

# **National Servers Program Implementation Plan National eResearch Collaboration Infrastructure Project**

**University of Melbourne**

**June 2010**

---

<b>1 OVERVIEW</b>	<b>2</b>
<b>2 POLICY DEVELOPMENT WITH THE SECTOR</b>	<b>3</b>
<b>3 SERVICE AND SUPPORT ENVIRONMENT DEVELOPMENT PLAN</b>	<b>5</b>
3.1 ITS positioning	5
3.2 Infrastructure overview	6
3.2.1 Virtual Machine configuration	6
3.2.2 Data storage infrastructure	6
3.2.3 Connectivity and Bandwidth	8
3.2.4 Failover and redundancy	8
3.2.5 Authorisation and Authentication	9
3.2.6 VM Capacity growth	9
3.2.7 Data centre development	9
3.3 Deployment plan	9
3.4 Access model	10
3.5 Roles and responsibilities	11
3.6 Support and the National Servers Program helpdesk	11

# 1 Overview

This document outlines a 12-month plan to develop and deploy a first installment of the National Servers Program as outlined in the *eResearch Collaboration Infrastructure EIF Interim Project Plan*. This program intends to provide server support for essential core eResearch services to run with reliability standards appropriate to sector needs. The plan focuses on the July 2010 to Jun 2011 period in order to provide initial infrastructure which core services can be deployed on, which are intended to support the three other Collaboration infrastructure components – Virtual Labs, eResearch Tools and the Research Cloud. EIF funds will be used to develop and deploy this infrastructure for publically funded researchers. This National Servers Program is distinct from the Research Cloud program which is intended to provide a National Research Applications Cloud that will allows researchers to develop, publish and access research enabling software tools, and operate those tools in secure and shared environments.

Servers that are deployed in the first twelve months may run core eResearch services such as:

- ANDS (Australian National Data Service) utilities,
- Critical identity management services required for the Australian Access Federation (AAF),
- Exposure of underlying storage capacity developed through the Storage Infrastructure plan as a national drop box service,
- THREDDS middleware (Thematic Realtime Environmental Distributed Data Services) to bridge data providers and data users,
- Servers required by the Virtual Laboratories and Research Cloud programs initial research developments, etc.

This initial deployment will begin with a three-month (Jul 2010 – Sep 2010) demand consultation process to ascertain how many core services need to be hosted as well as their server infrastructure requirements, with a go-live date of November 2010.

This implementation plan outlines policy, operation and infrastructure to be initially provisioned. The details of this will be reviewed through extensive consultation with the sector during Jul 2010 – Mar 2011 and the revised versions will be included in the Final Project Plan.

## 2 Policy Development with the Sector

The National Servers Program will establish policies covering:

- The criteria to be applied to admit services for virtual server support,
- The criteria for virtual server resource allocation,
- The pricing policy to apply to different categories of access and users,
- The characteristics of the virtual servers to be offered to service owners,
- The characteristics of the functional support to be offered to service owners.

The National Servers Program will also develop a standard form agreement that will define roles, responsibilities and contributions between the parties to the overall delivery of services operating from the support environment established by the program.

The full development of these policies and standard form agreements will occur during the lifetime of the Project.

- This Interim Plan will define these for the first year of the Project, July 2010 to Jun 2011,
- The intention is to develop and define a final version in the Final Project Plan, by March 2011,
- Further amendments will then be incorporated in subsequent Annual Business Plans.

The following interim policy structure will be applied for Virtual Server support (November 2010 – June 2011).

Policy	Details
<b>Admission Criteria</b>	<p>The Project will establish an expert panel to review and recommend services suitable for support by the program.</p> <p>The Panel will recommend final criteria to the Project Board by September 2010 addressing the following key factors:</p> <ul style="list-style-type: none"> <li>• A service 'owner' can be identified capable of entering into a contractual arrangement with UoM</li> <li>• The service underpins the integrated management or access to collaboration, data, compute or analysis resources and capabilities</li> <li>• The service is considered a core to other eResearch services and activities.</li> <li>• The service will support activities by researchers from multiple institutions</li> <li>• The service has suitable maintenance and user support resources</li> </ul>
<b>Resource Allocation</b>	<p>A standard configuration of virtual server resource unit (VSRU) will be defined and services will be allocated appropriate multiples of that VS Resource Unit by the Program.</p> <p>The proposed virtual server offerings and corresponding VSRUs are outlined in Section 3.2.1. These will be refined if required to better match needs over the lifetime of the Project.</p>
<b>Pricing</b>	<p>The Program will establish an annual capital and operating cost per VSRU.</p> <p>During 2010/11, the full costs will be absorbed by the Project.</p> <p>It is expected that in out years, arrangements will be concluded for the operating component of the infrastructure to be met through participant contributions to the Project.</p>

