

vScaler® launches its AI Reference Architecture

Submitted by: vScaler

Monday, 3 February 2020

The development of AI solutions is complex and varies greatly depending on the application. From data capture and ingestion, to preprocessing and algorithm development – the AI lifecycle involves many stages and technologies. The dynamic nature of these AI environments typically means that the datasets are constantly growing, and algorithms are being optimised for greater accuracy. An infrastructure that can handle growth seamlessly and grow as your requirements grow is crucial.

vScaler (https://www.vscaler.com/vscaler-ai-launch/?utm_source=vscaler&utm_medium=pr&utm_campaign=ai) – an optimised cloud platform built with AI and Deep Learning workloads in mind – provides you with a production ready environment with integrated Deep Learning application stacks, RDMA accelerated fabric and optimised NVMe storage, eliminating the administrative burden of setting up these complex AI environments manually.

The NVIDIA® (<http://nvidia.com>) DGX-2™ system (as well as previous generation DGX-1) have been integrated into vScaler as the building block for intensive workloads, providing the power of 16x V100 32GB GPUs. Featuring GPUDirect RDMA thanks to compatible, high performance RDMA fabric from Mellanox® (<http://www.mellanox.com>), the platform provides a significant decrease in GPU-GPU communication latency and completely offloads the CPU.

The AI optimised Architecture includes vScaler Cloud Storage, which centralizes all I/O and places active files on the fastest tier (Flash/NVMe) but also leverages the cost benefits of an erasure coded HDD tier and provides policy-based movement between these tiers— all within a single platform.

David Power, CTO at vScaler adds “With AI becoming more prevalent in almost every industry, it’s clear that a platform that can manage the complete AI development and production life cycle will become more and more critical for organisations wishing to adopt an AI strategy. Whether you are starting out on your AI journey or looking to scale an existing infrastructure, vScaler can help with simplifying the configuration and management of all software and storage in a cost-effective and easy to use environment”.

The vScaler platform features containerised stacks (including the RAPIDS suite of software libraries) which enable users to spin up application specific environments with the appropriate Deep Learning frameworks installed and ready for use, including TensorFlow, Caffe and Theano (and many more) orchestrated on individual VMs or in a scale-out manner using Kubernetes on bare metal. These frameworks are accelerated using the world’s fastest GPUs, purpose-built to dramatically reduce training time for Deep Learning and Machine Learning algorithms and AI simulations.

Organisations looking to gain an edge in business are turning to AI development to build their next generation of products and services. For a scalable and cost-effective alternative to legacy AI solutions, look to vScaler to start your AI cloud journey today.

To download the white paper in full click here

(http://www.vscaler.com/vscaler-ai-reference-architecture/?utm_source=vscaler&utm_medium=pr&utm_campaign=ai)

About vScaler

vScaler is a private Cloud platform built on Open Source technology that enables you to create a secure, scalable, cost-effective, flexible IT infrastructure. Driving out licensing costs, vScaler has been developed to reduce complexity, reduce risk & significantly reduce cost by creating a platform to deliver 'Anything as a Service'. Competitively priced, secure, fully supported and available immediately, vScaler is poised and ready to simplify your infrastructure and support your digital transformation.

www.vscaler.com

vScaler is a private Cloud platform built on Open Source technology that enables you to create a secure, scalable, cost-effective, flexible IT infrastructure.

For further information:

Orla Leland, CMO

vScaler

+44 (0)203 889 0662

orla.power@vscaler.com