of Functionality

Expand functionality of DaaS to meet the requirements of the broader research community as identified in part by WP5. Pre-established DaaS functionality extensions include: a Web interface for users to request, initialize, and manage the DaaS; incorporation of data storage models other than relational databases (i.e. XML-based and document-based); management and monitoring tools (i.e. storage capacity management, access control). Beyond this, the WP will provide a means of capturing metadata from DaaS for storage in one or more registries (whether local or national) and implement a system of single identifiable addresses for public access to databases in the DaaS. *[June 2011 – February 2012]*

- Implement online DaaS provisioning mechanism *(July 2011)*
- Develop interface for DaaS based on Restful Web Services *(September 2011)*
- Extend DaaS to support XML databases *(October 2011)*
- Adapt DaaS to interface with non-relational databases e.g. document-oriented databases such as CouchDB (assuming user demand) *(November 2011)*
- Integrate DaaS with local and national registries if available *(December 2011)*
- Implement additional priority functionality as identified by WP 5 *(December 2012)*
- Develop DaaS-specific management and monitoring tools for support staff *(January 2012)*

DaaS Developer (1 FTE)

WP 8 – Service Design and Initiation

Initiate a supported DaaS shared service. Specific development tasks will include: service-level planning and documentation; organizational preparation for service delivery; design of management and operational processes; implementation of the outputs from WP6 and WP7 onto production VI; invocation of operational processes and service management functions. *[June 2011 – January 2012]*

- Approve initial technical architectures *(July 2011)*
- Design metrics and targets *(July 2011)*
- Develop and agree SLD *(July 2011)*
- System test DaaS on VI *(August 2011)*
- Agree service governance model *(August 2011)*
- Design required operational processes and service-management tools *(September 2011)*
- Design management processes, including change control *(October 2011)*
- Launch DaaS in staging environment *(October 2011)*
- Design and test disaster recovery plan *(October 2011)*
- Launch DaaS to early adopters *(November 2011)*
- Launch shared DaaS service as production service *(December 2012)*

Lead Architect (0.5 FTE); DaaS Developer (0.2 FTE); VI Specialist (0.2 FTE)

WP9. Deliver DaaS to External VIs, and deliver proof-of-concept DataDrive

Deliver the DaaS as a supported production service to be hosted and run on virtual infrastructure, initially

- Deliver pilot DaaS on Eduserv VMware infrastructure *(September 2011)*
- Deliver to Southampton proof-of-concept DaaS on VMWare for their researchers *(September 2011)*

provided by Eduserv, but in time via the JANET brokering services. Investigate to see whether it can be operated on a public VI and furthermore explore the possibility of *bursting-out* from the Oxford VI to an external VI automatically. Southampton will operate DaaS on the Eduserv VI for their researchers as a proof-of-concept. Oxford will similarly deploy their DataDrive SaaS on both the Eduserv VI and Oxford VI as a proof of concept for Oxford's researchers. *[June 2011 – February 2012]*

- Test DaaS functionality and fix as required *(October 2011)*
- Deploy proof-of-concept Datadrive SaaS for use by Oxford's researchers *(November 2011?)*
- Deliver production DaaS on VMware infrastructure *(December 2011)*
- Explore through JANET brokerage service possibility of deploying DaaS on public VI *(December 2011)*
- Evaluate possibility of bursting out from Oxford VI to external VI automatically *(February 2012)*

Lead architect (0.4 FTE); DaaS Developer: (0.3 FTE); VI Specialist (0.5 FTE); Eduserv (0.4 FTE)

WP 10 – Documentation and Support

Produce documentation describing the DaaS and its management tools and how they may be installed on a virtual infrastructure. Separate sets of documentation will address end-users (researchers), software developers, database administrators, and systems administrators. Will include the production of an online user manual and contextual help within the software. Collaborate with DCC in developing material to ensure suitability for the HE community beyond Oxford.*[July 2011 – March 2012]*