<b>Server QoS levels</b>	<p>Two levels of Server QoS will be offered in 2010/11:</p> <ul style="list-style-type: none"> <li>• Production Servers (Daily backup, 1-day restore, non-replicated, no-failover, user-generated VM snapshots)</li> <li>• Production Redundant (Daily backup, 1-day restore, replicated, failover, user-generated VM snapshots).</li> </ul> <p>More detail is given in Section 3.2.3. These will be refined if required to better match needs over the lifetime of the Project.</p>
<b>Service Functional Support</b>	<p>Two levels of functional support will be provided to services supported by the program – unmanaged and managed (explained in detail in Section 3.5). In summary, <i>unmanaged</i> provides access to standard VM operations (Start, Shutdown, Reset, Snapshot) to a user-installed operating system, as well as performance reporting and weekly backup. Further to this, the <i>managed service</i> provides a fully configured virtual machine with a standard operating system (Windows, Solaris or Linux), installed, patched and maintained.</p>

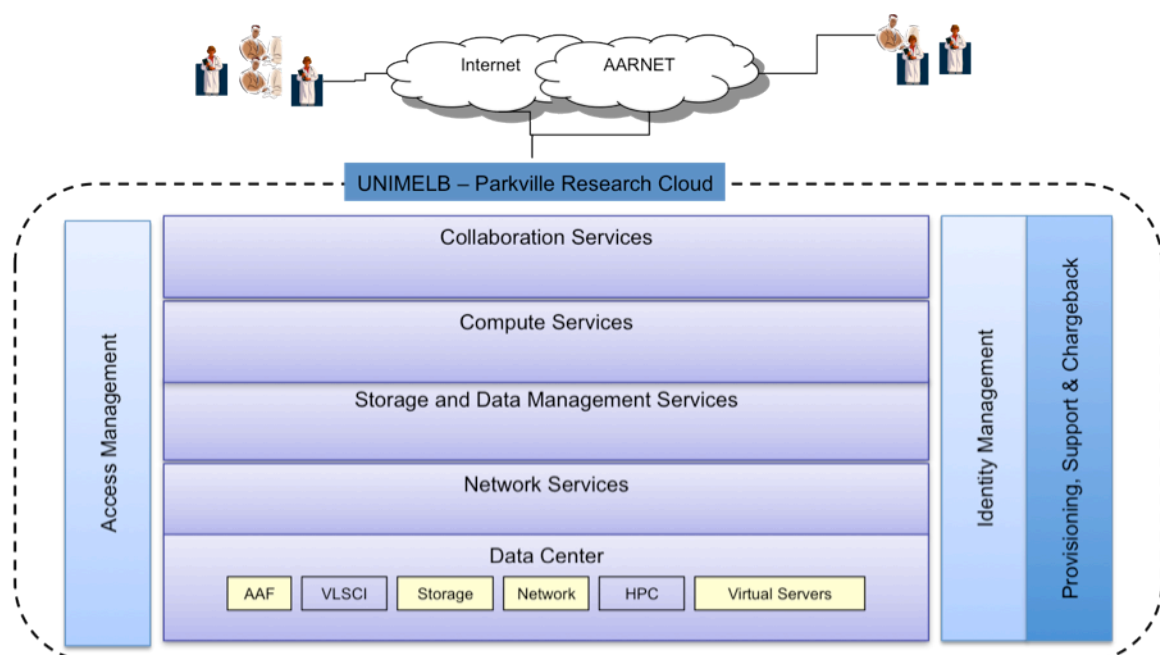
### 3 Service and support environment development plan

#### 3.1 ITS positioning

University of Melbourne Information Technology Services (ITS) is well advanced in establishing world-class core technology infrastructure to support Australian researchers. The strategic technology roadmaps cut across the complete infrastructure spectrum including Data Centre, Network, Storage, High-Performance and Server technologies leveraging virtualisation as far as possible. ITS are also implementing a new Identity Management platform to improve and expedite provisioning processes and offer “Identity as a Service” to both Melbourne and external/federated researchers and applications. Integrating with AAF infrastructure is a fundamental component of this project.

By extending the existing Parkville based data centre to host researcher infrastructure, including an IBM BlueGene computer, and augmenting this capacity with a separate facility at Noble Park ITS will offer a highly scalable and fault-tolerant environment to accommodate growing needs. The storage solution is utilising virtualisation technology to provide a highly scalable platform for SAN and NAS based storage. A dedicated storage team is already in place to support this facility. Virtual machine provisioning is already a service offered by ITS. The strategy will include development of a hybrid private/public cloud infrastructure-as-a-service offering to scale out on demand.

