

Dell EMC HPC Ready Architecture for AI and Data Analytics with DKube

An end-to-end MLOps platform for cost-effective on-premises deep learning

Architecture Guide

This paper provides high-level guidance for building a converged architecture capable of running HPC, AI and data analytics workloads on a single infrastructure with the DKube MLOps platform.

Updated February 2021

[Addendum to H18161, the HPC Ready Architecture for AI and Data Analytics](#)

Table of contents

Introduction	3
DKube MLOps platform overview	5
Dell EMC Ready Architecture for AI and Data Analytics overview	7
Summary	10
Learn more.....	10

Introduction

Solution overview

Converging HPC with AI and data analytics on a single system gives you the horsepower to run high-performance data analytics (HPDA), boost high-frequency trading, enhance risk analysis, improve fraud detection, collect and analyze data from Internet of Things (IoT), and speed motion picture animation and special effects cycles — just to name a few.

However, to optimize budgets along with performance, it's critical to match advanced computing resources to business requirements. System configuration can be a complex task, requiring a balance between workload requirements, performance targets, data center constraints and pricing. Many enterprises don't have the time to research, optimize and deploy advanced computing systems to deliver the required return on investment.

With the Dell EMC HPC Ready Architecture for AI and Data Analytics, Dell Technologies' engineers have done the heavy lifting, so you can quickly deploy a high performance computing solution that matches the needs of the business.

Now, with the One Convergence® DKube™ machine learning operations (MLOps) software platform, the solution can be optimized for deep learning. With DKube, engineers and data scientists can develop and optimize proofs of concept, train them on large data sets and monitor their performance from the unified Bright Computing® portal.

Because of the unique capabilities of the Dell EMC HPC Ready Architecture for AI and Data Analytics, DKube deep learning workloads can run alongside other data intensive workloads on the same system. Plus, system resources can be allocated flexibly, and easily moved between groups based on utilization. This flexibility allows rebalancing the system over time to meet the evolving needs of the organization.

Together, Dell Technologies and One Convergence provide the performance, availability and control of an on-premises solution for deep learning with cloud-like simplicity and agility. Because of the unique capabilities of the Dell EMC HPC Ready Architecture, DKube deep learning workloads can be run alongside other data-intensive workloads on the same system. Plus, system resources can be allocated flexibly, and easily moved between groups based on utilization. This flexibility allows rebalancing the system over time to meet the evolving needs of the organization.

Together, Dell Technologies and One Convergence provide the performance, availability and control of an on-premises solution for deep learning with cloud-like simplicity and agility.

Read the [Dell EMC HPC Ready Architecture for AI and Data Analytics Architecture and Deployment Guidance](#) white paper for more details on the underlying architecture.

Document purpose

This document is intended to provide high-level guidance on architecture design for a converged architecture. The document provides the following information:

- DKube MLOps platform overview
- Dell EMC Ready Architecture for AI and Data Analytics overview

Audience

This document is intended for solution architects and IT operations personnel who want to deploy a single environment to run multiple workloads for HPC, AI, and data analytics, including DKube.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Author: Ramesh Radhakrishnan, Luke Wilson, John Lockman

Note: For links to additional documentation for this solution, see the [Dell Technologies Solutions Info Hub](#).

DKube MLOps platform overview

The MLOps workflow

Deep learning consists of more than simply running model code through a notebook and identifying a few hyperparameters to optimize the output. Deep learning development must incorporate a full MLOps workflow, with the goal of deploying and maintaining an accurate model on a production server.

The following figure shows the full MLOps flow:

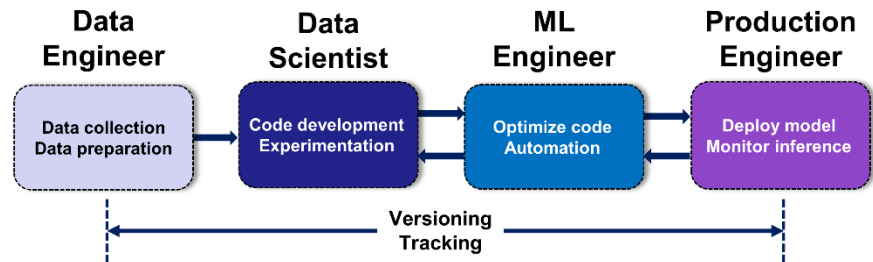


Figure 1. MLOps workflow

MLOps success factors

Any MLOps project should be thoughtfully planned with the following requirements in mind.

A secure, flexible, enterprise MLOps workflow:

- The process must be able to accommodate anything from a small team to a large organization that spans many people.
- The handoff between functions must be clear and formal; this is especially important for large organizations.
- Workflow governance, privacy and security are critical. Data must be tracked as it makes its way through the process and be easily reproducible for auditing.
- Models must provide complete insight into their provenance. The code, datasets, hyperparameters, and platform hardware/software details must be easily determined, and the metadata must follow the model through its life cycle.
- Since model development is iterative, versioning capability must integrate with the overall solution.