- Produce DaaS user manual 1st draft *(September 2011)*
- Complete contextual user help within software *(December 2011)*
- Fully document code to assist future software development *(documentation added throughout)*
- Produce installation guide for setting up DaaS and monitoring software on a virtual infrastructure *(December 2012)*
- Produce final version of user manual *(February 2012)*

Technical Author (0.8 FTE); DaaS Developer (0.1 FTE); VI Specialist (0.1 FTE)

WP 11 – Data Management Training

Produce and trial training material for researchers using the DaaS. Will include online and face-to-face material. A suite of training material will be developed from a basic introduction for those with no prior experience of database development, to specialist content relating to the use of particular data types such as geospatial data, plus using metadata to document databases developed using the DaaS. This work package will involve close cooperation with the DCC who will be expected to continue the training function after the completion of the project and will be part of their 'Training the Trainer' programme, and will include, for example, a module 'Creating structured data'. *[July 2011 – March 2012]*

Courses will be developed according to findings of WP 5). Sample training may include:

- Develop course: Quick introduction to the DaaS *(July 2011)*
- Develop course: Building a database with the DaaS *(September 2011)*
- Develop course: Using the DaaS for collaborative projects *(October 2011)*
- Develop course: Customizing the DaaS *(November 2011)*
- Develop course: Geospatial data and the DaaS *(January 2012)*
- Develop course: Structuring humanities data *(February 2012)*

Analyst (0.5 FTE); DCC (0.2 FTE)

WP12. Cost models and cost/benefit analysis

Consists of three aspects: an update of the ROI analysis to quantify cost and efficiency savings likely to be gained from the DaaS; a later more detailed costing of operating the VI and DaaS service at various levels of support; and a cost-recovery Business Plan which details the on-going support costs incurred by running the virtual infrastructure and SaaS at Eduserv and public VIs. The aim is to derive an evidence-based ROI model, and to create a viable Business Plan for continuation after the end of the VIDaaS project. Work will be undertaken in close collaboration with Eduserv and JANET. *[August 2011 – March 2012]*

- Update the ROI analysis submitted in the bid *(September 2011)*
- Deliver initial Business Plan for continuation beyond project *(December 2011)*
- Deliver final ROI evaluation based in experience through the project *(March 2012)*
- Deliver final Business Plan for operation of DaaS beyond the life of the VIDaaS project *(March 2012)*
- Work with DCC to develop a model for supporting SaaS on a shared VI to ensure that full savings can be made which result from the UMF investment *(March 2012)*
- Deliver plan to offer 'warranty' for DaaS after the completion of the project *(March 2012)*

Lead architect (0.1 FTE); Project Manager (0.2 FTE); DCC (0.1 FTE)

WP – 13 Communication and Dissemination

Participate in Programme-wide dissemination activities as well as pursuing effective communications opportunities both locally and nationally to ensure awareness and take-up of the services to be provided. Collaborate closely with DCC to develop training materials and disseminate news relating to the DaaS service. Prepare for service to be overseen by DCC in the long term. and for on-going training *[April 2011 – March 2012]*

- Produce communications plan after consultation with the DCC and JISC *(April 2011)*
- Update project website as outputs delivered *(on-going)*
- Post at least one blog entry each month, plus other material of interest to an external audience *(on-going)*
- Stage stakeholder workshop *(January 2012)*
- Participate in other events in collaboration with JISC *(on-going)*

Analyst (0.2 FTE); other project team members as required

WP – 14 User Acceptance Testing

System and user acceptance testing is contained within other WPs. This WP is concerned with the management of UAT, including process definition. The trialist group (WP 5) together with other interested parties will form an early adoption group, providing input to acceptance testing and user-led formative assessment. *[August 2011 – Jan 2012]*

- Form early adopters group consisting of project stakeholders and practitioners (*May 2011*)
- Define evaluation and testing plan (*June 2011*)
- Manage user acceptance testing (*ongoing, as per WPs*)

Project Manager (0.2 FTE); Analyst (0.2 FTE)

Project dependencies are towards the 'soft' end of the scale, allowing a relatively robust work plan that is able to absorb delays to some work packages without immediately interrupting others. No WP requires the completion of another prior to commencement, although individual milestones within a WP may depend upon milestones in other WPs being achieved – notably in work packages 8 and 9 where elements of testing cannot be carried out before the DaaS is available in the appropriate staging environment.