The above infrastructure capability is being supported by the implementation of appropriate provisioning processes, service designs and costing and pricing considerations. Analysing the take up and ongoing usage of these services will be critical in feeding back metrics to sponsors and stakeholders. The University data warehouse is a mature platform and will easily accommodate this information providing analytics and reporting via web-based interfaces.



## 3.2 Infrastructure overview

The National Servers Program is mainly comprised of scalable virtual server platform, dedicated storage infrastructure and network components that are key to deliver the required services. The proposed architecture, when in full service, will provide high availability, scalability and disaster recovery services effectively and efficiently. Further details are provided in the following sections.

### 3.2.1 Virtual Machine configuration

Oracle Sun Fire x86 servers will be the preferred virtualisation platform. The optimum hardware specification evolves over time and at present, units are purchased with 128GB RAM and 24 2.4 GHz processor cores. Each of these physical servers can accommodate 32 small guests, 8 large guests or 4 extra large guests, where:

- **Small VM** - 2GB virtual RAM, 1 virtual CPU, 40GB storage (1 VSRU)
- **Large VM** - 8GB virtual RAM, 4 virtual CPUs, 160GB storage (4 VSRUs)
- **Extra-large VM** – 16GB virtual RAM, 8 virtual CPUs, 1TB storage (8 VSRUs)

ITS at the University of Melbourne offers extended support services for a limited number of operating systems. The initial supported guest platforms will include:

- Windows Server 2008 R2
- Red Hat Enterprise Linux 5
- Oracle Solaris 10 for x86/x64

The demand consultation will finalise the operating systems to be supported in managed configurations and the means by which that support will be achieved. Researchers/service owners are able to install their own guest OS (OpenBSD, FreeBSD, Linux variants, etc.) however extended support services are not available for these platforms (see Section 3.5 for more detail).

### 3.2.2 Data storage infrastructure

Where needed, extra SAN-based storage can be deployed further to the standard 40 GB, 80 GB and 1 TB of local storage that are available to each VM. The storage service is comprised of the following scalable foundations:

- SAN Storage Arrays
- Fibre Channel SAN
- NAS Gateway Cluster
- Enterprise Tape Library

SAN storage arrays deliver three tiers of virtualised block mode storage within the datacenter. A range of price/performance points is offered to meet varied project requirements:

Tier	Type	Size	Rotational Speed (RPM)	RAID level	Relative cost
1	Fibre Channel	600GB	15K	5	High
2	SAS	600GB	15K	5	Medium
3	SATA	2TB	7.2K	6	Low

An additional tier, Tier 4 – Tape (via HSM) is planned for Q2 2011.

Storage arrays and the NAS Gateway cluster are interconnected via a 4Gbps Fibre Channel SAN, comprising enterprise class fibre channel directors configured to deliver dual, redundant fabrics. Additional switches will allow the SAN to scale from a dual core to dual core/edge topology as growth requires.

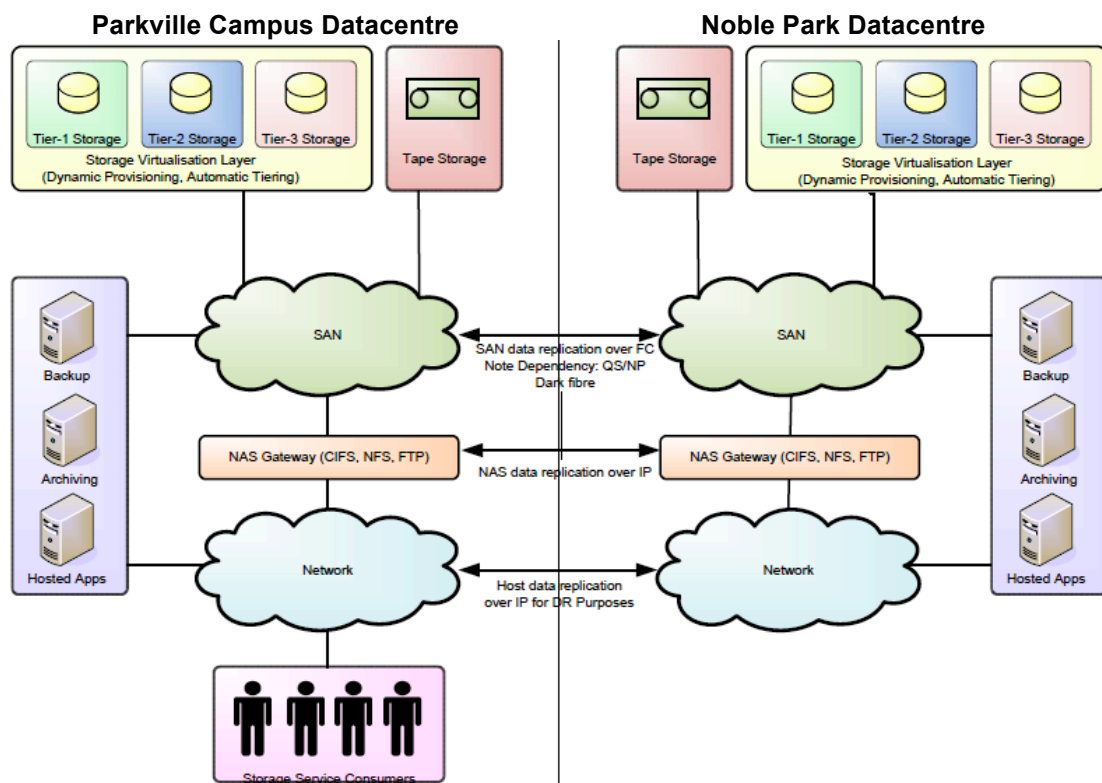
SAN host connectivity is via redundant fibre channel host bus adapters.

