

CASE STUDY: MULLARD SPACE SCIENCE LABORATORY AT UNIVERSITY COLLEGE LONDON (UCL)



“We had already experienced a couple of catastrophic data losses due to RAID6 failures on old arrays in the past which were dealt with by restoring backups. We were painfully aware that data volumes were increasing, so the risk of loss was higher and restore times would get longer. We needed better RAID.”

- Alan Brown,
*Networking and Unix Systems
Manager, MSSL*

THE STORAGE CHALLENGE

University College of London's Department of Space and Climate Physics, (Mullard Space Science Laboratory - MSSL), is a world-leading research organization and is the UK's largest university-based space research group. MSSL delivers a broad, cutting-edge science program, underpinned by a strong capability in space science instrumentation, space-domain engineering, space medicine, systems engineering and project management.

MSSL's scientific research ranges from cosmology and the study of extragalactic objects, to studies of the solar system, the Earth, and humans working and living in space. MSSL continues to research and develop the next generation of space instrumentation, and since its inception in 1966, MSSL has taken part in more than 35 satellite missions and over 200 rocket experiments.

MSSL demands the most from their storage and requirements range from handling tens of millions of surveys at 64KB to space probe datasets at 10GB+ each and just about everything in between.

MSSL needed to replace their aging clustered SAN file servers. The clustered SAN had been chosen for reliability and failover, but distributed lock management leads to major degradation in throughput and latency. In addition to user impact, inadequate performance had a major impact on MSSL's large backup requirements and carried a risk of a significant impact on business continuity should large scale data recovery be needed.

“The most compelling extra feature we found with TrueNAS was that it’s based on OpenZFS. Unlike traditional RAID systems, when TrueNAS detects an error from a hard drive it not only corrects it but simultaneously writes the correction back to the hard disk. This is a valuable feature that mostly goes unnoticed but is critically important in large storage systems.”

- Alan Brown

“TrueNAS uses OpenZFS, which is incredibly reliable. It will provide levels of performance that will make your head spin.”

- Alan Brown

WHY MSSL CHOSE TRUENAS®

When MSSL began to look for a new storage solution they were looking for better latency, higher performance, improved file-serving reliability, and a storage solution that scaled well with the rest of the file-serving system. Having significant experience with large disk setups, MSSL found that many of the types of solutions they had used in the past still had significant file-serving performance limitations that would be difficult to overcome; they were looking for a proven commercial storage solution that met their needs.

Alan Brown, the Networking and Unix Manager at MSSL for over 15 years, first encountered ZFS on Solaris, but became even more interested as ZFS was ported over to the Linux and BSD-based operating systems. Alan continued to follow the development of ZFS and OpenZFS (the Open Source variant of ZFS). He recalled previous experiments with ZFS where massive amounts of spinning disk storage were coupled with SSD read cache and backed by write caching leading to the seek load on mechanical drives being reduced dramatically. These features would be very helpful with some common MSSL processing patterns which repeatedly read the same files.

Alan decided to research how different storage solutions were leveraging the power of the OpenZFS file system and how it compared with ZFS. After browsing through products from many different vendors, Alan was happy to find that TrueNAS by iXsystems is built on OpenZFS and can process a massive level of read activity. TrueNAS caches reads of the same file, allowing repeated access with next to no performance degradation. This would be a huge improvement over the older non-cached system MSSL was using and would alleviate one major bottleneck. Alan also reviewed SAN-clustered file systems but found that they still typically have poor latency performance due to internodal voting requirements. Alan also found that OpenZFS would eliminate the periodic offline file system checking needed in Ext4, GFS2, etc. and the associated disk non-availability.

“In real-world use, the performance, stability, and reliability of our TrueNAS is simply another level beyond anything we’ve used in the past. As older storage units age out, it is highly likely we will again be looking at TrueNAS appliances.”

- Alan Brown

The OpenZFS file-system is well known among storage experts as one of the most resilient copy-on-write file systems available. It uses data integrity checksums and replication of important blocks to ensure that your data is safe and secure against various types of common failures. MSSL concluded that TrueNAS was a natural fit for their storage requirement because of the proven track record of the ZFS file system, and because of iXsystems’ reputation for producing quality engineered products.

CONCLUSION

MSSL deployed a TrueNAS storage array with High Availability Unified Storage Appliance with over 200 TB of storage. The TrueNAS storage array enables MSSL to grow to well over 1PB. As their data grows, MSSL can easily upgrade their storage controllers, enabling them to scale to a massive 5PB, while meeting their data bandwidth and IOPS requirements.

After implementing their TrueNAS storage array, MSSL has seen a real-world improvement in simultaneous data throughput by up to 50 times. Even with this improvement in throughput, MSSL is estimating that they have headroom for a ten-fold increase in traffic.

TrueNAS was a solid choice for MSSL because it provides a level of stability and resilience beyond what is commonly seen in the enterprise storage industry. TrueNAS allows MSSL to have centralized laboratory storage for millions of files that MSSL curates, increased reliability, and the scalability MSSL needs as their repository grows. iXsystems provides MSSL with an extra spares kit for their TrueNAS storage array as well as



next-day hardware support in the event that they need assistance.

ABOUT IXSYSTEMS

By leveraging decades of expertise in hardware design, its contributions to many Open Source software communities, and corporate stewardship of leading Open Source projects (FreeNAS and TrueOS), iXsystems has become an industry leader in building innovative storage solutions and superior enterprise servers for a global marketplace that relies on open technology.

Thousands of companies, universities, and government organizations have come to rely on iXsystems' storage, servers, and consultative approach to doing business. Headquartered in the heart of Silicon Valley since its founding in 1996, the dedication to white-glove customer service, industry-leading support, and transparent technological contributions has never wavered and continues to help lay the foundation for a new era powered by open technology.

CONSIDERATIONS

Special thanks to Frank Awuah of Storm Technologies, Alan Brown, and Paul Lamb of MSSL for their assistance in this case study.

TrueNAS is a registered trademark of iXsystems, Inc. All rights reserved.