2.3 Standards

Compatibility between the outputs of the VIDaaS Project and existing infrastructure used in the UK HE sector is an important goal of the project. It shall be ensured, therefore, that the infrastructural elements of development adhere to existing standards, and the layered approach to the use of open standards described in the JISC Standards catalogue is used. The extended DaaS will be built from, and released as, open source components, in order to enable other institutions to contribute and further develop the software to meet changing requirements in the future.

2.4 Intellectual Property Rights

Any IPR resulting from this project will remain the property of the organization generating it. Under the University of Oxford's policy on intellectual property (which covers all University employees and students), the University claims ownership of a range of intellectual property rights with commercial potential. The University does not assert any claim to the ownership of copyright in artistic works, books, articles or lectures, apart from those specifically commissioned by the University. Results arising from projects funded by the JISC at Oxford would therefore usually be owned in the first instance by the University as the employing institution. It is proposed that software outputs developed within this project will be released under an open source software licence.

2.5 Project Governance and Management

The project will be led by Oxford University Computing Services, with the support of a Steering Group comprising internal and external representatives from the key stakeholder communities (e.g. JISC, the DCC, Eduserv, JANET, Southampton, senior academic staff, and representatives of other projects funded under the same programme).

A smaller working group, comprising members of the project team, specialists and user representatives as required, will be established in order to help facilitate an agile and iterative approach to the development of the project outputs.

The Project Manager will have the responsibility of day-to-day coordination of the project. The Lead Architect role will be fulfilled by three key members of the project team, providing technical direction to the development of the hybrid VI, the virtualisation layer, and the DaaS. The project has overall direction from the Director of IT and the Head of the Infrastructure Systems and Services Group.

2.6 Risk Assessment

The risk analysis presented below is only an initial assessment of project risks, a more complete risk analysis will be completed for the overall project plan.

Risk	Probability (1-5)	Severity (1-5)	Score (PxS)	Mitigation
Staffing				
Failure to allocate appropriate staff for the project.	2	4	8	The project will use existing staff, either with experience of working on the Sudamih Project or on infrastructure development projects. Contractors will be used to backfill allocated staff or to assist with well-defined WPs.
Loss of key staff before end of	2	3	6	Ensure regular communication between all project

the project.				staff, so that processes and progress are clearly understood by the team, and staff may be redeployed to cover different work packages.
Organizational				
Lack of a sustainable business model.	2	4	8	Work within the Programme, taking advantage of expertise and support materials developed by JISC to ensure that all aspects of costs/benefits are considered.
Expectations mismatch between project and wider research community.	2	3	6	Ensure cross-disciplinary and cross-institutional academic representation on Steering Group. Ensure requirements gathering is broad and feeds through to technical development. Fully involve researchers in trialling software from an early stage.
Estimations of time required to complete technical work packages are inaccurate.	2	2	4	Allocated staff are familiar with the nature and scale of this type of project. Progress on each WP, and resource allocation will be closely monitored.
Outputs are difficult or demanding to support.	2	2	4	Ensure that DCC and Eduserv are closely engaged with technical development and support WPs.
Lack of coordination between project stakeholders.	1	3	3	Ensure clear reporting and communication lines; take advantage of existing institutional communication structures.
Technical				
Requirements of researchers outside humanities disciplines are very different, or demanding to meet during timescale.	2	3	6	Ensure speedy transmission of requirements to technical developers so that they may be integrated. Impose boundaries on scope of technical implementation if required to ensure completion of productionization phase within time-span.
Development of large-scale online database provisioning system proves to be too complex.	1	4	4	Ensure members of staff have requisite technical skills. Build upon existing infrastructure and DaaS development work, using same developers.

