

Dell EMC Ready Solution for AI

Deep Learning with Intel

June 2019

H17621

Design Guide

Abstract

This design guide describes the high-level architecture of the Ready Solution for AI, Deep Learning with Intel. This solution for deep learning training and inference uses the Kubernetes-based Nauta deep learning training platform to perform distributed, multinode model training.

Dell EMC Solutions

Copyright

The information in this publication is provided as is. Dell Inc. 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 © 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 06/19 Design Guide H17621.

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



Contents

- Executive summary.....4
- Architecture overview.....5
- Design considerations.....8
- Validation and use cases.....9
- Conclusion.....13
- References.....14

Executive summary

Overview

The Dell EMC Ready Solution for AI, Deep Learning with Intel is a scalable, flexible container-based solution for the development and testing of deep learning and machine learning models using Intel Xeon Scalable processors. The solution uses a combination of a high compute-density rackmount server chassis, Ethernet networking, expandable hybrid disk/flash storage, optimized libraries, and deep learning frameworks to enable customers to train their own artificial intelligence models quickly.

Audience

This guide is for solution architects, system administrators, and others who are interested in accelerating deep learning with advanced computing and data management solutions.

We value your feedback

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

Authors: Lucas A. Wilson, Ph.D. and Srinivas Varadharajan

Note: The [AI Info Hub for Ready Solutions](#) on the Dell EMC Communities website provides links to additional documentation for Ready Solutions for AI.

Architecture overview

Physical architecture

This Ready Solution consists of one master node, a Dell EMC PowerEdge R740xd 2U rack server, and 16 compute nodes that are housed in four PowerEdge C6000 2U compute chassis. The following figure shows this Ready Solution with an optional Dell EMC Isilon H600 storage array.



Figure 1. Ready Solution with optional Isilon storage

A Dell EMC Networking S3048-ON 1U rack-mounted 1 GbE top-of-rack (ToR) switch acts as the system's iDRAC and management network. The switch also provides a connection to the Internet and external components. A Dell EMC Networking S4128F-ON rack-mounted 10 GbE SFP+ switch provides higher bandwidth connectivity between the compute nodes for distributed training and data movement.

The optional Isilon H600 NAS system is attached through the S4128F-ON switch with four 40 GbE QSFP+ connections.

The following figure shows the Ready Solution physical architecture.

