Crusoe Builds a Climate-aligned AI Cloud With Power for the Future

Where Crusoe saw methane being flared in oil fields, they saw potential. Crusoe pioneered infrastructure that taps into wasted energy to power advanced computing systems for machine learning training, real-time inference, and other AI processes. Their Crusoe Cloud offers an AI platform that costs a fraction of hyperscale alternatives and lowers the environmental impact of running these compute-intensive workloads. To store their customer's massive datasets and sustain fast, direct access to their demanding AI/ML workloads required block storage that could deliver scalable high performance, consistent low latency, and streamlined operations as a service provider.

"We're building a platform. Lightbits has built a data platform that has enabled us to build our AI cloud platform. We want to continue to provide functionality at every level of the stack to meet our customers wherever they are."

THE CHALLENGE: PROCESSING MASSIVE DATASETS WITH SPEED AND EFFICIENCY

Training large AI models is compute-intensive, requiring expensive GPUs and even more electricity than mining cryptocurrency. This requires a balanced approach to power-efficiency, and high performance to keep GPUs busy. Along with the massive, specialized compute and storage requirements, the success of training AI models rests on high-speed, consistently rapid access to data. The more data the model has access to and the faster it can receive that data, the better the AI outcome. This quest for more precise modeling drives a never-ending pursuit of ever-higher performing storage to enable the iterative training that yields better results, faster.

Due to the size of these datasets, Crusoe needed a solution with a cost model that scaled linearly with capacity. As Mike McDonald, director of product management at Crusoe, explained, "We have to offer the right price/performance. Our customers strive to be at the cutting edge, generating new foundational models or running inference, but they have to do it in an economical way."

Because Crusoe Cloud is operated by a small team consisting mostly of software engineers, not storage specialists, they needed a software-defined solution, coupled with storage drives delivering a mix of power-efficiency and performance, that would be easy to operate. And it needed to be backed by resilient, trouble-free platforms and a high-quality, developer-focused support team when needed.

Industry

AI Cloud Service Providers (AIaaS)

Solution

The Lightbits data platform enabled Crusoe to speed up their high-performance compute cloud while minimizing costs with powerful block storage to manage their customers' massive data sets.

Business Requirements

- High Performance
- Highly Scalable
- Highly Available
- Cost-Efficient

Business Benefits

- **Performance at Scale:** Performance that rivals local flash with far greater scalability.
- **Efficiency.** Persistent storage that is disaggregated from compute resources.
- **Availability and Data Protection:** Clustered architecture for availability with built-in snapshotting
- **Cost-Efficiency.** Scalable, high performance with efficiency to compete with hyperscalers.

Environment & Workloads

Distributed LLM model training, AI inference and generative AI, high performance computing (HPC), ray tracing, rendering and other graphics workloads.

THE SOLUTION: A POWERFUL AI CLOUD THAT RISES ABOVE THE HYPERSCALERS

To help them find the right solution, Crusoe consulted with Alexey Stolyar, CEO International Computer Concepts (ICC), a leading system integrator with expertise in GPU-accelerated AI cloud computing and high-performance computing (HPC). After evaluating Crusoe's needs, Stolyar recommended they investigate Lightbits for its exceptional performance and full suite of enterprise data services, along with Micron® NVMe™ SSDs for storage.

Crusoe performed extensive performance tests across several block storage options, including against Ceph, a popular open-source, software-defined storage offering. "Lightbits on Micron SSDs blew everything else out of the water, with consistently higher throughput, IOPS, and lower latency across the board. We were impressed with the performance and consistency of Lightbits," McDonald shared.

Lightbits demonstrated up to a 4x performance advantage in terms of bandwidth. Lightbits also scales IOPS with increased load while maintaining low latencies. It outperformed the competition by consistently maintaining latencies under 0.5 milliseconds (ms) compared to the competition which exceeded 2.5ms under random access. With these results, Crusoe determined that Lightbits would perform at scale to meet their growth targets while maintaining sub-millisecond latency.

They also were pleased that Lightbits would meet their ambitious high availability and data protection goals through its fast, efficient snapshotting technology across multiple availability zones.

Hardware is another important decision even in SDS solutions like this. They chose Micron 7000 series SSDs for their reliability, low latency, and high density suited for the demanding AI/ML applications that run on Crusoe Cloud. The Micron 7000 series has a low, average, active power consumption of only 17 watts for more than 15TB of storage, which helps Crusoe Cloud save energy and reduce power costs.

THE RESULTS: A PLATFORM TO PUSH THE BOUNDARIES OF AI/ML FOR THE FUTURE

With Lightbits, Crusoe customers can now build a petabyte-scale file system using only a couple of instances with most of that data managed by Lightbits. That gives customers access to persistent, fault-tolerant, high-performance storage that was previously ephemeral. Crusoe can also efficiently store custom operating system images for their customers, enabling Crusoe to free up paused virtual machines and enable customers to resize their virtual machines.

Even at the petabyte scale, Crusoe customers are consistently seeing higher throughput and IOPS and lower latency across a wide variety of workloads. That's a critical factor for customers with mixed large and small read workloads. As McDonald explained, "With Lightbits, we're going to continue to push the boundaries of high-performance storage as we get into multi-modal workloads with AI models that have text, video and images. That's a wide gamut of storage workloads we need to optimize for performance and efficiency."

Lightbits also gives Crusoe a data platform that ensures availability and is easy for their team to maintain. Lightbits' advanced snapshot and backup capabilities provide crucial recovery options for Crusoe Cloud customers. Now customers can safeguard their data against unexpected loss by taking storage-efficient snapshots and rapidly restoring data when needed. Micron 7000 series SSDs enhance these advanced data management capabilities with 6.8GB of read throughput and up to 5.6GB of write throughput per SSD to ensure fast performance.

What started as wasted energy venting into the clouds is now harnessed to drive the world toward a more sustainable future with efficient storage and data management from Crusoe Cloud, Lightbits, and Micron.

For more information on Lightbits: www.lightbitslabs.com.
For more information on Micron SSDs: www.micron.com/SSD
For more information on Crusoe Cloud: www.crusoe.ai
For more information on International Computer Concepts: www.icc-usa.com

To get started contact your Lightbits or ICC representative today.

1830 The Alameda, San Jose, CA 95126 E: info@lightbitslabs.com www.lightbitslabs.com

LBCS05/2024/03

INTERNATIONAL COMPUTER CONCEPTS | Micron® | lightbits