

ZEBRA MEDICAL VISION USES ZADARA TO SCALE STORAGE WITHOUT LIMITS

Zebra Medical Vision provides radiologists with artificial intelligence (AI) tools that save lives. The company's deep learning algorithms empower radiologists to improve outcomes for acute conditions and identify at-risk patients faster than otherwise possible. Zebra Medical was recently recognized as one of the 10 most innovative health companies of 2020 by Fast Company.¹

When dealing with acute conditions, early discovery often means a life can be saved. Zebra Medical technologies spot symptoms that cannot be seen with eyes alone, making early diagnosis more likely.

Zebra Medical has grown rapidly since it was founded in 2014. The Zebra Medical AI1 ("All In One") solution is installed in hospitals globally, having secured FDA approvals and the CE mark. The company has continuously innovated, partnering with Oxford University Hospitals' osteoporosis team and developing a COVID-19 solution that analyses chest CT scans to diagnose, triage, and monitor patients.

This evolution, along with an array of data challenges faced by medical imaging companies generally, led Zebra Medical to seek a new storage solution that met growing capacity and performance needs.

Zadara was the company's chosen provider. Zebra Medical needed a solution that could deliver extreme throughput at virtually any scale. "That was the killer feature for us," explains Eyal Toledano, Co-Founder and CTO, Zebra Medical.

HIGHLIGHTS

- High performance at any scale
- Rapid capacity scaling without limits
- Cloud-like managed service, including 24/7 support and utility pricing
- AWS S3 integration with consistent tools across cloud and data center
- High throughput of 8GB/s read, 2 GB/s write enabled all teams' workloads to run concurrently

THE CHALLENGE

It takes a vast amount of data to create Zebra Medical solutions. The company develops algorithms that analyze medical images (such as CT scans, x-rays, and mammography scans) and detect medical conditions. Insights produced enable radiologists to care for more patients, with faster and more accurate diagnoses.

Algorithms are developed using deep learning (DL), a branch of AI. Each DL model needs to be trained and validated with many high-quality and diverse samples. This creates a unique set of challenges.

Scalability being one issue. The company started acquiring a lot more data, a lot faster than before. Medical scanners also continuously evolve to offer higher resolutions, which in turn generate larger files. Multiply this by the dataset size, typically hundreds of thousands or even millions of scans, and you quickly reach very high volumes.

Another challenge was **throughput**. The Zebra Medical team had started out with three DL researchers, then reached 20 and continues to grow. The company needed to increase how many studies it could ingest and consume, and how many inference jobs it could run, concurrently. Zebra's previous storage architecture could not accommodate these needs. And while data teams in other industries can slow storage growth by deleting old data, an active archive is the only option for medical imaging companies like Zebra Medical. Nothing can be deleted due to strict regulations related to record retention. If the team initially trains a model with half a million samples, for example, and then acquires a million new samples, the half-million from before gains new meaning. There is never a moment where data becomes obsolete and can be sent to a lower tier; it needs to remain hot and be available at all times.

Zebra Medical therefore needed a scalable, flexible and cost-effective data storage solution that could quickly adapt to meet growing and changing business needs.

SCALE WITHOUT LOSS OR LIMIT

The Zadara Storage Cloud provided a fully-managed service with on-premises storage and a simple 'pay-only-for-what-you-use' pricing model.

Critically for Zebra Medical, Zadara storage offers unlimited flexibility - including on-demand refreshes at no additional cost - and high performance that scales alongside capacity. Performance of 8 GB/s read and 2 GB/s write has been maintained, even as volumes continue to grow. The solution integrates seamlessly with AWS S3 and supports Zebra Medical's network upgrade from 10 Gigabit Ethernet to faster 100 GbE. With other solutions, performance claims were based on small quantities of data held in cache; on the other hand, Zadara was designed to handle the quantities the company needed out of the box.

Since first deploying the Zadara solution, Zebra Medical's storage capacity has nearly doubled and yet has achieved balanced, linear performance while able to ingest 200-300 TB of data at a moment's notice.

"The Zadara solution is built for the quantities we need. Everything is in the same tier in a cost-effective manner."

Eyal Toledano
Co-Founder, CTO, Zebra Medical

THE BENEFITS

NO NOISY NEIGHBOURS

For the clinical researchers at Zebra Medical, the Zadara solution has meant freedom to focus on studies instead of infrastructure. Everyone can run their workloads at the same time.

LIKE CLOUD ON-PREMISES

Having a managed service and pay-as-you-go pricing delivers the benefits of public cloud, with the advantages of on-premises hardware. Zebra Medical has access to 24/7 expert management and support. This means teams don't need to worry about issues that come up with the hardware. Operations teams have shifted their focus to core business priorities.

GPU PROXIMITY

Zebra Medical uses on-premises GPUs to train its deep-learning models. "Doing GPU in the cloud isn't yet cost-effective," explains Toledano. "So, the fact that Zadara gives us an on-prem solution to match our GPUs, with cloud-like service, is a good fit for us."

With the Zadara on-premises storage solution, data can be copied from storage to the GPU cache at much higher speeds than if they were copied from the cloud. The cost of downloading terabytes of data every day is also eliminated.

S3 OFF THE SHELF

Compliance with AWS S3 storage improves flexibility and efficiency further. "We use Zadara like we use S3, with off-the-shelf tools like AWS CLI and Boto3," explains Toledano. "It saves a lot of custom development time - most of the problems are already solved, so we only need to do monitoring."

The company can also overflow workloads to the cloud when its GPUs are full. Workloads are more flexible and not bound to the data center.

"It's great that we know the hardware is reliable and it's all taken care of, in a similar manner to if we'd gone with a cloud solution."

Eyal Toledano
Co-Founder, CTO,
Zebra Medical

"We use Zadara like we use S3, with off-the-shelf tools like AWS CLI and Boto3. It saves a lot of custom development time - most of the problems are already solved, so we only need to do monitoring."

Eyal Toledano
Co-Founder, CTO,
Zebra Medical

SCALING WITH AGILITY

When Zebra Medical has needed to increase storage capacity, Zadara has been quick to meet the company's new requirements. Capacity upgrades of half a petabyte happen seamlessly with no downtime for installation. There is never a need to sacrifice data, or any risk that the company will max out its storage capacity.

WHAT'S NEXT

As the company continues to grow and innovate, Zebra Medical aims to reach more patients in new countries and to develop algorithms that diagnose new medical conditions. Because every country is diverse in its prevalent healthcare, the company needs as much data variance as possible to make its algorithms effective anywhere in the world.

The company plans to use the object storage capacity and AWS S3 integration provided by Zadara to make data mining even more effective as it continues to improve patient care across the globe.

"With Zadara, whatever scale we need, whatever throughput we need, we can get to it. That's the killer feature for us."

Eyal Toledano
Co-Founder, CTO,
Zebra Medical

To learn more about Zebra Medical Vision, please visit www.zebra-med.com.

¹Fast Company. 2020. The 10 Most Innovative Health Companies Of 2020.

Available at: <https://www.fastcompany.com/90457848/health-most-innovative-companies-2020>. Accessed 23 June 2020.

zadara

Zadara is enterprise storage made easy.
Any data type. Any protocol. Any location.

To learn more about how Zadara can help your enterprise IT needs, you can visit:

www.zadara.com
support@zadara.com