Hewlett Packard Enterprise

HPE Reference Configuration for HPE Container Platform on HPE Apollo Gen10 Servers

A turnkey hybrid cloud solution for accelerating AI (ML, DL) and big data analytics at scale

Contents

Executive summary ... 3
Introduction ... 3
Solution components ... 4
HPE Container Platform ... 4
HPE Data Fabric ... 6
HPE Container Platform for Analytics ... 7
HPE ML Ops ... 9
HPE Apollo Servers for Analytics ... 11
HPE Container Platform on HPE Apollo Gen10 servers best practices ... 12
HPE Apollo servers configuration for HPE Container Platform ... 12
HPE ProLiant Servers and NVIDIA GPUs ... 13
HPE Container Platform CPU and memory ... 13
HPE Container Platform storage ... 14
HPE Container Platform networking ... 15
Summary ... 15
Version history ... 16
Resources and additional links ... 17

Executive summary

Enterprise organizations everywhere and across all industries are embarking on a hybrid cloud journey. Their key goals are speed, agility, simplicity, consistency, and cost effectiveness - to support digital transformation, business innovation, and accelerated growth.

However, current IT practices and various incompatible application deployment environments have created challenges for organizations to achieve their objectives. To name a few of those challenges:

- Difficult management of workloads across hybrid cloud environments
- Portability limitations and complexity due to bare metal or VM based deployments
- Lack of standards across different developer groups
- Slow to provision new environments. This significantly stifles innovation as teams must wait for environments to be setup and provisioned
- Vendor lock-in is a real concern and pertains to the use of proprietary technology, even the cloud adopters face the same issue
- Siloed infrastructure increases overhead costs including administrative overhead and the cost of additional infrastructure
- Landscape inconsistencies with different tool versions and packages across different deployments making it hard to reassign workloads to different infrastructure
- Release cycle is too time consuming for applications due to the monolithic/rigid nature of deployments
- Difficult and inefficient to scale infrastructure to fit application changes

In order to unleash business opportunities through digital transformation, enterprises need to overcome these restrictions with the new concepts and solutions of the next generation of IT practices.

Hewlett Packard Enterprise recognizes the next wave of technology innovation will be built on a more agile application development platform using modern cloud-native architectures, with portability across any infrastructure from edge to core to cloud. Containers and Kubernetes are central to this new approach, providing key enabling technology to ensure the organizations success in their data analytics journey across the hybrid cloud.

The HPE Container Platform leverages the unique functionality and data analytics experience from BlueData and MapR, together with 100% open source Kubernetes for container orchestration. The HPE Container Platform is designed to run both microservices-architected cloud-native and monolithic non-cloud-native applications with persistent data for hybrid cloud deployments spanning on-premises, public cloud, and edge environments.

The HPE Container Platform provides an enterprise-class solution to deploy and manage containerized environments at scale with Kubernetes. Developers have secured on-demand access to their environments, so that they can develop apps and release code faster, with the portability of containers to build once and deploy anywhere. IT teams can manage multiple Kubernetes clusters with multi-tenant container isolation and data access, for any workload from edge to core to cloud. Also, they can extend the benefits of containers beyond cloud-native microservices-architected stateless applications, providing the ability to containerize monolithic stateful applications with persistent data.

Target audience: This Reference Configuration is for IT decision makers and architects who have a directive to invest in the cloud, modernize their analytics environment, and to reduce their data center footprint. It is also intended to assist platform engineers and cloud architects in defining and implementing a container and/or hybrid cloud strategy. This information is intended for DevOps managers, infrastructure engineers, and LOB managers to improve their operational efficiencies in either containers, hybrid cloud, or in accelerating the time to value for developing and deploying analytics.

Document purpose: This Reference Configuration provides an overview of the deployment of the HPE Container Platform on the HPE Apollo Servers for Analytics.

This solution testing performed in January 2020.

Introduction

It is imperative for enterprises to accelerate their digital transformation strategy in order to become competitive and innovative in their organization, however, there are many headwinds due to traditional IT practices and different application deployment environments.

New cloud-native applications are developed from the ground up using containers and a microservices-architecture, leveraging a DevOps and continuous integration/continuous delivery (CI/CD) approach. However, enterprises also have a wide variety of existing legacy applications that are not cloud-native. Many of these legacy analytics applications have a traditional monolithic three-tier, client-server software architecture. Organizations today also want to modernize these existing monolithic enterprise applications by either rewriting them in a modern programming language, re-architecting them as cloud-native applications, or porting them over to a modern computing platform. The goal is to reduce the cost of operations and bring agility and modern DevOps efficiencies to these legacy applications. But these modernization efforts are time-consuming and expensive.

Running these stateful analytic applications in containers provides an easier way to bring the modernization benefits of agility and efficiency, without the need to go through costly refactoring of applications. However, most containers today run on virtualized infrastructure due to concerns around the security, isolation, and manageability, and live migration of containers. Most of this is due to the ubiquity of VMs because it is the standard way to provision infrastructure since the early 2000s.

Hewlett Packard Enterprise is fundamentally changing that equation by delivering a container platform for a broader range of application architectures, with portability across on-premises, public cloud, and edge environments. HPE Container Platform also eliminates the complexity and expense of a virtual machine layer - with bare metal containerization which eliminates the VM "performance tax" while ensuring enterprise-grade security, bare metal performance, scalability, and reliability.

According to Gartner, adoption of containers and Kubernetes is increasing for cloud-native applications with 75% of global organizations expected to be running containerized applications by 2022[1]. At the same time, containerized workloads spanning hybrid environments is also set to increase over the next few years. As containers use continues to grow rapidly, vendors are investing heavily in bringing this to the enterprise market as customers are preparing to invest in containers as part of their digital transformation initiatives[2].

The HPE Container Platform combines BlueData, the MapR File System, and 100% open source Kubernetes with HPE Apollo Gen10 servers in a combined solution that delivers a turn-key solution platform, which provides high performance and flexible architecture to rapidly deploy containers supporting new application frameworks. Ultimately, this results in faster digital transformation for the business. With services from HPE Pointnext and HPE GreenLake, you can decide whether to purchase hardware upfront, or move to a pay-as-you-go consumption model.

Solution components

This solution is built by deploying the HPE Container Platform on HPE Apollo Gen10 servers, which provides high density compute and data store and elasticity required in large enterprises. The major Hewlett Packard Enterprise products and technologies used in this solution are:

- HPE Container Platform
- HPE ML Ops
- HPE Apollo Gen10 servers

These components are briefed in the following sections.

The solution of HPE Container Platform on HPE Apollo servers for analytics is being delivered as a set of Reference Configurations and Reference Architectures highlighting the key features of the HPE Container Platform.

HPE Container Platform

The HPE Container Platform is a unified container platform built and designed for both cloud-native applications and stateful analytics applications running on any infrastructure either on-premises, in multiple public clouds, in a hybrid model, or at the edge.

[1] 6 Best Practices for Creating a Container Platform Strategy," Gartner, Inc., 20192 "Worldwide Container Infrastructure Software Market Shares"
[2] 2017: Containers Poised for Growth," Doc # US43408418, IDC, Dec 2018

Figure 1 shows the high level HPE Container Platform architecture.

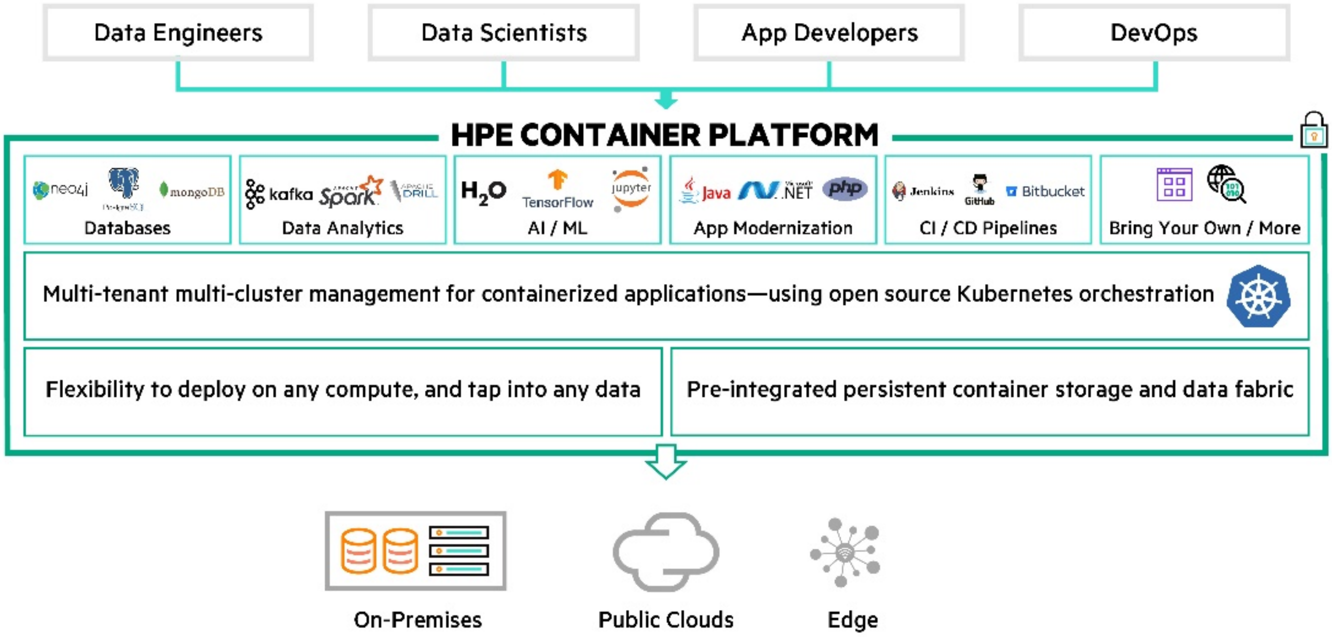

Figure 1. HPE Container Platform architecture

Key features

- **Multi-Tenant management:** Fast, easy deployment, management, and monitoring of multiple clusters with out-of-the-box configuration of networking, load balancing, and storage
- **Hybrid deployments:** Ability to deploy on any infrastructure- on-premises, in multiple public clouds, or at the edge
- **Pre-integrated with persistent container storage:** Fully managed, pre-integrated, scale-out, edge-ready persistent storage with the HPE Data Fabric plus DataTap and FSMount provide connectivity to external data without copying data locally
- **100% open-source CNCF Kubernetes:** With innovations such as KubeDirector, an open-source Kubernetes-based controller to deploy non-cloud-native, stateful apps. The HPE Container Platform is a CNCF certified Kubernetes distribution.
- **Enterprise-grade security and control:** Integrations into enterprise security and authentication services with support for high availability, fault tolerance, and resiliency for mission-critical enterprise applications

Key benefits

- Greater flexibility: A unified platform for orchestration of cloud-native and non-cloud-native applications on-premises, in any public cloud, and at the edge
- Boost productivity: Delivers a self-service experience through a curated App Store of prebuilt application images. Streamlined deployment and management for a wide range of analytics use cases including application modernization, AI/ML, big data analytics, IoT, and CI/CD
- Reduced risk: Enterprise-class security with integrations into enterprise security and authentication services. In-place access to enterprise data sources without creating data copies
- Reduced cost: Lower total cost of ownership with reduced admin overhead and elimination of virtualization tax with bare-metal containerization
- Improved ROI: Improves utilization of hardware resources and provides a cloud-like experience for non-cloud-native monolithic applications, which increases the return on hardware investment

Figure 2 shows block diagram for the HPE Container Platform Control Plane layout.

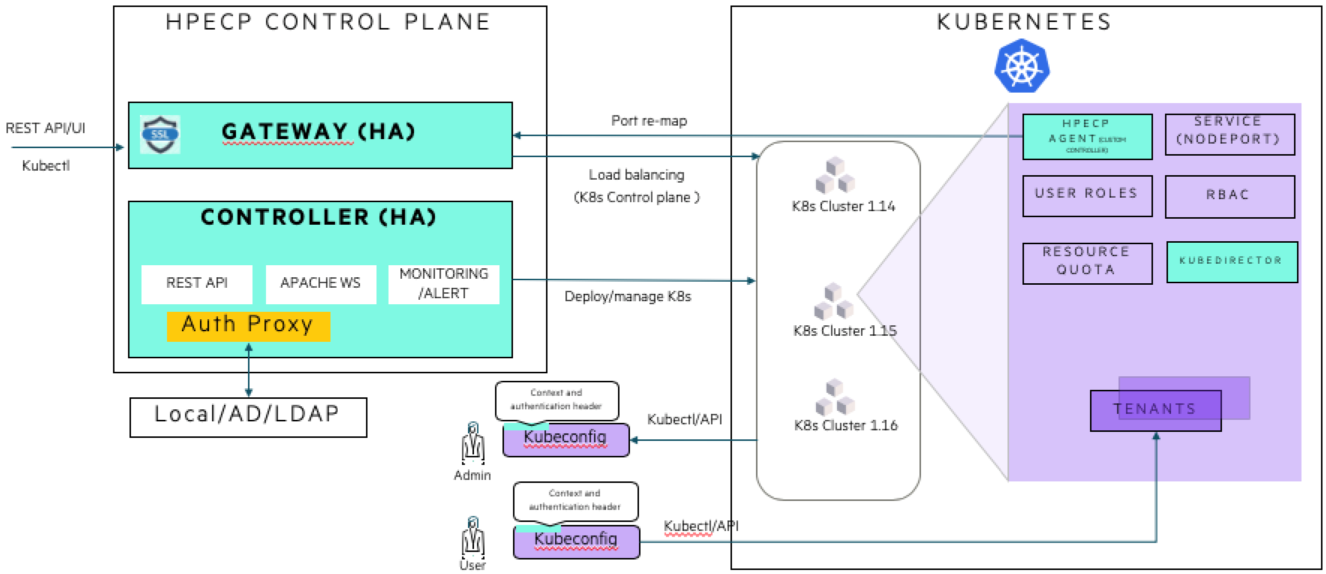

Figure 2. HPE Container Platform Control Plane

HPE Data Fabric

The HPE Container Platform will install and configure the HPE Data Fabric for AI and analytics. This HPE Data Fabric as shown in below figure is a distributed file and object store from MapR that manages both structured and unstructured data. It is designed to store data at exabyte scale, support trillions of files, and combine analytics and operations into a single platform. It supports industry standard protocols and APIs, including POSIX, NFS, S3, and HDFS. With production-ready capabilities like policy-based data tiering, consistent snapshots, and mirroring, the HPE Data Fabric serves as the enterprise standard for meeting stringent storage and processing SLAs across on-premises, hybrid cloud, and edge deployments.

All applications running in containers will be able to natively access data across the fabric through DataTap as well as through FSMount. Persistent volumes will be seamlessly available across clusters from this HPE Data Fabric. The HPE Data Fabric provides pre-integrated, scale-out, edge-ready persistent storage along with HPE Data Fabric for data services. The unique features provided by HPE Data Fabric are:

- Enterprise data persistence with fast, flexible, and consistent data access for multiple tenants leveraging a global name space and supporting multi-protocol access
- Auto-tiering that enables effortless data scale with seamless hot, warm, cold data tiering across hybrid cloud environments
- Secure and portable data access and controlled data and app mobility from core to cloud to edge using a common security and governance model
- Deployed clusters viewed as a single, logical, local cluster that are run globally
- Distributed metadata service and support of limitless scale (billions of files, PBs of data) with no single point of failure
- Ability to bring your AI/ML tool of choice without having to create another copy/silo of data
- HA/Resiliency/DR capabilities for mission-critical deployments through automatic services failover, container re-replication and mirroring
- Single distributed HPE Data Fabric that can store files, tables, and message topics with data portability across nine (9) industry standard APIs
- Multi-tenancy to support a range of application types in one platform

The HPE Data Fabric components are described at a high level in Figure 3

Figure 3 shows the HPE Data Fabric.

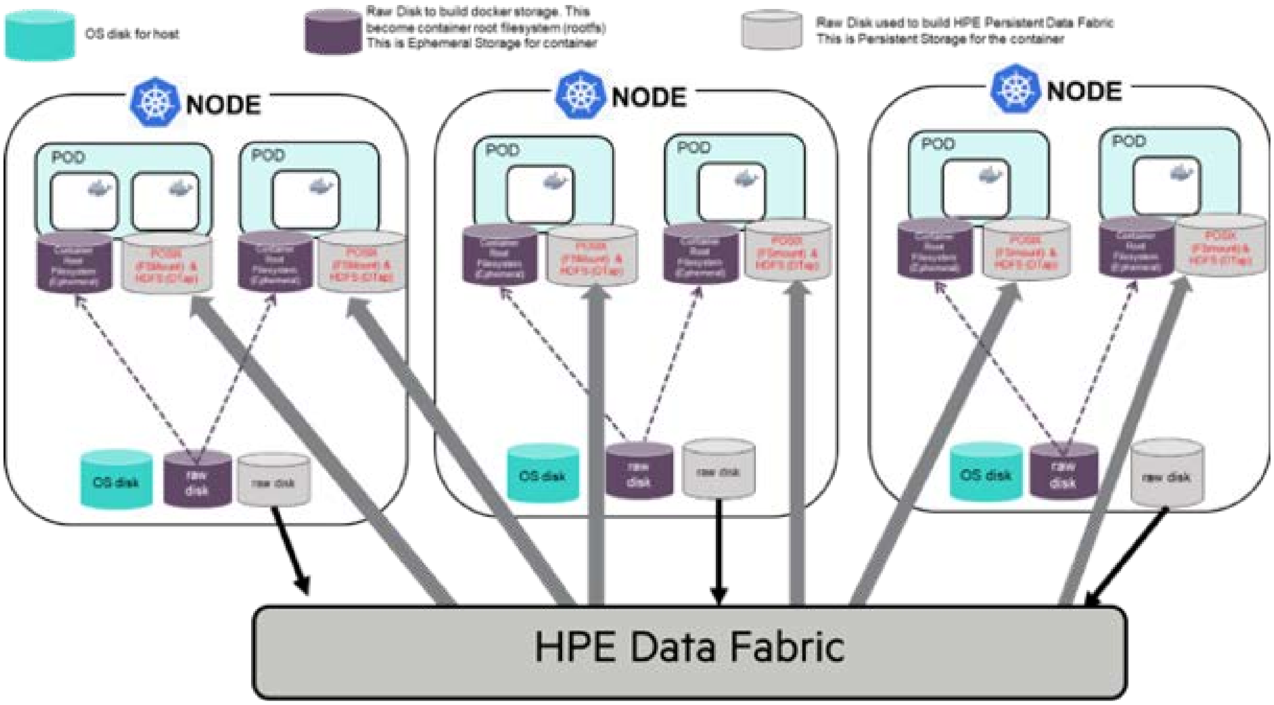

Figure 3. HPE Data Fabric

HPE Container Platform for Analytics

The HPE Container Platform leverages the Docker container technology and patented innovations to deliver self-service, speed, and efficiency for AI, ML, DL, and big data analytics.

The HPE Container Platform can be deployed on any infrastructure – whether on-premises, in the public cloud (e.g. AWS, Azure, and GCP), or a hybrid architecture. The HPE Container Platform installs as a software layer between the underlying server infrastructure and the big data distributions, AI/ML libraries and applications. The use of Docker is completely transparent and the HPE Container Platform customers benefit from greater agility and performance due to the lightweight nature of containers. They can leverage the flexibility of Docker to simplify development for AI/ML and big data applications, and the portability of containers to support both on-premises and public cloud deployments.

The key capabilities of HPE Container Platform for Analytics include:

- **Integrated platform for ML/DL and big data analytics:** The HPE Container Platform is an infrastructure platform purpose-built for big data and/or AI applications—including data science, analytics, machine learning (ML), and deep learning (DL). HPE Container Platform is designed to have enterprise-grade security, networking, and supports a variety of local and remote storage options.
- **Runs on-premises and/or on public cloud virtual machines (VMs):** The HPE Container Platform can be deployed on-premises, in the public cloud, or in a hybrid environment that includes both public cloud and on-premises resources.
- **Create virtual clusters:** The HPE Container Platform uses containers to replicate the functionality of physical clusters while adding flexibility and scalability at reduced cost. With HPE Container Platform, the containerized clusters can be created, modified, re-prioritized, and removed on demand in response to ever-changing needs within individual business units/departments. This reduces time-to-value from months to hours.
- **Multi-tenancy and enterprise-grade security model:** The HPE Container Platform integrates with the enterprise LDAP and Active Directory authentication systems. Administrators can create groupings of users and resources that restrict access to jobs, data, or clusters based on department membership and/or roles. The result is an integrated, secure, multi-tenant infrastructure.

- **Self-service portal:** The HPE Container Platform includes a self-service web portal that allow users to create and manage clusters, create and manage nodes, run jobs, and view monitoring statistics. A users visibility into resource utilization and their ability to take various actions on the platform is controlled by each users role and tenant membership, in accordance with the existing enterprise security policies. For example, department administrators can use the portal to provision nodes/clusters to run their own big data and AI applications without impacting nodes/clusters that are assigned to different departments and without having to manage the physical infrastructure.
- **GPU-As-A-Service (GPUaaS):** The HPE Container Platform provides significant advantages for running containers on bare-metal, ensuring optimal performance for AI / ML / DL applications and reducing overall TCO. Specifically for GPU-accelerated workloads, customers can improve GPU utilization and efficiency with our GPU-as-a-Service solution that allows organizations to virtualize and pool their on-premises GPU resources and share them with multiple teams and applications to create a cloud-like experience and elastic consumption model with the performance of bare metal. Now customers can develop and deploy their AI workloads on containers to dramatically reduce cost and complexity, while exploiting the full value and performance of GPUs to deliver the fastest results and highest throughput.
- **RESTful API:** The HPE Container Platform supports a RESTful API that surfaces programmable access to the same capabilities available via the self-service portal.
- **Superior performance:** The HPE Container Platform provides storage I/O optimizations to deliver bare metal performance to big data and AI applications without the penalties commonly associated with virtualization. The CPU cores and RAM in each host are pooled and then partitioned into virtual resource groups based on tenant requirements.
- **Works with existing infrastructure:** The HPE Container Platform enables your enterprise to repurpose its existing infrastructure investments for big data and AI deployments. HPE Container Platform can run on your physical and virtualized infrastructure, including CPUs and GPUs, as well as on all three major public clouds (Amazon Web Services, Google™ Cloud Platform, and Microsoft® Azure). Existing storage protocols are also supported (HDFS, HDFS with Kerberos, and NFS).
- **Reduced IT overhead:** The HPE Container Platform streamline the operations and reduces IT cost by automating provisioning, unifying management, and supporting push-button upgrades.
- **Increases utilization while lowering costs:** The HPE Container Platform saves the hardware and operational cost while simultaneously eliminating the complexity of managing multiple physical clusters.
- **High availability:** The HPE Container Platform supports three levels of high availability (at the platform, virtual cluster, and/or gateway node level) to provide redundancy and protection.
- **Compute and storage separation:** The HPE Container Platform enable the decoupling of analytical processing from data storage, providing the ability to independently scale compute and storage capacity instantly on an as-needed basis.
- **In-place access to both on-premises enterprise storage and cloud storage:** The HPE Container Platform can run jobs directly against both existing enterprise-class storage systems and cloud storage systems. The separation of compute and storage provided by the HPE Container Platform avoids the need to move or duplicate data before running analytics.
- **Support for custom apps and versions:** The HPE Container Platform includes an “App Store” with pre-configured Docker images for common AI/ML tools, big data frameworks, and applications. The platform provides an App Workbench that enables administrators to easily modify and update the pre-configured Docker images in their App Store (e.g., with a different version of the application or distribution) – or create new images for other applications and tools. Each HPE Container Platform customer has its own App Store populated with the applications and tools that their users need, thus providing the ultimate flexibility and configurability.

Figure 4 is a partial sample of the applications available in the App Store.

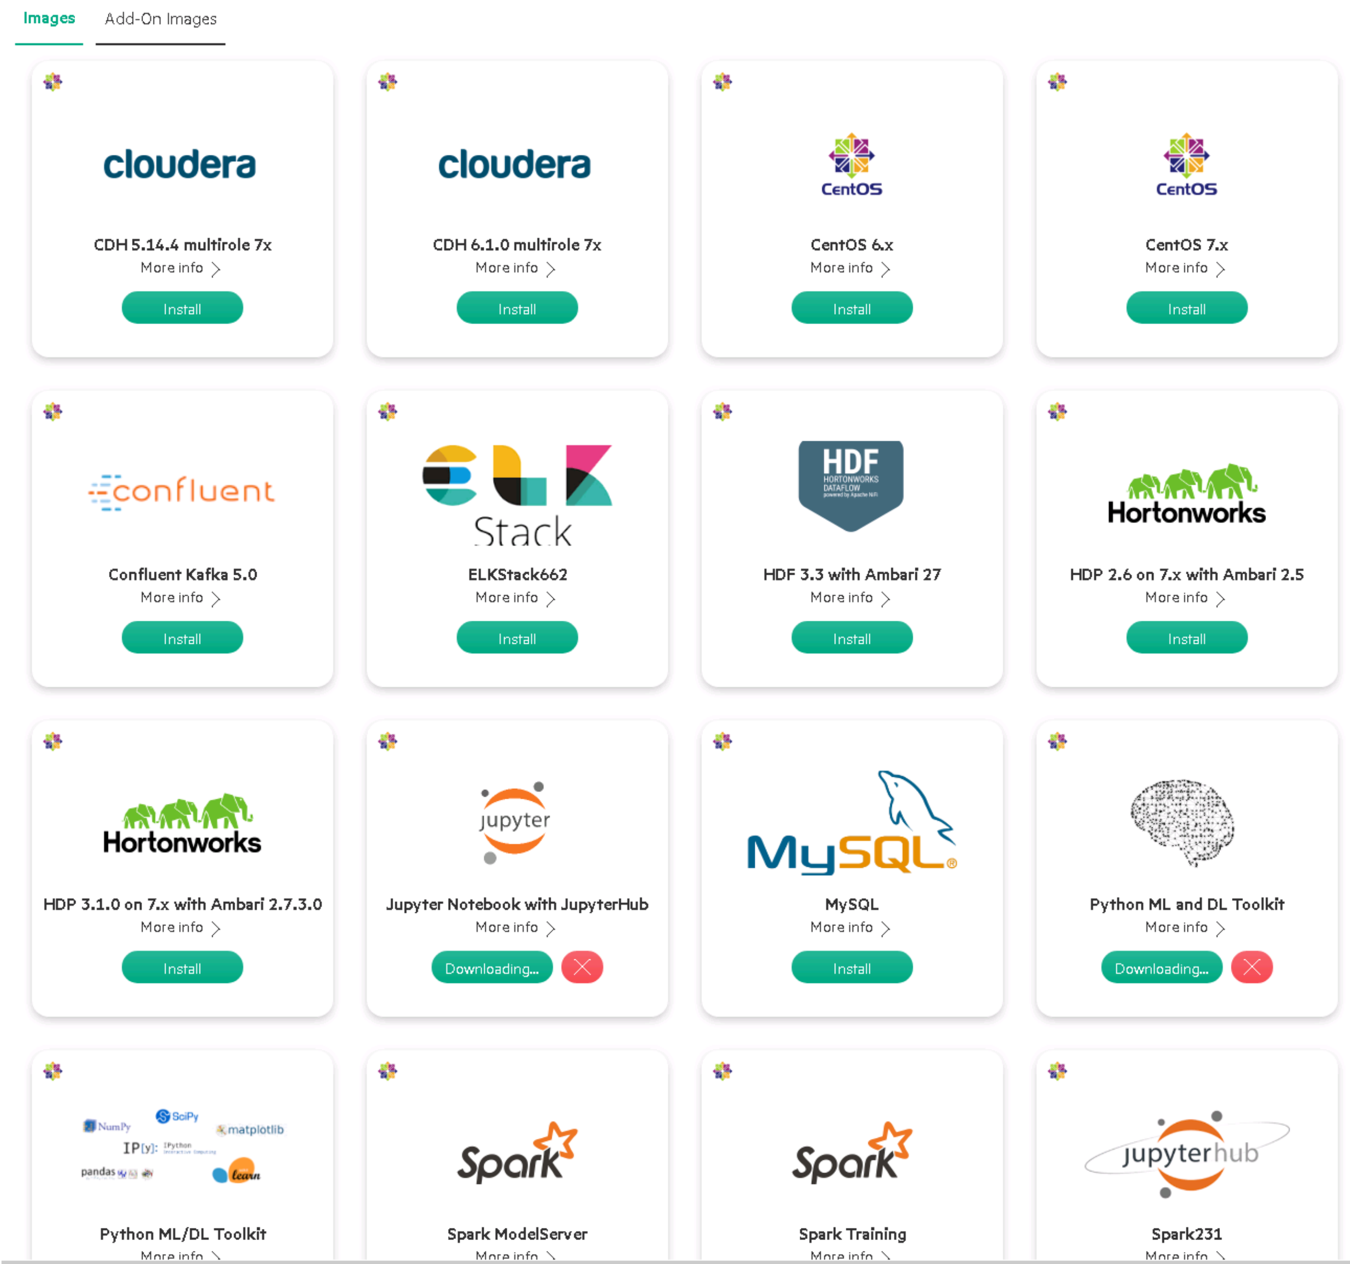

Figure 4. HPE Container Platform App Store

HPE ML Ops

HPE Machine Learning (ML) Ops is an optional software add-on to the HPE Container Platform that extends the speed and agility of the Container Platform by bringing the power of containers to the entire machine learning lifecycle — data preparation, model build, model training, model deployment, collaboration, and monitoring. HPE ML Ops is an end-to-end data science solution with the flexibility to run on-premises, in multiple public clouds, or in a hybrid model and respond to dynamic business requirements in a variety of use cases.

Figure 5 shows the HPE ML Ops solution.

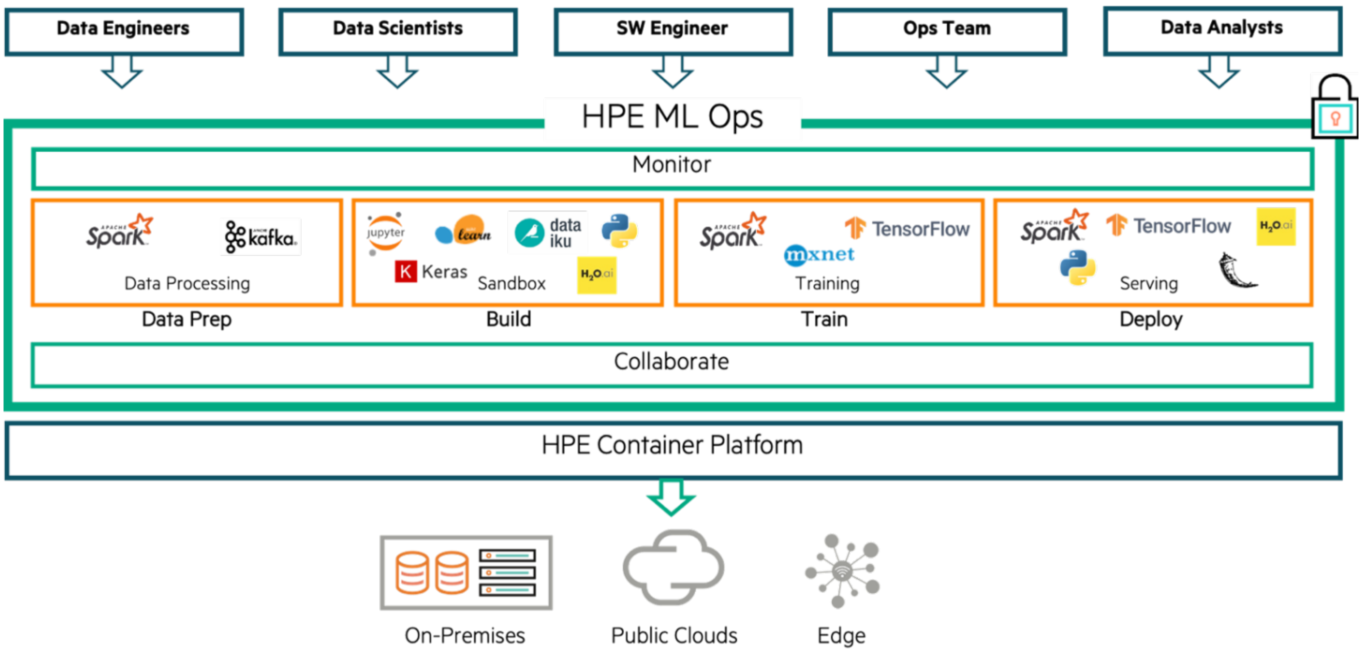

Figure 5. HPE Machine Learning Ops

With the HPE ML Ops solution, data science teams involved in the ML/DL lifecycle of building and deploying ML models can benefit from the industry's most comprehensive operationalization and lifecycle management solution for enterprise AI. The solution contains:

- **Model Building:** Pre-packaged, self-service sandbox environments for ML tools and data science notebooks such as TensorFlow, Keras, PyTorch, and more.
- **Model Training:** Scalable training environments with secure access to big data. single node or distributed multimode clusters—for development, test, or production workloads. Patented innovations provide high performing training environments, with compute and storage separation, that can securely access the shared enterprise data sources on-premises or cloud-based storage.
- **Model Deployment:** Flexible, scalable, and can deploy the model's native runtime image, such as Python, R, H2O into a secure, highly available, load balanced and containerized endpoint. An integrated model registry enables version tracking and seamless updates to models in production. Auto scaling from HPE ML Ops dynamically scales nodes for the scoring engines.
- **Model Monitoring:** End-to-end visibility across the ML model lifecycle. Complete visibility into runtime resource usage such as GPU, CPU, and memory utilization. Ability to track, measure, and report model performance along with third party integrations track accuracy and interpretability.
- **Collaboration:** Enable CI/CD workflows with code, model, and project repositories. Project repository and GitHub integration of HPE ML Ops provide source control, ease collaboration, and enable lineage tracking for improved auditability. The model registry stores multiple models, including multiple versions with metadata for various runtime engines.
- **Security and Control:** Secure multi-tenancy with integration to enterprise authentication mechanisms. HPE ML Ops software provides multitenancy and data isolation to ensure logical separation between each project, group, or department within the organization. HPE ML Ops integrates with the enterprise security and authentication mechanisms such as LDAP, Active Directory, and Kerberos.
- **Hybrid Deployment:** Support for on-premises, public cloud, or hybrid cloud. HPE ML Ops runs on-premises on any infrastructure, on multiple public clouds, or in a hybrid model, providing effective utilization of resources and lower operation costs.

Figure 6 shows the ML/DL lifecycle.

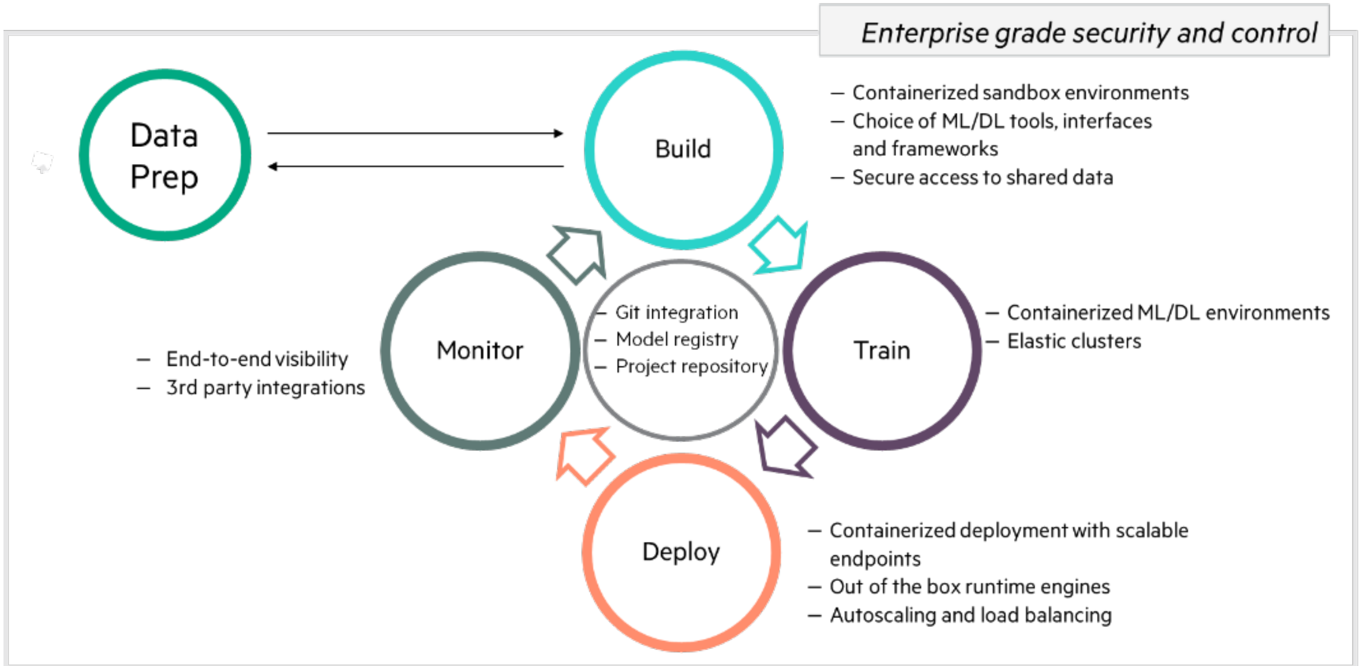

Figure 6. ML/DL lifecycle

HPE Apollo Servers for Analytics

The family of high performance HPE Apollo servers are one of the primary server platforms used in the family of the HPE Elastic Platform for Analytics (EPA) offering. It is a premier modular infrastructure foundation to accelerate business insights, enabling organizations to rapidly deploy, efficiently scale, and securely manage the explosive growth in volume, speed, and variety of big data and ML/DL workloads. EPA harnesses the power of faster Ethernet networks which enables a building block approach to independently scale compute and data store allowing you to scale your compute and data store at different rates. The base HPE EPA system uses the HPE Apollo 4200 as a storage block and the HPE Apollo 2000 as a compute block. By leveraging a building block approach, customers can simplify the underlying infrastructure needed to address a myriad of different business initiatives around data warehouse modernization, analytics, ML/DL, and build large-scale data lakes with diverse sets of data. As the workloads and data storage requirements change (often uncorrelated to each other), the HPE EPA system enables customers to easily scale by adding compute and storage blocks independently from each other, maximizing the infrastructure requirements for the workload demands.

Figure 7 highlights the different building blocks that are part of the HPE Elastic Architecture.

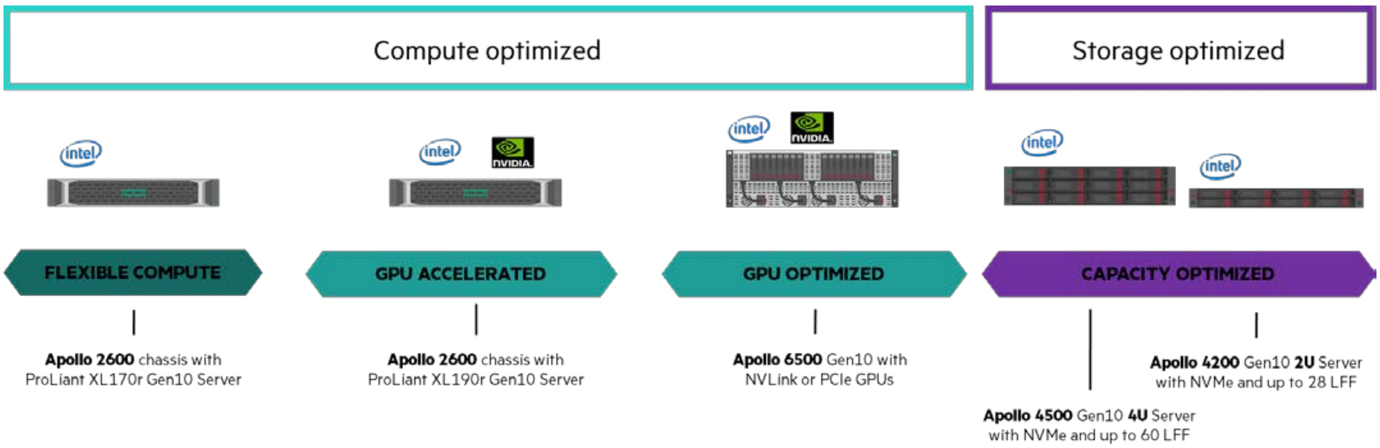

Figure 7. HPE Elastic Architecture

HPE Container Platform on HPE Apollo Gen10 servers best practices

Deploying the HPE Container Platform on the HPE Apollo Gen10 servers provides greater flexibility in deploying your workloads and managing your resource growth, by decoupling storage from compute. This section is intended to provide the high-level guidance and best practices for deploying the HPE Container Platform and the HPE ML Ops solution on the HPE Apollo Gen10 servers.

HPE Apollo servers configuration for HPE Container Platform

The HPE Container Platform uses four distinct host types, as shown in Table 1, with the recommended HPE Apollo server model.

Table 1. HPE Container Platform host types

Host Type	HPE Server Model
HPE Container Platform - Controller/Shadow Controller/Arbiter	Apollo 2000 with up to 4x HPE ProLiant XL170r
HPE Container Platform – Compute	Apollo 2000 with up to 4x HPE ProLiant XL170r
HPE Container Platform – Compute with GPUs	Apollo 2000 with up to 2x HPE ProLiant XL190r each with 1 or 2 GPUs Apollo 6500 with HPE ProLiant XL270d with 4 or 8 GPUs
HPE Container Platform - Gateway	Apollo 2000 with up to 4x HPE ProLiant XL170r
HDFS Storage	Apollo 4200 with up to 28x LFF HDD drives

HPE Apollo 2000 Compute Servers

HPE Apollo 2000 systems deliver a scalable, high-density workload-optimized compute module, supporting HPE ProLiant XL170r and HPE ProLiant XL190r Gen10 Servers. The HPE Apollo 2000 system can support dedicated and hybrid configurations that host both XL170r and XL190r servers.

- **HPE ProLiant XL170r Gen10 Server:** For compute intensive workloads, HPE ProLiant XL170r delivers four servers in a single 2U chassis. Each HPE ProLiant XL170r server is serviced individually without impacting the operation of other servers sharing the same chassis to provide increased server uptime. Each server harnesses the performance of up to 2666MHz memory (16 DIMM slots per server) and two Intel processors in a very efficient solution that shares both power and cooling infrastructure.
- **HPE ProLiant XL190r Gen10 Server:** For GPU workloads, the HPE ProLiant XL190r Gen10 server delivers two servers in a single 2U chassis. Each HPE ProLiant XL190r Gen10 server is serviced individually without impacting the operation of the other server sharing the same chassis to provide increased server uptime. Each server supports two NVIDIA® graphic cards and harnesses the performance of up to 2666MHz memory (16 DIMM slots per server) and two Intel processors in a very efficient solution that shares both power and cooling infrastructure.

For more information, refer to HPE Apollo 2000 servers.

HPE Apollo 6500 GPU system

The HPE Apollo 6500 Gen10 system is an ideal High-Performance Computing (HPC) and deep learning platform providing unprecedented performance with industry leading GPUs, a fast GPU interconnect, high bandwidth fabric, and a configurable GPU topology to match your workloads. Driving these requirements is the ability of computers to autonomously learn, predict, and adapt using massive data sets which is driving innovation and competitive advantage across many industries and applications. This system with rock-solid reliability, availability, and serviceability features include up to eight GPUs per server, NVLink for fast GPU-to-GPU communication, Intel® Xeon® Scalable processors support, choice of high-speed / low latency fabric, and workload enhanced using flexible configuration capabilities. While aimed at deep learning workloads, the system is suitable for complex simulation and modeling workloads.

For more detailed information, refer to HPE Apollo 6500 Gen10 system.

HPE Apollo 4200 Gen10 Storage Servers

HPE Apollo 4200 Gen10 Servers make up the HDFS storage layer, providing a single repository for big data. The HPE Apollo 4200 enables you to save valuable data center space through its unique density optimized 2U form factor which holds up to 28 LFF disks or 54 SFF disks and with the capacity for up to 392 TB per server. The rear drive cage can also support up to 6 SFF NVMe SSDs. It can grow your big data solutions with an infrastructure that is ready to scale. Another benefit is that the HPE Apollo 4200 fits easily into standard racks with a depth of 32 inch per server – no special racks are required.

The storage controllers in the HPE Apollo 4200 support HPE Secure Encryption, an HPE Smart Array controller-based data encryption solution that provides encryption for data at rest.

For more detailed information, refer to HPE Apollo 4200 systems.

HPE ProLiant Servers and NVIDIA GPUs

HPE ProLiant servers offer NVIDIA accelerators for high performance computation for deep learning, high-performance computing (HPC) workloads, or graphics. The NVIDIA accelerators for HPE ProLiant servers seamlessly integrate GPU computing with selected HPE server families. Designed for power-efficient, supercomputing, NVIDIA accelerators deliver dramatically higher application acceleration than a CPU-only approach for a range of deep learning, scientific, and commercial applications. The thousands of NVIDIA CUDA® cores of each accelerator allow it to divide large computing or graphics tasks into thousands of smaller tasks that can be run concurrently, thus enabling much faster simulations and improved graphics fidelity for extremely demanding 3D models.

For more detailed information, refer to HPE AI and deep learning.

HPE Container Platform CPU and memory

The CPU and memory requirements presented in Table 2 provide some general guidelines. However, your actual needs will depend on several factors, including:

- For optimal performance, the total number of CPU cores on each host should equal the flavor(s) you will be using multiplied by the number of concurrent instances of that flavor that will be in use on that host. For example, the HPE Container Platform default small flavor uses four (4) vCPU cores per virtual node. If you plan to run four virtual nodes of this flavor on each host, then each host should have 16 CPU cores. You should also consider having at least one extra CPU core available for management and other overhead. This level of sizing ensures one physical CPU core per vCPU; however, the HPE Container Platform does allow you to over-provision the CPU cores.
- The amount of RAM on each host depends on the flavors used multiplied by the number of concurrent instances of that flavor that will be in use on that host. The HPE Container Platform requires an additional 1/16 of this total plus another 5 GB. For example, the HPE Container Platform-default medium flavor uses 12 GB of RAM per virtual node. With four virtual nodes of this flavor on each host, each host should have 56 GB of RAM or more.
  - Medium flavor (12 GB RAM) multiplied by four virtual nodes equals 48 GB of RAM
  - 1/16 of 48 GB is 3 GB, and 48 GB plus 3 GB is 51 GB
  - Add the additional 5 GB to obtain 56 GB total

Table 2. HPE Container Platform host CPU and memory recommendations

Host Type	Deployment Size	Memory	Processor
HPE Container Platform – Controller, Shadow Controller, Arbiter	• Starter	192 GB	(2) Intel Cascade Lake 5215 - 10C 2.5 GHz
	• Medium	384-768 GB	(2) Intel Cascade Lake 6226 - 12C 2.8 GHz
	• Large	384-768 GB	(2) Intel Cascade Lake 6242 - 16C 2.8 GHz
HPE Container Platform – Compute	• Starter	192 GB	(2) Intel Cascade Lake 5215 - 10C 2.5 GHz
	• Medium	384-768 GB	(2) Intel Cascade Lake 6226 - 12C 2.8 GHz
	• Large	384-768 GB	(2) Intel Cascade Lake 6242 - 16C 2.8 GHz
HPE Container Platform – GPU	• Starter	384 GB	(2) Intel Cascade Lake 5215 - 10C 2.5 GHz
	• Medium	384-768 GB	(2) Intel Cascade Lake 6226 - 12C 2.8 GHz
	• Large	384-768 GB	(2) Intel Cascade Lake 6242 - 16C 2.8 GHz
HPE Container Platform – Gateway	All	192 GB	(2) Intel Cascade Lake 5215 - 10C 2.5 GHz

To assist in sizing an HPE Container Platform cluster, Hewlett Packard Enterprise has developed a sizing tool that can be found at https://solutionsizers.ext.hpe.com/EPASizer/.

Note
HPE Gen10 systems support a variety of flexible memory configurations. For optimal performance, it is recommended to balance the total memory capacity across all installed processors and making use of all six memory channels per CPU with up to two DIMM slots per channel.

HPE Container Platform storage

This section describes the various storage requirements for the HPE Container Platform and how the platform makes datasets available to the containerized clusters.

Container local data storage

When a virtual Hadoop cluster is created, HDFS is provisioned within the Docker containers that comprise a virtual Hadoop cluster. The underlying storage for the HDFS data nodes in the containers resides on local disks in the physical servers hosting those containers. HPE Container Platform refers to the set of local disks as node storage. When using HDFS storage in a virtual cluster, the data does not persist beyond the life of the virtual cluster.

Ephemeral storage

Ephemeral storage is built from the local storage in each host of the HPE Container Platform deployment and used for the disk volumes, that back the local storage for each virtual node. Installing an HPE Container Platform host reserves a subset of the local disks on that host for node storage. The platform creates Linux® physical volumes on those disks, and then uses those physical volumes to create a Linux volume group. From this Linux volume group, a Linux logical volume is created. This Linux logical volume is assigned to the Linux Docker subsystem, which then uses the Linux device mapper functionality to allocate portions of the logical volume to the containers running on that host, for use as local storage within those containers.

Persistent storage

HPE Container Platform supports deploying an HPE Data Fabric on the local disks within the servers running the HPE Container Platform services. This local storage can serve as HDFS storage or as Persistent volumes for Kubernetes clusters. Persistent volumes for Kubernetes stateful clusters are seamlessly available either from the native data fabric or HPE Nimble Storage using the HPE Container Storage Interface Driver that is deployed during the cluster creation.

Remote data storage

Getting the maximum flexibility from a container-based solution such as HPE Container Platform requires being able to independently scale compute and storage resources. It is also essential to support the persistence of big data datasets beyond the lifespan of a big data compute cluster. The HPE Container Platform DataTap and IOBoost technologies allow big data clusters running on HPE Container Platform to access remote data, regardless of location or format.

DataTap creates a logical data lake overlay that enables access to shared data in the enterprise storage devices. This allows users to run big data and ML/DL jobs using the existing enterprise storage, without the need to make time-consuming copies or transfers of data to local disks. IOBoost augments DataTap's flexibility by adding an application-aware data caching and tiering server to ensure high-speed remote data delivery.

Operating system storage

For all host types, the recommended storage for the operating system is two 960 GB SSD's in a RAID 1 configuration.

HPE Container Platform storage recommendations

Table 3 lists the recommended minimum storage configuration for each HPE Container Platform host type.

Table 3. HPE Container Platform storage recommendations for HPE ProLiant XL170r, HPE ProLiant XL190r, and HPE ProLiant XL270d

Host Type	HPE Container Platform – Storage Type	Storage Recommendation
HPE Container Platform – Controller, Shadow Controller, Arbiter	• OS • Ephemeral Storage • Persistent Storage	• 2 x 960 GB SSD configured as RAID 1 • 3 x 6.4 TB mixed use SSD. • 0-1 x 2 TB SATA 7.2 SFF HDD.
HPE Container Platform – Compute	• OS • Ephemeral Storage • Persistent Storage	• 2 x 960 GB SSD configured as RAID 1 • 3 x 6.4 TB mixed use SSD. • 0-1 x 2 TB SATA 7.2 SFF HDD.
HPE Container Platform – Gateway	OS	2 x 960 GB SSD configured as RAID 1

Note

The HPE ProLiant XL170r and HPE ProLiant XL190r have six drive bays. Two are used for the OS drives and four are available for ephemeral storage and persistent storage. Choose the appropriate size and number of disks based upon your ephemeral storage and persistent storage space requirements.

HPE Container Platform networking

Hewlett Packard Enterprise recommends deploying 25Gbps or greater Ethernet adapters on all HPE Container Platform hosts. Table 4 shows the recommended networking hardware for the HPE Container Platform clusters.

Table 4. HPE Container Platform networking recommendations

Host Type	Network Recommendation
HPE Container Platform – All host types	HPE Eth 10/25Gb 2P 640FLR-SFP28 Ethernet Adapter

Summary

Companies are driving digital transformation and investing in innovation to better compete. They are looking to simplify production environments in a hybrid cloud environment. In order to do so, they may have a mandate to move their application portfolio to containers. They are struggling to migrate non-cloud-native enterprise applications running on-premises into containers due to lack of time and expertise to re-architect or refactor all these applications to cloud-native microservices-based applications.

With the HPE Container Platform, enterprise organizations now have a unified Kubernetes-based software solution for both traditional analytics workloads and cloud-native application workloads, streamlining deployment and operation with consistent orchestration and management. The platform acts as the control plane for container management and provides persistent container storage across multiple versions of open source Kubernetes for container orchestration.

This next-generation container platform dramatically reduces cost and complexity by running containers on bare-metal – while providing the flexibility to deploy on virtual machines and cloud instances in a hybrid or multi-cloud model. The solution provides customers with greater efficiency, higher utilization, and bare-metal performance by "collapsing the stack" and eliminating the need for virtualization. Developers have secured on-demand access to their environments, so that they can develop apps and release code faster, with the portability of containers to build once and deploy anywhere. IT teams can manage multiple Kubernetes clusters with multi-tenant container isolation and data access for any workload from edge to core to cloud. And they can extend the benefits of containers beyond cloud-native microservices-architected stateless applications, providing the ability to containerize monolithic stateful analytic applications with persistent data.

When deployed on HPE Apollo servers, enterprises can get started by delivering a simpler, more scalable approach to modernize analytics applications on container architectures. This helps organizations to drastically increase the velocity of application development and accelerate digital transformation.

This Reference Configuration provides an enterprise grade solution and enables the organization to increase agility, simplify operations and deliver a cloud-like experience and better return on investment.

Version history

Document version	Date	Description of changes
1.0	03/20/2020	Initial Publication

Resources and additional links

HPE Container Platform, https://www.hpe.com/us/en/solutions/container-platform.html

BlueData, https://www.hpe.com/info/bluedata

MapR, https://www.hpe.com/us/en/solutions/mapr.html

HPE big data solutions, http://www.hpe.com/bigdata

HPE Reference Architectures, http://www.hpe.com/info/ra

HPE Servers, http://www.hpe.com/servers

HPE Storage, http://www.hpe.com/storage

HPE Apollo 2000 systems, https://buy.hpe.com/b2c/us/en/servers/apollo-systems/apollo-2000-system/apollo-2000-system/hpe-apollo-2000-system/p/1010192759

HPE Apollo 4000 systems, https://buy.hpe.com/b2c/us/en/servers/apollo-systems/apollo-4000-system/apollo-4200-server/hpe-apollo-4200-gen10-server/p/1011147097

HPE Apollo 6500 systems, https://buy.hpe.com/b2c/us/en/servers/apollo-systems/apollo-6500-system/apollo-6500-system/hpe-apollo-6500-gen10-system/p/1010742495

HPE Synergy, https://www.hpe.com/info/synergy

HPE Nimble Storage, https://www.hpe.com/us/en/storage/nimble.html

HPE Networking, https://www.hpe.com/networking

HPE Technology Consulting Services, http://www.hpe.com/us/en/services/consulting.html

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Intel, Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. NVIDIA and NVIDIA CUDA are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Google © 2012 Google Inc. All rights reserved. Google and the Google Logo are registered trademarks of Google Inc.

a50001145enw, version 1.0, March 2020