

October 2024

High-Performance Storage and AI Driving Animal Conservation at the Zoological Society of London



Intro

At the intersection of cutting-edge technology and conservation efforts, the Zoological Society of London (ZSL) and PEAK:AIO have developed an innovative approach to data management for wildlife studies. This collaboration, featuring the deployment of high-density Solidigm storage and advanced AI, significantly impacts the monitoring, protecting, and understanding of animal populations. At the core of this project is the PEAK:AIO solution, driven by Solidigm's D5-P5336 61.44 TB QLC drives, providing the infrastructure for rapid data storage and analysis—enabling researchers to make faster, more accurate decisions that ultimately benefit conservation efforts.

We recently had the opportunity to visit the London Zoo and explore how PEAK: AIO's data storage systems and cutting-edge AI-driven platforms are used by the ZSL to process vast amounts of data generated from wildlife monitoring projects, ranging from urban conservation for hedgehogs to global species observation. The zoo can tackle data management challenges inherent in modern conservation efforts by integrating cutting-edge storage and AI technologies. We also had the privilege of discussing these innovations with Mark Klarzynski, CEO and Founder of PEAK:AIO, [Professor Chris Carbone](#), a professor of zoological conservation, and others involved in this groundbreaking initiative.



From Hedgehogs to Global Species Monitoring

One of the most fascinating projects at the ZSL involves urban wildlife conservation, specifically focusing on hedgehog populations in London. Chris Carbone shared with us the challenges hedgehogs face, a species whose habitat is increasingly fragmented by human activity. “They struggle to disperse between populations,” Carbone says, “which leads to issues like inbreeding.” This problem is common across many species globally as urbanization encroaches on natural habitats.



Conservationists are working to combat this by creating pathways or “green corridors,” allowing animals like our hedgehogs to move freely between isolated habitats. However, awareness and public engagement are critical components in ensuring the success of these initiatives. The ZSL is also focused on making people

aware of local wildlife populations and encourages them to help maintain these green spaces, taking small steps to enable the animals to thrive in otherwise challenging environments.

But where do high-capacity Solidigm SSDs and NVIDIA DGX Supercomputers come into play? According to Carbone, they’ve become critical tools. Modern wildlife monitoring technologies have evolved rapidly over the past two decades, allowing researchers to collect vast data. However, this data comes with its challenges—storage and analysis.

The Explosion of Data: From Camera Traps to Genomics

In the early days of wildlife monitoring, camera traps were revolutionary but had limitations. Chris reminisced about the days when film-based cameras could only take 36 images, often leading to missed opportunities. “A troop of monkeys might use up all the film on the first day, leaving you nothing useful.” Fast forward to today, and modern digital cameras can capture tens of thousands of images, giving researchers a much more complete picture of wildlife behavior.

The amount of data generated has grown exponentially, especially with global conservation efforts using camera traps across multiple countries. For instance, India’s camera traps are some of the largest in the world, deploying over 10,000 cameras and generating millions of images. The [London HogWatch](#) project, too, has accumulated over 15 million images. With data collection scaling at this rate, managing and processing such large datasets has become a significant challenge.



Genomics is another area that contributes to data deluge. Researchers previously analyzed small fragments of animal genomes to understand population structures. However, with the increasing affordability and accessibility of whole-genome sequencing, data sets have expanded to terabytes. [Environmental DNA \(eDNA\) projects](#) further contribute to this growth, using DNA traces in environments to survey species populations.

Thanks to the cutting-edge density of the Solidigm QLC SSDs, the teams working on these projects have plenty of room to do work currently and have capacity available for the foreseeable future. This allowed the project to be scoped and built so that it would not require constant upgrades to meet data needs.

Building a Supercomputer in a Zoo

One of the more unexpected challenges the team faced while installing the new edge data center at the zoo involved the outdoor cooler unit and a group of Chinese water deer. The cooling system, essential for maintaining the proper temperature for the high-performance computing equipment, was located near an area inhabited by these sensitive animals. While upgrading the cooler to a larger unit, the noise and activity from the installation disturbed the deer, requiring the team to carefully adjust their plans to ensure the animals' comfort and safety.



The resolution was to temporarily relocate the water deer to a more secure and peaceful area within the zoo while the work was completed. The situation added a layer of complexity that most data center installations wouldn't encounter, highlighting the unique challenges of deploying technology in environments shared with wildlife. Once the cooler was installed and the animals returned to their habitat, the team continued the project with minimal disruption to the zoo's daily operations.

Edge Computing at London Zoo

Given the volume of data, ZSL required a powerful yet compact data center to handle these massive datasets. Enter PEAK:AIO and their high-performance storage solution deployed at the edge. This solution supports the zoo's imaging and genomics workloads. Inside the edge data center at ZSL's London Zoo offices, two DGX platforms from NVIDIA, combined with PEAK:AIO systems, provide a remarkable 1.2 petabytes of redundant and secure high-speed storage space.