Focus on deep learning:

- The model must be deployable on a variety of frameworks. Different requirements for performance, cost, load balancing, and proximity mean that the serving capability must be flexible.
- The application must be easily installed on the target platform, and the workflow must be intuitive.
- The hardware resources must be automatically identified and configured for use. Allocation of resources must be flexible, such that they are securely shared to maximize utilization.

DKube: an end-to-end MLOps platform

The DKube deep learning application provides an end-to-end MLOps platform. Data engineers, data scientists, machine learning engineers and production engineers can coordinate to seamlessly deliver powerful models.

Go from bare metal to MLOps in hours

One Convergence handles the difficult task of integrating the required standards, enhancing the flow and capability with added value software components, enabling them to work together in a GUI-based production environment.

Move between on-premises and the cloud

Although an on-premises platform is most appropriate for large datasets and compute-intensive workloads, there may be situations when a hybrid cloud environment makes sense. DKube allows seamless migration between on-premises systems and the cloud. The workflow, look and feel are identical across platforms.

Track versioning and lineage

Complex models require substantial experimentation and iteration to get them right. To keep all that activity manageable, DKube provides full versioning. In addition, a run or a model can be traced back to show all the inputs used to create it. Your code and dataset can be traced forward to understand where it has been used.

Compare models to find the best fit

DKube allows you to compare model metrics in order to identify which combination of inputs provides the best outcome.

Test and deploy models for production serving

Once the optimal model has been identified, its inference capability can be tested locally, then deployed for production serving.

Leverage multi-tenancy and collaboration

Resources — both technical and human — are expensive. Resources and data on the system can be allocated and shared among groups of users.

Use popular standards

Based on Kubeflow & Kubernetes, DKube supports the most popular frameworks, including JupyterLab®, RStudio®, TensorFlow™, Keras, PyTorch® and scikit-learn. Authentication and code management are supported for GitHub® LDAP, and Bitbucket. Automation is supported through Kubeflow Pipelines, Katib, and GitOps.

Scale with your workload

Different workloads have different resource requirements. Model development and experimentation might need only CPUs on a workstation, while training will make use of powerful GPUs on scalable clusters. DKube automatically identifies resources and makes them available for use.

Enjoy multi-node resilience

DKube supports multi-node resilience to address critical, high availability environments.

Choose the package that works for you

Organizations have differing needs. Individuals and smaller groups can choose the Lite package that offers an affordable mix of popular features. Larger organizations that require a full MLOps capability can move up to the Enterprise package.

Dell EMC Ready Architecture for AI and Data Analytics overview

The scale of HPC meets the flexibility of containers

Engineering teams from Dell Technologies and One Convergence have worked together to tune the Dell EMC Ready Architecture for AI and Data Analytics with the DKube deep learning MLOps application. The solution can be tailored for a wide variety of HPC, AI and data analytics workloads.

The architecture's resource pool can be dynamically assigned and managed by an HPC resource manager. Or resources can be assigned to containerized AI/data analytics workloads, which are orchestrated by the open-source Kubernetes® container orchestration system. This flexibility allows IT teams to change the balance of the system to meet the evolving computational needs of users.

Compute building blocks

The Ready Architecture leverages different servers for different roles. The architecture is defined by node types: management, compute, high-performance acceleration and dense acceleration.

Management and compute nodes

[Dell EMC PowerEdge R740/xd Servers](#) are recommended for the management and compute nodes. The R740/xd can be configured as a compute-only platform, or as an accelerated platform which supports NVIDIA® A100, V100 or T4 Tensor Core GPUs, AMD® Instinct™ Mi100 GPUs or Intel® Performance Accelerator Cards (PACs) with Intel Arria® 10GX Field Programmable Gate Arrays (FPGAs).

[Dell EMC PowerEdge R7525 Servers](#) have more AMD processing cores and faster data transfer speeds with PCIe Gen 4. Each two-socket server has up to 2TB memory for scale-out environments and up to 24x direct-connect NVMe. It also accommodates up to three double-wide or six single-wide accelerators.

[Dell EMC PowerEdge C6420 Servers](#) can serve as dense management and compute nodes with up to four dual-socket servers and high capacity storage in a 2U rack. Maximizing density, scalability and energy efficiency per U, the PowerEdge C6420 comes with multiple backplane and storage combinations.

[Dell EMC PowerEdge C6525 Servers](#) are compute-dense servers designed to boost data center performance with more processing cores and fast data transfer speeds with PCIe Gen 4. With up to four dual-socket servers and up to 144TB of storage, this server maximizes density, scalability and energy efficiency.

