

Modern Apps

FlexPod

NetApp March 25, 2024

This PDF was generated from https://docs.netapp.com/us-en/flexpod/modern-apps/fpd-ai-dl-design.html on March 25, 2024. Always check docs.netapp.com for the latest.

Table of Contents

odern Apps
FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Design
Deploy NetApp Trident CSI plug-in on Cisco Container Platform with FlexPod
FlexPod Datacenter for OpenShift Container Platform 4 - Deployment
FlexPod Datacenter with Docker Enterprise Edition for Container Management
FlexPod Datacenter for OpenShift Container Platform 4 - Design
FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Deployment 3
3D graphics visualization with VMware and NVIDIA on Cisco UCS - White paper
3D graphics visualization with Citrix and NVIDIA - White paper

Modern Apps

FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Design

Haseeb Niazi, Cisco Arvind Ramakrishnan, NetApp

This document provides design details around the integration of the Cisco UCS C480 ML M5 platform into the FlexPod Datacenter solution to deliver a unified approach for providing AI and ML capabilities within the converged infrastructure. By providing customers the ability to manage the servers with combined AI and ML capabilities with the familiar tools they use to administer traditional FlexPod systems, the administrative overhead as well as the cost of deploying deep learning platform is greatly reduced. The design presented in this CVD also includes other Cisco UCS platforms such as the C220 M5 server with two NVIDIA T4 GPUs and the C240 M5 server equipped with two NVIDIA V100 32GB PCIe cards as additional options for handling concurrent AI and ML workloads.

FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for Deep Learning - Design

Deploy NetApp Trident CSI plug-in on Cisco Container Platform with FlexPod

This document provides step-by-step procedures for deploying the NetApp Trident Container Storage Interface (CSI) plug-in on a Cisco Container Platform Kubernetes tenant cluster in a FlexPod solution.

Deploy NetApp Trident CSI plug-in on Cisco Container Platform with FlexPod

FlexPod Datacenter for OpenShift Container Platform 4 - Deployment

Haseeb Niazi, Cisco Alan Cowles, NetApp

Red Hat OpenShift is an enterprise-ready Kubernetes container platform to manage hybrid cloud and multi-cloud deployments. Red Hat OpenShift Container Platform includes everything needed for hybrid cloud, enterprise container, and Kubernetes development and deployments. It includes an enterprise-grade Linux operating system, container runtime, networking, monitoring, container registry, authentication, and authorization solutions.

Combining Red Hat OpenShift with the FlexPod Datacenter solution can simplify the deployment and the management of the container infrastructure. Customers can benefit from improved efficiency, better data protection, lower risk, and the flexibility to scale this highly available enterprise-grade infrastructure stack to

accommodate new business requirements. The pre-validated converged solution approach helps organizations achieve the speed, flexibility, and scale required for all of their application modernization and digital transformation initiatives.

FlexPod Datacenter for OpenShift Container Platform 4 - Deployment

FlexPod Datacenter with Docker Enterprise Edition for Container Management

Muhammad Afzal, Cisco John George, Cisco Amit Borulkar, NetApp Uday Shetty, Docker

Docker is the world's leading software container platform for developers and IT operations to build, ship, and run distributed applications anywhere. With microservices architecture shaping the next generation of IT, enterprises with large investments in monolithic applications are finding ways to adopt Docker as a strategy for modernizing their application architectures and keeping the organization competitive and cost effective. Containerization provides the agility, control, and portability that developers and IT operations require to build and deploy applications across any infrastructure. The Docker platform allows distributed applications to be easily composed into a lightweight application container that can change dynamically yet non-disruptively. This capability makes the applications portable across development, test, and production environments running on physical or virtual machines locally, in data centres, and across the networks of different cloud service providers.

FlexPod Datacenter with Docker Enterprise Edition for Container Management

FlexPod Datacenter for OpenShift Container Platform 4 - Design

Haseeb Niazi, Cisco Alan Cowles, NetApp

Cisco and NetApp have partnered to deliver a series of FlexPod solutions that enable strategic data center platforms. The FlexPod solution delivers an integrated architecture that incorporates best practices for computing, storage, and network design, thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance, and support that can be used in various stages (planning, designing and implementation) of a deployment.

FlexPod Datacenter for OpenShift Container Platform 4 - Design

FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Deployment

Haseeb Niazi, Cisco Arvind Ramakrishnan, NetApp

This document provides deployment details and guidance around the integration of the Cisco UCS C480 ML M5 platform into the FlexPod data center solution to deliver a unified approach for providing AI and ML capabilities within the converged infrastructure. This document also explains the NVIDIA GPUs configuration on Cisco UCS C220 and C240 platforms. For a detailed design discussion about the platforms and technologies used in this solution, refer to the FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning design.

FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Deployment

3D graphics visualization with VMware and NVIDIA on Cisco UCS - White paper

This document describes the performance of the VMware ESXi hypervisor and VMware Horizon with NVIDIA Tesla P4, P6, and P40 solution on Cisco UCS C240 M5 Rack Servers and B200 M5 Blade Servers.

3D graphics visualization with VMware and NVIDIA on Cisco UCS - White paper

3D graphics visualization with Citrix and NVIDIA - White paper

This document describes the performance of Citrix XenDesktop on Citrix XenServer with NVIDIA Tesla P4, P6, and P40 cards on Cisco UCS C240 M5 and B200 M5 servers with SPECviewperf 13.

3D graphics visualization with Citrix and NVIDIA - White paper

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.