2.7 Exit and Sustainability Plans

By the end of the project, a full production-level virtual infrastructure supporting DaaS will have been prepared and deployed at the VIs in Oxford and Eduserv, and possibly on the commercial cloud and on VIs in other HEIs (WPs 3, 7 and 7). A costing model will have been developed (WP12) and mechanisms put in place to measure usage and charge accordingly (WP4). The software outputs will be released under an open source licence and made available to other interested parties along with the associated user support and training materials. OSS Watch will assist the VIDaaS project in creating a community around the open source code and advise on establishing the project in a code repository with the appropriate management tools and processes for long-term sustainability.

Oxford and the DCC have agreed to build a DaaS business plan (WP12) (with contributions sought from Eduserv where appropriate) that will include a cost-recovery funding model, and is *designed to enable the service to be developed further and be fully supported for use by the national higher education community from 1 April 2012*. This will be crucially important. It is anticipated that DCC will: be the 'point of provision' for future HE services which have been developed and delivered through the UMF programme, set future service level definitions, and contract units to develop services further and support the national services which result. DCC has agreed to continue to train researchers nationally to use DaaS and to manage their research data. Within the DaaS Business Plan, the DCC and Oxford, in discussion with OSS Watch, will recommend a sustainable model by which the software can continue to be developed (for example through the Software Sustainability Institute). DaaS will continue to be delivered as a production service for the HE community on a cost recovery basis beyond the life of the VIDaaS project.

3. Engagement with the Community

3.1 Stakeholder and Practitioner Engagement

Project stakeholders and practitioners will be engaged throughout the life of the project via the Steering Group and Project Working Group. Academic researchers will be involved throughout the software development via the trialist group. The project will work closely with Eduserv, the DCC, OSS Watch, and colleagues at the University of Southampton to ensure the compatibility and long-term sustainability of outputs. The project website and the planned workshop will also help ensure

on-going communication between the project and stakeholders.

Stakeholder	Interest / Stake	Importance
Academic researchers at UK HE institutions	As the ultimate users of the DaaS, academic researchers are a key stakeholder. The VI-hosted DaaS will improve research practice and increase the impact and value of research data, by providing mechanisms for data management, collaboration, discovery, and dissemination.	High
JISC	VIDaaS will act as an exemplar shared services/cloud project. The good management and curation of research data is a JISC priority, and the DaaS is recognized as a tool that is likely to be of broad application in raising standards of research data management.	High
University of Oxford	The University considers the project an important component in its strategy for the cost-effective and sustainable management of research data. The project contributes to the provision of shared infrastructure services to the collegiate University and beyond.	High
Digital Curation Centre (DCC)	The development of a national digital curation infrastructure is of one of the central challenges of the DCC. Their interest in the VIDaaS project will be particularly significant, as they will be the 'point of provision' for DaaS for the HE community after the project completes, and will work with Oxford to deliver research data and DaaS training.	High
Higher Education Funding Council for England (HEFCE)	At a national governmental level, the VIDaaS project is one element in a programme of university funding ultimately intended to bring cost efficiency savings to the HE sector. HEFCE will see evidence of cost savings and efficiencies as a result of the project.	Medium
Other projects funded by the Universities Modernisation Fund	As the other projects funded by the UMF will all be aiming to create a national infrastructure, it is important that the projects work together to ensure compatibility and avoid unnecessary duplication of activity.	Medium
Universities and Colleges Information Systems Association (UCISA)	UCISA represents UK HE in the provision and development of academic, management and administrative information systems. They have a particular interest in shared IT services in HE and should find the hybrid VI aspect of the project attractive.	Medium
Research funding agencies	Given that the DaaS and associated infrastructure will provide a useful service for researchers in terms of managing their data and opening it up to other researchers in the future, the funding agencies may wish to recommend its use in future funding bids.	Low