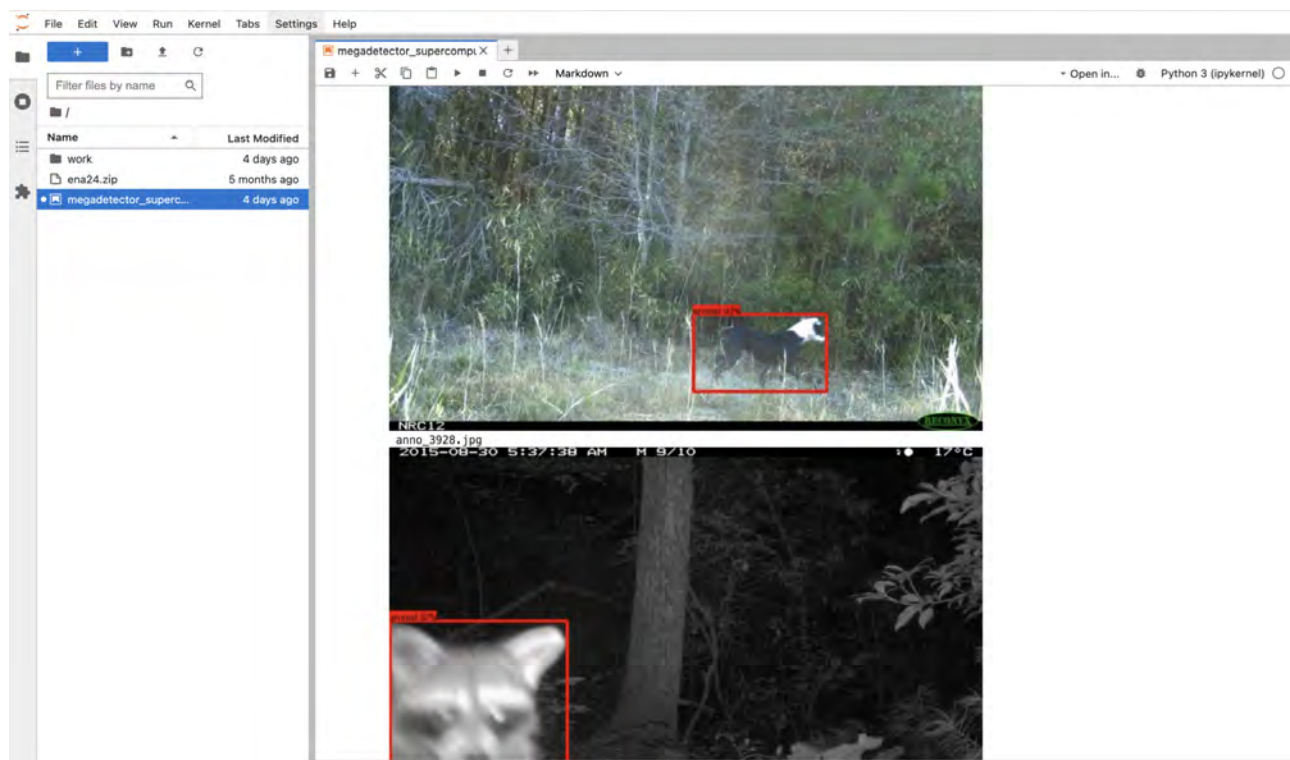
Mark Klarzynski, CEO & Founder of PEAK:AIO, explains that these storage systems are designed to keep up with the demands of AI workloads and wildlife conservation needs. The DGX systems enable fast AI processing, which can analyze millions of images in less time than traditional computing setups. “We are pushing the boundaries of performance in such a small space,” Mark said. “It’s not just about storing the data but also about processing it as quickly and effectively as possible.”

With RDMA over NFS and NVMe over Fabric, these storage systems can provide high-throughput access to

