



LONDON
SCHOOL of
HYGIENE
& TROPICAL
MEDICINE



Case Study

Software-Defined Storage
IT Infrastructure Management

LSHTM: SUSE Simplifies the Flow of Disease- Fighting Research

The London School of Hygiene and Tropical Medicine (LSHTM) is a global leader in public health research, helping develop new ways to fight disease worldwide. However, collecting data from thousands of clinical studies over decades was creating significant storage challenges for the school. With SUSE Enterprise Storage running on SUSE Linux Enterprise Server, LSHTM was able to consolidate data across numerous medical projects, resulting in greater transparency, more efficient allocation of scarce resources, and a better experience for researchers around the world.

London School of Hygiene & Tropical Medicine at a Glance:

The London School of Hygiene & Tropical Medicine (LSHTM) is a world-leading center for research and education that is committed to improving health and health equity in the UK and worldwide.

Industry and Location

- Education, London, England

Product and Services

- SUSE Linux Enterprise Server
- SUSE Enterprise Storage
- SUSE Manager

Results

- + Cuts time spent updating servers from several weeks to 15 minutes
- + Provides scalable, long-term storage for important medical research data
- + Saves one month's workload per year on storage management

Overview

LSHTM is a world-leading center for research and postgraduate education in public and global health with an annual research income of more than £180 million per year. The expertise of its staff and students encompasses many disciplines, and it is one of the highest-rated research institutions in the UK.

LSHTM has 4,000 students and more than 3,000 staff working all around the world with core hubs in London and at Medical Research Council Units in The Gambia and Uganda. In 2016 it was named University of the Year in the Times Higher Education awards. LSHTM has an international presence, collaborative ethos, and is uniquely placed to help shape health policy and translate research findings into tangible impact.

Challenge

FIGHTING THE BATTLE AGAINST DISEASE

In a world of increasingly complex and demanding public health issues, LSHTM's work in researching and combatting disease is more vital than ever. To do so effectively, large and complex research projects (including some clinical trials that span decades) must be able to coordinate, share, and store data in a secure and centralized way.

“At the end of the day, as the IT department, we are here to facilitate LSHTM's important work. With SUSE technology, we can provide researchers with the tools they need, and empower them to get on with their job.”

Steven Whitbread
Systems Manager
LSHTM

“With SUSE Manager, we can essentially patch multiple servers at once and schedule them for particular times on a given day - meaning we can now achieve in 15 minutes what used to take several weeks. This easily saves several days of work a month.”

Steven Whitbread
Systems Manager
LSHTM

“Our storage landscape has historically been somewhat disorganized,” says Steven Whitbread, Systems Manager at LSHTM. “People here work on a large number of data sets - some pre-existing, and some collected by themselves out in the field. A lot of our researchers conduct telephone and video interviews to gather survey results. We also have specialized medical data types, including microscopy, medical imaging and so on.”

Examples of LSHTM’s long-term projects include: a nine-year study of **drug effectiveness and public health strategy** in The Gambia and Kenya; and a twelve-year project to evaluate the **impact of cancer control policy** on English patients. These projects involved extensive interviews, consultations, strategic planning and raw data gathering. With such large and varied data sets to keep track of, data storage and management were frequently a challenge.

“I’d often have researchers come to me almost out of nowhere and say, ‘We need eight terabytes,’ and I’d have to produce that at fairly short notice,” describes Whitbread. “So finding enough storage was one problem, but the biggest issue is that we have research projects that run for decades. We’re now storing data for projects that have been running for over thirty years, and our researchers often need to easily access and review that data. For example, they might interview members of a local population and then go back ten years later to see what’s changed. That means we often have to keep multiple versions of the same basic data. Remote access is also very important for some individuals who are conducting research all over the world and rarely come back

to LSHTM’s main site in London. So not only do we need to store a large amount of data, we also need to make sure it is accessible from all over the world.”

He adds: “Until recently, there was a lot of independent storage around LSHTM. People would often keep work and archive datasets on departmental NAS devices, which created many issues. Different teams weren’t able to collaborate on this siloed data.

“It was also a challenge to keep up with auditing and compliance as an organization because it wasn’t clear what each research team had got squirreled away. We knew we needed to implement a new central data storage solution to make sure we, and our researchers, could do our jobs more effectively.”

AUGMENTING THE INFRASTRUCTURE

LSHTM recognized that its legacy enterprise storage system needed an upgrade in the face of the growing size and complexity of its data environment. Whitbread explains: “We have a high-performance clustered primary storage solution, which is still core to our storage needs, but it just didn’t have the archiving capabilities we needed to support researcher needs. So, we began looking for something to augment it.

“We wanted an enterprise storage system with a lower cost-per-terabyte, partly to make archiving more cost-effective and to allow us to bring in all of the disparate storage on personal and departmental drives. Our existing archive storage system was a bit of a nightmare to manage, especially when or if we chose to migrate. We wanted a scalable solution that could grow with us - that we could organically

“With SUSE Enterprise Storage we can simply add new nodes to the storage infrastructure, and the migration is done for us. This means that our migration processes, which we run every three to five years as old hardware gets replaced, can now be achieved much faster – saving the equivalent of one month’s workload every year in maintaining older data.”

Steven Whitbread
Systems Manager
LSHTM

update as we went along. It was important to us that we avoid rip-and-replace.”