3.2 Dissemination Plans

Internal and external dissemination activities will run throughout the project, in close collaboration with JISC and the DCC. The two most important aims of dissemination in the VIDaaS project are to publicize the intended service to potential users and to ensure that the development community are aware of what we are doing (to avoid duplication of effort as well as to attract ideas). Dissemination will comprise multiple channels of communication, including:

- A project website, maintained with regular updates.
- At least one workshop, to stimulate discussion amongst stakeholders. Internal dissemination will be facilitated through additional less formal events as well as through the proposed stakeholder engagement structure.
- Knowledge transfer and demonstrations of the VIDaaS service at relevant events.
- Full engagement with DCC and JISC, in particular through the DCC Research Data Management Forum and the JISC Managing Research Data Programme.

A more complete dissemination plan will be created with the overall project plan (WP1).

4. Project Budget

WITHHELD

The VIDaaS project is aligned with, and will contribute to, Oxford University's strategy to deploy shared infrastructure services befitting the devolved IT environment within the collegiate University. The new University Shared Data Centre provides the launch-pad for a suite of virtualised infrastructure and software services, with the potential for seamless bi-directional integration with HE virtualised infrastructure. The VIDaaS project contributes to the entire infrastructure stack, including development of transferable knowledge relating to the building, maintenance, and business planning of both IaaS and production quality SaaS.

5. Project Team

Staff will be allocated to the project who have expertise in the management and development of IT infrastructure, virtualisation, and the Sudamih DaaS.

1. **Professor Paul W Jeffreys** – Providing overall direction and leadership for the project. Additionally, ensuring the project is well integrated within, and supported by, the University. Paul is Director of IT, at the University of Oxford and is PI for the Sudamih project.
2. **Dr Michael Fraser** – Providing direction to the project (which will be led from within the OUCS Infrastructure Systems and Services Group). Michael is head of the OUCS Infrastructure Systems and Services Group and has led a number of projects related to research support infrastructure.
3. **Dr James A J Wilson** – Providing project management. James is currently project manager for Sudamih and has managed a number of other JISC-funded projects in the past including Eidcsr, IJDDiP and Intute Arts and Humanities. (Full time)
4. **Mr Jon Hutchings** – Providing virtual infrastructure specialist skills. Jon has internationally-recognised expertise in the development of VMware-based virtual infrastructure.
5. **Mr Asif Akram** – Providing technical development expertise for the development of the DaaS. Asif is the lead developer for the Sudamih DaaS. Before joining Oxford Asif worked on a number of projects relating to medical research at Imperial College London. (Full time)
6. **Dr Meriel Patrick** – Providing business analyses and user requirements skills. Meriel is the analyst for the Sudamih project. She has previously worked on user analysis for Intute, and is an experienced teacher and trainer. (50% time)
7. **Mr John Ireland** – Providing systems integration and management expertise. John is head of the OUCS Systems Development and Support team.
8. **Mr Peter Jones** – Providing technical leadership and coordination, especially for the development and deployment of shared infrastructure services. Peter is head of the OUCS Shared Infrastructure Services team.
9. **Ms Elena Blanco** – providing technical and user documentation expertise. Elena is a member of the OUCS Information and Support Group, and of the OSS Watch team.
10. **NSMS –** NSMS, part of OUCS, has extensive expertise in VMware having adopted the ESX platform in 2005. Several members of NSMS are VMware Certified Professionals (VCP) who have presented at international VMware Conferences, as well as being active community members. NSMS have excellent links with the VMware organisation at many levels, enabling access to beta programs and favourable access terms to research and development solutions for other groups within Oxford and the HE community.
11. **Sysdev** – The Systems Development and Support team within OUCS are responsible for many of the core services developed and provided to members of the University. They will contribute expertise in the deployment and integration of identity and access management services; large-scale server configuration and management; and the development and support of enterprise standards-based online applications.

The DI staff effort will be provided by the named Project Team.