


DEEP LEARNING IN THE CLOUD

Deep Learning Frameworks in the cloud
powered by GPUs



Deep learning is redefining what's possible. From early-stage startups, scientific research centers to large web service providers, Deep Learning has become the fundamental building block in delivering revolutionary solutions that define and shape our world.

Today's leading Deep Learning models typically take days or weeks to train, forcing analysts to make compromises between accuracy and time to deployment. In order to speed up training times, GPUs are typically deployed - but can incur complex management challenges as well as increased costs if not utilised efficiently.

vScaler - an optimised cloud platform built with Deep Learning workloads in mind - enables anyone to quickly deploy scalable, production-ready deep learning environments via an optimised private cloud appliance.

Spin up application specific environments with the appropriate Deep Learning frameworks installed and ready for use, including TensorFlow, Caffe and Theano*. 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.

Your vScaler Deep Learning appliance can be hosted on your own premises or in one of our renowned Data Centre facilities and comes with optional modular managed services add-ons.

vScaler provides you with a production ready environment with integrated Deep Learning application stacks and optimised NVMe storage, eliminating the administrative burden of setting up these complex environments manually.

Faster Training & Inference

Accelerated using the world's fastest GPUs, purpose built to dramatically reduce training time.

GPU and vGPU

Options for dedicated GPUs for intensive workloads or vGPU for development environments included.

Scalable

Scale-up or down with the touch of button, under a single customisable management portal.

 **SCALER**

WHAT DO I GET FROM VSCALER?

Self Service Deep Learning Environments

vScaler empowers your end users to set up the environments they need for their work or research. With instant access to resources on-demand, our platform eliminates the need for system administration skills and allows researchers to concentrate on the task at hand.

NVMe Accelerated Storage

Modern GPUs used in AI and ML have an amazing appetite for data - up to 16GB/s per GPU. Starving that appetite with slow storage, or wasting time copying data back and forth is a waste of GPU cycles, which is why vScaler incorporates NVMe accelerated storage to ensure the most efficient use of our GPU resource.

Multiple Hosting Options

Deploy your vScaler cloud either on premise, in our DCs or across both in a hybrid model and scale up or down as your requirements, change through a single management interface.

In-House Deep Learning Expertise

Our team of experts offer hands-on training in the latest AI and accelerated computing methods used to solve real-world problems. Designed for developers, data scientists, and researchers, our training can be delivered online or onsite via instructor-led courses.

HOW CAN I MOVE TO VSCALER?

The switch from your old system can be done in stages with little to no impact on your daily operations. vScaler offers full service on-boarding and service off-boarding assistance, as well as a modular Managed Services options, with no-vendor lock-in.

FRAMEWORKS

Tensorflow

Caffe

Theano

Amber

Gromax

Namd

BrytlytDB

Intelligent Voice

Torch

(*Others available on request)

TECH SPECIFICATIONS

Our Hyper-converged 2U private cloud appliance is designed using class leading technologies, the latest processor family, 2666Mhz DDR4 memory, 25GB interconnects and all SSD storage delivering superior performance with our optimised stack.

APPLIANCE (PER NODE)	
MODEL	VS-H-2824
SERVER COMPUTE	Intel 6138 20C/40T Xeon Scalable Skylake
STORAGE CAPACITY	6x SSD [480 GB, 800 GB, 960 GB, 1.2 TB, 1.6 TB, or 1.9 TB],
MEMORY	384GB DDR 2666MHz
NETWORK CONNECTIONS	2x 1 GbE, 1x 1GbE RJ45 (IPMI)
CONNECTIONS	Dual-Port 25 GbE RDMA Mellanox
GPU	Up to 4x NVIDIA Tesla V100 per appliance

For More information,
Contact vScaler

info@vscaler.com | www.vscaler.com | +44 (0)20 3889 0662