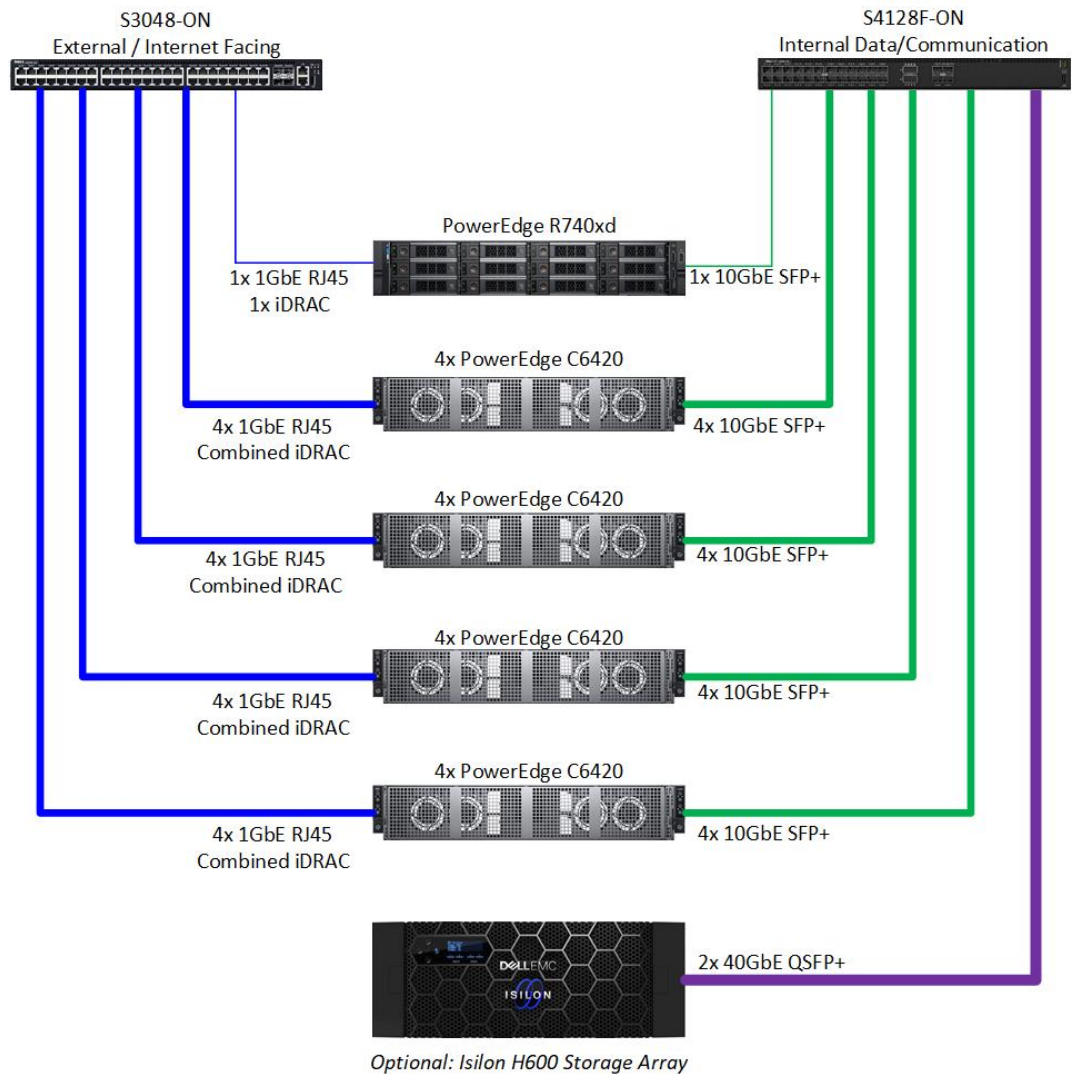


Figure 2. Ready Solution architecture

Logical architecture

This Ready Solution is a container-native distributed platform that uses Kubernetes for container orchestration and resource management. The R740xd 2U rack server acts as a:

- Master node that handles all Kubernetes scheduling and orchestration tasks
- Network File System (NFS) host that serves a shared file system to the rest of the servers
- User login node that provides users access to the system and to the Nauta data science platform, which uses Kubernetes

Each of the 16 PowerEdge C6420 compute servers acts as a worker node in the Kubernetes environment. Each worker accepts pod scheduling requests from the Kubernetes master and runs user and system jobs that are encapsulated in Docker containers.

Hardware and software components

The following tables list the hardware components and R740xd RAID configuration for this Ready Solution.

Table 1. PowerEdge R740xd hardware components

Component	Quantity	Description
Processor	2	Intel Xeon Scalable Gold 6230
Memory	12	32 GB 2933 MHz DDR4
Storage	12	12 TB 7.5 K RPM HDD
Network adapter	1	Intel X710 10 Gb SFP+ Ethernet adapter

Table 2. RAID configuration for the PowerEdge R740xd master node

Virtual disk	Configuration	RAID	Mounted as
VD0	2 x 12 TB HDD (11 TB usable)	RAID 1	/(root)
VD1	10 x 12 TB HDD (97 TB usable)	RAID 5	/data (NFS and ETCD)

Table 3. PowerEdge C6420 hardware components

Component	Quantity	Description
Processor	2	Intel Xeon Scalable Gold 6230
Memory	12	16 GB 2933 MHz DDR4
Storage	1	250 Gb M.2 boot optimized SSD storage
Network adapter	1	Intel X710 10 Gb SFP+ Ethernet adapter

The following table lists the software components for this Ready Solution.

