Position on hold on MIT Careers site - Not currently accepting additional applications

ORCD Office of Research Computing and Data

MIT Position Description

Job Title: Mgmt 3, IT Generalists	**Position Title:** Lead Storage Engineer
Pay Grade: 11	**% Effort or Wkly Hrs:** 100%
Department: Office of Research Computing and Data	**Reports to:** Director of Platforms, Infrastructure, and Data Services
Prepared by: Renée Hellenbrecht	**Date:** April 3, 2025

Position Overview

The Office of Research Computing and Data (ORCD) is seeking a hands-on **Lead Storage Engineer** to design, implement, and manage high-performance storage infrastructure in collaboration with a diverse team of systems engineers. This role will play a pivotal part in the strategic transformation of storage architecture, ensuring scalability, reliability, and efficiency to support ORCD's continued growth.

The successful candidate will possess deep expertise in database or storage infrastructure, including distributed storage, high-performance parallel file systems, object storage, and data lifecycle management. Experience in large-scale data platforms, cloud storage, and HPC computing environments is essential. The role requires a strong foundation in Linux system tuning, storage optimization, data migration, and performance monitoring to support research workloads effectively.

Additionally, the ideal candidate will exhibit strong analytical and structured problem-solving skills, coupled with excellent communication and interpersonal abilities. This position will be instrumental in fostering cross-functional collaborations between engineering and operations teams across MIT, ensuring alignment with institutional objectives and long-term strategic initiatives.

The Office of Research Computing and Data at MIT serves every part of the Institute. It supports research missions ranging from breakthroughs in fusion power to cyber defense research and works closely with research at all five of MIT’s schools and the MIT College of Computing. This creates a uniquely rewarding environment in which new research workloads leveraging

advanced research computing tools appear regularly. ORCD is often at the heart of delivering services to support evolving needs as efficiently as possible.

Principal Duties and Responsibilities (Essential Functions**):

- Design, implement, and optimize large-scale storage infrastructure to support research computing and high-performance workloads.
- Architect and maintain high-availability storage solutions, including SAN, NAS, object storage, and parallel file systems.
- Develop and implement data lifecycle management strategies, ensuring efficient tiering, archiving, and retrieval mechanisms.
- Monitor and enhance the performance, capacity, and reliability of storage systems, identifying and resolving bottlenecks.
- Automate storage provisioning, monitoring, and disaster recovery processes to improve operational efficiency.
- Collaborate with systems engineers to integrate storage solutions with compute and network infrastructure.
- Lead and own incident investigation and resolution related to storage, maintain response playbooks, and mitigate risks through proactive measures.
- Work with ORCD leadership, senior engineers, and research community facilitation to develop and transform IT infrastructure into compelling storage solutions.
- Build and maintain roadmaps for next-generation storage services provided by ORCD.
- Supervise, mentor, and direct staff and contractors in storage-related initiatives.
- Performs other duties as assigned.

Supervision Received:

This position reports to ORCD's Director of Platforms, Infrastructure, and Data Services.

Supervision Exercised:

This position supervises, mentors, and directs staff and contractors.

Position on hold on MIT Careers site - Not currently accepting additional applications

ORCD Office of Research Computing and Data

Qualifications & Skills:

MINIMUM REQUIRED EDUCATION AND EXPERIENCE:

Education

- Bachelor's degree in engineering, computer science, related field or equivalent industry experience

Experience/Skills

- At least 7 years of experience in storage engineering, data infrastructure management, and/or open-source storage systems or database storage architecture.
- Experience working within a computing environment (Linux, Unix) supporting HPC or research/life sciences.
- Deep understanding of storage technologies, including parallel file systems (Lustre, GPFS), object storage (Ceph, MinIO), and enterprise SAN/NAS solutions (NetApp, Dell EMC, IBM Flash SAN).
- Expertise in Linux-based storage environments, including RHEL tuning, kernel optimizations, system performance tuning.
- Strong background in storage networking concepts (InfiniBand, SAS, SCSI, Multipath, NAS, SAN block storage).
- Advanced experience with RAID configurations, data replication, backup (TSM, RMAN), and disaster recovery solutions.
- Strong knowledge of automation tools and scripting (Ansible, Python, Bash) for storage management.
- Proven experience in migrating petabytes of data across storage platforms.
- Ability to analyze storage performance metrics and optimize I/O patterns for research computing workloads.
- Experience with monitoring platforms and storage-related alerting systems.
- Troubleshooting complex storage-related issues in coordination with internal teams and vendors.
- Comfortable working with wiki and ticketing tools such as Confluence, Jira, ServiceNow, or equivalent.

ORCD Office of Research Computing and Data

- Ability to communicate and collaborate effectively with both technical teams and non-technical clients and partners.

PREFERRED EDUCATION AND EXPERIENCE:

- Has worked within a DevOps, cloud infrastructure, or database service management environment.
- Experience with storage security best practices, encryption, and governance frameworks.
- Familiarity with software-defined storage solutions and hybrid cloud-based architectures.
- Experience with database clustering, replication, and high-performance data access strategies (Oracle RAC, PostgreSQL, MySQL, Vertica, Cassandra, CouchDB).
- Experience with cloud storage services (AWS S3, EBS, FSx for Lustre, EFS).
- Knowledge of storage hardware management (iDRAC, IPMI, Pacemaker, Corosync).
- Hands-on experience with data integration, ETL/ELT processes, and storage-based data pipeline automation.
- Red Hat Certified System Administrator (RHCSA) certification.
- Ability to compile and troubleshoot Linux source code and optimize packages for performance.

Hours, Location, Salary

This is a full-time 40 hours per week position. A hybrid schedule, with both remote and on-campus work, is available, with a minimum of 3 days each week on-campus. This role requires occasional travel to our data center in Holyoke, MA (typically every other month).

The role is an MIT salary grade 11, with a salary range of $140,000 – $165,000 annually.

How to Apply

Employment is contingent upon the completion of a satisfactory background check. We are unable to sponsor a visa for this role, so applicants must already be authorized to work in the United States.

Position on hold on MIT Careers site - Not currently accepting additional applications

ORCD Office of Research Computing and Data

Apply through the MIT Careers website (Job ID: 24910). Please include both a resume and a cover letter with your application.