Hewlett Packard Enterprise

Technical white paper

HPE EZMERAL RUNTIME ENTERPRISE

Software architecture overview

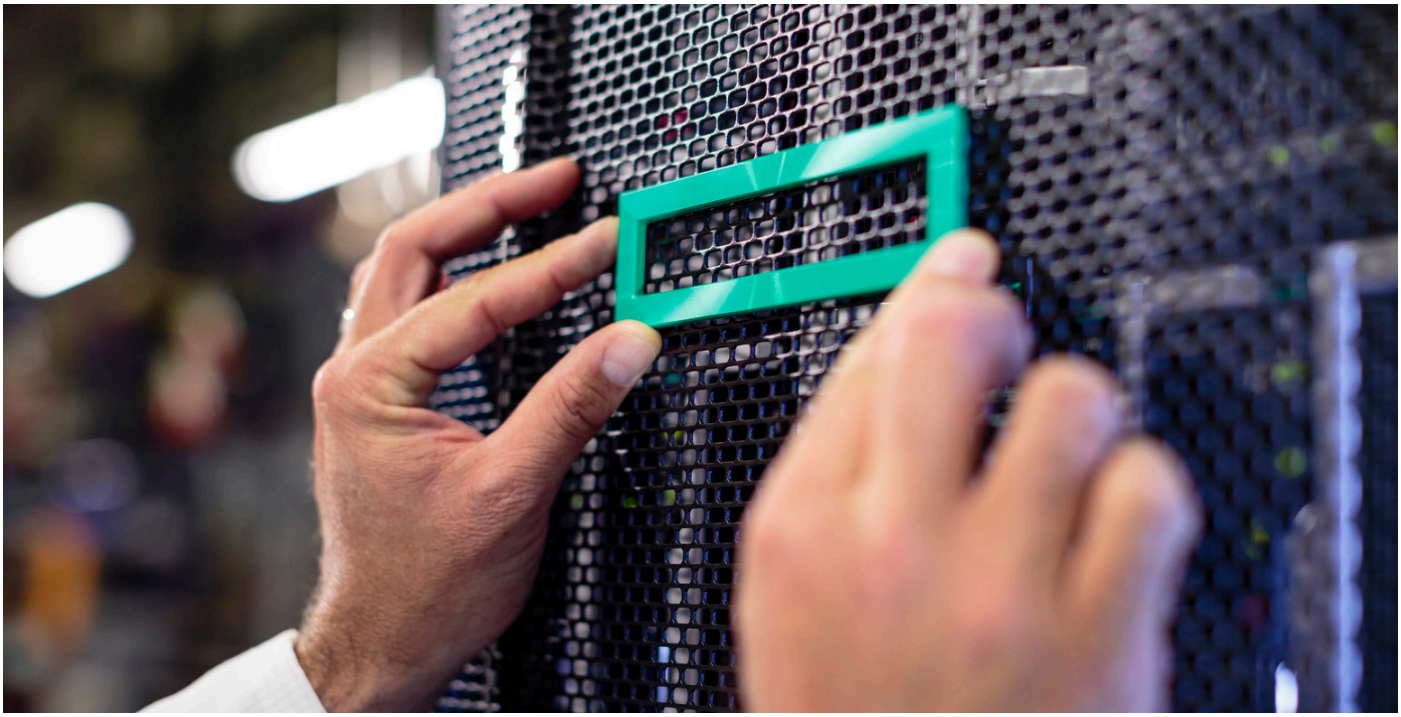

CONTENTS

ABSTRACT

In this paper, you will learn about the architecture and technologies driving HPE Ezmeral software. HPE Ezmeral is software for your future-state architecture delivering digital and AI transformation. With HPE Ezmeral, you can manage data simply, deliver data science quickly, operationalize models routinely, and align models and digital applications. Hewlett Packet Enterprise has a control plane with Kubernetes orchestration, a data plane with a global namespace for data from core to edge, and we secure workloads and applications with zero trust.

Read on to find out more about how HPE Ezmeral Runtime Enterprise enables key personas in the machine learning (ML) and data analytics space. HPE Ezmeral Runtime Analytics for Spark can be added to HPE Ezmeral Runtime Enterprise to support multiple versions of Spark on Kubernetes. In addition, you will learn about the HPE Ezmeral ML Ops features supporting end-to-end data science pipelines.

Also included in this paper is a discussion of the HPE Ezmeral Data Fabric File & Object Store. HPE offers a global, large-scale, and multiprotocol data fabric built to support data-intensive applications such as AI and ML applications and a huge array of data science applications that may require streaming, database, file, or object storage.

