



DEEP LEARNING

AI startup revs up its business with PowerEdge servers and NVIDIA Tesla GPUs



ZIFF Inc. | Artificial Intelligence | United States

Business needs

ZIFF, Inc. needed a highly reliable, enterprise-class server to train algorithms for artificial intelligence.

Solutions at a glance

- Dell EMC™ Ready Solutions for AI, Machine and Deep Learning
- Dell EMC™ PowerEdge™ C4140 servers
- NVIDIA® Tesla® V100 Tensor Core GPUs

Business results

- Increased productivity with extremely reliable computing infrastructure
- Eliminated continual problems with server crashes
- Accelerated training for models with tens of millions of images

ZIFF can train models with tens of millions of images in

LESS THAN A DAY



ZIFF traded constant server crashes for a server that

JUST WORKS



The rise of a new AI database

ZIFF Inc. is a small company with a big vision for the future of artificial intelligence (AI). The Utah-based startup pioneered the development of a groundbreaking AI database for unstructured image, audio and video data. This database arms the company's customers with built-in indexing, search, training and inference capabilities for the type of data that doesn't fit into the tidy rows and columns of traditional databases. The ZIFF database anticipates data consumption by automatically building out AI models using the image, audio and video data and metadata fed into it.

"When we talk about unstructured data — like image, audio and video — we're talking about data that doesn't fit neatly in an Excel spreadsheet," says ZIFF co-founder David Gonzalez, who serves as the company's chief executive officer and chief technology officer. "It's not numbers. It's not categories. It's not transactions. I like to say that unstructured data is the real data behind the real work that your real people do."

Unstructured data can stem from phone calls fielded at a help desk, photos from a manufacturing floor, images used in medical diagnostics and countless other audio, video and image files that are generated in the course of the business day. It's the work of the business, Gonzalez says. And the deep-learning AI brain inside the ZIFF unstructured database can unlock the insights buried within it.

A case in point

Let's consider an example offered up by Ben Taylor, chief AI officer for ZIFF. One of ZIFF's customers, Chatbooks, is a scrapbooking company that builds physical scrapbooks from photos and information that users post on social media

sites. The idea is to help people hold on to their memories by offering an easy way to turn digital images into photo books. For this work, the company leverages ZIFF's AI database and its built-in deep learning capabilities to automate what would otherwise be a labor-intensive process.

"Using deep learning, they were able to build AI workloads that can organize an album and then actually predict which photos are most likely to succeed, or which photos are best," Taylor says.

That's a big advantage for Chatbooks and its customers, according to Nate Quigley, the company's CEO and co-founder. "Thanks to ZIFF, we're able to make our apps so much more intuitive, smarter and user-friendly," Quigley says. "ZIFF has been helping our users quickly turn mountains of photos into beautiful books containing the images they most want to remember, and almost effortlessly."¹

Getting started — on a bumpy road

Like many AI startups, ZIFF started out using a consumer gaming machine, and soon after realized that the gaming machine couldn't handle the volumes of data required for deep learning. It crashed often, and when it was up and running it couldn't scale effectively to mine through datasets that can encompass millions or tens of millions of images and audio files.

"We struggled, like most startup companies," Gonzalez says. "We built our own system based on gaming GPUs and a big gaming motherboard . . . The thing would crash once a week, sometimes once a day. It just wasn't going to work for us."

¹ ZIFF, Inc. news release, "[ZIFF, Inc. Announces Unstructured AI Database Product Release](#)," April 24, 2018.



Nurturing next-gen AI leaders

ZIFF is taking advantage of the resources made available to startups via the [NVIDIA Inception](#) program. NVIDIA Inception nurtures dedicated and exceptional startups that are revolutionizing industries with advances in AI and data science. The program functions as a virtual accelerator that helps startups during critical stages of product development, prototyping and deployment. Every startup in the program gets a custom set of ongoing benefits, from hardware grants and marketing support to training with deep learning experts.

To cure their never-ending server headaches, and meet the needs of a business poised for growth, Gonzalez and Taylor set their sights on an enterprise-class system built for the rigors of deep learning – the Dell EMC PowerEdge™ C4140 server.

A server built for AI

The [Dell EMC PowerEdge C4140](#) is designed for the most demanding workloads in cognitive and technical computing. This high-density, accelerator-optimized server speeds up applications with four NVIDIA Tesla V100 Tensor Core GPUs in a 1U, two-socket package.

Under the hood, ZIFF's PowerEdge C4140 servers have four [NVIDIA® Tesla® V100 GPUs](#). It pairs NVIDIA® CUDA® and Tensor Cores to deliver the performance of an AI supercomputer in a GPU.²

“The Volta architecture processes up to 8,000 images a second, Gonzalez says. “And that means that we are solving tens of millions of image problems for people in a day or less. On more standard hardware – or if you have your own system that you’ve built out – you’re looking at weeks to months to solve that problem. We can turn that around on the Dell EMC system inside a day or two.”

Better still, the PowerEdge C4140 server has proven to be rock-solid reliable – so much so that the ZIFF team doesn't

² NVIDIA website, “[NVIDIA Volta, The New GPU Architecture, Designed to Bring AI to Every Industry.](#)”



need to pay any attention to it.

“The best compliment you can give a hardware company is that they're forgettable,” Taylor says. “And we feel like, on the hardware side, we don't think about it anymore. We think about the innovation. We think about the work that we're doing.”

Gonzalez echoes those sentiments. “The PowerEdge C4140 has been a godsend for what we do,” he says. “The thing just sits in our data center and works.”

IT simplicity with Ready Solutions

The PowerEdge C4140 server is at the core of the [Dell EMC Ready Solution for Deep Learning with NVIDIA](#) so when ZIFF is ready to scale out, Dell EMC has done the testing and validation to take the guesswork – and a great deal of the risk – out of deploying new IT solutions.

“The nice thing that we love about using Dell EMC is we don't care about what configurations we go with,” Taylor says. “We just want something that's very, very fast and that works – top of the line, it works.”

Prior to deploying the PowerEdge C4140 server, the ZIFF team had to make its best guesses about hardware configurations and then spend time wrestling with system configuration issues and component incompatibilities. And that's not where any company wants to spend its time.

“From a business perspective, we don't care about any of that,” Taylor says. “We hate all of that. We just want to do stuff. We just want to run our software. We don't want to be configuring servers and figuring out what we should put in them. We just want best practices baked in. Dell EMC did that very well. Dell EMC has preconfigured systems that just work for our application.”



Learn more about Dell EMC Ready Solutions for AI
DellEMC.com/readyforai



View all customer stories at
DellEMC.com/PowerEdge-Stories



[Share this story](#)

Copyright © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. NVIDIA and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. This case study is for informational purposes only. The contents and positions of staff mentioned in this case study were accurate at the point of the interview conducted in June 2018. Dell and EMC make no warranties – express or implied – in this case study.

