Annexure A: Scale-Out File Storage System for TW Kambule Laboratories

for Mathematical Sciences Teaching and Examinations

Platform Type:	Scale-out file storage system
Primary Use:	Ubuntu Linux desktop home directories
Target Users:	At least 3,000 users (long-term planning: at least 6,000)
Concurrent Clients:	1,000 Ubuntu desktops
Initial Capacity:	200 TB usable space
OEM Support:	Five years, extendable to seven years
Hardware SLA:	Failed parts replaced within 24 hours

Mathematical Sciences Support (MSS)

October 2026

Contents

1 Executive Summary 3
2 Background and Need 4

1 Executive Summary

TW Kambule Laboratories requires a scale-out file storage platform. Its primary purpose is to provide centrally managed Ubuntu Linux home directories for at least 3,000 users. Each user must have an initial allocation of 50 GB, and the platform must provide at least 200 TB of usable filesystem capacity after all required data-protection overheads, plus a separate snapshot reserve of at least 20% of usable capacity (at least 40 TB at 200 TB usable). The platform must support up to 1,000 simultaneous desktop clients and accommodate login storms of up to 1,000 desktops within five minutes. The accepted performance baseline is at least 64,000 total operations/s, 24,000 metadata operations/s and 320 MB/s read throughput under an agreed MSS test method.

The requirement is driven by planned and approved growth in the Mathematical Sciences, greater reliance on persistent user environments, high-concurrency teaching and assessment, and software-development workloads that generate intensive small-file and metadata activity. A scale-out architecture provides a single global namespace and allows capacity and aggregate performance to grow by adding nodes.

The platform may use either an all SSD or a hybrid SSD/HDD architecture. It must protect against individual disk failures, and provide node-level data and service redundancy from day one so that failure of any single storage node does not cause data loss or loss of storage service. It must include redundant power and networking. Backend/inter-node storage traffic must use a redundant storage switching fabric with at least two backend switches, or a functionally equivalent redundant architecture, so failure of any single backend switch does not interrupt storage service. It must support FreeIPA and Kerberos-secured NFS, SMB, per-user quotas, snapshots and future replication. Block and object storage protocols are optional and must not substitute for the mandatory file-storage capabilities. Client-facing storage uplinks must connect to the University's planned Cisco C9500-48Y4C Catalyst switch using 25 GbE SFP28 optical links from delivery while remaining compatible with the current 10 GbE client environment. All required optics at both ends and fibre patch cabling must be included. The C9500 is not part of the required redundant backend storage switching fabric.

The core scale-out storage nodes/appliances and filesystem software must be supported by one OEM. The solution must be supplied by that OEM or through an OEM-authorised partner. OEM-backed support must be provided for five years and be extendable to seven years, with failed hardware parts replaced within 24 hours. Physical installation is mandatory and must be performed by the OEM or an OEM-authorised partner. Before final acceptance, the supplier must perform the temporary acceptance-test configuration; after final acceptance, MSS will perform the final production software configuration and environment-specific integration.

2 Background and Need

Networked home directories are a foundational teaching service. A storage outage can prevent large numbers of students from logging in, accessing their work or participating effectively in lectures, practicals, tutorials, tests and examinations. As user numbers and concurrent desktops grow, a standalone file server creates increasing risks in performance, maintenance, migration, recovery and availability.

The proposed platform is intended to provide a sustainable service rather than a short-term capacity increase. It must expand without dividing users across separate servers, mount points or storage silos. Adding nodes should increase both usable capacity and aggregate performance while preserving a single namespace.

A scale-out platform also reduces lifecycle risk. Snapshots, automated data placement and non-disruptive expansion reduce the need for large manual migration exercises and make the service easier to maintain as it grows.

The platform must be supplied as an integrated solution. Normal expansion by adding nodes must increase capacity and aggregate performance without data migration or disruptive reconfiguration and must not create separate user-facing filesystems or namespaces. Nodes purchased in later years must be able to join the existing system, and older nodes must be able to be retired, without disruption to the user-facing namespace.

Failed hardware parts require replacement within 24 hours, 365 days per year.