Further, you will learn how the open-source KubeDirector customer controller architecture drives our user-friendly application store in the web UI or through the fully automated REST API interface from HPE. This application deployment technology dramatically simplifies the deployment of large-scale, stateful, and non-cloud-native applications requiring persistent data.

Next, you will learn how everything is delivered with **enterprise-grade security** throughout the platform with features such as integration with your existing Active Directory (AD) / LDAP and Kerberos user authentication systems, role-based access control (RBAC), Kubernetes native security features, data fabric security features, and using SPIRE to implement Secure Production Identity Framework for Everyone (SPIFFE) zero-trust security for software services.

Finally, we will explore how high availability is woven into all aspects of HPE Ezmeral software.

ABOUT THIS DOCUMENT

The discussion is organized into five subjects:

- **HPE Ezmeral Runtime Enterprise**
  - Learn about the architecture and key features of the platform
- **Data orchestration with HPE Ezmeral Data Fabric File & Object Store**
  - Learn about the architecture of the HPE Ezmeral Data Fabric File & Object Store—a data orchestration layer specifically designed for modern, large-scale applications with massive data requirements
- **Simplify application modernization**
  - Learn how you can modernize your stateful, large-scale applications with HPE Ezmeral Runtime Enterprise
- **Security**
  - Learn how security is deeply integrated into every aspect of HPE Ezmeral Runtime Enterprise, HPE Ezmeral Data Fabric File & Object Store, and every modern application you deploy with the platform
- **High availability**
  - Learn how high availability is deeply rooted in all HPE Ezmeral software components

HPE EZMERAL RUNTIME ENTERPRISE: AN ENTERPRISE-GRADE SOFTWARE ORCHESTRATION PLATFORM DESIGNED TO DEPLOY MODERN APPLICATIONS

There is a lot of talk about containers and Kubernetes in the IT world today.

Kubernetes is a great way to orchestrate applications that have been containerized. The architecture of Kubernetes supports massive scale, and it is robust and secure. As explained in the Kubernetes.io documentation, Kubernetes approaches security using the 4Cs of cloud-native security: cloud, clusters, containers, and code.

Although a powerful and well-accepted standard method for orchestrating containerized workloads, most users of Kubernetes will describe it as a **steep learning curve** and not exactly **user friendly**. HPE Ezmeral Runtime Enterprise solves that problem.

HPE Ezmeral Runtime Enterprise is a software-defined container orchestration management control plane that allows users to easily organize compute and storage resources located anywhere and quickly create Kubernetes clusters on those resources. Multiple versions of Kubernetes can be running at the same time under HPE Ezmeral Runtime Enterprise management.

Much of the complexity and configuration details are simplified and presented as API commands, or a user can interact with HPE Ezmeral Runtime Enterprise using the graphical web user interface (web UI). This is like the experience you would have using a public cloud. The difference is that once you use HPE Ezmeral Runtime Enterprise you have total control of where your data is being stored, where your server (host) resources are coming from, and how everything is accessed and secured.

With HPE Ezmeral Runtime Enterprise, you also have the option to import clusters built on any public cloud platform and manage them alongside your local or on-premises clusters or clusters built on resources at the edge.

Key features

- **HPE Ezmeral Runtime Analytics for Apache Spark:** An add-on license to HPE Ezmeral Runtime Enterprise that gets data scientists and data engineers up and running quickly. With this license, data engineers can leverage Apache Spark's Delta Lake support for lakehouse infrastructure.
- **HPE Ezmeral ML Ops:** A full-service, end-to-end suite of tools integrated to enable large-scale, ML. A data scientist can utilize open source tools such as Jupyter Notebooks, Kubeflow, and Kale integrated into dynamic, scalable Kubernetes clusters and namespaces.
- **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. It allows you to build once and run anywhere.
- **Public cloud cluster import:** Unified control plane makes it easy to import external Kubernetes clusters and includes support for importing clusters from cloud vendors such as Amazon Elastic Kubernetes Service (Amazon EKS), Google™ Kubernetes Engine (GKE™), and Azure Kubernetes Service (AKS).
- **Create standards-based Kubernetes orchestrated clusters to simplify app development:** Takes the complexity out of using Kubernetes. Using simple API or web UI forms, users can take full advantage of the massive scale provided by Kubernetes orchestration. This open approach enables modern application development using containerized microservices, as well as the rich open-source ecosystem for CI/CD and DevOps.
- **Multicluster, multi-tenant Kubernetes management:** Fast, easy deployment, management, and monitoring of Kubernetes clusters both on-prem and off-prem for single-pane-of-glass management and visibility across environments.
- **Enterprise-grade security:** Built-in security controls to 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; Argo CD 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:** Cross-server GPU sharing and NVIDIA® multi-instance GPU fractionalization improve collaboration and GPU utilization across on-premises, hybrid- and multi-cloud environments at enterprise scale.

