DELL Technologies

Solution brief

Optimized Kubeflow implementation with One Convergence DKube

Quickly build deep-learning platforms for detecting disease in X-rays with GPU clusters on-premises

Customer Results

2 hours vs. 9 months
to run analysis[1]

218% ROI
over 3 years[2]

20 million
images used to train a deep neural network[3]

One of the most exciting areas of interest for deep learning is medical imaging. Radiology is a particularly good fit, since it combines visual data with complicated, hard-to-define outcomes. The neural networks that make up a deep learning model can be trained with massive data sets, eventually superseding human capabilities for identifying disease in X-rays.

Public cloud is often used for deep learning, but it has limitations for radiology because data sets can be quite large and the information is subject to data privacy laws. Many times, the best choice is to move compute resources to where the data resides. This reduces the cost associated with moving large data sets in and out of public cloud and addresses many regulatory requirements for data privacy.

However, while on-premises multi-node cloud-native platforms for deep learning address these concerns, they come with a different set of challenges. Creating on-premises Kubernetes for deep learning applications can be difficult and time-consuming, taking data scientists away from more valuable activities. That's because:

- Pulling together the solution components and integrating them into an IT environment can be complicated.
- Ensuring algorithms function reliably in a production environment is often outside the ability of most organizations.
- Enabling the application to scale effectively is difficult, even for seasoned IT veterans.

Dell Technologies and One Convergence® have joined forces to provide integrated, scalable and cost-effective deep learning systems that combine Dell's leadership family of Dell EMC PowerEdge servers with the One Convergence DKube™ cloud-native machine learning operations (MLOps) solution. By joining forces, Dell and One Convergence simplify on-premises deployment of GPU clusters on-premises for detecting disease in X-rays.

PowerEdge server GPU-optimized portfolio for AI

Dell Technologies provides an extensive portfolio of high performance computing technologies and solutions for deep learning. You can choose from a range of PowerEdge servers optimized with GPUs available as PCIe cards that fit into server PCIe slots, and as SXM2 modules mounted to the server motherboard. With years of experience and an ecosystem of curated technology and service partners, Dell Technologies provides workstations, servers, networking, storage and services that reduce complexity and provide the high performance computing required for data analytics and AI.

[1] Dell EMC Case Study, Caterpillar Autonomous Mining, August 2017.

[2] Forrester Study commissioned by Dell EMC, The Total Economic Impact ofDell EMC Ready Solutions for AI, Machine Learning with Hadoop, August 2018.

[3] Dell EMC Video Case Study, AI startup ZIFF.ai revs up its business with Dell EMC, June 2018.

Learn more

- Dell Technologies and One Convergence Partner to Demonstrate Sample Models to Detect Disease in X-Rays
- DKube User Guide
- kubeflow.org
- hpcatdell.com
- delltechnologies.com/ai
- delltechnologies.com/servers

DKube complements Dell Technologies HPC/AI offerings with a comprehensive deep learning development engine that includes:

Rapid prototyping	• Integrated development and debug • Jupyter® and TensorBoard™ support • Dashboards, metrics and logs	**Repeatability**	• Standard pipeline creation and maintenance • Code and data set versioning • Simple, file-based data set and configuration management
Automation	• Native support for Kubeflow pipelines • Automated experiments and runs • Container-based management of pipeline elements	**Productization**	• Model packaging for deployment • Export to popular inference platforms • Evaluation and validation

Dell Technologies has an extensive PowerEdge server portfolio that can be paired with DKube to build a deep learning solution that meets your specific requirements. These on-premises GPU-optimized solutions bring together qualified software and hardware components for a robust yet flexible AI experience.

Server	PowerEdge T640	**PowerEdge** R740/R7425	PowerEdge 940XA	DSS 8440	PowerEdge C4140
Max GPUs	4x NVIDIA® V100	3x NVIDIA V100 6x NVIDIA T4	4x NVIDIA V100	10x NVIDIA V100	4x NVIDIA V100
Target workloads	VDI, ML/DL training and inference, database/ analytics	VDI, ML/DL training and inference, HPC, database/ analytics	ML/DL training, database/ analytics	ML/DL training and inference	ML/DL training, HPC

One Convergence and Dell Technologies

One Convergence and Dell have joined forces to create an industry leading flexible, scalable and intuitive deep learning solution, focused on radiology. DKube from One Convergence is a cloud-native MLOps solution that simplifies deep learning workflows, and its power can be applied to the most challenging problems in machine learning. DKube is based on open standards and is compatible with TensorFlow, Jupyter, Scikit-Learn, Katib-based hyperparameter tuning, and Kubeflow Pipelines.

With Dell Technologies and One Convergence, demanding data science professionals can be up and running on a Dell production-ready system in a few hours, scale with their requirements, and focus on deep learning rather than on infrastructure.