



FOR IMMEDIATE RELEASE

Contact:

Mike LaPan
Cirrascale Cloud Services
(888) 942-3800
mike.lapan@cirrascale.com

**CIRRASCALE CLOUD SERVICES NOW OFFERING AMD EPYC™ 7002 SERIES
PROCESSOR-BASED SERVERS IN ITS DEDICATED, MULTI-GPU DEEP LEARNING CLOUD**

*The 2nd Gen AMD EPYC™ processor is the world's first x86 7nm data center CPU,
bringing breakthrough performance for various deep learning workflows.*

San Diego, CA -- August 8, 2019 -- Cirrascale Cloud Services®, a premier provider of multi-GPU deep learning cloud solutions, today announced it is now offering AMD EPYC™ 7002 Series Processors as part of its dedicated, multi-GPU cloud platform. AMD EPYC 7002 Series Processors are expected to deliver up to 2X the performance-per-socket and 4X peak FLOPS per-socket over the previous generation AMD EPYC processors. With up to 64 cores, the 2nd Gen AMD EPYC Processors are perfectly suited to drive performance and reduce bottlenecks in autonomous vehicle, medical imaging, and other deep learning application workflows.

"In addressing the growing demands of today's deep learning workloads, we're excited about offering our customers the latest generation of AMD EPYC Processors as a solution in our dedicated, multi-GPU cloud platform," said Mike LaPan, vice president, Cirrascale Cloud Services. "The higher core per GPU ratios will help our customers improve overall price/performance for their multi-GPU cloud server deployments."

"We designed the 2nd Gen AMD EPYC Processor Series to give our customers exactly what they said they needed," said Scott Aylor, corporate vice president and general manager, Datacenter Solutions Group, AMD. "With leadership in architecture, performance, and security, the 2nd Gen AMD EPYC Processors are a superior choice and set a new standard for the modern datacenter."

The 2nd Gen AMD EPYC Processors power the most demanding workloads with breakthrough performance providing up to 64 cores, 128 threads, 8 memory channels with up to 4TB of memory per socket, and 128 PCIe® 4.0 lanes coupled with embedded security features to help defend the CPU, applications and data. Cirrascale targets customers that want to maximize GPU utilization by removing network and storage bottlenecks within its cloud and offers its services at a flat weekly and monthly rate to customers with no egress charges for data transfers. Additionally, with Cirrascale Cloud Services, the customer experience isn't virtualized, so they receive dedicated access, allowing them to gain the absolute raw horsepower of their cloud server.

Customers can gain access to the new AMD EPYC configurations by visiting and signing up for the Cirrascale Cloud Services platforms at <http://www.cirrascale.com> or by calling (888) 942-3800.

About Cirrascale Cloud Services

Cirrascale Cloud Services is a premier provider of public and private dedicated, multi-GPU cloud solutions enabling deep learning. The company offers cloud-based solutions for large-scale deep learning operators, service providers, as well as HPC users. To learn more about Cirrascale Cloud Services and its unique dedicated, multi-GPU cloud solutions, please visit <http://www.cirrascale.com> or call (888) 942-3800.

AMD, the AMD Arrow logo, EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. Cirrascale Cloud Services, Cirrascale and the Cirrascale logo are trademarks or registered trademarks of Cirrascale Cloud Services LLC. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

###

¹Projections as of July 3, 2019 for AMD EPYC processors using computer modeling of preproduction parts and SPECrate®2017_int_base internal testing results. Results may vary with production silicon testing. EPYC 7601 results as of June 2019 <http://spec.org/cpu2017/results/res2019q2/cpu2017-20190411-11817.pdf>. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. ROM-23

²Estimated generational increase based upon AMD internal design specifications for "Zen 2" compared to "Zen 1". "Zen 2" has 2X the core density of "Zen 1", and when multiplied by 2X peak FLOPs per core, at the same frequency, results in 4X the FLOPs in throughput. Actual results with production silicon may vary. ROM-04