- **Frictionless data access:** HPE Ezmeral Data Fabric, dataTap, and FSMount let you connect 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 through natively integrated Istio™ service mesh.
- **Self-service web portal:** Lets users create and manage clusters, create, and manage nodes, run jobs, and view monitoring statistics. RBACs are designed for every action a user can or cannot take, which also makes it simpler for a user to get to just the tools they need. For example, department administrators can use the portal to provision nodes/clusters whereas a non-admin user can log in and go directly to the application they need to work with and not bother with cluster administration.
- **RESTful API:** Supports a RESTful API that surfaces programmable access to the same capabilities available through the self-service portal.
- **1-click provisioning:** 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 applications through KubeDirector and App Workbench.
- **Leverage existing storage and compute resources:** Repurpose existing, large-scale data deployments. Multiple storage protocols are supported (NFS, HDFS, HDFS with Kerberos, and others).

Key benefits

- **Greater choice:** Hybrid architecture provides the foundation for next-gen lakehouse analytics, delivering greater freedom and choice to data engineers, analysts, and scientists.
- **Greater flexibility:** A unified platform for orchestration of cloud-native and non-cloud-native applications on-premises, in any cloud, and at the edge.
- **Superior performance:** Provides storage I/O optimizations to deliver data to Big Data and AI applications without the penalties commonly associated with virtualization and containerization or the separation of compute and storage resources. The CPU cores and RAM in each host are pooled and then partitioned into virtual resource groups based on tenant requirements.
- **Improved productivity:** Containerized applications deliver faster results and higher throughput on bare metal.
- **Reduced risk:** Extensive policy-based privilege management and control let you easily define and tailor access, trust levels, and privileges for people, teams, and data spaces.
- **No lock-in:** Using 100% open-source components and open APIs, you can pick up your application and data and move it without needing to refactor.
- **Improved ROI:** Utilization of hardware resources is improved by sharing a common data infrastructure across teams and workloads and providing a cloud-like experience for non-cloud-native monolithic applications, increasing the return on hardware investment.
- **Reduced IT and administrative overhead:** Streamline operations and reduce IT costs by automating provisioning, unifying management, and supporting push-button upgrades.
- **Increases utilization while lowering costs:** Delivers hardware and operational cost savings while simultaneously eliminating the complexity of managing multiple physical clusters; also allows clusters to be paused/unpaused at will, meaning that you only use the resources that you need when you need them.
- **High availability:** Supports three levels of high availability to provide redundancy and protection.
- **Compute and storage separation:** Decouples analytical processing from data storage, giving you the ability to independently scale compute and storage capacity instantly on an as-needed basis. This independent scaling enables more effective utilization of infrastructure resources and reduces overall costs.

ARCHITECTURE OVERVIEW

Figure 1 provides a comprehensive view of how to conceptualize everything you get when you install HPE Ezmeral Runtime Enterprise.