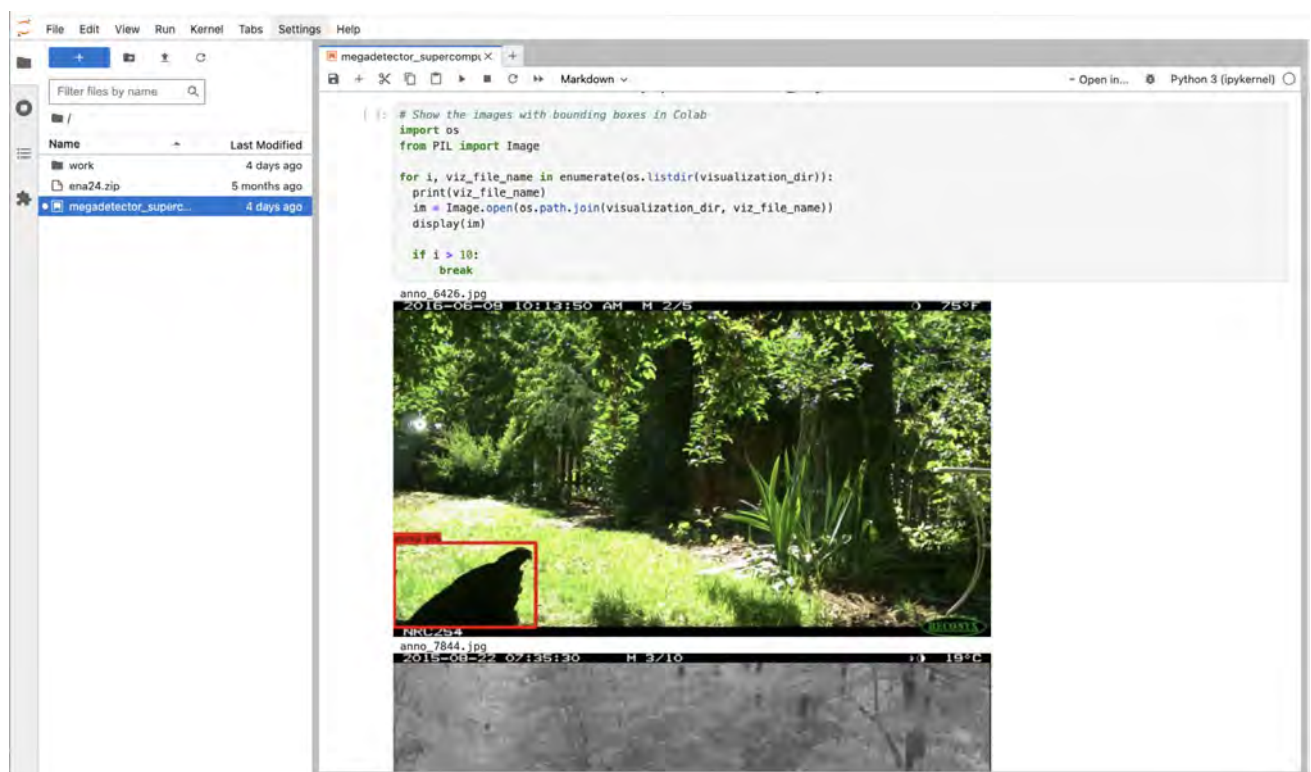
data, which is necessary when processing data at the scale ZSL requires. GPU-accelerated AI workloads demand this level of performance, especially when dealing with tasks like species classification and object detection in images.

The Role of AI in Conservation

One of the standout AI tools in use at London Zoo is Microsoft's MegaDetector. This AI model is instrumental in automating image preprocessing by detecting and filtering out non-relevant items like humans or cars and isolating animals for further analysis. Previously, a standard laptop might only process a few images per minute, but with the new infrastructure, thousands of images can be processed in the same timeframe.



