

Greater workload flexibility and better economics with AMD EPYC processors

CUSTOMER



INDUSTRY

Global Collaboration Platform

CHALLENGES

Deliver reduced costs over a vastly varied range of compute workloads

SOLUTION

Deploy AMD 1st and 2nd Generation EPYC processors in data centers

RESULTS

20 percent reduction in hardware costs in data centers

AMD TECHNOLOGY AT A GLANCE

1st Generation AMD EPYC CPUs AMD EPYC 7642 CPU with 48 cores Dropbox needs little introduction. The company has become a household name since the company was founded in 2007. It hit over 15 million paying users by 2020, and over 600 million users of all types by the beginning of 2020. These users upload a massive amount of files per day.

Although Dropbox is best known for its cloudbased storage, there are a lot of very clever features built into its service that enhance its functionality. The ability to preview Microsoft Word files and images in many formats or play back video are just a few, and all require integrated applications that have subtly differing needs. To deliver their services Dropbox primarily utilizes four types of servers to handle its wide range of requirements -Compute, Storage, Database, and Apache Hadoop[®] Distributed File System

(HDFS). The compute servers are the ones that are tasked with running all the integrated apps and utilities. This vast array of tasks necessitates the most flexible servers available, and this led Dropbox towards AMD EPYC technology. With the massive number of clients and files, efficiency of compute as well as price performance are critical for Dropbox. The selection of AMD EPYC CPU-powered servers has resulted not only in greater flexibility, but a 20 percent cost reduction for them as well.

Varying needs of Dropbox workloads

"We were looking for an evolutionary design that was a natural progression from our prior generations of compute servers," explains Steven Hart, former CPU Staff Engineer at Dropbox. However, the large and diverse workload requirements made the optimum platform choice far from straightforward. "It's hard to enumerate all the workloads that we deploy on compute servers. It's anything from video trans-coding, to distributed caches, to orchestration software. It makes that class of systems challenging to design for, because we can't optimize for one workload, which is the case with other systems, for example a storage system or a database system. We must consider the whole field of applications. There are around 10 that we can optimize for, but we

"The system that 2nd Generation AMD EPYC processors allowed us to build, gave us the best cost per transaction to meet the performance requirements."

Steven Hart, CPU Staff Engineer, Dropbox can't give short shrift to any application. We have to deploy something that benefits all our workloads and serves in multiple roles with a variety of applications whose performance profiles and resource demands may differ vastly. Some are network intensive, some are CPU intensive, some are memory intensive, so we're looking for that golden median."

When AMD EPYC technology became available, it had clear potential to fulfil these needs. "We were fortunate enough to have somebody on the team that had worked at AMD, who brought the 'Zen' Architecture to our attention," says Refugio Fernandez, Manager, Supply Chain, Dropbox. "The single socket capability was appealing. We started meeting with AMD and our technical team."

Dropbox began testing AMD EPYC 1st Generation processors. "Our benchmarking process starts out with publicly available tests, for example SPECrate[®], to look at the improvements in core count between different options," says Hart. "As we mature the product and we can integrate it more into our infrastructure then we can run canary workloads [ed. various model workloads to test for functionality] or internal versions of our production workloads on the system to see the actual benefit that we realize in production."



In their testing, Dropbox was able to identify quantifiable benefits from moving to a single-core platform with AMD EPYC CPUs. "Any time you have more devices," adds Hart, "you complicate the design and if we can find a single-socket system that can perform in that same role, that introduces simplicities that benefit us from a verification perspective and a system design perspective."

Increasing rack density with AMD EPYC processors

The positive test results and potential for platform simplification led to Dropbox deploying thousands of AMD EPYC 7001 CPUs in 2018 and in 2019. "With the capabilities of a single socket platform, we were able to see that there were cost savings. We were able to reduce our Bill of Materials Cost, and ultimately impact and optimize our TCO," explains Jesse Lee, Commodity Manager, Dropbox. "The option to select a number of sockets that suited our workload and requirements allowed us to benefit from lowered power consumption, as well as lowered capital expenditure."