Table 4. Software components

Component	Description
BIOS	Dell BIOS 2.1.8
Operating system	Red Hat Enterprise Linux (RHEL) 7.6
Docker	Community Edition 18.06
Kubernetes	1.10.11
Helm	2.9.1
Nauta	1.0 Enterprise Support
Remote Desktop Protocol (XRDP)	0.9.9-1

Design considerations

This Ready Solution is a flexible, scalable solution for developing deep learning-based models by using a distributed learning approach. It takes advantage of Intel Xeon Scalable processors to run wide-width Single Instruction, Multiple Data (SIMD) operations in various numerical formats (single-precision FP32, double-precision FP64, 8-bit integer, and 16-bit integer). This Ready Solution uses optimized linear algebra and neural network libraries and primitives to improve neural network training performance. Because it is based on Intel Xeon Scalable processors, the solution provides greater memory footprints for training large neural network architectures and large format datasets that might not be suitable for limited memory-footprint accelerators.

This Ready Solution uses 10 GbE connectivity between nodes for both data motion and model training. The 10 GbE connections simplify solution scaling and enable easy integration into an organization's computing and networking infrastructure.

Optional scale-out Isilon hybrid storage provides customers with a scalable storage solution if the on-node storage—approximately 97 TB, as shown in Table 2—is insufficient.

Validation and use cases

Dell EMC has validated the functionality of the Nauta deep learning platform on this Ready Solution, including the following capabilities:

- Starting a Jupyter notebook
- Launching the Nauta web user interface and the Kubernetes Dashboard
- Submitting a TensorFlow training experiment
- Submitting a distributed TensorFlow training experiment
- Submitting a distributed TensorFlow + Horovod parallelized training experiment
- Submitting a parameter sweep
- Canceling an experiment
- Submitting a batch prediction experiment
- Submitting a streaming prediction experiment
- Launching a TensorBoard session

Dell EMC has also tested multiple use cases on this Ready Solution. Each use case examines a different neural network topology and ensures that the solution can train models for many different domains, such as image classification, language translation, and retail product recommendation.

For each of the three use cases that we used for validation, we performed scale-out tests on PowerEdge C6420 compute nodes—from 1 node up to 16 nodes. We submitted all tests through the Nauta deep learning platform. We ran the tests by using Docker containers that were orchestrated by Kubernetes. Each validation test ran a single compute pod per node, each containing four Docker container instances (four processes per node).

**Use case 1:
Image
classification**

For the image classification use case, we used frontal chest x-ray data from the [US National Institutes of Health \(NIH\) ChestXray14](#) dataset. We validated runs by using a ResNet-50 topology that is pretrained on the [ImageNet Large Scale Visual Recognition Challenge 2012 \(ILSVRC2012\)](#) dataset. We performed scale-out testing by using the TensorFlow and Horovod libraries on various quantities of PowerEdge C6420 compute nodes, from 1 to 16, as shown in the following figure.




Figure 3. Image classification use-case chart

Use case 2: Language translation

For the language translation use case, we used Google's [Tensor2Tensor](#) library's "Big Transformer" neural network architecture to train an English to German translation model by using the [WMT16 English to German](#) dataset. We performed scale-out testing by using TensorFlow and Horovod libraries on various quantities of PowerEdge C6420 compute nodes, from 1 to 16, as shown in the following figure.