Solution

FINDING SUSE

After considering several commercial and open source alternatives, LSHTM decided to deploy SUSE Enterprise Storage, a software-defined storage solution based on Ceph open source technology.

By selecting an open source solution, LSHTM knew they could reap the benefits inherent in community-supported technologies — benefits like greater flexibility, cutting-edge innovation, and robust security. And SUSE’s early adoption of Ceph-based storage, with its pioneering approach to both file and object storage-based systems, stood out as key differentiators.

Designed as a distributed storage cluster, SUSE Enterprise Storage provides an intelligent and highly resilient software-defined storage solution that separates the storage management from the underlying hardware. All capacity is combined across multiple storage arrays and can be grouped into one or several storage pools according to the specific requirements. By expanding a cluster with additional nodes when extra capacity is needed, the virtual storage pools can be scaled essentially without limit. Automatically redistributing data, the solution makes best use of the capacity and input/output performance of the additional nodes, facilitating optimal communication between the storage array and the host.

“In the end, it was cost-effectiveness that sold us on the SUSE solution, combined with the exceptional



level of support that SUSE offers,” says Whitbread. “We did look at several competitors that offered similar features, but we realized that they would have cost us more in the long term than our existing solution. Because the SUSE solution is completely agnostic from the hardware perspective, we also avoid having to invest in costly high-end storage devices. We were also impressed by the number of references for SUSE Enterprise Storage – it was reassuring to know that so many other organizations had already made this work.

“What’s more, we have been running SUSE Linux Enterprise Server for over a decade now, so adopting SUSE Enterprise Storage made sense with our existing infrastructure.”

DEPLOYING THE SOLUTION

LSHTM decided to deploy the SUSE Enterprise Storage 5 on a single initial cluster for bulk storage and for writing backup and archival data to the cluster.

“The SUSE team was extremely helpful in getting the solution set up,” remarks Whitbread. “They’ve also been very informative about the scalability of the solution, and we are looking to expand the cluster in the future.”

A large part of the appeal of SUSE Enterprise Storage 5 was its unique graphical user interface (GUI) – based on openATTIC – which gives managers an appealing suite of visualization tools.

“The visualization tools simplify data management, especially for our managers with less IT experience. It also makes it easier to explain and justify decisions,” describes Whitbread. “The tools enable us to report on data trends effectively and plan our future growth accordingly.”

“In addition, the SUSE technology integrates well with the traditional storage access methods most of our teams were already using. It has also enabled us to provision much larger data volumes. Our previous system had an upper limit on how much data could be stored, but with SUSE Enterprise Storage 5 we no longer have to split data across multiple volumes like we did before. The whole storage landscape is better unified and runs far more smoothly.”

SUSE Enterprise Storage 5 has been widely adopted by researchers at LSHTM and has become an accepted part of its software culture. “Our researchers have really embraced the new SUSE technology,” reflects Whitbread. “Our whole storage landscape has become more transparent, and it is much easier to manage and track data.”

MANAGING EXPECTATIONS

In addition to SUSE Enterprise Storage, LSHTM recently deployed SUSE Manager to handle its network of 30 physical servers and more than 100 virtual servers more effectively.

Whitbread says: “SUSE Manager has enabled us to effectively centralize the way we manage, update, and patch our server network, which has been a revelation.”

“Previously, when we needed to update our servers, we had to manage each one individually, as well as run the installations after hours. With SUSE Manager, we can essentially patch multiple servers at once and schedule them for particular times on a given day. In other words, we can now achieve in

15 minutes what used to take several weeks. This easily saves several days of work a month, giving us more time to be proactive in the way we manage and maintain our server network.”

Results

SMOOTHING THE TRANSITIONS

In their effort to support excellence in public and global health research, LSHTM currently operates a SUSE Enterprise Storage cluster of 177 terabytes, of which 77 terabytes are currently in use, storing over 50 terabytes of medical research data. This figure is only set to grow as research continues and the School onboards more legacy systems. They already plan to expand the cluster in the future.

“SUSE Enterprise Storage makes the whole migration process so much easier for us,” asserts Whitbread. “Previously, if we needed to replace part of our storage infrastructure, we had to manually migrate up to 80 terabytes at a time.”

“Furthermore, SUSE Enterprise Storage makes backup and archiving much more straightforward and less time-consuming than they were previously. From the IT team’s point of view, that makes our jobs easier, and from our researchers’ point of view, it means they always have access to any data they need. We can avoid downtime as we migrate data from old hardware to new, even when we are in the process of upgrading.”

A TRANSPARENT FUTURE

LSHTM is very pleased with how its SUSE solutions have helped them provide better support for researchers and their IT department.

Whitbread states: “The whole implementation has been about making storage easier to manage from an IT perspective while ensuring that data is backed up and available for our research staff. I think IT in general should be transparent. It should be about providing a service that just works for its users, and that is something that SUSE clearly understands.”

“At the end of the day, as the IT department, we are here to facilitate LSHTM’s important work. With SUSE technology, we can provide researchers with the tools they need, and empower them to get on with their job.”



SUSE

The Power Of Many

SUSE
Maxfeldstrasse
90409 Nuremberg
www.suse.com

For more information, contact SUSE at:
+1 800 796 3700 (U.S./Canada)
+49 (0)911-740 53-0 (Worldwide)