NAS file level services (CIFS, NFS v3, v4) are delivered via a high performance, scalable, clustered NAS gateway within the data centre, campus or remote sites. Connectivity within the datacenter or across campus leverages the 10GbE enterprise/campus network.

Both SAN & NAS storage offerings provide the option of

- Seamless, policy driven storage tier migration,
- Remote data centre replication,
- Thin provisioning - delivering storage oversubscription and allowing seamless client volume growth from actively managed shared storage pools.

Data protection is provided via a network backup service to a central enterprise tape library.



### 3.2.3 Connectivity and Bandwidth

The National Servers Program platform will be connected to AARNet by two 10 Gb links. This will be generally shared with University of Melbourne users, however dedicated bandwidth to services can be provided via QoS. To enable appropriate fair sharing of available network resources, VMs will have published network I/O limits implemented. These will be adjustable according to the application requirements.

If required, external traffic can be directed into the SAN (a possible scenario in a national-shared drop box implementation), however this is not enabled by default.

Depending on the service security requirements, the VM can be hosted outside the firewall or inside the firewall, and in either case each VM can configure their own host-based firewall. If needed, demilitarized zones (DMZ) can be configured to more easily enable appropriate group firewall configurations, etc.

Through the sector consultation process research requirements as well as enabling network resources will be detailed.

### 3.2.4 Failover and redundancy

The virtualisation platform is designed to support an “active/active” model to deliver business continuity of virtualised applications and services. The virtualisation platform will offer multiple levels of failover.

<b>Planned downtime</b>	The impact of scheduled downtime required for hardware maintenance will be largely eliminated by the ability to dynamically move guests to another server.
<b>Snapshots</b>	Users will have the ability to snapshot their virtual machines and return to a known state at some future time.
<b>Unplanned downtime</b>	The virtualisation platform will be built on high-end servers with fault tolerant features including multiple power supplies, highly redundant shared SAN storage with multi-pathing, and multiple teamed NICs. In the event of a hardware failure, users have the option of having their guest servers automatically restarted on another machine.
<b>Fault tolerance</b>	VMware Fault Tolerance is a feature that supports a second virtual machine that runs on another physical server and mirrors the state of a primary virtual machine, running in parallel and processing the same I/O. In the event of loss of the primary machine, transparent failover to the second guest occurs with no data loss and no loss of network connectivity. This facility will be available to users by request.
<b>Software replication</b>	Software replication mechanisms will be put in place to allow virtual machine images to be copied between the data centres. This will allow users to fall back to an alternative data centre with an attractive RPO (recovery point objective) and RTO (recovery time objective).
<b>Hardware replication</b>	As the number of virtual machines increases, and with it the demand for guest replication, responsibility for replication will be delegated to the storage layer.



### **3.2.5 Authorisation and Authentication**

For authorisation and authentication purposes, there are two general classes of users.

- *Service owners* – those responsible for maintaining their VM instance and the application(s) deployed within that instance
- *Service users* – the general research community that will use the services running on the VMs.

Service owners will be authenticated against AAF and once their credentials are established they will be able to access National Servers Program resources to manage their VM instance (start, stop, restart, generate snapshot, etc.). Separately, the service owners can configure their hosted applications to use AAF authentication/authorisation, etc., according to their requirements.