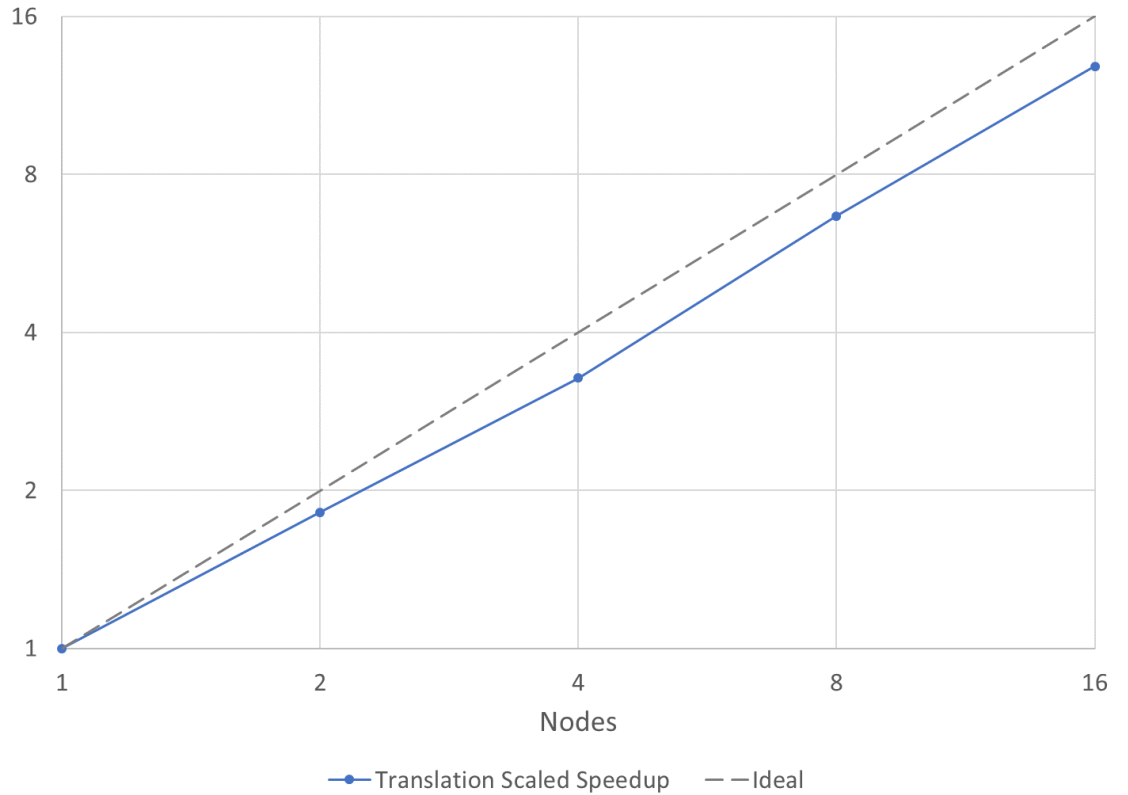


Figure 4. Language translation use-case chart

**Use case 3:
Product
recommender**

The recommendation engine use case trains a Restricted Boltzmann Machine (RBM) for ratings prediction by using the [MovieLens 20M](#) stable dataset. We used TensorFlow and Horovod libraries to train the RBM. We performed scale-out testing on various quantities of PowerEdge C6420 compute nodes, from 1 up to 16, as shown in the following figure.




Figure 5. Product recommender use-case chart

Conclusion

Summary

The combination of optimized libraries and deep learning frameworks with the processing power of Dell EMC PowerEdge servers with Second Generation Intel Xeon Scalable processors with Intel Deep Learning Boost enables customers to develop, test, and deploy artificial intelligence models. The Dell EMC Ready Solution for AI, Deep Learning with Intel has a scalable design that improves neural network training performance and supports greater memory footprints. This Ready Solution simplifies solution scaling and integration and enables quick model training for varied domains.

We validated the Dell EMC Ready Solution for AI, Deep Learning with Intel for deep learning training and inference tasks through use cases addressing the broad categories of image classification, language translation, and product recommendation. In each of the use cases, the solution demonstrated near-linear scaling across all 16 compute nodes. The solution can continue to scale efficiently to greater node counts to accommodate system expansion.

References

The following documentation provides additional and relevant information:

- [ChestXray14 Dataset](#)
- [ImageNet 2012 Dataset](#)
- [Google Tensor2Tensor Library](#)
- [WMT16 Language Translation Datasets](#)
- [GroupLens Public MovieLens Dataset](#)