NEWS RELEASE

Penguin Computing™ Receives Americas 2017 NVIDIA Partner Network High Performance Computing Partner of the Year Award

3/27/2018

SAN JOSE, CA – GPU Technology Conference – March 27, 2018 – Penguin Computing™, provider of high performance computing, artificial intelligence (AI), enterprise data center and cloud solutions, today announced that it has been recognized with the Americas 2017 NVIDIA Partner Network (NPN) High Performance Computing (HPC) Partner of the Year Award at the 2018 NVIDIA GPU Technology Conference (GTC).

The NVIDIA Partner Network chose Penguin Computing for this honor because of its success in the HPC industry, having grown its NVIDIA GPU business 778% over the past four years. For example, as an early contributor and Open Compute Project (OCP) solution provider, Penguin Computing was recently awarded a contract with the U.S. Department of Energy's National Nuclear Security Administration (NNSA) tri-laboratory program to provide NVIDIA Tesla Volta-based GPU systems.

Today, the company's experience with NVIDIA GPU architectures will be on display as Penguin Computing Senior Director of Advanced Solutions, Kevin Tubbs, Ph.D., speaks at GTC in the session "Best Practices in Designing and Deploying End-to-End HPC and AI Solutions." Understanding the challenges faced by organizations looking to build AI systems, Penguin Computing will provide insights into the design principles and technologies that have proven successful in their AI deployments for customers in the Top500. Penguin Computing helps customers designing and deploying large-scale GPU clusters including datacenter and environmental challenges, network performance and optimization, data pipeline and storage challenges as well as workload orchestration and optimization.

"We are proud to be able to combine Penguin Computing's customized, Linux high-performance computing (HPC) capabilities with the latest Volta GPU architecture from NVIDIA and enable data scientists, researchers and engineers to tackle AI and computing challenges that were once impossible," said Tom Coull, CEO at Penguin

Computing. “This award is a direct result of our partnership with NVIDIA and our commitment to our customers success through advanced GPU-based solutions. We look forward to working with NVIDIA and helping bring even more GPU-based customer projects to fruition in the future.”

“With more than 550 optimized applications, NVIDIA Tesla is the most adopted platform for accelerating HPC and AI,” said Craig Weinstein, Vice President, Americas Partner Organization, NVIDIA. “Penguin Computing has made high performance computing a priority designing GPU solutions into vertical markets ranging from oil and gas and higher education to government/defense and healthcare.”

About Penguin Computing
Penguin Computing, Inc. is a 20-year-old, U.S.-based global provider of high-performance computing (HPC), artificial intelligence (AI), and data center solutions with more than 2,500 customers in 40 countries, across eight major vertical markets. Penguin Computing offers a comprehensive portfolio of hardware, software and services, including solutions based on the Open Compute Project (OCP), as well as financing and top-rated customer support. Penguin Computing products include Linux-based servers; software; integrated, turn-key clusters; enterprise-grade storage and bare metal HPC on cloud via Penguin Computing® On-Demand™ (POD).
Visit **www.penguincomputing.com/** to learn more about the company.

For further information, contact: Rachel Shatz Karbo Communications for Penguin Computing 650-270-1097 penguin@karbocom.com