In practice, the University's gateways will provide SSL termination, firewalls and an in-house identity and access management system (IDAM) to manage all identity and access related activities. IDAM will have the ability to link to the AAF via OpenSSO in order to feed relevant roles and permissions. This will simply be an additional component of the IDAM service; the Communications gateway will query the IDAM via a radius component (e.g. Juniper UAC radius controller).

### **3.2.6 VM Capacity growth**

Provisioning for the first 12 months of the infrastructure is based on the deployment of 25 to 50 virtual servers with minimum of 40GB of local storage each.

Over the following three years the capability will exist to support an order of magnitude growth in requirements for services, i.e. 250 – 500 virtual servers. Although these figures are subject to change based on demand-mapping consultation exercise and the wider sector consultation process, the scalable architecture presented here means it is possible to provide more computing and storage as the demand increases. The deployment plan provided below in Section 3.3 provides more details in this regard.

### **3.2.7 Data centre development**

Initial deployment for phase 1 will be at the Noble Park (NP) Data Centre. It is envisaged that the National Servers Program will be spread across NP and Queensberry Street Data Centres as a single logical data centre through further roll out during full server deployment.

## **3.3 Deployment plan**

The following figure outlines the deployment plan in three main phases – demand consultation, basic access and full service.

Phase	Consultation	Basic	Full Service
Timeline	July 2010 – Mar 2011	Nov 2010 – June 2011	July 2011 – Dec 2013
Services	Requirements gathering Demand projections Architecture development Provisioning	Go-live Manual provisioning of VM Limited backup Limited disaster recovery Limited high availability	Automated provisioning of VM Full backup and archiving Built-in disaster recovery Built-in high availability
VM		VMs =10 to 50 at each site	VMs =250 to 500 at each site

\* **One VM** = 1 vCPU, 2 GB virtual RAM, 40 GB local Storage

Note: Additional computing resources, storage and VM upgrade will be available on demand.

### 3.4 Access model

It is expected that during 2010 to 2011, at least 100 VSRUs will be available to support core eResearch services. In the case that demand exceeds available virtual resources, various criteria will be used to prioritise services for deployment as well as a review process.

The management of the service will report and review the quality of the support for the service offering, so that poorly maintained or increasingly inoperable services may be removed from the infrastructure.

- An expert allocation committee will be established during the demand consultation process. This will be a subcommittee of the governance board and consist of 4 members.
- Priority will be given to core services that support other eResearch services, research activities, disciplines and communities.
- The fundamental model will provide a quantum of access of 1 VSRU reassessed annually. Associated storage and extra capabilities (backup, failover, redundancy) can be requested as well.
- The allocation committee will review each service every 12 months. This review process will take into consideration the continuing demand for the service, its reliability and usage statistics reported from the infrastructure and the level of the service-owner response to application issues and user-queries relating to that service.
- The access model will be refined and reviewed in the Final Project Plan according to the sector consultation process.

### **3.5 Roles and responsibilities**

In this initial deployment the University of Melbourne will be responsible for:

#### **Unmanaged:**

- Management of host hardware
- Power, cooling and environmental management for the host hardware
- Network infrastructure within the data centre, cabling infrastructure and storage
- Networking configuration
- Network security and physical security
- Managed virtual server infrastructure
- Researcher access to relevant virtual server guest management components
- Researcher access to virtual server guests
- Monitoring virtual server infrastructure
- Access to standard operations (Start, Shutdown, Reset, Snapshot) guest snapshot based technologies at the virtualisation layer

#### **Managed:**

- Management of the virtual server guest OS within the virtual server
- Upgrade and patching of the virtual server guest OS
- Detailed application backup of content within the virtual server guest
- Access to the infrastructure beyond the customer's virtual server guest
- Security of virtual guest operating system
- Management of storage within the virtual server

Extra services that will be considered/implemented according to specific requirements:

- Regular reporting usage of the VM resources at the operating system level
- Backup of VM and associated storage (where needed)
- Redundancy and failover (where needed)

The service owner / researcher will be responsible for:

- Application level reporting
- Application level patching and upgrading

Further considerations and future developments that might be useful for the community by the VM service provider include:

- Automated testing and reporting of server/application functionality
- Service-level automated usage reporting

### **3.6 Support and the National Servers Program helpdesk**

A helpdesk service will be developed at the University of Melbourne which will be triaged through ITS. Responsibility for hardware, virtual machine and operating system specific problems will be directed to the appropriate owner in ITS. Application-specific problems will be directed to the service owner.