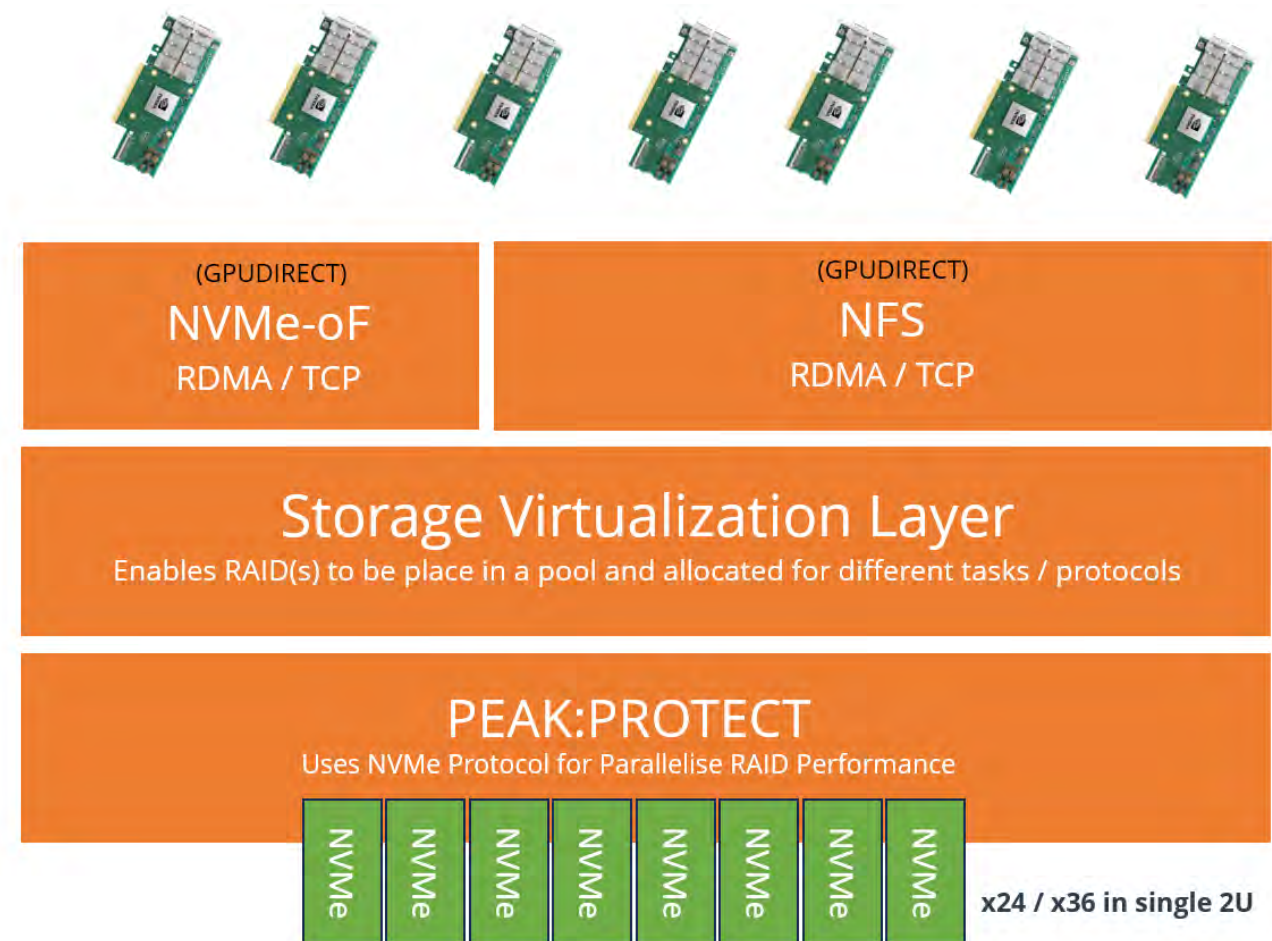
This is just the first step in a more complex AI-driven process. After the initial filtering, AI models are employed for species recognition—a much more challenging task. “You’d think we’d be better at it by now,” Carbone muses, “but it’s incredibly difficult.” Vegetation often obscures animals, making it difficult for AI to classify species correctly. However, new techniques, such as using bounding boxes to focus AI attention, improve accuracy.



Once the images have been tagged and analyzed, the data can be fed into more complex models. These models provide conservationists with critical insights, such as population estimates and species interaction patterns. The results of these analyses help shape conservation strategies in London and around the world.

Real-Time Backup and Data Security

Data integrity is crucial when dealing with invaluable conservation data. PEAK:AIO has built-in redundant backup systems to ensure that none of this data is lost. The 1.2 PB of Solidigm storage is backed up in near real-time using NVMe over RDMA, and the data can be recovered almost instantly without needing a traditional restore process. Klarzynski emphasizes, “You don’t actually recover; you just use it.”



For added security, the backup process ensures that data is immutable, meaning it cannot be tampered with—a crucial feature, given the sensitive nature of conservation data. Should something go wrong, the data is immediately accessible, preserving years of research.

The Future of Conservation Technology

As conservation efforts continue to scale, so will the demand for advanced technology solutions. The work performed at London Zoo serves as a model for what is possible when you combine AI, high-performance computing, and edge-based storage systems. Mark from PEAK:AIO points out that while they've achieved incredible things already, they only scratch the surface. "There are things that AI can do that we don't even know about yet," he says.

Global collaboration will likely become more critical in the future. Researchers envision global camera trap networks, similar to how weather stations operate, providing real-time data on species and ecosystems across the globe. As the world faces increasing environmental challenges, these technologies will become indispensable tools in the fight to protect our planet's wildlife.

Conclusion

The collaboration between ZSL, PEAK:AIO, and Solidigm at the London Zoo demonstrates how technology can drive real-world change in conservation. By leveraging high-performance storage and AI, conservationists can process massive amounts of data in record time, providing new insights into animal behavior, population dynamics, and ecosystem health. The future of wildlife conservation is undeniably digital, and projects like these are shaping how we protect our planet's most vulnerable species.

Jordan Ranous, AUTHOR

AI Specialist; navigating you through the world of Enterprise AI. Writer and Analyst for Storage Review, coming from a background of Financial Big Data Analytics, Datacenter Ops/DevOps, and CX Analytics. Pilot, Astrophotographer, LTO Tape Guru, and Battery/Solar Enthusiast.



This report is sponsored by Solidigm. All views and opinions expressed in this report are based on our unbiased view of the product(s) under consideration.