High-performance acceleration node

The [Dell EMC PowerEdge C4140 Server](#) provides a high-performance GPU building block for HPC simulation and deep learning training. This two-processor server supports up to four NVIDIA V100 GPUs with NVLink™ in just 1U. It features a patented interleaved GPU design to optimize both space and airflow for maximum compute performance.

Dense acceleration node

The [Dell EMC DSS 8440 machine learning server](#) supports up to 16 accelerators, a high-performance switched PCIe fabric for rapid I/O, and up to 10 local NVMe and SATA drives for optimized data access. The DSS 8440 has an open architecture, based on industry-standard PCIe fabric, allowing for customization of accelerators, storage options and network cards.

Network fabric building blocks

The Dell EMC HPC Ready Architecture for AI and Data Analytics can have one or more network fabrics, ranging from 25 Gigabit Ethernet to 100 Gigabit Mellanox® InfiniBand® low-latency fabric, which is tailored specifically for HPC environments.

Storage and file system building blocks

Multiple storage and file system choices are available to best address your workload needs.

[Dell EMC Isilon F800](#) is an all-flash system that provides a high-performance Network File System (NFS) storage building block that is well suited for large-scale HPC simulation workloads, deep learning training, machine learning training, and inferencing. Isilon F800 all-flash storage is powered by the Dell EMC OneFS operating system. The Isilon F800 provides a powerful yet simple scale-out NFS storage architecture to speed access to massive amounts of unstructured data while dramatically reducing cost and complexity.

The [Dell EMC Ready Solution for HPC BeeGFS® Storage](#) transparently spread user data across multiple servers. File system performance and capacity can be scaled out to the desired level by simply increasing the number of servers and disks in the system. You can scale seamlessly from small clusters up to enterprise-class systems with thousands of nodes — on-premises or in the cloud.

[Dell EMC Ready Solutions for HPC PixStor™ Storage](#) give you high performance, highly scalable, enterprise-class software-defined storage that empowers you to search, manage, securely isolate and protect your data, collaborate and share across distances, and run in the cloud. The Dell EMC engineering-validated design for PixStor storage delivers high performance with limitless scale at lower cost than traditional legacy solutions.

System management

Using Bright Cluster Manager® from Bright Computing, the Dell EMC Ready Architecture for AI and Data Analytics with DKube brings all these capabilities together and lets you manage them from a single pane of glass. Bright View integrates directly with the Dell EMC iDRAC management system, allowing changes to BIOS settings and system firmware from the same console.

Resource balancing

Workload needs change over time. You can avoid idle resources and excessive wait times by moving compute resources from HPC scheduling to Kubernetes and back again as needed. Plus, you can independently scale your shared storage to create the right balance of compute and storage for your needs.

Summary

The Dell EMC HPC Ready Architecture for AI and Data Analytics with DKube is a converged infrastructure solution. With this solution, you can design a unified architecture with multipurpose, balanced nodes to support all your HPC, AI and data analytics workloads.

The solution employs a building block approach—supporting multiple servers, networking, storage and software environment options—so you can design the architecture that best meets your needs.

If you require assistance with your project, Dell Technologies has a team of AI experts dedicated to staying on the cutting edge, to help you keep pace with this constantly evolving landscape. You can collaborate with Dell Technologies engineering teams at one of the worldwide [Customer Solution Centers](#), tap into the resources of one of our nine [HPC & AI Centers of Excellence](#), or test and tune systems at the [HPC & AI Innovation Lab](#).

From design and implementation to support and systems management, Dell Technologies offers a comprehensive services portfolio for AI, including [consulting](#), [education](#), [deployment](#) and [support](#). And a wealth of leasing and financing options from [Dell Technologies Financial Services](#) can help you find opportunities to obtain new IT while optimizing cash flow.

Learn more

- delltechnologies.com/referencearchitectures
- delltechnologies.com/hpc
- delltechnologies.com/ai

The information in this publication is provided as is. Dell Technologies makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020-1 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.

Published in the USA 02/21 White paper DELL-WP-HPC-RA for-AI-DA-DKube-USLET-101

Intel® and Xeon® are registered trademarks of Intel Corporation in the U.S. and other countries. NVIDIA® and NVLink™ are trademarks or registered trademarks of NVIDIA Corporation. Kubernetes® is a registered trademark of The Linux Foundation. PixStor™ is a trademark of Arcapix Holdings. Mellanox® and InfiniBand® are registered trademarks of Mellanox Technologies, Ltd. BeeGFS® is a registered trademark of Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung E.V. Bright Computing® is a trademark of Bright Computing, Inc. One Convergence® and DKube™ are trademarks or registered trademarks of One Convergence, Inc. TensorFlow™ is a trademark of Google Inc. The Jupyter trademark is registered with the U.S. Patent & Trademark Office. RStudio® is a registered trademark of RStudio, Inc. GitHub® is an exclusive trademark registered in the United States by GitHub, Inc.

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice