Hewlett Packard Enterprise

Reference Configuration

HPE REFERENCE CONFIGURATION FOR HPE EZMERAL RUNTIME ENTERPRISE 5.6 AND HPE EZMERAL DATA FABRIC ON HPE PROLIANT DL325 GEN11 AND HPE PROLIANT DL385 GEN11 SERVERS

HPE Ezmeral Runtime Enterprise (ERE) version 5.6

CONTENTS

EXECUTIVE SUMMARY

Enterprise organizations across all industries are embarking on a hybrid cloud journey. With a hybrid cloud, companies can place each workload in its optimal environment, which has the effect of speeding deployments, reducing costs, and allowing the right security measures to be put into place to meet compliance standards for sensitive data. Integrating public cloud with private systems allows organizations to take advantage of pay-as-you-go public cloud pricing models to handle high-volume batch computing jobs and sudden traffic spikes while maintaining control and predictability with owned infrastructure.

The current IT practices and various incompatible application deployment environments have created challenges for organizations to achieve these objectives. Some of the key challenges are as follows:

- Modernizing legacy apps to take advantage of the latest agile cloud-native and DevOps innovations are difficult and time-consuming.
- Managing workloads in automated platforms are spanning across hybrid cloud environments is challenging.
- Provisioning a new environment is a slow process and can significantly stifle innovation as teams must wait for the environment to be available.
- Vendor lock-in is a genuine concern and pertains to the using any proprietary technology, on-premises, or public clouds.
- Siloed infrastructure increases overhead costs including administrative overhead and the price of additional infrastructure.

To unleash business opportunities through digital transformation, enterprises must overcome these restrictions and adapt to the new concepts and solutions of the next generation IT practices.

Hewlett Packard Enterprise believes the next wave of technology innovation will be built on a more agile application development platform and will use modern, cloud-native architecture with portability across any infrastructure from the edge to the cloud. Application containers and open-source Kubernetes (K8s) are central to this new approach, providing key features and technologies to ensure organizations succeed in their digital transformation journey.

HPE Ezmeral Runtime Enterprise is the industry's first enterprise-grade container platform designed to deploy cloud-native and non-cloud-native applications using open - source Kubernetes – running on bare-metal or virtualized infrastructure, on any public cloud, and at the edge. HPE Ezmeral Runtime Enterprise helps enterprise organizations accelerate their data-driven digital transformation efforts – to enhance customer and user experiences, grow new revenue streams, and deliver faster business innovation. With the HPE Ezmeral Runtime Enterprise solutions, we provide our customers with:

- A unified multi-cluster control plane - To deploy, run, and manage multi-Kubernetes clusters.
- A platform to modernize both cloud-native and non-cloud-native monolithic apps - With differentiated capabilities to containerize monolithic and stateful applications, without the need to refactor code.
- Ready-to-run workload-optimized solutions - To build containerized applications for a wide range of use cases, with their choice of open-source or ISV tools, and deploy them anywhere – on-premises, in any public cloud, or at the edge.

HPE Ezmeral Runtime Enterprise includes a pre-integrated HPE Ezmeral Data Fabric for persistent container storage – and it provides the foundation for our HPE Ezmeral ML Ops solution. Our container platform is the cornerstone of our HPE Ezmeral software portfolio and a key element of our overall as-a-service strategy with HPE GreenLake (including our new container service).

The HPE Ezmeral Runtime Enterprise on HPE ProLiant DL325 Gen11 and HPE ProLiant DL385 Gen11 servers, along with automation scripts help deploy a production-ready environment with software-defined storage spanning multiple clusters and data centers.

Target audience: This Reference Configuration is for IT decision-makers and architects who are directive to invest in the cloud, modernize their legacy applications, and reduce their data center footprint. The intent is to assist platform engineers and cloud architects in defining and implementing a container, or hybrid cloud strategy. DevOps managers, infrastructure engineers, and Line-of-Business (LOB) managers will benefit from this paper and improve their operational efficiencies in containers or hybrid cloud or accelerate the time to value for developing and deploying analytics.

Document purpose: This Reference Configuration provides an overview of the deployment of HPE Ezmeral Runtime Enterprise and HPE Ezmeral Data Fabric on Bare Metal running on HPE ProLiant DL325 Gen11 and HPE ProLiant DL385 Gen11 servers using Local Disks for persistent volumes.

INTRODUCTION

Today's economic market conditions make it imperative for enterprises to accelerate their digital transformation strategy to become competitive and innovative in their industries. However, traditional IT practices and different application deployment environments present daunting challenges that one must overcome. According to International Data Corporation (IDC), the adoption of containers and Kubernetes is increasing for cloud-native applications. Around 2.8 billion container instances are expected by 2025, and 15% of enterprise applications are containerized. At the same time, containerized workloads spanning hybrid environments will also increase over the next few years. Container use is proliferating, and vendors are investing heavily to bring containerized workloads to the enterprise market as customers prepare to invest in containers as part of their digital transformation initiatives.

New cloud-native applications are developed from the ground up leveraging containers and microservices architecture, using a DevOps and continuous integration/continuous delivery (CI/CD) approach. Beyond these new cloud-native architectures, enterprises also use and own a variety of existing legacy applications that are not cloud-native. Many legacy applications have a traditional monolithic, three-tier, client-server software architecture. Organizations want to modernize these existing monolithic enterprise applications, a few options for performing this modernization effort include rewriting the legacy applications in a modern programming language, re-architecting them as cloud-native applications, or porting them over to a new computing platform. The goal is to reduce the cost of operations while bringing agility and modern DevOps efficiencies to these legacy applications. These modernization efforts are time-consuming and expensive. Running some of the non-cloud-native, monolithic applications in containers provides an easier way to bring the modernization benefits of agility and efficiency without the need to do a costly refactoring of code. When performance and isolation are the key drivers of outcome for containerized applications, the most efficient approach is to deploy the containerized infrastructure on Bare Metal computes.

Hewlett Packard Enterprise is fundamentally changing that equation by delivering a container platform for a broader range of application architecture with portability across on-premises, public cloud, and edge environments. HPE Ezmeral Runtime Enterprise is uniquely positioned to help customers experience these challenging times and prepare for this rapidly changing digital future—leveraging AI and data-driven innovation. HPE Ezmeral Runtime Enterprise eliminates the complexity and expense of a virtual machine layer with Bare Metal containerization that eliminates the expense of the hypervisor layer. This ensures enterprise-grade security, performance, scalability, and reliability.

The HPE Ezmeral portfolio allows you to:

- Run containers and Kubernetes at scale to modernize apps, from the edge to the cloud.
- Manage your apps, data, and ops – leveraging AI and analytics for faster time-to-insights.
- Ensure control for governance, compliance, and lower costs.
- Provide enterprise-grade security and authentication to reduce risk.

This Reference Configuration provides information on deploying of HPE Ezmeral Runtime Enterprise and HPE Ezmeral Data Fabric architecture on HPE ProLiant DL325 Gen11 server and HPE ProLiant DL385 Gen11 server a general-purpose workload-optimized platform to serve the needs of DevOps teams, CI/CD workflow integration, application modernization, and hybrid cloud solutions for the enterprise. This solution provides a cloud-like experience to our customers from edge to core to the cloud. The HPE ProLiant DL325 Gen11 and HPE ProLiant DL385 Gen11 a general-purpose built server, storage, and networking hardware are the foundation for an infrastructure that provides both rapid deployment and scaling while delivering the highest levels of performance, quality, and availability. Scale computing, storage, networking, and scale-up and scale-out solutions for your big data analytics and machine learning. This solution also showcases how to modernize a legacy application using KubeDirector.

The combination of HPE Ezmeral Runtime Enterprise (with pre-integrated HPE Ezmeral Data Fabric), HPE ProLiant DL infra delivers a business-focused architecture that rapidly deploys modern containers supporting the new application framework. 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.

Key benefits includes the following:

- Flexible unified platform for cloud-native and non-cloud-native apps
- Enterprise-class security
- Fast, easy deployment and management of Kubernetes clusters provides rapid time to value
- Flexible multi-cluster, multi-tenant control plane facilitates the ability to deploy on-premises, public clouds, or at the edge
- Curated App Store allows the rapid creation of solutions for various use cases. For example, CI/CD, app modernization, NoSQL databases, data analytics, and AI/ML

Key differentiators of HPE Ezmeral Runtime Enterprise and HPE Ezmeral Data Fabric on HPE ProLiant DL Gen11 servers

The key differentiators are as follows:

- Supporting the Kubernetes guidelines by disaggregating compute and storage at the platform layer, in line with the fundamental tenet of modern data center design.
- Deployment from power-on to developer ready takes only hours, not weeks or months with the automation of OS deployment, and installation of the HPE Ezmeral Runtime Enterprise.
- Simplifies the user experience with a self-service portal and one-click application deployment.
- HPE Ezmeral Runtime Enterprise connects to external data without copying data locally by leveraging the HPE Ezmeral Data Fabric.
- HPE Ezmeral Data Fabric can also be leveraged for scale-out, edge-ready persistent storage along with data platform for data services.

HPE innovations such as KubeDirector offers a controller to deploy non-cloud-native apps. Comprehensive security models deliver as-a-service with end-to-end security from HPE ProLiant DL Gen11 server's silicon root-of-trust to HPE Ezmeral Data Fabric's data encryption and HPE Ezmeral Runtime Enterprise integration with security services such as Active Directory. This Reference Configuration provides an overview of deploying and managing multiple HPE Ezmeral Runtime Enterprise delivering containerized clusters on top of the best-in- HPE ProLiant DL Gen10 servers and HPE Ezmeral Data Fabric.

SOLUTION OVERVIEW

The solution leverages the HPE ProLiant DL Gen11 AMD servers Infrastructure along with major HPE products and technologies as follows:

- HPE Ezmeral Runtime Enterprise
- HPE Ezmeral Data Fabric
- HPE ProLiant DL325 Gen11 compute
- HPE ProLiant DL385 Gen11 compute

These components are discussed in detail in the following sections:

HPE Ezmeral Runtime Enterprise

HPE Ezmeral Runtime Enterprise provides you with an enterprise-grade platform to deploy Kubernetes at scale for a wide range of use cases on Bare Metal or virtualized infrastructure. It can be run on premises, in hybrid and multi-cloud environments, and at the edge. HPE Ezmeral Runtime Enterprise is also the industry's first container platform designed to run modern applications (both cloud-native and non-cloud-native) with persistent data, making it easier for enterprises to manage their apps with containerized application deployments.

HPE Ezmeral Runtime Enterprise is the full-featured compute, storage, and container-management foundation that supports workload-specific solutions such as HPE Ezmeral ML Ops and HPE Ezmeral Unified Analytics.

Figure 1 shows the overview of the HPE Ezmeral software portfolio.

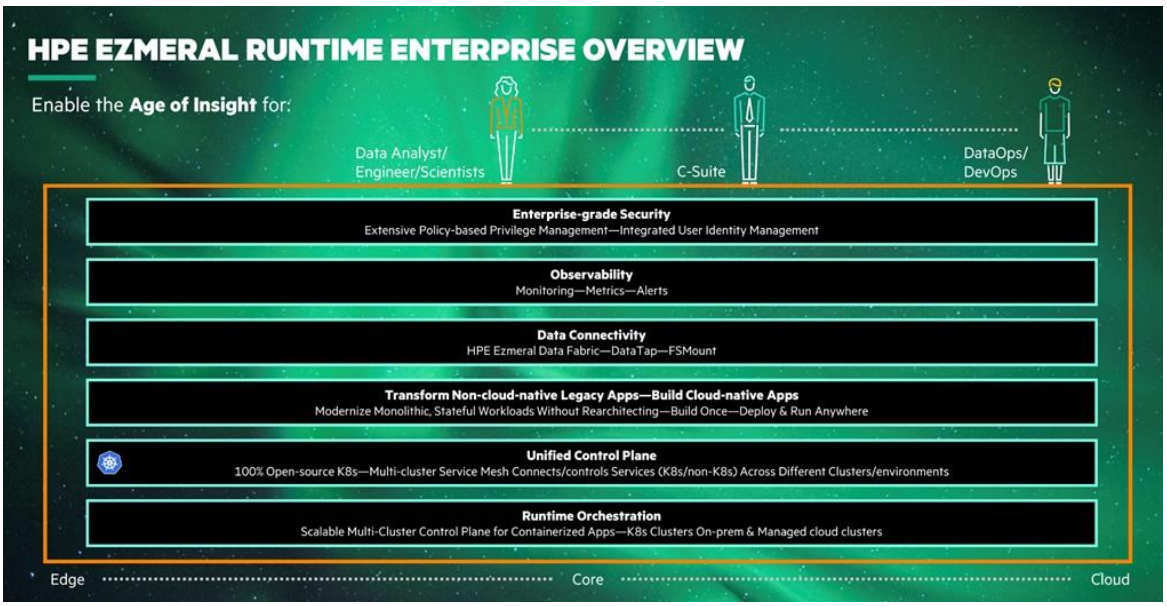

FIGURE 1. HPE Ezmeral software portfolio overview

Key features

- **Edge to cloud:** The industry's first and only 100%, open-source, Kubernetes, hybrid analytics platform spanning edge to cloud helps enterprises modernize their apps with containerized application deployments on Bare Metal or VMs spanning on-premises, multiple clouds, and at the edge; allows you to build once, run anywhere.
- **Multi-cluster, multi-tenant Kubernetes management**: Fast, easy deployment, management, and monitoring of Kubernetes clusters both on-prem and off-prem from a single pane of glass.
- **Enterprise-grade security:** Built-in security controls integrate with identity providers such as AD/LDAP; single sign-on; SAML integration; role-based access controls for secure access to the platform; Falco container runtime security for proactive threat detection and alerting.
- **GitOps-based centralized policy management and drift management**: Seamless and fleet management of clusters; ArgoCD leveraged to ensure clusters are consistent and immutable for continuous compliance.
- **Turnkey solution**: Easily containerize cloud-native and non-cloud-native apps; KubeDirector—an open-source custom Kubernetes controller—allows you to deploy non-cloud-native apps without rearchitecting or refactoring.
- **Accelerated analytics**: GPU sharing by using NVIDIA Multi-instance GPU fractionalization improves collaboration and GPU utilization.
- **Frictionless data access**: HPE Ezmeral Data Fabric, DataTap and FSMount let you connect to and manage data wherever it is located.

- **Built-in Service Mesh and observability**: For intelligent traffic shaping, load balancing, canary rollouts, and A/B testing of application microservices; visualize tenant-granular workload traffic for rapid troubleshooting and analysis via natively integrated Istio Service Mesh.
- **One-click provisioning**: An App store of curated, prebuilt, ready-to-run solutions for a wide range of applications including AI/ML, DataOps, analytics, CI/CD, DevOps apps and services, with the ability to BYO application via KubeDirector and App Workbench.
- **Available via HPE GreenLake**: Cloud services for HPE Ezmeral Runtime Enterprise are available through HPE GreenLake to deliver a preconfigured platform designed for multi-cluster, multi-tenant Kubernetes deployment.
- **License Information**: Information about the features included with an HPE Ezmeral Runtime Enterprise license, with a comparison to other HPE Ezmeral Runtime Enterprise product licenses, is provided in the product QuickSpecs.

HPE Ezmeral Runtime Enterprise 5.6

HPE Ezmeral Runtime Enterprise 5.6 includes the following new features:

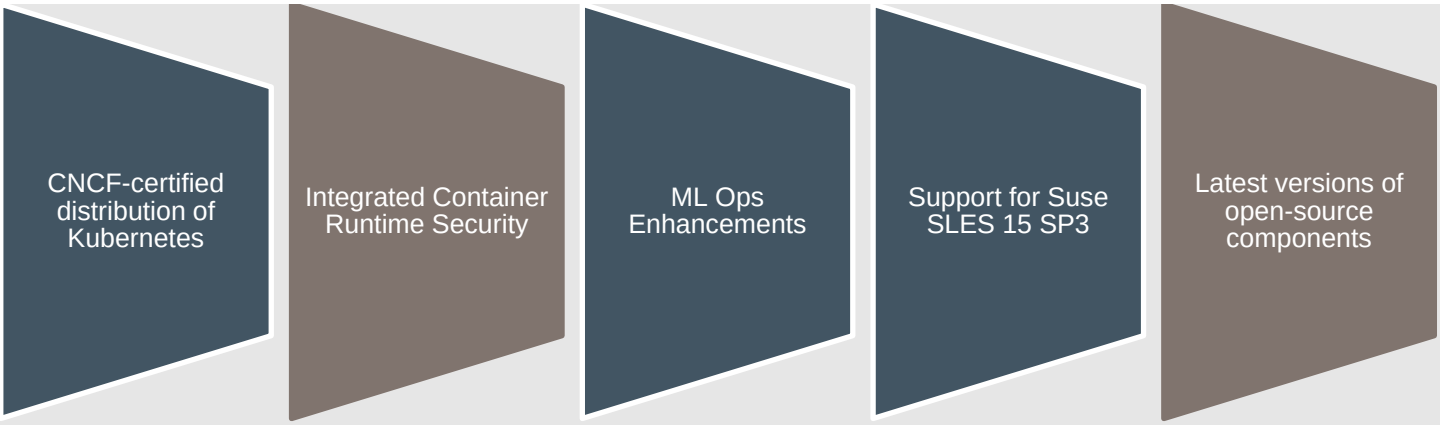

FIGURE 2. HPE Ezmeral Runtime Enterprise 5.6 What's New

OS Support

HPE Ezmeral Runtime Enterprise 5.6.0 adds support for SLES 15 SP3.

Kubernetes Distribution and Version Support

- Beginning with HPE Ezmeral Runtime Enterprise 5.6.0, Hewlett Packard Enterprise provides a CNCF-certified distribution of Kubernetes. The suffix -hpe<number>labels the Hewlett Packard Enterprise distributions of Kubernetes.
- The Kubernetes versions that were released with HPE Ezmeral Runtime Enterprise 5.6.0 include the following:
  - 1.22.15-hpe3
  - 1.23.14-hpe2
  - 1.24.8-hpe2
- With the introduction of Kubernetes bundles, later versions of Kubernetes can be added to compatible releases of HPE Ezmeral Runtime Enterprise. For information about supported versions of Kubernetes and Kubernetes bundles, see Support Matrixes.
- For deployments that are upgraded to 5.6.x versions of HPE Ezmeral Runtime Enterprise, existing Kubernetes clusters with compatible Kubernetes versions are permitted for a limited time but must be migrated to a Hewlett Packard Enterprise distribution of Kubernetes manually. For more information, contact Hewlett Packard Enterprise Support.

Updated Versions of Open Source Components

HPE Ezmeral Runtime Enterprise 5.6.0 supports the latest versions of open-source components. For more information about supported add-ons, see Support Matrixes.

HPE Ezmeral Data Fabric on Bare Metal Support

HPE Ezmeral Runtime Enterprise 5.6.0 provides enhanced support for registering HPE Ezmeral Data Fabric on Bare Metal as tenant storage as follows:

- Adding support for registering HPE Ezmeral Data Fabric on Bare Metal 7.0 as tenant storage. Registering HPE Ezmeral Data Fabric on Bare Metal 7.0 as tenant storage continues to be supported.
- Simplifying the registration process.
- Adding support for registering the same HPE Ezmeral Data Fabric on Bare Metal instance as tenant storage for multiple deployments of HPE Ezmeral Runtime Enterprise.

HPE Ezmeral Data Fabric on Kubernetes in Non-Production environments only

- HPE Ezmeral Data Fabric on Kubernetes is an implementation of HPE Ezmeral Data Fabric cluster.
- Beginning with HPE Ezmeral Runtime Enterprise 5.5.0, HPE Ezmeral Data Fabric on Kubernetes is available for use in non-production deployments of HPE Ezmeral Runtime Enterprise, but it is not supported for production environments.

Embedded Data Fabric is discontinued

- Embedded Data Fabric is an implementation of HPE Ezmeral Data Fabric that is locally embedded and runs on HPE Ezmeral Runtime Enterprise hosts.
- Embedded Data Fabric is not supported on 5.5.0 and later versions of HPE Ezmeral Runtime Enterprise.
- Hewlett Packard Enterprise recommends migrating pre-5.6.0 deployments of Embedded Data Fabric to HPE Ezmeral Data Fabric on Bare Metal for production or non-production use cases. Alternatively, HPE Ezmeral Data Fabric on Kubernetes can also be considered for non-production use cases only.

HPE Ezmeral ML Ops Enhancements

HPE Ezmeral Runtime Enterprise 5.6.0 includes the following enhancements and changes to the HPE Ezmeral ML Ops features:

Support for Kubeflow 1.6

Users can take advantage of the latest innovations made in the Kubeflow open-source end-to-end Machine Learning platform which is installed and configured with a single click of a button.

Secure Model Management in an automated tenant environment

- Built-in support for experiment tracking and model management.
- Multiple tenants can use a shared model registry and experiment tracker; end users are authenticated and provided controlled access to their own model and model metadata.

Enhanced User Experience

- Streamlined access to data sources, notebooks, experiment tracking, and model registry in an intuitive UI experience for data scientists.
- Improved data scientist productivity to build and operationalize machine learning (ML) models quickly.

Deprecation of KubeDirector training and deployment clusters

- The KubeDirector-based training and deployment functions have been deprecated in favor of Kubeflow-based training operator and model serving frameworks.

- Links to the **Training** and **Model Serving** pages no longer appear in the **Tenant view** UI, and Training magic functions include an annotation that KubeDirector is depreciated.
- Training magic functions are still available for use. However, a note of warning appears with information on KubeDirector deprecation and links to transition from KubeDirector to Kubeflow Training and Deployment.
- For information on transitioning applications from KubeDirector to Kubeflow Training, see Tutorial: Transition from KubeDirector to Kubeflow Training.
- If you want to continue to use the KubeDirector-based applications, you must initialize a Custom Resource (CR) for those applications. After they are provisioned, training clusters continue to appear as part of the KubeDirector Notebook cluster creation process. For information about initializing KubeDirector applications in HPE Ezmeral Runtime Enterprise 5.6.0, see Tutorial: KubeDirector Training and Serving.

Support for Apache Spark 3.2 compatible with HPE Ezmeral Data Fabric on Bare Metal 7.0

Includes secure access to read and write data from HPE Ezmeral Data Fabric based on the user identity.

Support for bring-your-own (BYO) open-source Apache Spark 3.2 and 3.3 images

- The Spark operator has been enhanced to allow users to submit Spark applications referencing any open-source Spark images.
- The Spark operator wraps enterprise security features, such as authenticating users against their AD/LDAP credentials and checking user Role-based access control (RBACs), to allow or deny executing Spark jobs on the cluster.
- For more information, see Spark Operator.

Accelerated GPU support through the NVIDIA RAPIDs plugin

- Users can now provision and consume GPUs to accelerate Spark SQL workloads.
- Users can bring their own open-source Spark images with built-in RAPIDS support or use the Apache Spark image with the built-in NVIDIA RAPIDS plugin for GPU, provided by Hewlett Packard Enterprise.
- For more information about Spark and GPU support, see Nvidia Spark-RAPIDS Accelerator for Spark.

Kubernetes Component Upgrades without Platform Software Upgrades

Beginning with HPE Ezmeral Runtime Enterprise 5.5.0 release, Kubernetes bundles enable Platform Administrators to upgrade the following Kubernetes-related components without being required to upgrade the HPE Ezmeral Runtime Enterprise platform software:

- The versions of Kubernetes that are available on the deployment.
- The versions of Kubernetes add-on components.
- The following processes are not **c**hanged:
  - There is no change in the upgrade procedure of an existing Kubernetes cluster.
  - If the Kubernetes version update depends on changes to the HPE Ezmeral Runtime Enterprise software, the Platform Administrator must upgrade HPE Ezmeral Runtime Enterprise software to a compatible version.
  - Beginning with HPE Ezmeral Runtime Enterprise 5.5.0 release Kubernetes bundles package the software to support newer Kubernetes versions, updated add-ons, and software fixes. For more information about Kubernetes bundles, see Kubernetes Bundles.

UI Changes for Platform Administration Tasks

Beginning with Ezmeral Runtime Enterprise 5.5.0, UI screens that are used for platform administration have been consolidated and are grouped in the Global Settings section of the main menu:

- The **Controllers/Upgrade** screen is now the **Controllers & HA s**creen. From the **Controllers & HA** screen, Platform Administrators can enable Platform High Availability (HA) and add the hosts that will become the Shadow Controller and Arbiter to the deployment.

- The **System Settings** has an **Updates** tab from which you manage the HPE Ezmeral Runtime Enterprise software and the Kubernetes bundles. Platform Administrators can view information about the installed software and initiate and manage the installation of controller software updates.
- The **User Management** screen now has a **Site Admin** tab, which is for managing users with Platform Administrator rights (the Site Admin role).

For more information on the new features of HPE Ezmeral Runtime Enterprise 5.6, see What's New in Version 5.6.x. https://docs.containerplatform.hpe.com/56/reference/Whats_New_in_Version.html.

Figure 3 represents a block diagram of the Kubernetes architecture within the HPE Ezmeral Runtime Enterprise.

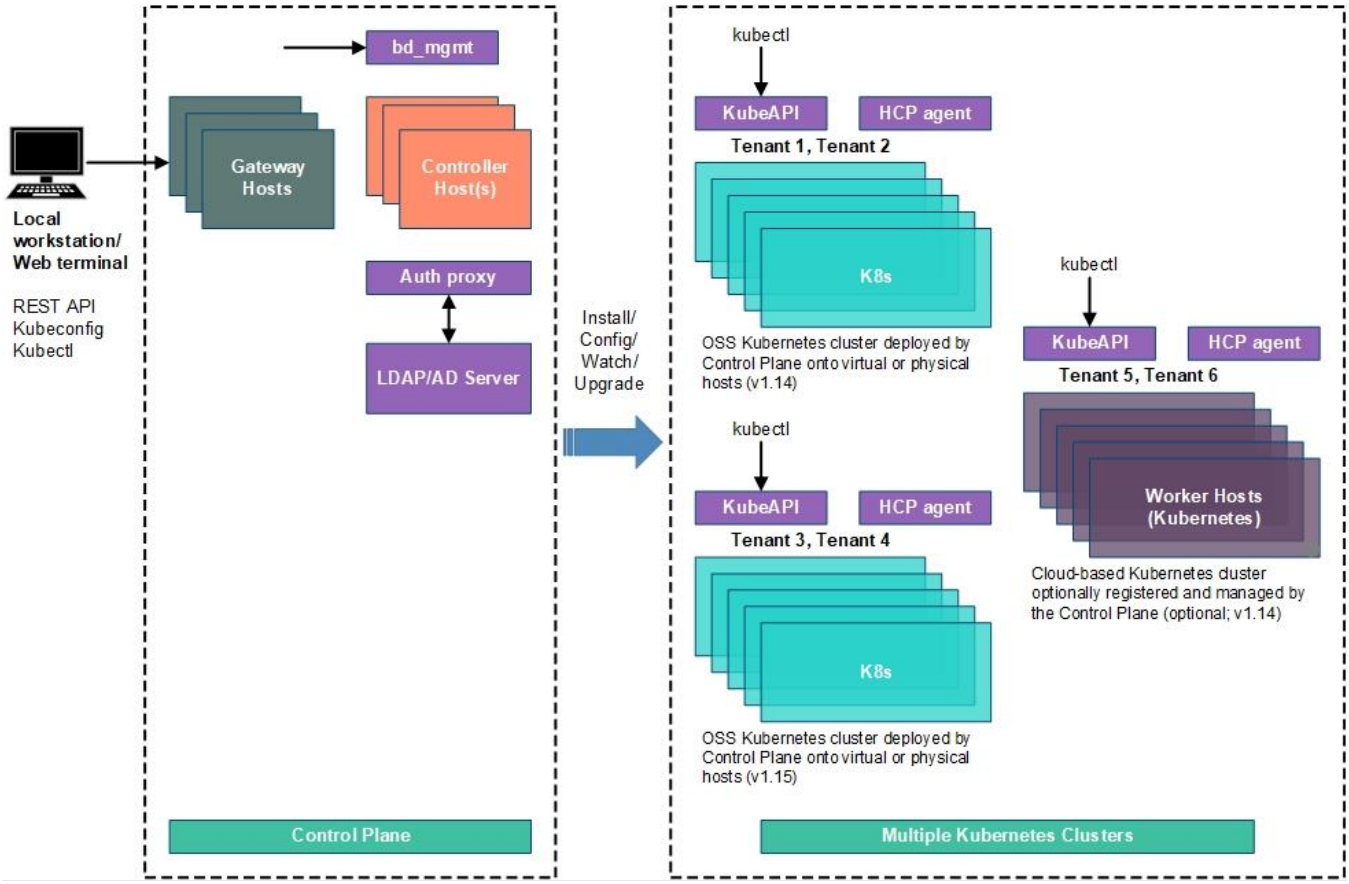

FIGURE 3. Kubernetes architecture within the HPE Ezmeral Runtime Enterprise

HPE Ezmeral Data Fabric

By default, the HPE Ezmeral Runtime Enterprise installs and configure the persistent data fabric for AI and analytics and K8s workload. This data fabric is a distributed file and object store that manages structured and unstructured data. It is designed to store data at an Exabyte scale, support multiple files, and combine analytics and operations into a single platform. It supports industry-standard protocols and application programming interfaces (APIs), including Portable Operating System Interface (POSIX), Network File System (NFS), S3, and Hadoop Distributed File System (HDFS). With production-ready capabilities like policy-based data tiering, consistent snapshots, and mirroring, the persistent data fabric serves as the enterprise standard for meeting stringent storage and processing service-level agreements (SLAs) across on-premises, hybrid cloud, and edge deployment.

All applications running in containers will be able to natively access data across the fabric through DataTaps and FSMount. K8s persistent volumes will be seamlessly available across clusters from this constant data fabric.

The data fabric provides pre-integrated, scale-out, and edge-ready persistent storage along with data services. The unique features provided by HPE Ezmeral Data Fabric are as follows:

- Enterprise data persistence with fast, flexible, and consistent data access for multiple tenants leveraging a global namespace and supports multi-protocol access.
- Auto-tiering enables an effortless data scale and data tiering across a hybrid cloud environment.
- Secure and portable data access, controlled data, and app mobility from the core to cloud to edge using a common security and governance model.
- Deployed clusters across geographies can be accessed via a single, global namespace.
- 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 creating another copy/silo of data.
- High Availability/Resiliency/Disaster Recovery capabilities for mission-critical deployment through automatic services failover, container re-replication, and mirroring.
- Single distributed data fabric that can store files, tables, and message topics with data portability across nine (9) industry-standard APIs.
- Multitenancy to support a range of application types on a single platform.

HPE Ezmeral Data Fabric on Bare Metal

HPE Ezmeral Data Fabric on Bare Metal is an implementation of HPE Ezmeral Data Fabric on physical or virtual machines that are not part of the HPE Ezmeral Runtime Enterprise deployment. HPE Ezmeral Data Fabric on Bare Metal is the supported implementation of HPE Ezmeral Data Fabric for HPE Ezmeral Runtime Enterprise production deployments. An HPE Ezmeral Data Fabric on Bare Metal cluster is external to the HPE Ezmeral Runtime Enterprise installation. After you have installed HPE Ezmeral Runtime Enterprise 5.6.0 or later, you can register the HPE Ezmeral Data Fabric on Bare Metal cluster.

Figure 4 shows the high-level components of the HPE Ezmeral Data Fabric on Bare Metal.

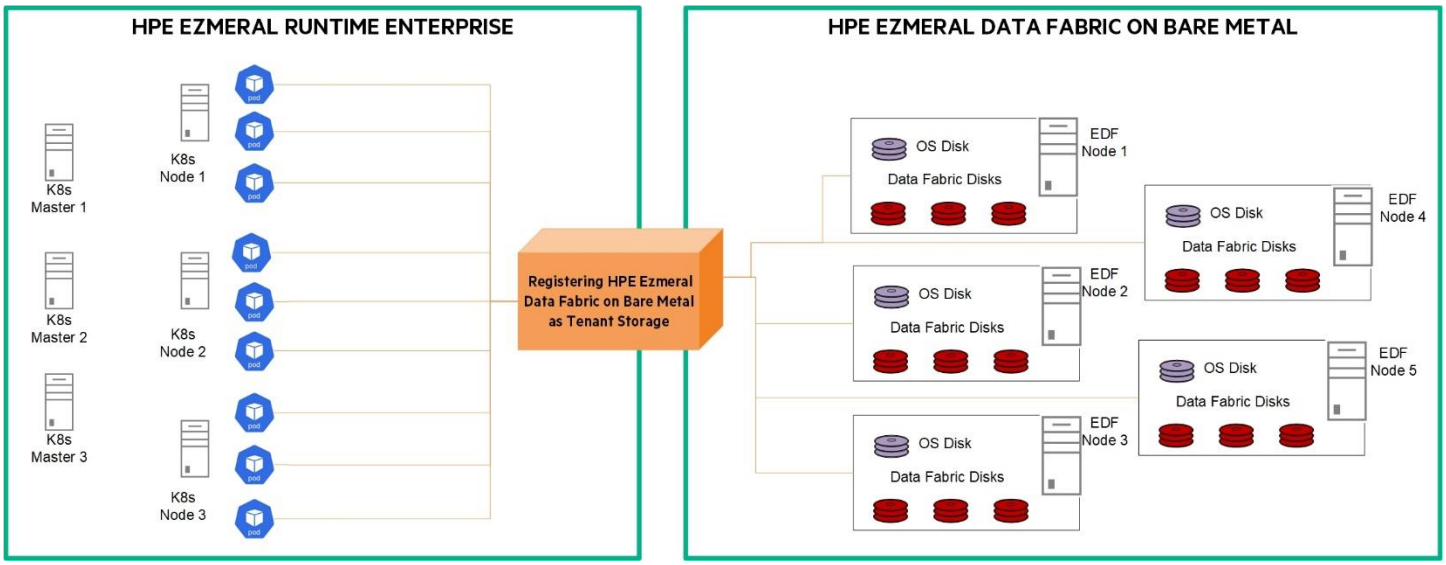

FIGURE 4. High-level architecture of HPE Ezmeral Data Fabric on Bare Metal

KubeDirector

The HPE Ezmeral Runtime Enterprise uses standard Kubernetes (K8s) facilities of custom resources and API extensions to implement a stateful scale-out of application clusters. This approach enables transparent integration with K8s user or resource management and existing K8s clients and tools.

In broad terms, KubeDirector is a custom controller (deployed into K8s) that waits for custom resources of a given type to be created or modified within some K8s namespace(s). On such an event, KubeDirector uses K8s APIs to create or update the resources and configuration of a cluster to follow the specification defined in the custom resource.

Unlike some other custom controller implementations, KubeDirector does not tie a custom resource definition to an application or contain hardcoded application-specific logic within the controller. The application characteristics are instead defined by metadata and an associated package of configuration artifacts.

Figure 5 shows the diagram representation of the KubeDirector custom controller.

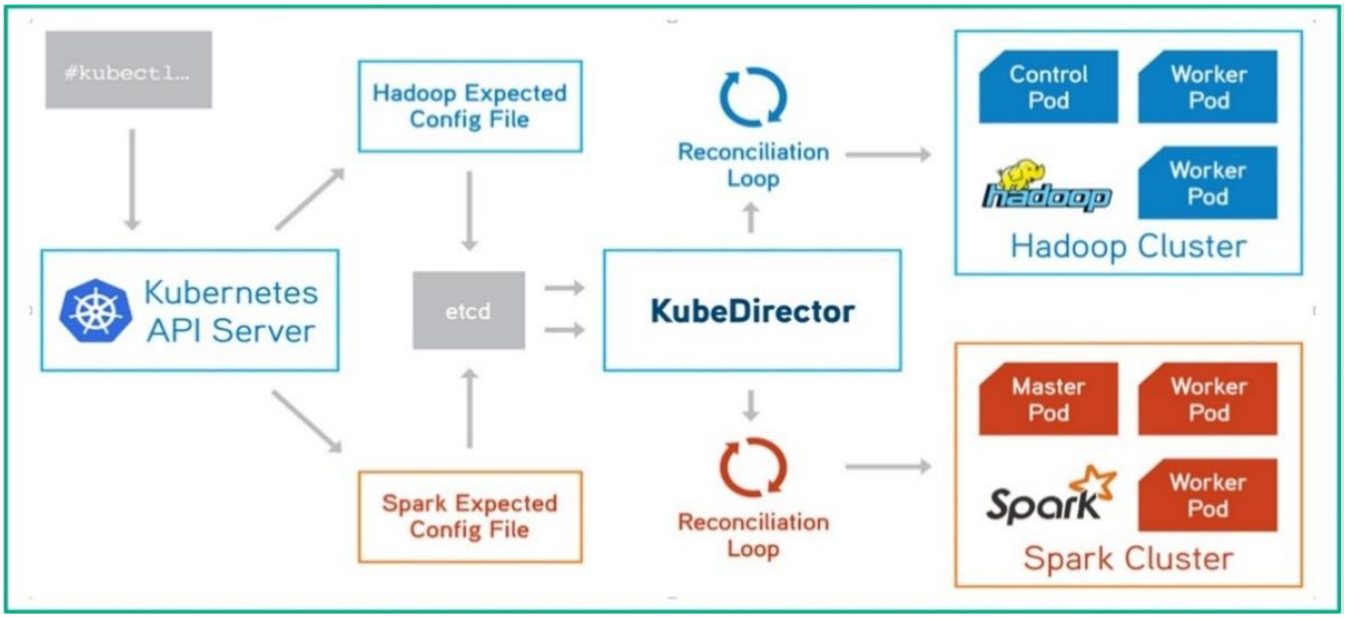

FIGURE 5. KubeDirector custom controller

The characteristics of the separation of responsibilities include the following:

- Application experts within or outside the organization, running the KubeDirector can enable application deployment without writing Go code or understanding the operation of custom controllers. This includes making easy incremental changes to adopt new versions of an application or tweak the setup choices exposed to the end-user.
- Site administrators can easily manage the application types and versions available within an organization. This is done without undergoing a custom controller code upgrade that could potentially disrupt operations.
- End-users can launch and reconfigure clusters using familiar K8s tools, selecting from application-specific choices provided to them by the experts.

The preconfigured KubeDirector operator deploys KubeDirector kind of apps. This is a generic operator that allows onboarding a variety of applications by providing roles, images per role, services, and ports to be exposed in a simple JSON format.

This solution describes the way of modernizing a legacy application using KubeDirector.

HPE ProLiant DL Gen11 servers

HPE ProLiant Compute engineered for your hybrid world to unlock more value from your data, power insights and innovation across the edge to the cloud.

HPE ProLiant DL325 Gen11 server

The HPE ProLiant DL325 Gen11 server is a low-cost 1U 1P solution that delivers exceptional value balancing compute, memory and network bandwidth at 1P economics. Powered by 4th Generation AMD EPYC™ Processors with up to 96 cores, increased memory bandwidth (up to 3 TB), high-speed PCIe Gen5 I/O and EDSFF storage, and supporting up to 2 GPUs at the front, this server is a superb low-cost, 1U 1P, performance solution for your virtualized workloads. The silicon root of trust anchors the server firmware, creating a fingerprint for the AMD Secure Processor that must be matched precisely before the server will boot. The HPE ProLiant DL325 Gen11 server is an excellent choice for virtualized workloads such as software-defined compute, CDN, VDI, and secure edge apps that require balancing processor, memory, and network bandwidth.

FIGURE 6. HPE ProLiant DL325 Gen11 server

HPE ProLiant DL385 Gen11 server

The new HPE ProLiant DL385 Gen11 server is an accelerator-optimized 2U 2P solution that delivers exceptional compute performance, upgraded high-speed data transfer rate and memory depth at 2P compute capability. Powered by 4th Generation AMD EPYC™ 9004 Series Processors with up to 96 cores, increased memory bandwidth (up to 6TB), high-speed PCIe Gen5 I/O, Gen5 EDSFF storage and the newly designed chassis supporting 8 single wide (SW) or 4 double wide (DW) GPUs The HPE ProLiant DL385 Gen11 server is a perfect accelerator-optimized 2U 2P solution.

FIGURE 7. HPE ProLiant DL385 Gen11 server

For more information on automation script, see https://github.com/hewlettpackard/hpe-solutions-hpecp.

SOLUTION ARCHITECTURE

HPE Ezmeral Runtime Enterprise can be used to create and manage multiple containerized clusters (Kubernetes-based or HPE Ezmeral Runtime Enterprise-based), either on-premises or in the public cloud. The external K8s cluster management feature is available for the HPE Ezmeral Runtime Enterprise. The on-premises cluster includes highly available HPE Ezmeral Runtime Enterprise controllers, HPE Ezmeral Runtime Enterprise gateway servers, and Kubernetes clusters with master and worker nodes. The persistent volume for the Kubernetes clusters is provided using HPE Ezmeral Data Fabric on Bare Metal.

The environment described in this Reference Configuration includes the following main functional components (the details are described in the subsequent sections of this document).

- **HPE Ezmeral Runtime Enterprise Controller:** Installed and configured on 3x HPE ProLiant DL325 Gen11 servers with multiple local disks set up as a raid for platform-level high availability functionality that protects HPE Ezmeral Runtime Enterprise against the failure of the Controller host. Platform-level high availability requires two designated Ezmeral worker hosts, such as Shadow Controller and Arbiter.
- **HPE Ezmeral Data Fabric on Bare Metal:** Consists of a Bare Metal 5x HPE ProLiant DL385 Gen11 Server Compute Module as Ezmeral Data Fabric nodes the deployment of Kubernetes clusters is achieved using the proven HPE Ezmeral Runtime Enterprise software.
- **Persistent Storage:** Persistent volume for the cluster is available from the native HPE Ezmeral Data Fabric on Bare Metal.
- **HPE Ezmeral Runtime Enterprise Gateway:** Gateway plays an important role as it provides a connection to HPE Ezmeral Runtime Enterprise- managed Kubernetes cluster services. All public service endpoints in KubeDirector- managed Kubernetes clusters will be exposed via HPE Ezmeral Runtime Enterprise Gateway re-mapped ports. Redundancy for Gateway hosts is provided by creating two sets of gateway servers. This is done by mapping multiple Gateway host IP addresses to a single hostname. This ensures that there is no single point of failure for the Gateway host. When this is done, either the DNS server or an external load balancer will load-balance requests to the hostname among all the Gateway hosts on a round-robin, least connection, etc. basis.

Figure 8 shows the layout of the HPE Ezmeral Runtime Enterprise deployment on HPE ProLiant DL Gen11 servers – without the K8s Compute Cluster nodes.

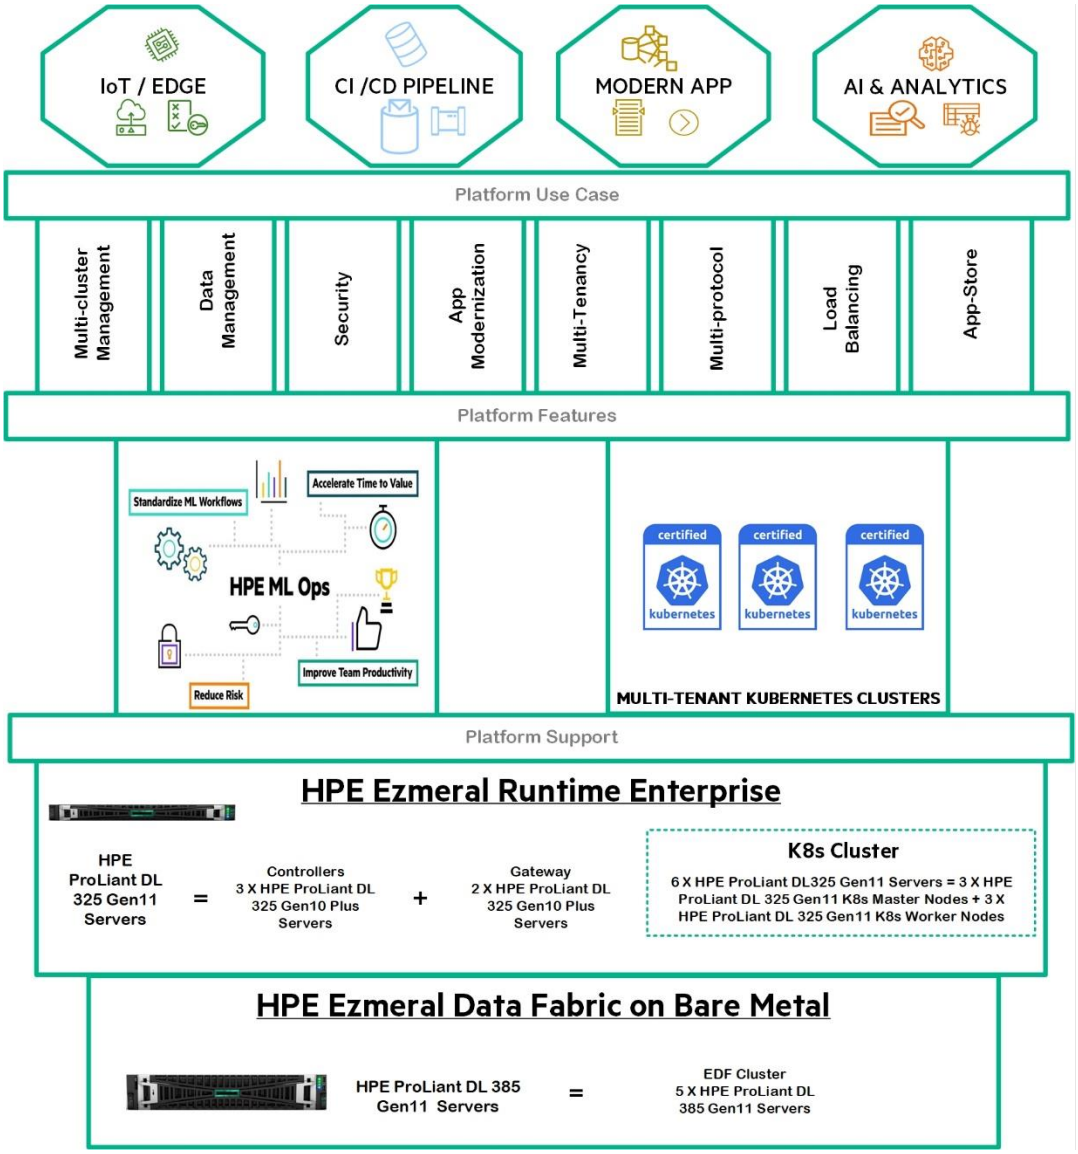

FIGURE 8. Solution layout of HPE Ezmeral Runtime Enterprise on HPE ProLiant DL Gen11 servers

NOTE

The HPE Ezmeral Runtime Enterprise Gateway discussed in this Reference Configuration is Bare Metal deployment.

The HPE Ezmeral Runtime Enterprise environment is built around HPE ProLiant DL Gen11 servers using a deployment setup consisting of Sixteen (16) HPE ProLiant DL Gen11 servers: three (3) HPE ProLiant DL325 Gen11 servers are deployed as HPE Ezmeral Runtime Enterprise Controllers in a highly available configuration, two (2) HPE ProLiant DL325 Gen11 servers are deployed as gateway Load Balancer server, three (3) HPE ProLiant DL325 Gen11 servers along with five (5) HPE ProLiant DL385 Gen11 servers for the EDF cluster nodes that are used to create the HPE Ezmeral Data Fabric on Bar Metal and finally six (6) HPE ProLiant DL325 Gen11 servers are deployed as Kubernetes compute cluster worker nodes and master nodes. Local Disks on the servers are used for providing storage for Ephemeral Disks and Persistent Volume in Kubernetes Clusters.

Figure 9 shows the physical layout of the server, storage, and network used for the solution.

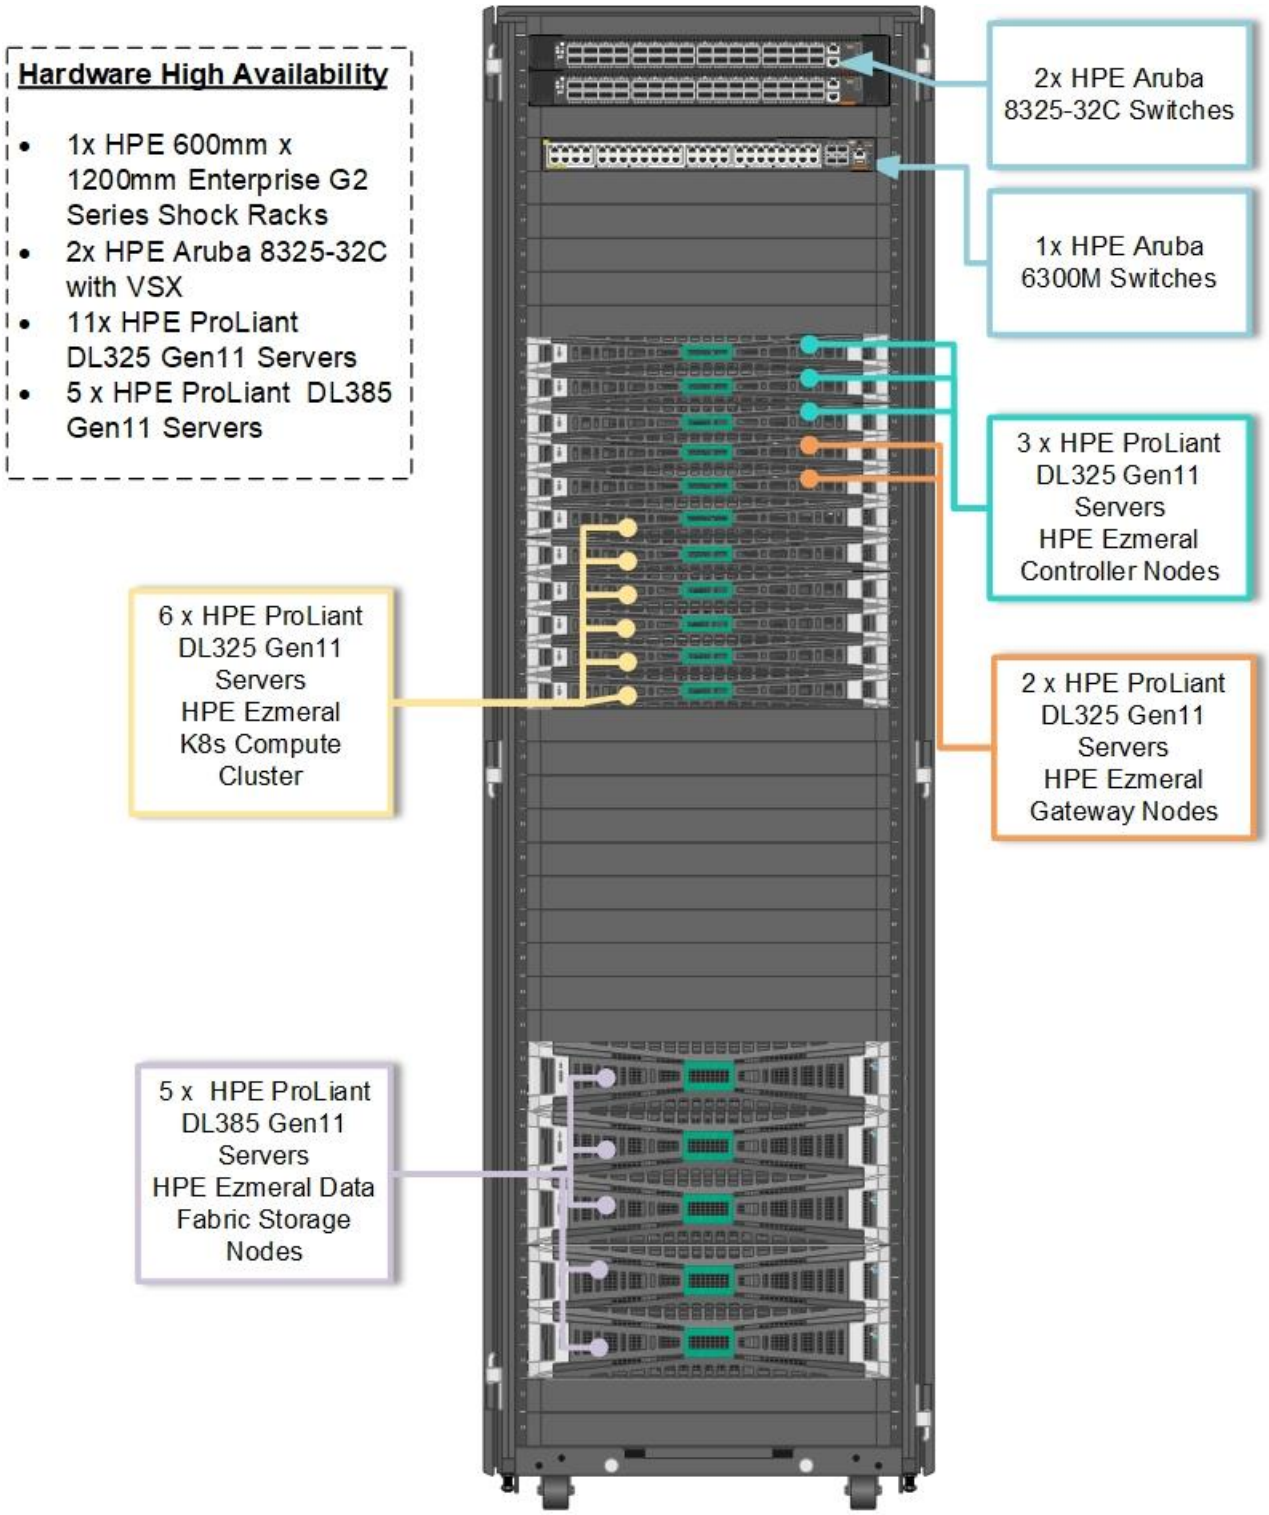

FIGURE 9. Front view of 16x HPE ProLiant DL325 Gen11 and HPE ProLiant DL385 Gen11 servers in the solution

SOLUTION COMPONENTS

This solution describes the deployment of the HPE Ezmeral Runtime Enterprise, the creation of a Kubernetes cluster with the native HPE Ezmeral Data Fabric on Bare Metal, and an app modernization use case using KubeDirector. HPE Ezmeral Runtime Enterprise Controllers, Gateway servers, and Kubernetes cluster nodes are configured to achieve high availability and are described as part of this Reference Configuration. The HPE Container Storage integration driver gets installed in an automated fashion as a part of the cluster installation.

Hardware

Table 1 lists the hardware components utilized in the design of this solution.

TABLE 1. Components utilized in this solution

Component	Qty	Description
HPE ProLiant DL325 Gen11 servers	11	Three (3) for HPE Ezmeral Runtime Enterprise Controller + Two (2) for HPE Ezmeral Runtime Enterprise Gateway + Six (6) for HPE Ezmeral Runtime Enterprise K8s Compute Cluster
HPE ProLiant DL385 Gen11 servers	5	Five (5) HPE Ezmeral Data Fabric K8 Worker Nodes
HPE Aruba 6300M	1	OOBM Switch
HPE Aruba 8325-32C BF	2	Each switch contains either 40/100GbE (QSFP+/QSFP28) or 1/10/25GbE (SFP/SFP+/SFP28)

Software

Table 2 shows the software that is used in this solution configuration.

TABLE 2. Software used in this solution

Component	Version	Description
HPE Ezmeral Runtime Enterprise	5.6	Control Plane deployed on three server nodes
SUSE Linux Enterprise Server	15 SP3	OS for all compute modules with SP1 update minimum
HPE Gateway Load Balancer	5.6	Configured automatically by HPE Ezmeral Runtime Enterprise using two server nodes
Kubernetes	1..22,123	Kubernetes version installed as a part of the solution
Docker	19.03.5	Docker server and client version
HPE Ezmeral Data Fabric on Bare Metal	7.0	HPE Ezmeral Data Fabric on Bare Metal for tenant persistent volume

Operating system installation

Installs the SUSE Linux Enterprise Server 15 SP4 and ensures that the system has appropriate licensing in place using SUSE subscription. Installs the latest patches using zypper and ensures that the prerequisites have been met for each server.

Disk requirement

A minimum of two (2) raw disks of 500GB or higher in addition to the OS disk is required. Minimum one (1) 1TB raw disk is for Docker storage (ephemeral storage). These disks are provided as a local disk.

The operating system installation and prerequisites for the deployment are automated as part of this solution. To view the playbooks to deploy operating systems and the prerequisites, see https://github.com/hewlettpackard/hpe-solutions-hpecp.

Switching and cabling

Figure 10 describes the cabling configuration of nineteen (19) HPE ProLiant DL Gen10 servers as well as the HPE Aruba 8325-32C BF switches and Aruba VSX, which takes advantage of the ArubaOS-CX modern architecture and delivers best in class high availability within the context of this solution.

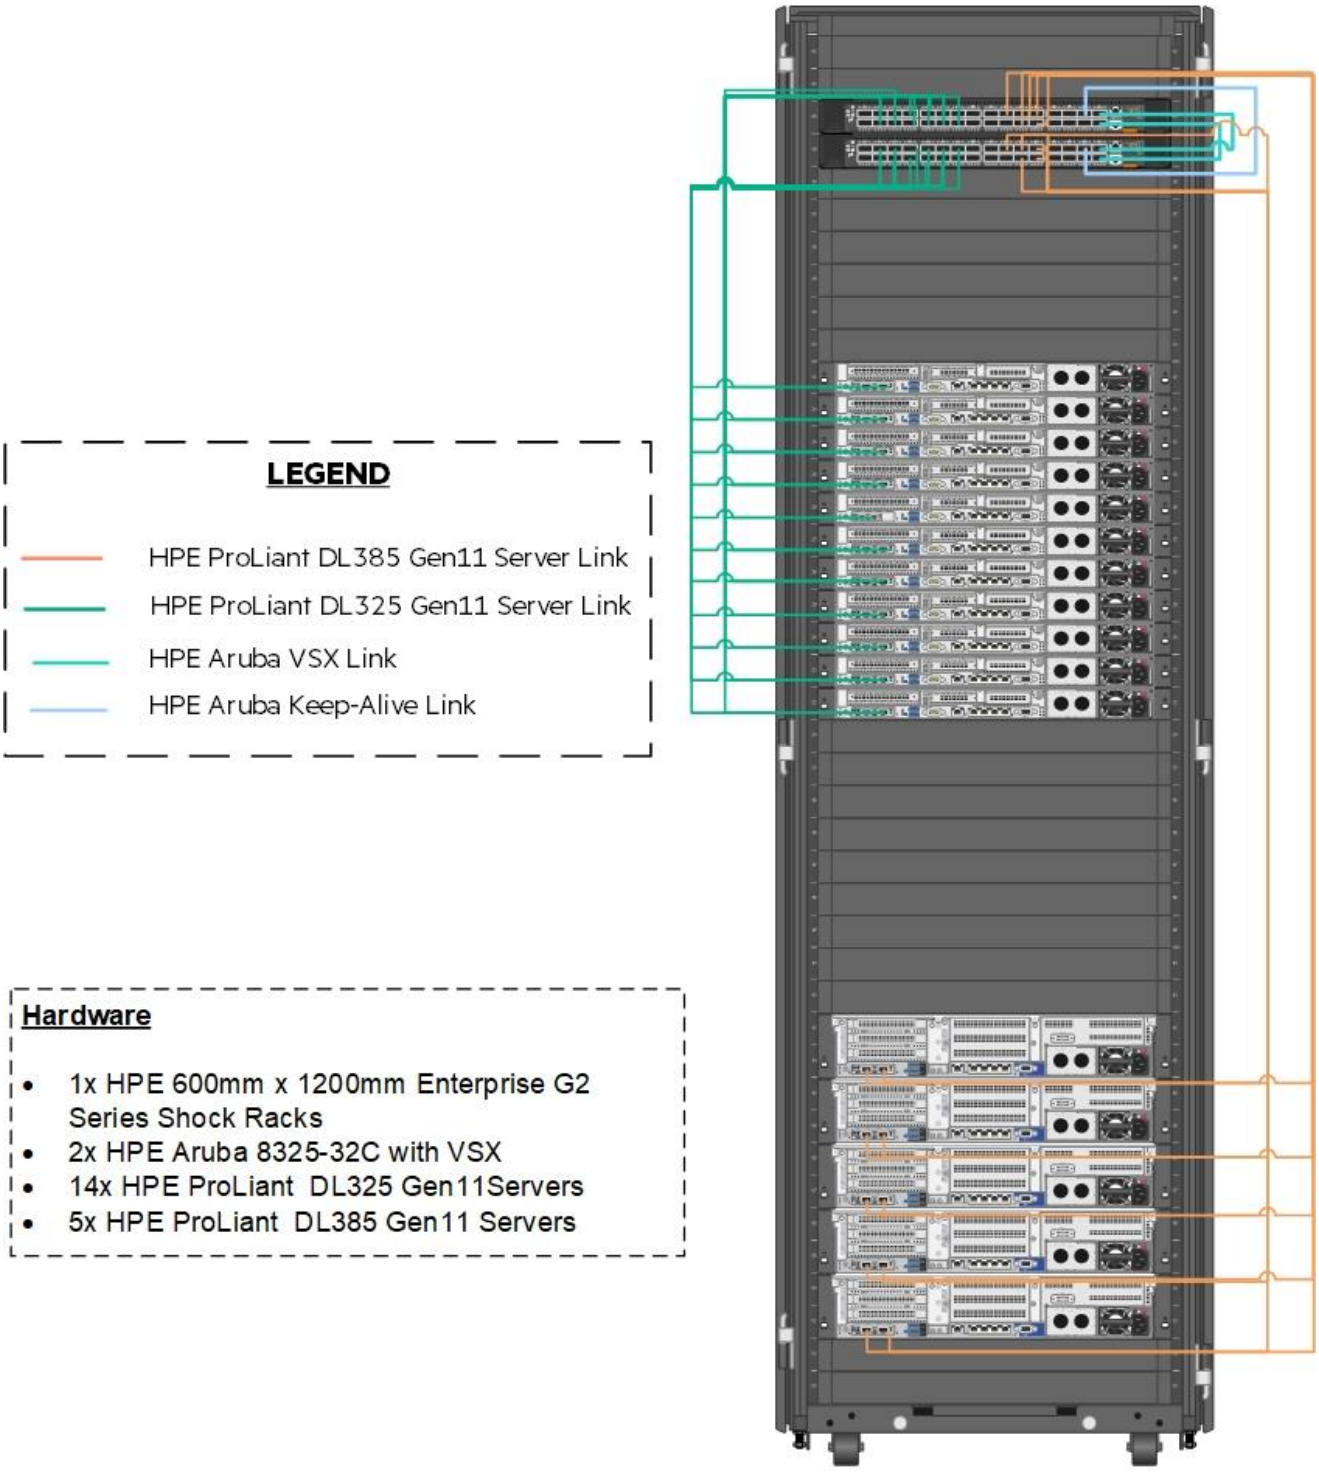

FIGURE 10. Frame and switch cabling within the solution

Storage

Kubernetes guidelines state managing storage is a distinct problem from managing compute instances. This is in line with the fundamental tenet of modern data center design that compute and storage are best considered separately from a planning and deployment perspective.

Based on the required performance either solid-state disks or hard disk drives can be used, where it is consumed by HPE Ezmeral Runtime Enterprise ephemeral disk for Docker and optionally by compute nodes as boot devices. Kubernetes Storage Cluster running on HPE ProLiant DL385 Gen11 server's local disk provides Persistent Volume (PV) for containers using the HPE EDF on Bare Metal.

Table 3 describes recommendations for the HPE Ezmeral Runtime Enterprise storage.

TABLE 3. HPE Ezmeral Runtime Enterprise storage recommendations for HPE ProLiant DL Gen10 servers for the RC

Host Type	Storage Type	Minimum Storage Recommendation for the RC
HPE Ezmeral Runtime Enterprise—Controller Shadow controller Arbiter controller	OS Node Storage	2 x 50 GB HDD configured as RAID 1 1 x 150 TB HDD configured as RAID 0
HPE Ezmeral Runtime Enterprise—Compute	OS Node Storage	2 x 50 GB HDD configured as RAID 1 1 x 1 TB HDD configured as RAID 0
HPE Ezmeral Runtime Enterprise—Gateway	OS	2 x 50 GB HDD configured as RAID 1

For more information, see Storage Partition Requirements.

Figure 11 describes the logical storage layout used in this solution.

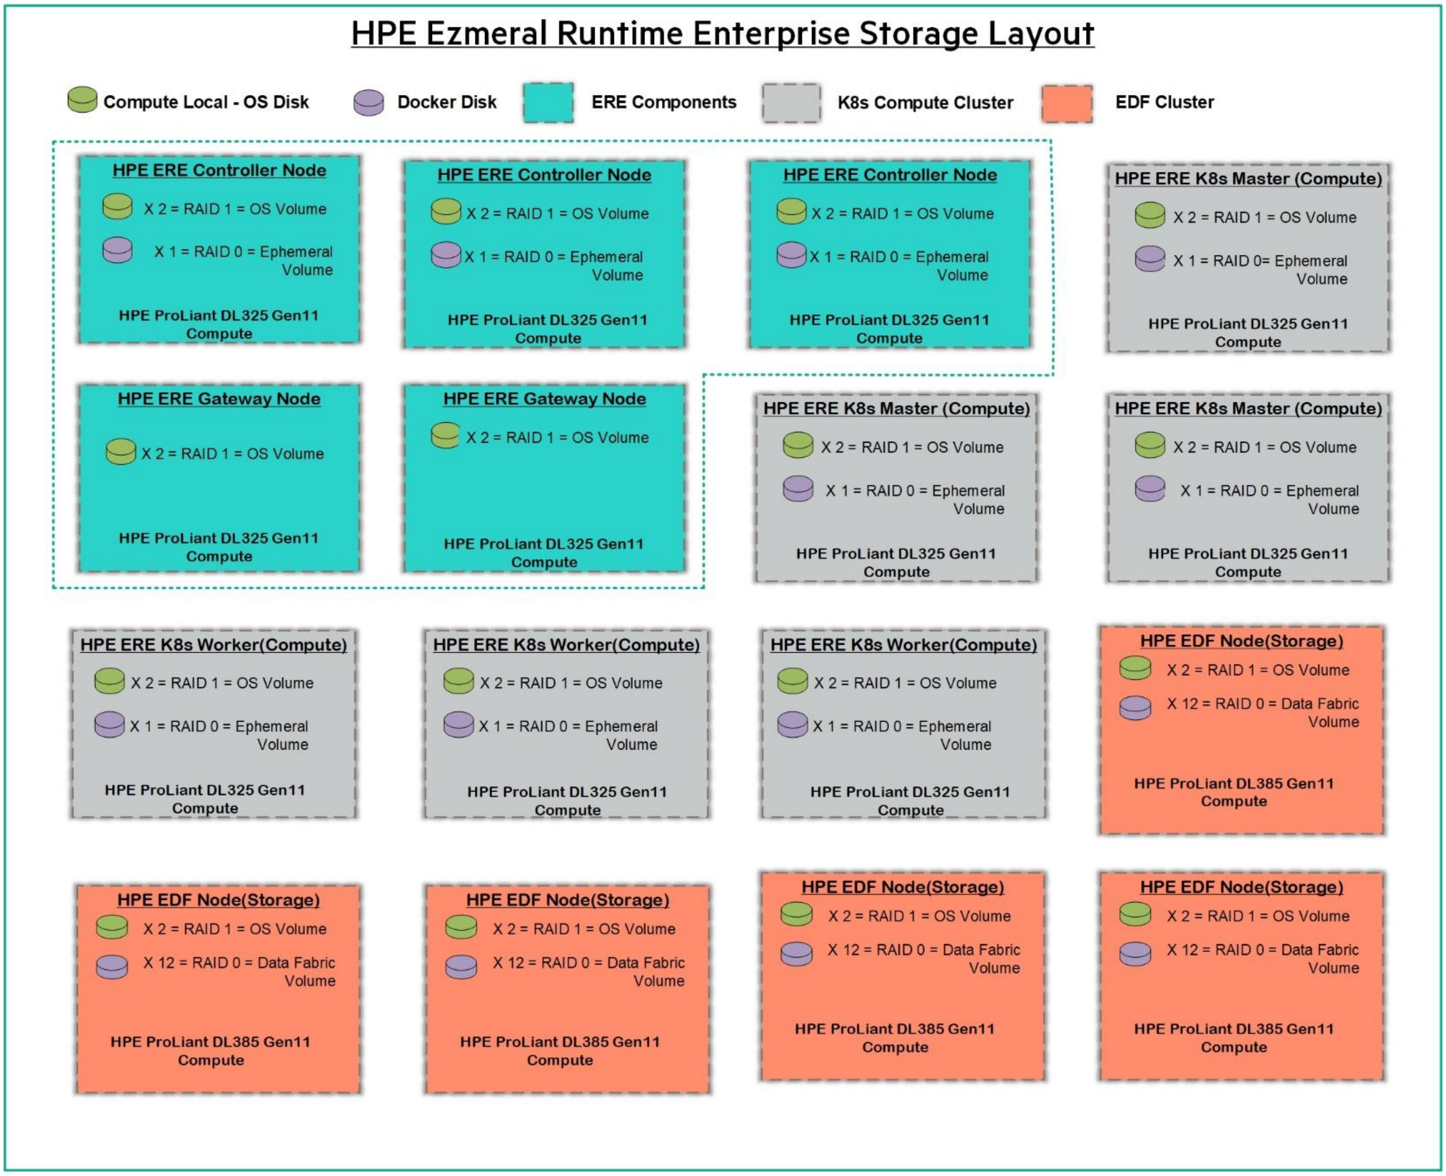

FIGURE 11. Logical storage layout used in HPE Ezmeral Runtime Enterprise

CAPACITY AND SIZING

Based on the solution deployed as part of this Reference Configuration, this section recommends the capacity for HPE Ezmeral Runtime Enterprise deployment that is dependent on the expected load in the environment. The following values are considered for defining the load:

- Number of Tenants supported
- Number of Clusters per tenant
- Number of Containers per clusters
- Resources around CPU, memory, and disk of the server

The required amounts of physical CPU, memory, and disk storage space can be computed from these values. After these are determined, they can be converted to specific server configurations for deployment.

For more information on HPE Ezmeral Runtime Enterprise Sizing, see Online Sizing Tool for Compute Capacity, HPE Sizing Tool for HPE Ezmeral Runtime Enterprise.

Hewlett Packard Enterprise has created a unique Sizer tool to assist in deploying the HPE Ezmeral Runtime Enterprise on HPE servers. The tool helps to size the user requirements quickly and easily. Users can create a customized sizer recommendation based on the vast options available on the tool's user-friendly web interface. This helps them to pick and choose various components and get an approximate cost of the configuration or BOM. To access the Sizer tool, see the following website: https://solutionsizers.ext.hpe.com/EzmeralCPSizer/.

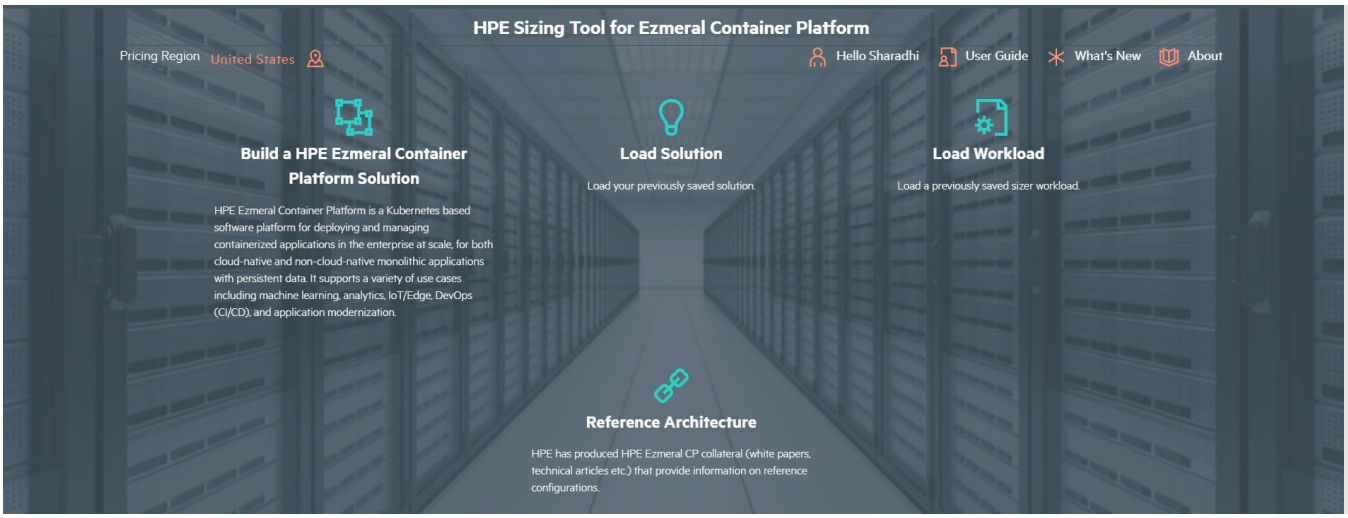

FIGURE 12. HPE Sizing Tool for HPE Ezmeral Runtime Enterprise

The HPE online sizer for HPE Ezmeral Runtime Enterprise sizes the compute tier where the HPE Ezmeral Runtime Enterprise software is running. This allows the Clusters to be deployed on servers optimized for running HPE Ezmeral Containerized service. Furthermore, the compute capacity can easily be scaled by adding more compute servers as new tenants and services are bought online without incurring the cost of scaling storage when not required.

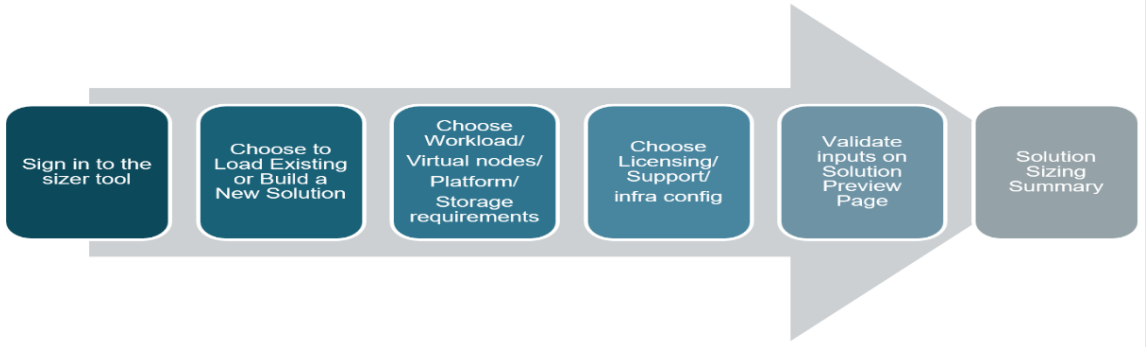

FIGURE 13. High-level workflow of HPE Sizing Tool for HPE Ezmeral Runtime Enterprise

CONNECTED ENVIRONMENT SOLUTION CONFIGURATION

This section provides high-level steps involved in the configuration of the solution for servers that have access to the internet and can download required packages and update directly.

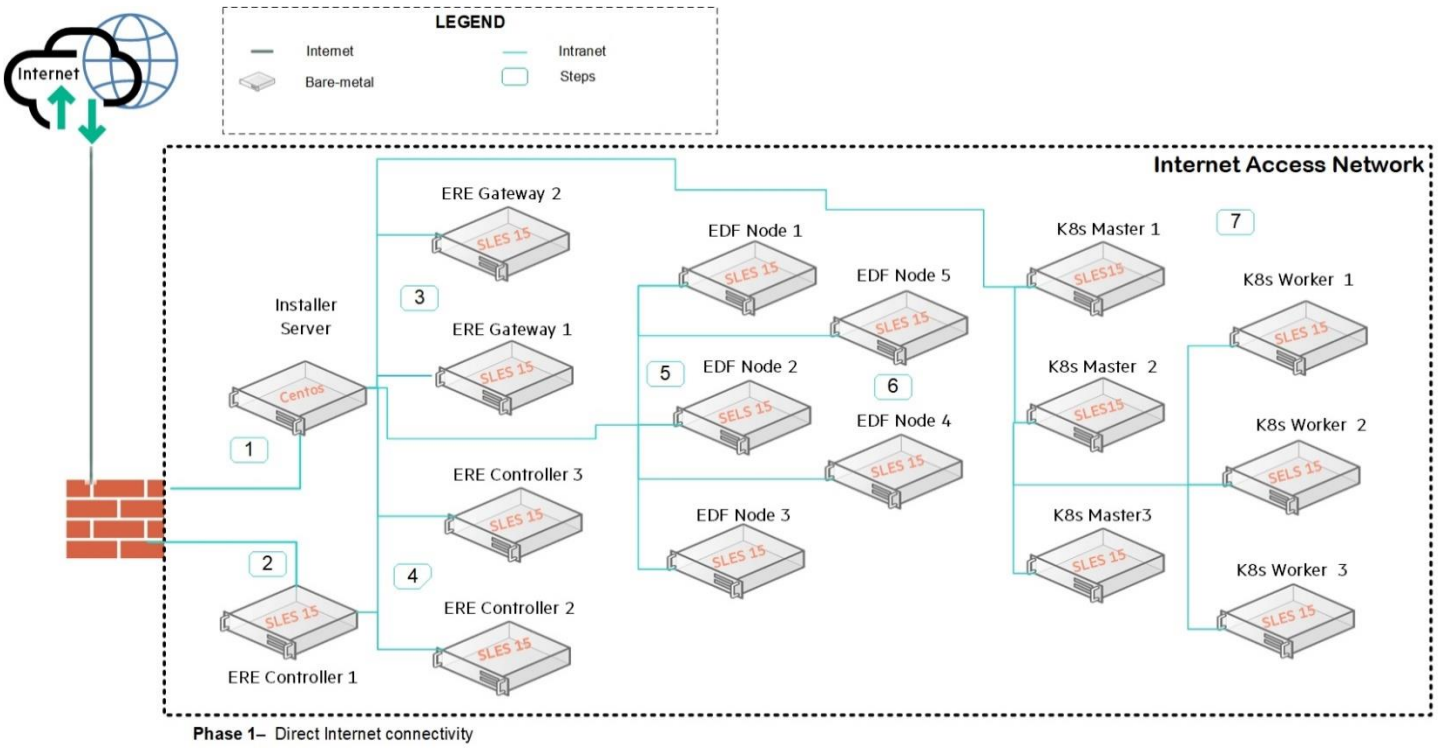

FIGURE 14. Connected environment solution layout

Pre-deployment environment configuration

Pre-deployment environment configuration is a high level two-step process that includes the following:

FIGURE 15. Pre-deployment environment configuration

To enhance the experience of deploying HPE Ezmeral Runtime Enterprise 5.6 using automation scripts even further, the Automated Installation (Automated) enables the administrator to deploy HPE Ezmeral Runtime Enterprise 5.6 using fewer playbooks. The prerequisites for the automated deployment method are dependent on the availability of iLO access to the servers. The server is connected to the application network and has the required storage available on the computes to deploy HPE Ezmeral Runtime Enterprise 5.6 successfully along with the Kubernetes cluster running HPE Ezmeral Data Fabric.

The high-level flow for the automated deployment method is as follows:

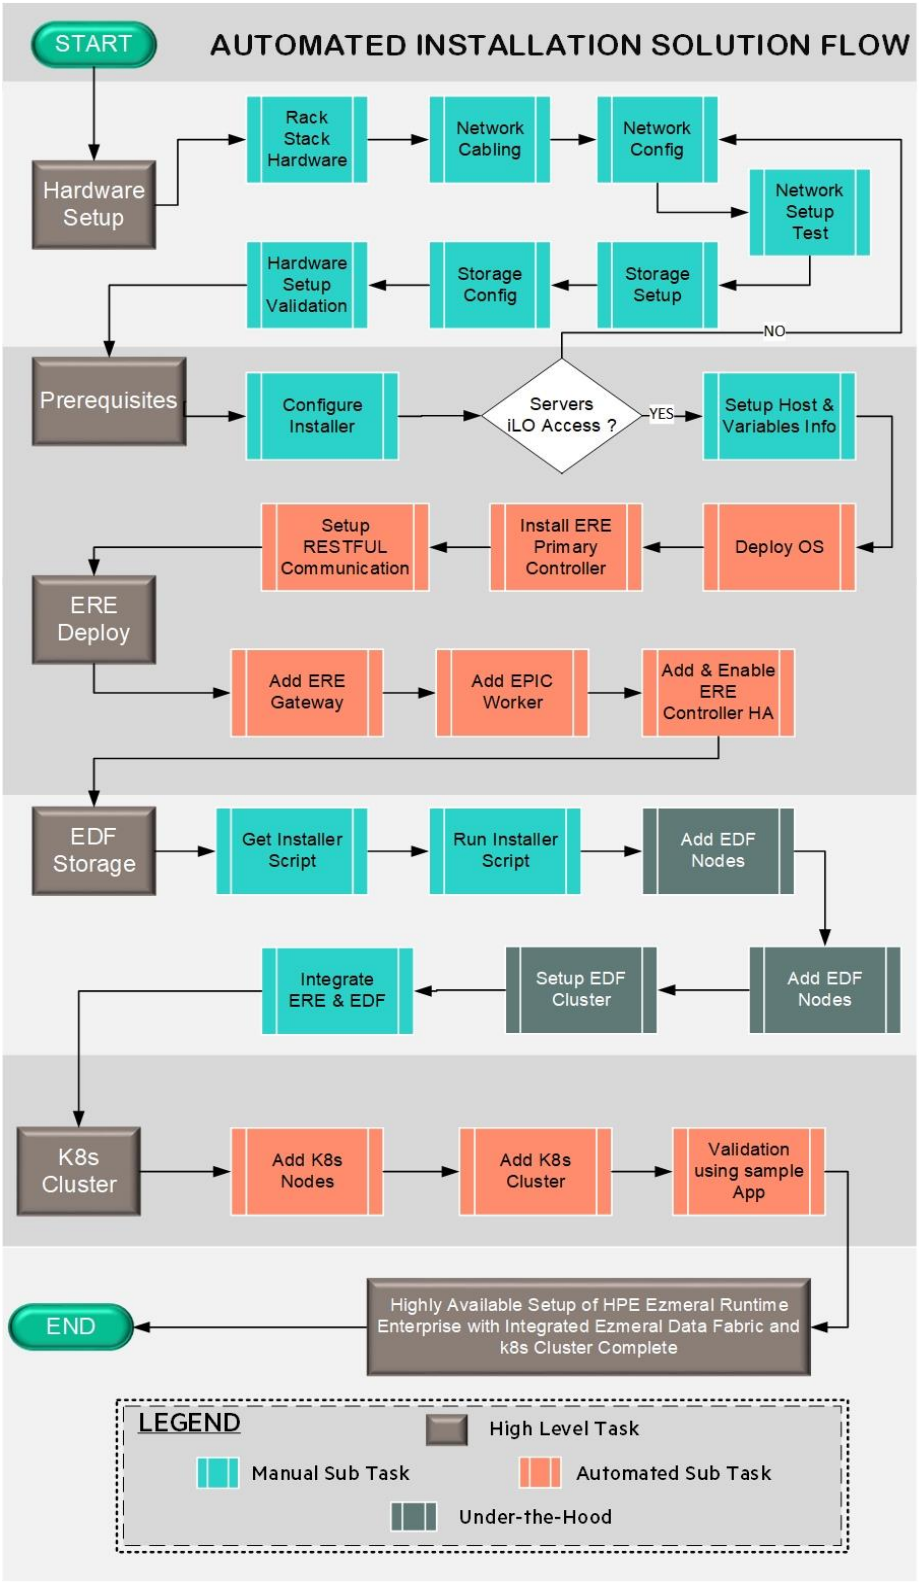

FIGURE 16. HPE Ezmeral Runtime Enterprise 5.6 Automation installation flow

The end-to-end deployment of HPE Ezmeral Runtime Enterprise 5.6 starts with the Operating System deployment on compute to install External HPE Ezmeral Data Fabric storage and registering it with the Platform for K8s Automated compute cluster usage. For more information on HPE Ezmeral Runtime Enterprise 5.6 automated method and scripts, see the following website: https://github.com/hewlettpackard/hpe-solutions-hpecp.

DISCONNECTED ENVIRONMENT SOLUTION CONFIGURATION (AIR-GAPPED INSTALL)

This section provides high-level steps involved in the configuration of the solution for servers that do not have access to the internet and download required packages and update from a repository server.

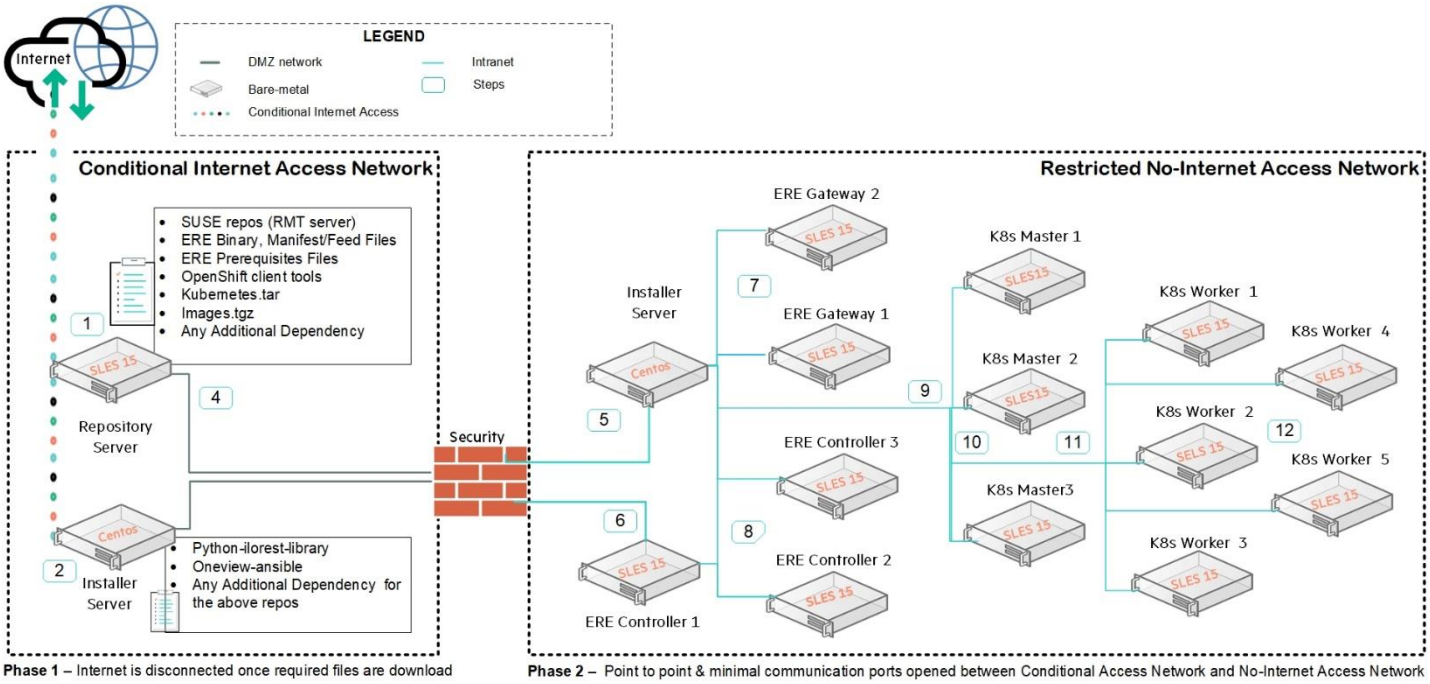

FIGURE 17. Disconnected Environment Solution layout –Airgap

Deployment configuration

Deployment environment configuration is a step-by-step process that includes the following:

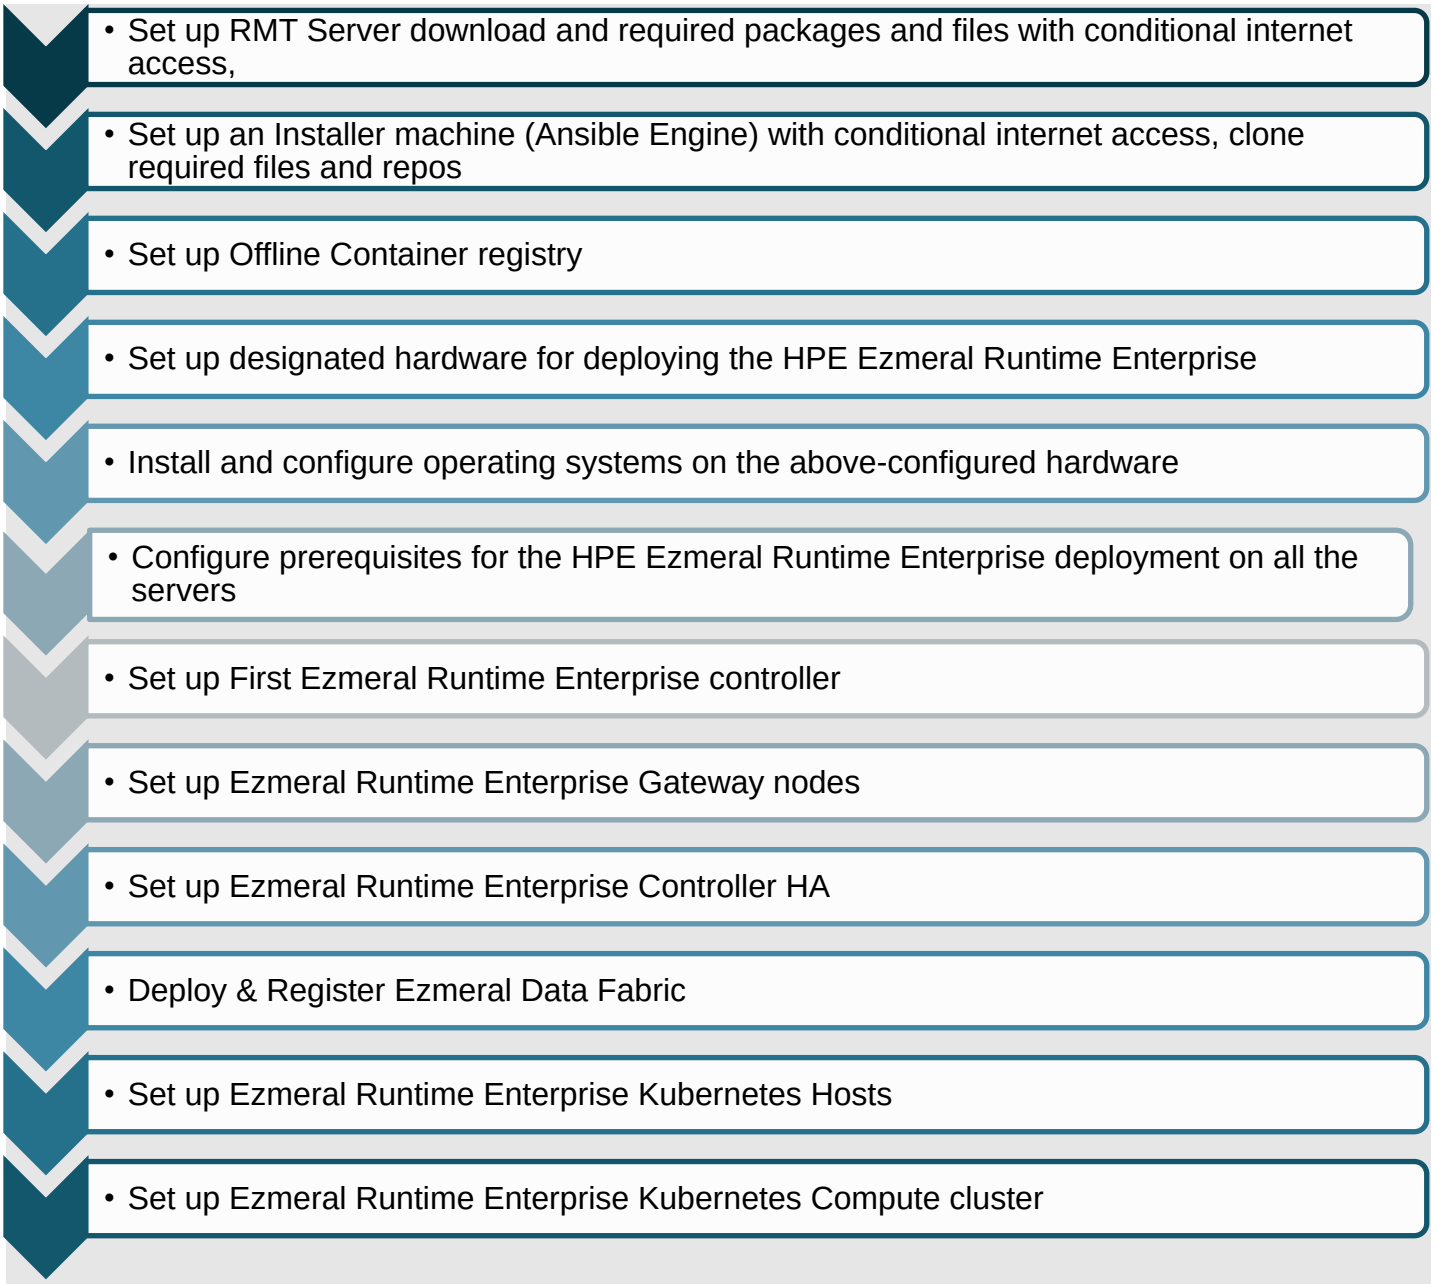

FIGURE 18. Deployment configuration

The playbooks to deploy the server profile and operating systems are available at https://github.com/hewlettpackard/hpe-solutions-hpecp. For more information, see the Deployment Guide.

GATEWAY LOAD BALANCER

Gateway servers map the services running in Ezmeral nodes or containers to ports and balance service loads as described in Gateway Hosts and Load Balancing. Add the designated gateway server(s) to the HPE Ezmeral Runtime Enterprise.

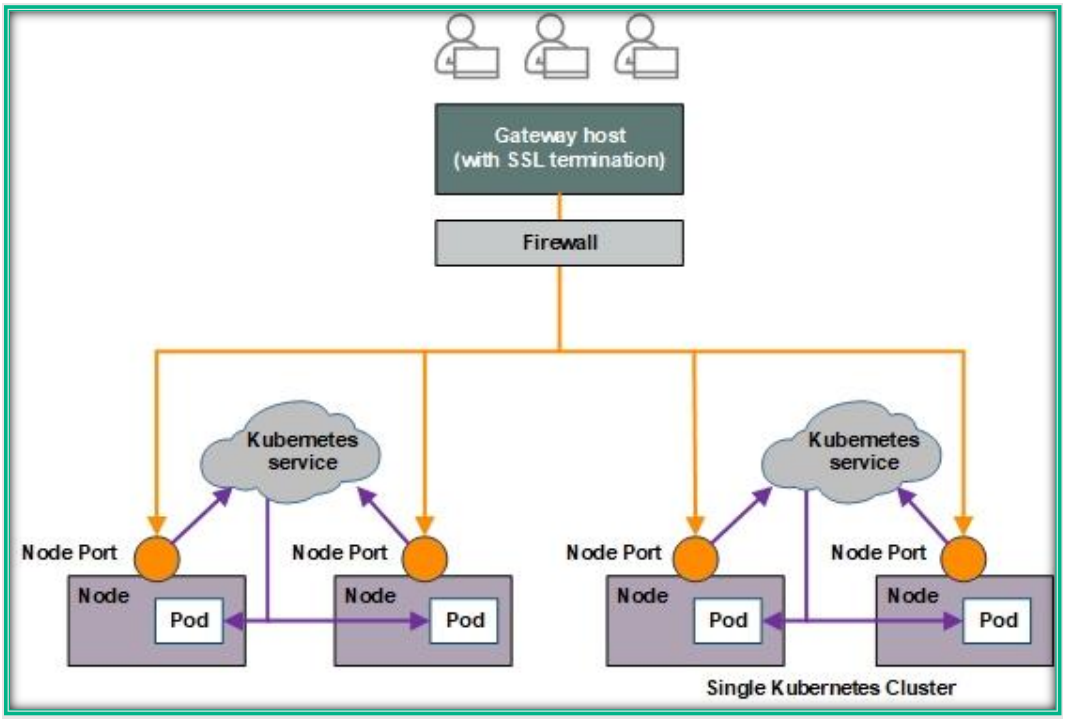

FIGURE 19. HPE Ezmeral Runtime Enterprise gateway Kubernetes orchestrator

Kubernetes orchestrator

- A simple, secure, and fully managed control path for end-users and admins alike to access Kubernetes API servers (such as when handling kubectl commands) as well as the service endpoints of multiple Kubernetes clusters. HPE Ezmeral Runtime Enterprise dynamically manages endpoints, including kubeconfig contents, as clusters and services are created, deleted, or updated.
- Automated load balancing for Kubernetes masters and services. Kube API traffic (via kubectl etc.) is load-balanced to both multi-masters—highly available Kubernetes clusters and multi-replica NodePort services.
- SSL termination to container service access points.

Gateway host high availability

Redundancy for multiple Gateway hosts can be achieved by mapping multiple Gateway host IP addresses to a single hostname. When this is done, then either the DNS server or an external load balancer will load-balance requests to the hostname among each of the Gateway hosts on a round-robin basis. This ensures that there is no single point of failure for the Gateway host.

If required existing load balancing solution such as NGINX can be used to further enhance the capability provided by the Gateway such as least connection load balancing or SSL termination.

The existing load balancing solution is shown in Figure 20 as an external load balancer.

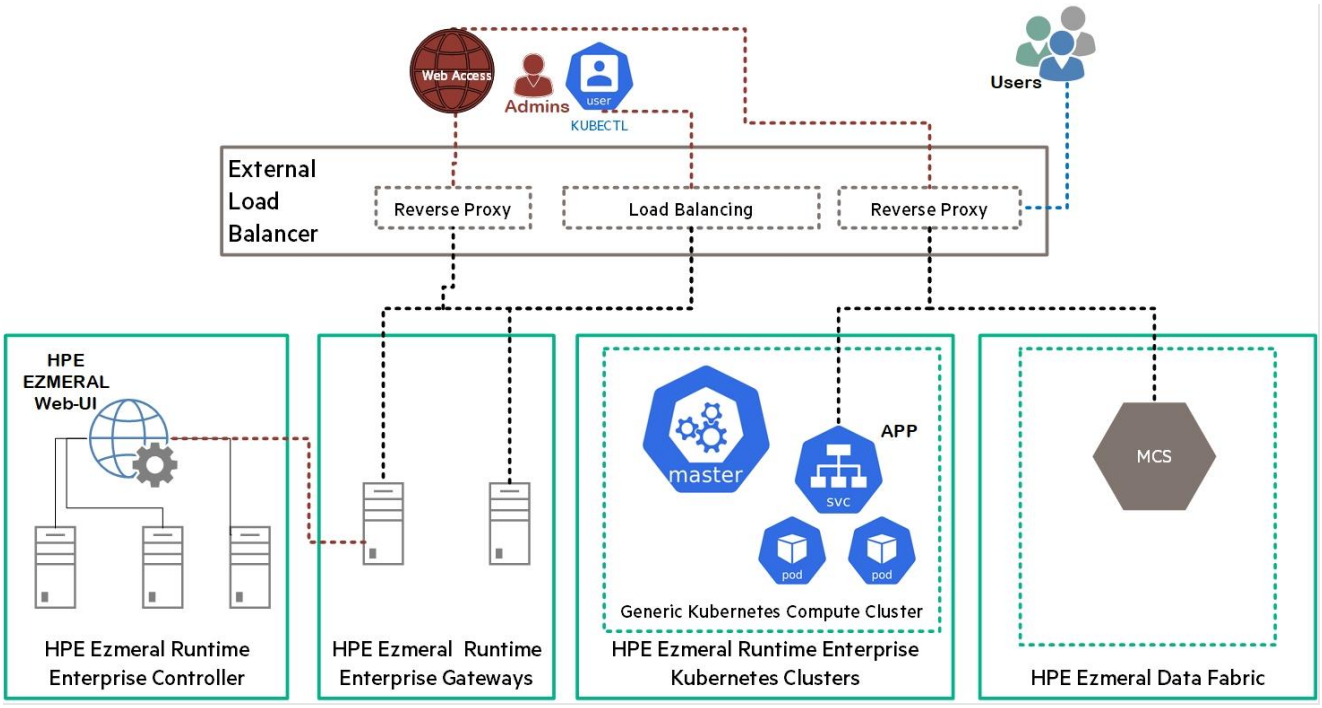

FIGURE 20. External Load Balancer for HPE Ezmeral Runtime Enterprise

TABLE 4. External Load Balancing reference for HPE Ezmeral Runtime Enterprise

Load balancer Service	TCP port/range	Description
HPE Ezmeral Runtime Enterprise Controller	80	This is for HPE Ezmeral Runtime Enterprise https Web UI access
HPE Ezmeral Runtime Enterprise Controller	443	This is for HPE Ezmeral Runtime Enterprise https Web UI access
HPE Ezmeral Runtime Enterprise Controller	8080	This is for accessing the Kubernetes Web Interface via the controller
HPE Ezmeral Data Fabric	8443	This is access to the MCS for the HPE Ezmeral Data Fabric via https
Kubernetes services	6443-10000	This is for Kubernetes Administrator management and functioning
Application services	10000-50000	This is for an end-user to access services running within the Kubernetes cluster

Load Balancing with NGINX

NGINX offers load balancing across multiple application ports and it is a commonly used technique for optimizing resource utilization, maximizing throughput, reducing latency, and ensuring fault-tolerant configurations. NGINX can be used to further enhance the capability of the HPE Ezmeral Runtime Enterprise Gateways after the gateways are configured using a common Hostname.

For more information on how to configure NGINX as an External Load balancer for HPE Ezmeral Runtime Enterprise, see. HPE EZMERAL CONTAINER PLATFORM 5.3 on HPE DL servers deployment guide

SOLUTION DIFFERENTIATOR

This Reference document is accompanied with automation playbooks to deploy the operating systems and the solution are available at https://github.com/hewlettpackard/hpe-solutions-hpecp and a deployment guide at https://hewlettpackard.github.io/hpe-solutions-hpecp/5.6-DL/. The automation script help saves time while fostering accuracy by decreasing the number of steps involved in setting up the solution. This in turn improves business productivity and promotes an Idea Economy, where success is defined by the ability to turn ideas into value faster than the competition.

Figure 21 shows the high-level reduction in steps & time when using the automation playbooks to deploy the solution. The graph depicts the graphical steps difference in forming a manual vs automated deployment of HPE Ezmeral Runtime Enterprise 5.6 environment on Bare Metal servers using scripts mentioned in this document.

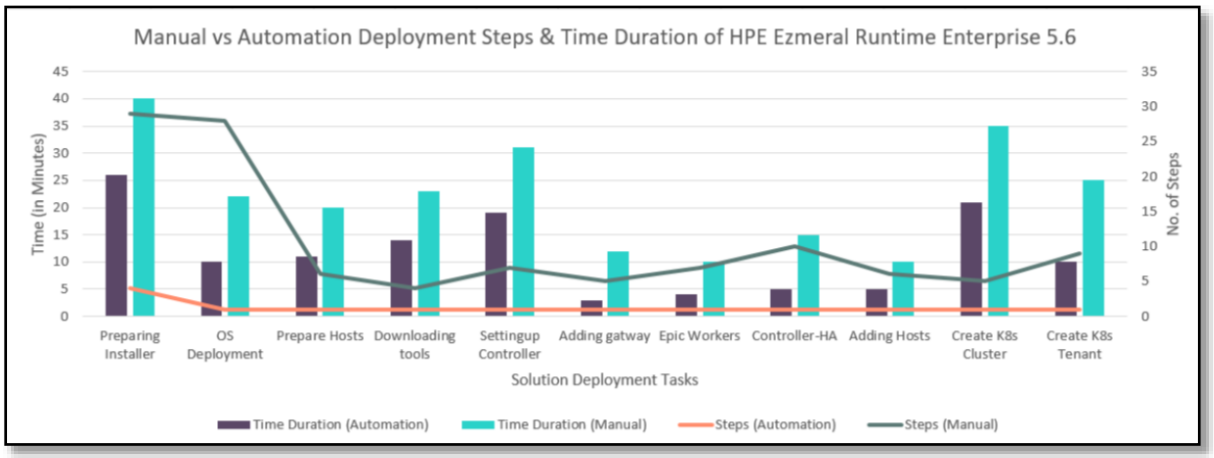

FIGURE 21. Comparison manual vs automated deployment steps of HPE Ezmeral Runtime Enterprise 5.6 on Bare Metal servers

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

This section describes the prerequisites for the solution and the best practices followed while implementing this solution.

NOTE
The solution assumes that the following infrastructure services and components are installed, configured, and function properly in the deployment environment i.e., DHCP, DNS, and NTP.

SECURITY

Role-based access control

Role-based access control (RBAC), also known as role-based security, is a framework that restricts system access. It does so by setting permissions and privileges to enable access to authorized users. Most organizations use role-based access control to provide their employees with varying levels of access based on their roles and responsibilities. This protects sensitive data and ensures employees can only access information and perform actions they need to do their duties.

Users and Roles in HPE Ezmeral Runtime Enterprise

The HPE Ezmeral Runtime Enterprise provides a multi-tenant hybrid cloud platform. This multi-tenant platform is based on role-based access control. Some of the user- and tenant-related considerations to be resolved when planning/maintaining your deployment include:

- **Tenants**: The number of tenants and the function(s) each tenant performs will determine how many users with the Tenant Administrator role will be needed and, by extension, the number of users with the Tenant Member role will be needed for each tenant. The reverse is also true since the number and functions of users needing to run jobs can influence how you define tenants. For example, different levels of confidentiality may mandate separate tenants.
- **Job functions**: The specific work performed by a given user will directly impact the role they receive. For example, a small organization may designate a single user as the Tenant Administrator for multiple tenants, while a large organization may designate multiple Tenant Administrators per tenant.
- **Security clearances**: You may need to restrict access to information based on a user's security clearance. This can impact both the tenant(s) a user has access to, and the role configured for that user within a given tenant.

Table 5 lists the role-specific acronyms that are used within the HPE Ezmeral Runtime Enterprise.

TABLE 5. HPE Ezmeral Runtime Enterprise role-specific acronyms

Acronym	Description
PLADMIN	PLADMIN stands for the Platform Administrator role.
KA	KA stands for the Kubernetes Tenant Administrator role.
KC	KC stands for the Kubernetes Cluster Administrator role.
KM.	KM stands for the Kubernetes Tenant Member (non-administrative) role.
TA	TA stands for the Big Data Tenant Administrator role.
TM	TM stands for the Big Data Tenant Member (non-administrative) role.
PA	PA stands for the AI/ML Project Administrator role.
PM	PM stands for the AI/ML Project Member (non-administrative) role.
SV	SV stands for the Site Viewer role. This is an internal role that is only visible to HPE Technical Support personnel.

For more information on the specific functions that can be performed by the above roles, see Users and Roles.

Platform security and monitoring

Sysdig Security and monitoring

To address the security challenges that exist in a containerized environment, this solution leverages the Sysdig SaaS Platform to secure and monitor the HPE Ezmeral Runtime Enterprise, an enterprise-ready Kubernetes platform that is installed and configured on HPE ProLiant DL Gen10 servers. After the HPE Ezmeral Runtime Enterprise Kubernetes cluster is deployed, the cluster access is granted to the Sysdig SaaS Platform. The Sysdig SaaS Platform is a cloud-based service where security and monitoring services will be available to the user based on the choice of their subscription. For security and monitoring of the HPE Ezmeral Runtime Enterprise, it is required to install the Sysdig Agent on the Kubernetes cluster. Sysdig Agents are lightweight entities that are installed within each node in the HPE Ezmeral Runtime Enterprise Kubernetes cluster. These agents run as a daemon to enable Sysdig Monitor and Sysdig Secure functionality. Sysdig Monitor provides deep, process-level visibility into a dynamic, distributed production environment. Sysdig Secure provides image scanning, run-time protection, and forensics to identify vulnerabilities, block threats, enforce compliance, and audit activity across an HPE Ezmeral Runtime Enterprise Kubernetes cluster.

The key benefits are as follows:

- Faster incident resolution using Sysdig Monitor for the cluster.
- Simplified compliance for the entire solution.
- Service-based access control for container security and monitoring.
- Less time spent on managing platforms, containers, and vulnerabilities.

The implementation of Sysdig in this solution uses the Software as a Service (SaaS) deployment method. The playbooks deploy Sysdig Agent software to capture the data from every node in the HPE Ezmeral Runtime Enterprise Kubernetes deployment. The captured data is relayed back to your Sysdig SaaS Cloud portal. The deployment provides access to a 90-day try-and-buy, fully featured version of the Sysdig software. For more information about Sysdig Agent deployment in the HPE Ezmeral Runtime Enterprise Kubernetes setup, see the HPE Ezmeral Runtime Enterprise on HPE ProLiant DL Gen10 servers Deployment guide.

NOTE
The Sysdig functionality is not part of this solution.

USE CASE

Application modernization with HPE Ezmeral Runtime Enterprise

Legacy application modernization enables organizations to create new business value from existing and aging applications. This is achieved by updating them with modern features and capabilities. Different strategies such as refactoring, re-architecting, rebuilding, and so on are adopted to modernize an application. Hewlett Packard Enterprise recommends starting with a lift and shift approach where the application is first containerized with as minimal changes as possible. After the application is containerized, it is easier and faster to upgrade to better infrastructure, refactor to microservices, and dynamically scale deployment. This approach enables regression testing of the application before further modifications are done.

In a lift and shift scenario, the same versions of OS, libraries, or components and applications are retained. Containers provide the application with an environment very similar to the existing environment. This paves the way to commence modernizing the application deployment while making minimal or no changes to the application itself. In this section, Hewlett Packard Enterprise describes the best practices for modernizing applications by containerizing the legacy application and deploying it on the HPE Ezmeral Runtime Enterprise using KubeDirector.

Traditional Kubernetes approach

This section describes the traditional way of deploying a stateful application on Kubernetes. The web application Methodize-Productivity is a simple productivity manager and to-do tool that uses ReactJS, ExpressJS, NodeJS, and MongoDB. The Node.js front, backend, and the MongoDB database use Persistent Volume (PV) and Persistent Volume Claims (PVCs) to store data. PV and PVCs are independent of the pod lifecycle and data are preserved even during events such as restarting, rescheduling, and deletion of pods. A Persistent Volume claim will be created at the deployment step and uses HPE Ezmeral Data Fabric as the default persistent volume.

Perform following steps to deploy a stateful application on Kubernetes:

1. Create the docker files for the three tiers.
2. Create a secret to store the credentials, if any.
3. Create the Service, PVC, and Deployment for the NodeJS frontend, backend, and MongoDB database respectively.

NOTE

The pods deployed in this approach do not make any explicit requests for a specific StorageClassName. The default MapR StorageClass is used. The NodeJS frontend YAML file is configured to use the NodePort Service instead of the load balancer service to enable the HPE Ezmeral Runtime Enterprise to re-map the ports.

Applying the YAML file creates the required resources and the description of the created service that helps you to access the endpoint for the application.

For a detailed description of the approach, see Kubernetes Cluster Usage Examples.

Deploying a stateful application using KubeDirector

KubeDirector is an open-source project designed to make it easy to run complex stateful scale-out application clusters on Kubernetes. KubeDirector is built using the custom resource definition (CRD) framework and leverages the native Kubernetes API extensions and design philosophy. This enables transparent integration with Kubernetes user or resource management as well as existing clients and tools.

KubeDirector provides the following capabilities:

- The ability to run non-cloud native stateful applications on Kubernetes without modifying the code. In other words, it is unnecessary to decompose the existing applications to fit a microservices design pattern.
- Native support for preserving application-specific configuration and state.
- An application-agnostic deployment pattern, minimizing the time to onboard new stateful applications to Kubernetes.

KubeDirector is a custom controller (deployed into K8s) that watches for custom resources of a given type to be created or modified within some K8s namespace(s). On such an event, KubeDirector uses K8s APIs to create or update the resources and configuration of a cluster to bring it into accordance with the specification defined in that custom resource. KubeDirector uses the following custom resources name, namely KubeDirectorApp and KubeDirectorCluster.

The app resource definition is used to create an application template that has the required metadata of the application such as the image used, the ports exposed, the various roles, and persistent volume. The cluster resource definition is used to instantiate the application template by providing the required resources such as CPU, memory, storage, and replicas. The application created this way has its endpoint exposed which is accessed using the service gateway endpoint mapping.

Deploying Methodize-Productivity web application using KubeDirector

Methodize-Productivity is a simple productivity manager and to-do tool that uses ReactJS, ExpressJS, NodeJS, and MongoDB.

Figure 22 shows the architecture of the legacy application.

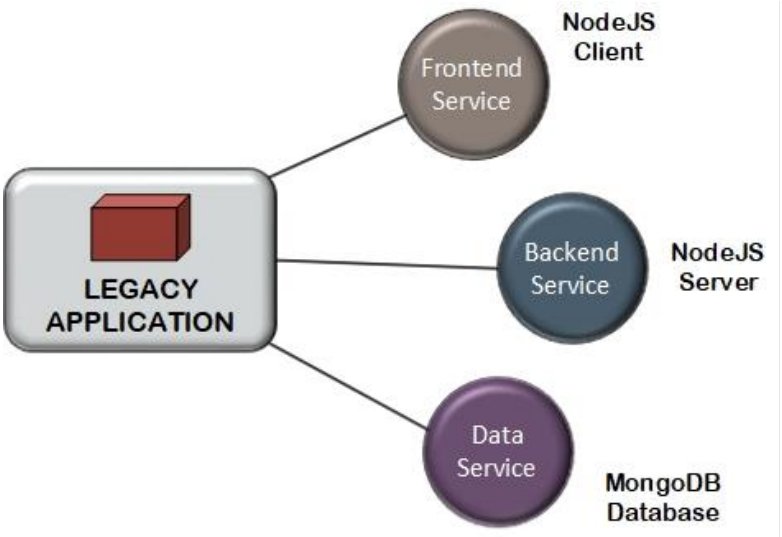

FIGURE 22. Legacy application architecture

Figure 23 shows the workflow steps to deploy the Methodize-Productivity application using KubeDirector.

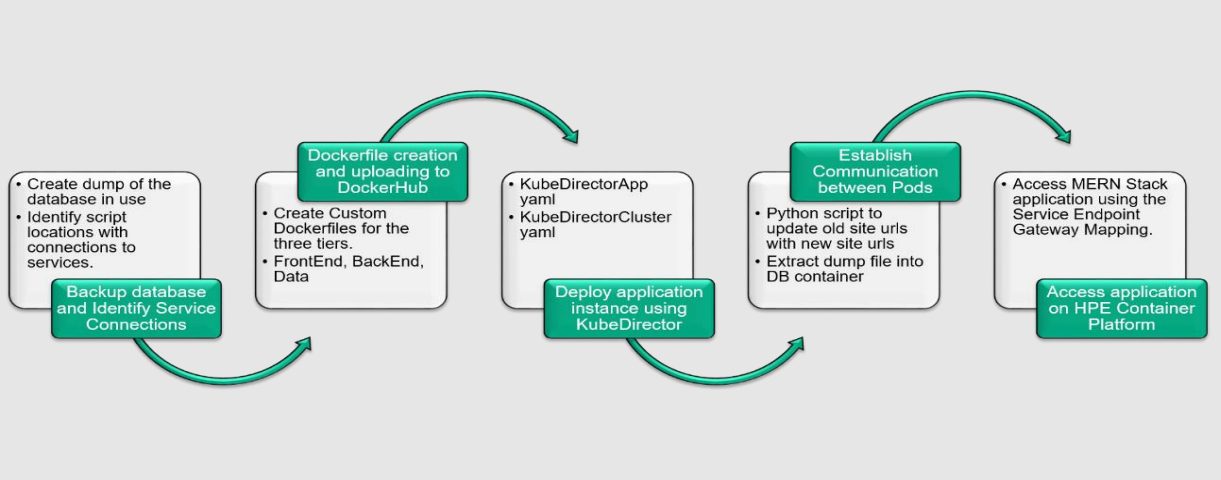

FIGURE 23. Workflow steps to deploy the Methodize-Productivity application using KubeDirector

For more information, see the Deployment Guide.

Business Continuity using Velero

Velero is an open source tool to safely back up, recover, and migrate Kubernetes clusters and persistent volumes. It works both on premises and in a public cloud. Velero consists of a server process running as a deployment in your Kubernetes cluster and a command-line interface (CLI). Velero uses the Kubernetes API to capture the state of cluster resources and to restore them when necessary. Backups will capture subsets of the cluster's resources, filtering by namespace, resource type, and/or label selector, providing a high degree of flexibility around what's backed up and restored.

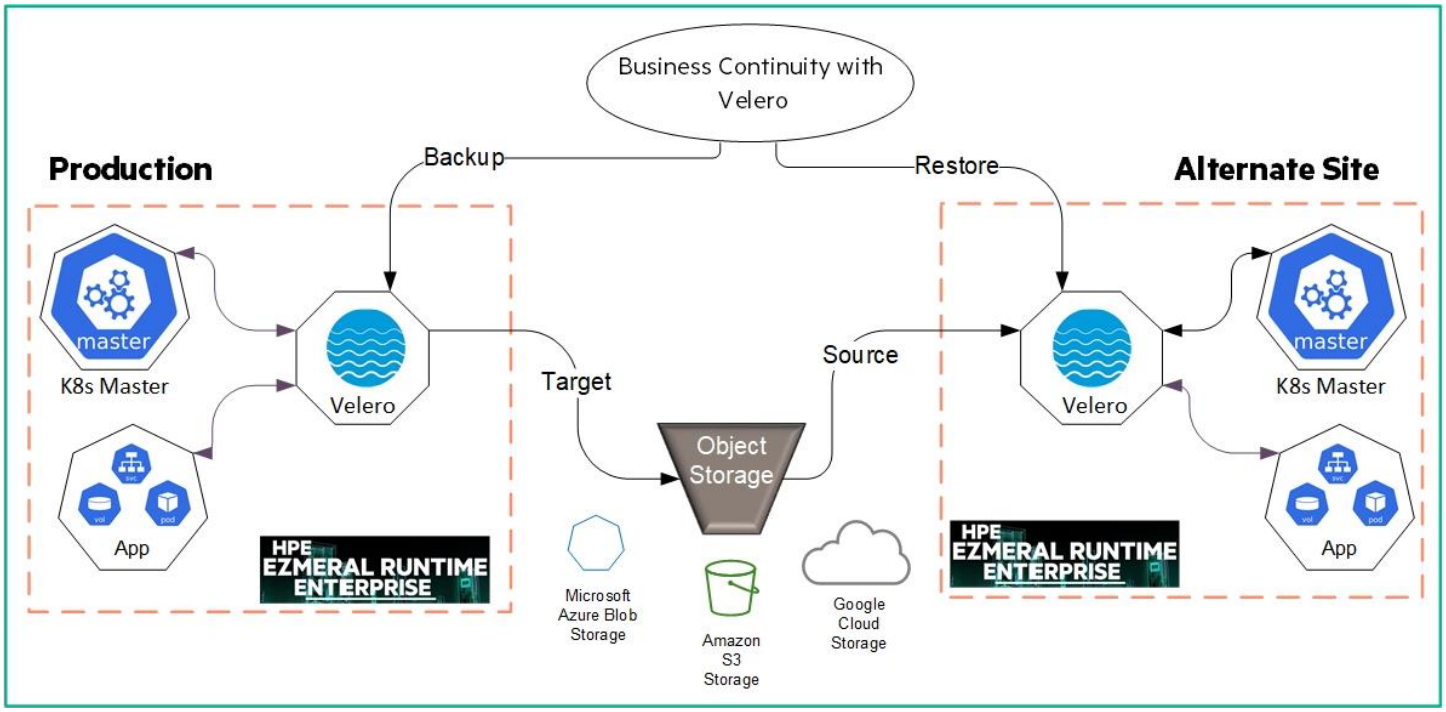

FIGURE 24. High-level overview of Velero and HPE Ezmeral Runtime Enterprise

Observability, security, and reliability

Microservices architecture brings a set of challenges in terms of security, network traffic control, and application telemetry. These challenges are overcome with the usage of a Service Mesh. A Service Mesh is a network of microservices that consists of applications and interactions between those applications. Istio, an open-source Service Mesh product developed by Google, IBM, and Lyft. It is one of the more feature-rich and complex options available in Service Mesh today. Istio works by deploying a special sidecar proxy that intercepts all network communication between microservices, using Istio's control plane functionality. Kiali works with Istio to visualize the Service Mesh topology. Kiali is a management console for Istio. Istio Service Mesh can be enabled while creating or editing any Kubernetes Cluster deployed by HPE Ezmeral Runtime Enterprise. The administrator can enable or disable Istio Service Mesh and also enable mTLS for each tenant within the cluster.

High-Level deployment flow

The prerequisite for deploying Istio is access to the HPE Ezmeral Runtime Enterprise Web UI. The high-level flow for setting up or enabling Istio for HPE Ezmeral Runtime Enterprise Kubernetes cluster is as follows:

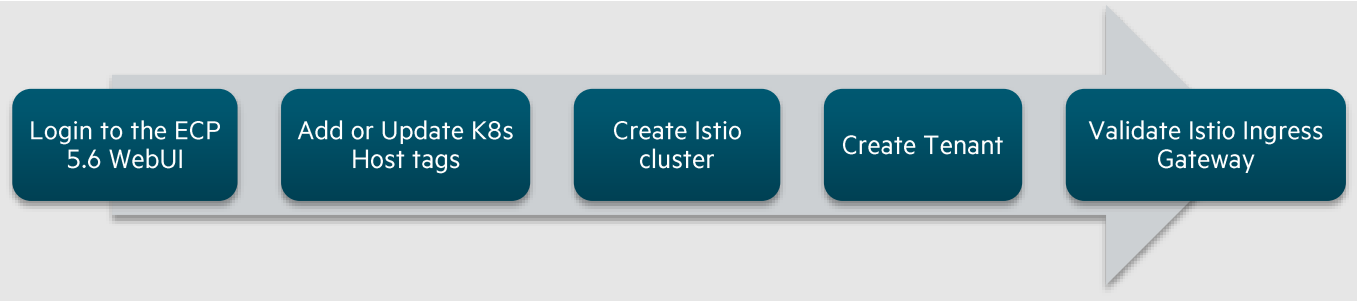

FIGURE 25. Workflow to install Istio Service Mesh in HPE Ezmeral Runtime Enterprise

Figure 26 shows the Istio service mesh and Kiali visualized Dashboard showcasing service observability

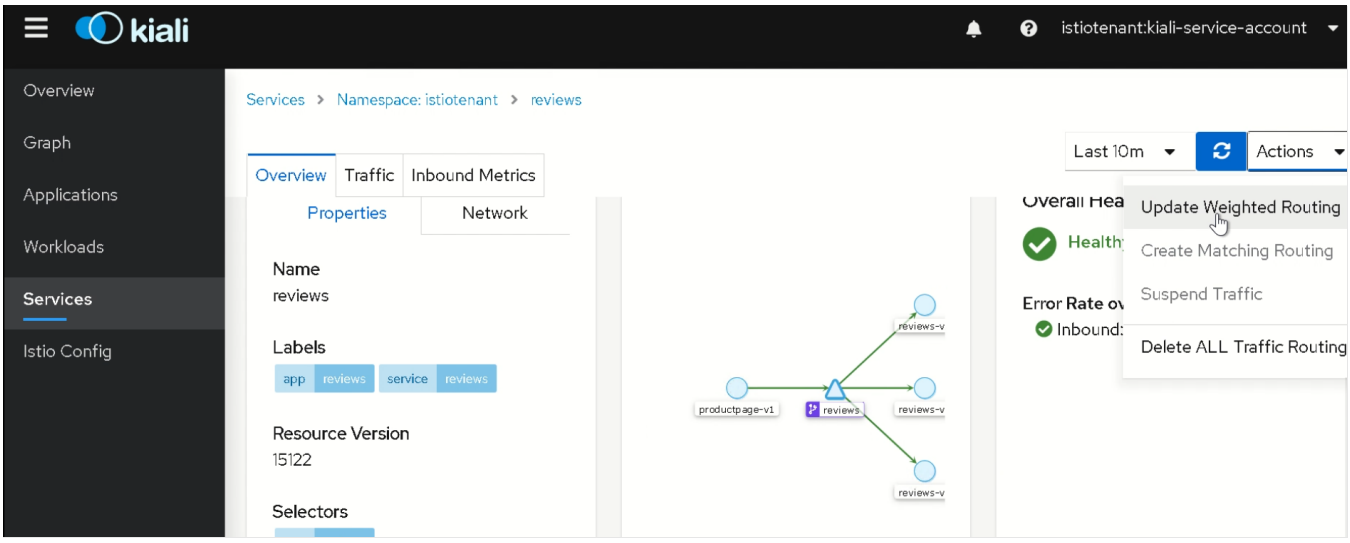

FIGURE 26. Istio service mesh and Kiali visualized Dashboard

Kubernetes cluster elasticity

To meet the on-demand growth and shrinking of the application instances, there might be times when the Kubernetes cluster nodes might be added to scale up resources. Expanding a Kubernetes cluster adds more resources and high availability to your cluster, for example, expanding from 1 master host to 3 master hosts adds high availability protection to the Kubernetes cluster. Shrinking a Kubernetes cluster is also preferred when existing cluster resources are needed for use elsewhere to manage effective usage of on-demand resources.

In HPE Ezmeral Runtime Enterprise 5.6 you expand and shrink your Internal Kubernetes Clusters. This Reference Configuration talks about an automated approach to expand or shrink your existing Kubernetes cluster. The high-level flow of the automated process is as follows:

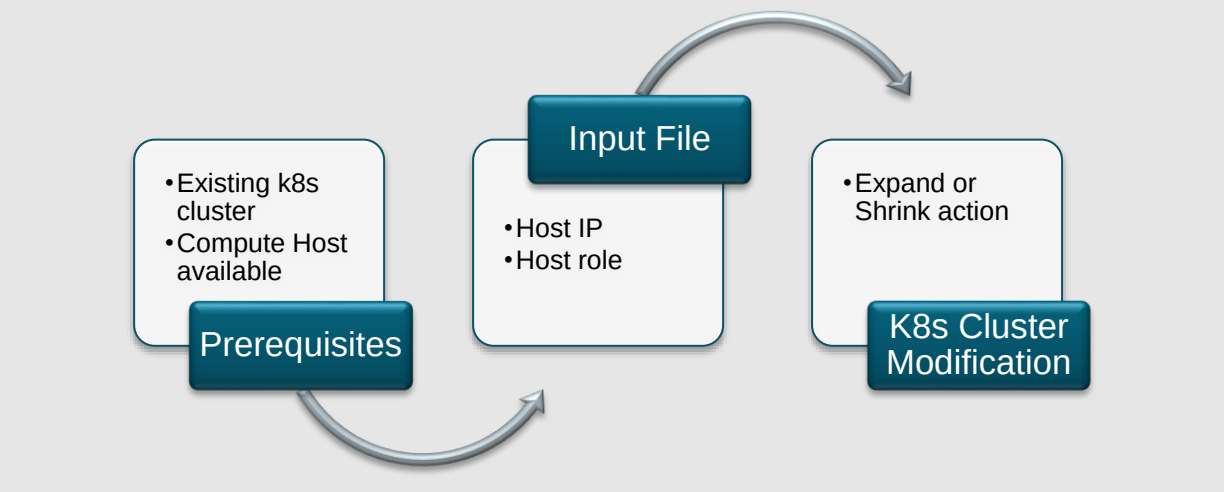

FIGURE 27. Kubernetes Cluster Elasticity flow

For more information on Kubernetes Cluster Elasticity, see the Deployment Guide.

SUMMARY

Companies are driving digital transformation and investing in innovation to compete. They want to deploy modern apps faster and simplify the production environment in a hybrid cloud architecture. To do so, they may have the mandate to move their application portfolio to the cloud and/or containers or both. Many organizations still struggle to achieve these goals due to a lack of time and expertise. This Reference Configuration showcases the lift-and-shift application modernization use case. It allows organizations to accelerate time-to-value by flawlessly building up a workable infrastructure the first time, every time.

With the HPE Ezmeral Runtime Enterprise, enterprises now have a unified Kubernetes-based software solution for DevOps, CI/CD workflow, application modernization across hybrid cloud architecture, streamlining deployment, air-gapped (disconnected) deployment and operations 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.

The solution provides customers a simpler, more scalable approach to modernizing applications 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 environment. So, 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 multitenant container isolation and data access, for any workload from edge to core to cloud. The benefits of containers beyond cloud-native microservices-architected stateless applications can be extended by providing the ability to containerize monolithic stateful analytic applications with persistent data.

The combination of HPE Ezmeral Runtime Enterprise paired with HPE ProLiant DL Gen11 server infrastructure and HPE Ezmeral Data Fabric deliver a modular architecture that can rapidly deploy modern containers supporting the new application framework. Ultimately, this results in faster digital transformation for businesses by helping organizations to increase the velocity of application development and accelerate innovation. This Reference Configuration overviews an enterprise-grade solution that helps organizations increase agility, simplify operations, and deliver a cloud-like experience while offering a compelling return on investment.

APPENDIX A: BILL OF MATERIALS

The following bill of materials (BOM) contains the core components utilized in the creation of this solution. Services, support, and software are not included in the BOM, and power distribution should be customized based on customer needs.

NOTE

Part numbers are at the time of testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult your Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative. For more information, see hpe.com/us/en/services/consulting.html.

TABLE A1. Bill of materials

Part number	Qty	Description
		Rack
P9K10A	1	HPE 42U 600mmx1200mm G2 Kitted Advanced Shock Rack with Side Panels and Baying
P9K10A 001	1	HPE Factory Express Base Racking Service
H6J85A	1	HPE Rack Hardware Kit
BW932A	1	HPE 600mm Rack Stabilizer Kit
BW932A B01	1	HPE 600mm Rack include with Complete System Stabilizer Kit
		HPE Ezmeral Runtime Enterprise HPE ProLiant DL325 Gen11 nodes
P54199-B21	11	HPE ProLiant DL325 Gen11 8SFF Configure-to-order Server
P54199-B21 ABA	11	HPE DL325 Gen11 8SFF CTO Svr
P53703-B21	11	AMD EPYC 9554P 3.1GHz 64-core 360W Processor for HPE
P53703-B21 0D1	11	Factory Integrated
P50312-B21	132	HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
P50312-B21 0D1	132	Factory Integrated
P54999-B21	11	HPE ProLiant DL325 Gen11 8SFF x1 Tri-Mode U.3 Backplane Kit
P54999-B21 0D1	11	Factory Integrated
P40510-B21	22	HPE 960GB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD
P40510-B21 0D1	22	Factory Integrated
P49049-B21	22	HPE 1.6TB SAS 24G Mixed Use SFF BC Multi Vendor SSD
P49049-B21 0D1	22	Factory Integrated
P25960-B21	11	Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE
P25960-B21 0D1	11	Factory Integrated
P02377-B21	11	HPE Smart Storage Hybrid Capacitor with 145mm Cable Kit
P02377-B21 0D1	11	Factory Integrated
P58335-B21	11	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
P58335-B21 0D1	11	Factory Integrated
P59668-B21	77	HPE ProLiant DL325 Gen11 Liquid Cooling Fan Kit
P59668-B21 0D1	77	Factory Integrated
P38997-B21	22	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
P38997-B21 0D1	22	Factory Integrated
BD505A	11	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	11	Factory Integrated

Part number	Qty	Description
P59619-B21	11	HPE ProLiant DL325 Gen11 8SFF x1 OCP2 Tri-Mode Cable Kit
P59619-B21 0D1	11	Factory Integrated
P48922-B21	11	HPE ProLiant DL3XX Gen11 Intrusion Cable Kit
P48922-B21 0D1	11	Factory Integrated
P08040-B21	11	HPE iLO Common Password FIO Setting
P52351-B21	11	HPE DL3XX Gen11 Easy Install Rail 2 Kit
P52351-B21 0D1	11	Factory Integrated
P58463-B21	11	HPE ProLiant DL325 Gen11 Liquid Cooling Heat Sink Kit
R6Z90AAE	11	HPE GreenLake for Compute Ops Management Standard 5-year Upfront ProLiant SaaS
		HPE Ezmeral Data Fabric HPE ProLiant DL385 Gen11 nodes
P53921-B21	5	HPE ProLiant DL385 Gen11 8SFF Configure-to-order Server
P53921-B21 ABA	5	HPEDL385Gen118SFFCTOServer
P53701-B21	10	AMD EPYC 9354 3.25GHz 32-core 280W Processor for HPE
P53701-B21 0D1	10	Factory Integrated
P50312-B21	60	HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
P50312-B21 0D1	60	Factory Integrated
P55082-B21	5	HPE ProLiant DL385 Gen11 8SFF Tri-Mode U.3 x1 BC Backplane Kit
P55082-B21 0D1	5	Factory Integrated
P40510-B21	10	HPE 960GB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD
P40510-B21 0D1	10	Factory Integrated
P49049-B21	15	HPE 1.6TB SAS 24G Mixed Use SFF BC Multi Vendor SSD
P49049-B21 0D1	15	Factory Integrated
P25960-B21	5	Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE
P25960-B21 0D1	5	Factory Integrated
P02381-B21	5	HPE Smart Storage Hybrid Capacitor with 260mm Cable Kit
P02381-B21 0D1	5	Factory Integrated
P57884-B21	5	HPE ProLiant DL3X5 Gen11 Smart Storage Battery 2P 96W Cable Kit
P57884-B21 0D1	5	Factory Integrated
P58335-B21	5	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
P58335-B21 0D1	5	Factory Integrated
P51181-B21	5	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE
P51181-B21 0D1	5	Factory Integrated
P38997-B21	10	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
P38997-B21 0D1	10	Factory Integrated
P57845-B21	5	HPE ProLiant DL385 Gen11 SFF Backplane Power Cable Kit
P57845-B21 0D1	5	Factory Integrated
P57847-B21	5	HPE ProLiant DL385 Gen11 8SFF OROC x1 SAS/SATA Cable Kit
P57847-B21 0D1	5	Factory Integrated
P57886-B21	5	HPE ProLiant DL385 Gen11 2U Standard/Performance FIO Air Baffle Kit
P58465-B21	30	HPE ProLiant DL3X5 Gen11 2U Performance Fan Kit
P58465-B21 0D1	30	Factory Integrated

Part number	Qty	Description
P50400-B21	5	HPE Gen11 2U Bezel Kit
P50400-B21 0D1	5	Factory Integrated
P52351-B21	5	HPE DL3XX Gen11 Easy Install Rail 2 Kit
P52351-B21 0D1	5	Factory Integrated
P58459-B21	10	HPE ProLiant DL3X5 Gen11 Performance 2U Heat Sink Kit
P58459-B21 0D1	10	Factory Integrated
R6Z89AAE	5	HPE GreenLake for Compute Ops Management Standard 3-year Upfront ProLiant SaaS
R9F63A B2B	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE PDU
HA454A1-021	2	HPE FE Strg and Ntwking Pkg 4 SVC
R9G06A	1	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE
R9G06A B01	2	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE
R9F61A	1	HPE Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit
R9F61A B2B	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE PDU
R9F61A 0D1	2	Factory Integrated
R9F57A	1	Aruba 1U Universal 4-post Rack Mount Kit for HPE
R9F57A 0D1	2	Factory Integrated
R9F59A	2	Aruba 4-post Rack Kit for HPE
R9F59A 0D1	2	Factory Integrated
		HPE Aruba 8325-32C 32-port 100G Switching
R9F67A	2	Aruba 8325-32C Power to Port Airflow 6 Fans 2 Power Supply Units Bundle for HPE
R9F67A 0D1	2	Factory Integrated
R9F67A B2B	2	Aruba 8325-32C Power to Port Airflow 6 Fans 2 Power Supply Units Bundle for HPE PDU
HA454A1-021	2	HPE FE Strg and Ntwking Pkg 4 SVC
R9F77A	2	Aruba 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable for HPE
R9F77A B01	2	Aruba 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable for HPE
C7535A	24	HPE RJ45 to RJ45 Cat5e Black M/M 7.6ft 1-pack Data Cable
C7535A 0D1	24	Factory Integrated
		HPE Aruba 6300M Switching
R9F63A	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE
R9F63A 0D1	1	Factory Integrated
R9F63A B2B	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE PDU
HA454A1-021	1	HPE FE Strg and Ntwking Pkg 4 SVC
R9G06A	1	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE
R9G06A B01	1	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE
R9F61A	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE
R9F61A B2B	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE PDU
R9F61A 0D1	1	Factory Integrated
		HPE Ezmeral Runtime Enterprise
R4T58AAE	608	HPE Cntr Plat 3yr Sel Sub 24x7 E-LTU (based on CPU core)
		SUSE Linux Enterprise Server

Part number	Qty	Description
N7F54AAE	19	SUSE Linux Enterprise Server 1-2 Sockets or 1-2 VM 1 Year Subscription 24x7 Support E-LTU

RESOURCES AND ADDITIONAL LINKS

HPE Reference Architectures, hpe.com/info/ra

HPE External Demo Portal https://hpedemoportal.ext.hpe.com/login

HPE servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE GreenLake Advisory and Professional Services, https://www.hpe.com/us/en/services/consulting.html

HPE Ezmeral Runtime Enterprise, https://www.hpe.com/us/en/solutions/container-platform.html

HPE Ezmeral Data Fabric https://www.hpe.com/us/en/software/data-fabric.htmll

Sysdig security, compliance, and monitoring for Kubernetes, https://sysdig.com/partners/hpe/

NGINX Load Balancing, https://www.nginx.com/partners/hp-enterprise/

HPE Ezmeral Runtime Enterprise Sizer tool https://solutionsizers.ext.hpe.com/EzmeralCPSizer/

HPE Ezmeral Runtime Enterprise 5.6 documentation https://docs.containerplatform.hpe.com/53/index.html

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

a50008368enw, version 1.0, April 2023

Hewlett Packard Enterprise