When 2nd Generation EPYC processors arrived, the potential benefits improved dramatically. "With this generation of AMD EPYC processors, we went to a higher core count CPU, the 48-core 7642," says Hart. "We saw we could easily scale production workloads, with very few issues. With the 7642, we were able to build the ideal balanced system. We could populate all the eight memory channels; we could drive them up to take advantage of DDR4 up to 3200MHz."

With the advent of AMD EPYC 2nd Gen processors, Dropbox dramatically increased their EPYC powered server deployment in 2020. "We raised our sights and were looking to address cost per unit performance," says Hart, "trying to drive that down and increase the amount of rack-level density that we have. So those were much more aggressive goals for the 2nd Generation CPU than we had for the 1st Generation." Dropbox was aiming to get the most density it could out of existing rack space. "There are cooling and power limitations on what we can put into the rack. Our goal is really to arrive at the densest rack-level solution that we can, that provides the best ratio of cost per transaction. AMD EPYC 7002 CPUs really hit the mark there."

Reducing costs per transaction

The high core counts of AMD EPYC CPUs really paid dividends for Dropbox. The raw performance allowed more efficient deployment of resources.

"When we've built a system that scales up to a factor of 3x [the application density], we have to take advantage of that," says Hart. "We do that in two ways. We have a set of applications that can scale with core count very well, and these applications receive enough traffic to justify scaling them up. We also have a significant number of applications that, if put on a single machine, would under-utilize the machine. That represents a lot of stranded performance resources on a single system. So along with a move to a higher core count processor, we are introducing multi containerization.

We're able to consolidate a lot of these workloads onto a single system, so the benefit is larger than 3x."

Around ten of Dropbox's applications fit into the first category and can be scaled with core counts and represent 70 percent of Dropbox's server capacity. The other 30 percent includes around 100 applications that can benefit from containerization into a shared environment. Dropbox uses Apache Mesos® orchestration for this. "Our overriding goal was to reduce the cost per transaction," says Lee. "That was the ultimate yardstick. The system that 2nd Generation AMD EPYC processors allowed us to build gave us the best cost per transaction to meet the performance requirements. It drove that cost per transaction

down." Lee estimates that "the Dropbox hardware cost saving exceeds 20 percent" over its previous solution.

"One of the challenges that our hardware design is running into is that we have existing infrastructure," says Hart. "We have limitations in cooling and power. For the foreseeable future, EPYC products are going to enable us to best utilize the physical resources that we have available in the data center, without resorting to exotic means. We can put more compute units in the rack, further optimizing our cost-performance metric."

"My message to our peers in our industry would be that this is real," concludes Fernandez. "AMD has come through, there's credibility behind it, it has worked for us and we can believe in it."

WANT TO LEARN HOW AMD EPYC[™] PROCESSORS MIGHT WORK FOR YOU?

Sign up to receive our data center content amd.com/epycsignup



"For the foreseeable

future, EPYC products

are going to enable us

to best utilize the

physical resources that

we have available in

the data center."

Steven Hart, CPU Staff

Engineer, Dropbox

About Dropbox

Dropbox is the one place to keep life organized and keep work moving. With more than 600 million registered users across 180 countries, we're on a mission to design a more enlightened way of working. Dropbox is headquartered in San Francisco, CA, and has offices around the world. For more information on our mission and products, visit <u>dropbox.com</u>.

About AMD

For 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit <u>amd.com/EPYC</u>.

©2021 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Apache Hadoop and Apache Mesos are registered trademarks or trademarks of the Apache Software Foundation in the United States and/or other countries. SPEC[®] and the benchmark SPECrate[®] are registered trademarks of the Standard Performance Evaluation Corporation. Learn more at www.spec.org. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.