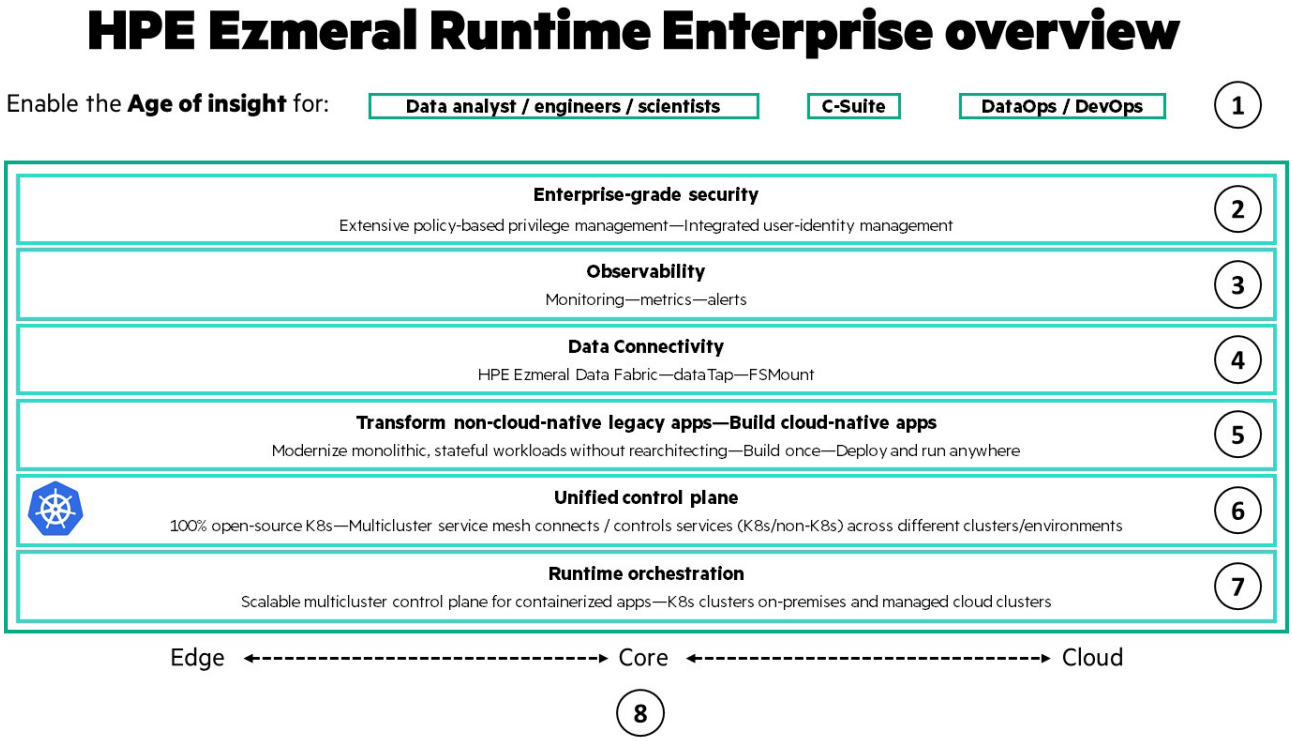

FIGURE 1. HPE Ezmeral Runtime Enterprise architecture

In Figure 1, there are circles next to each of the architectural layers. Here is a brief description corresponding to each number.

1. **Self-service:** Multiple different teams can help themselves with tools and infrastructure.
2. **Enterprise-grade security:** Multiple levels of security include built-in security controls to 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; and Kerberos authentication for your data. HPE Ezmeral Runtime Enterprise clusters and applications can be further secured by using strongly attested identities provided by SPIRE for authentication.
3. **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.
4. **Frictionless data access:** HPE Ezmeral Data Fabric File & Object Store, dataTap, and FSMount let you connect and manage data wherever it is located.
5. **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 (that is, legacy) apps without rearchitecting or refactoring.
6. **Unified control plane:** Makes it easy to import external Kubernetes clusters; includes support for importing clusters from cloud vendors such as Amazon EKS, GKE, and AKS.
7. **Multicluster, multi-tenant Kubernetes management:** Fast, easy deployment, management, and monitoring of Kubernetes clusters both on-prem and off-prem for single-pane-of-glass management and visibility across environments.
8. **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. It allows you to build once and run anywhere.

HPE EZMERAL RUNTIME ANALYTICS FOR APACHE SPARK

The HPE Ezmeral Runtime Analytics add-on license offers a centralized solution to the problem of complex, costly, and disjoint point solutions. Businesses that are working to be more data driven currently must work with multiple data formats and even different data architectures such as data warehouses and data lakes. In addition, an analytics workflow requires many different applications, and sometimes even multiple versions of the same application such as supporting Spark 2.4.x and Spark 3.1.x simultaneously.

HPE Ezmeral Runtime Analytics offers a tightly integrated solution where a tenant, cluster or storage administrator, data scientist, and data engineer can all benefit in many ways.

Features that support tenant, cluster, and storage administrators

- Point-and-click installation of multiple Spark versions into a Kubernetes namespace/tenant
- GPU resource optimization in multi-tenant deployments
- Simplified connectivity to existing data lakes and data warehouses utilizing dataTaps and NFS mounts
- Enterprise support and professional services to help you build and maintain it

Features that support data scientists, data analysts, data engineers, and developers

- Pre-built integration of Notebooks with REST interface to submit Spark jobs
- Airflow integration to design and run workflows
- Multiple Spark versions, including custom Spark images with required packages
- Delta Lake support within Spark applications
- Install Spark Thrift and History Servers on demand to your Kubernetes namespace

A Kubernetes cluster or tenant administrator can easily configure Kubernetes clusters to include the Spark Operator. The Spark Operator enables running Spark on Kubernetes clusters and can be added with the click of a button to any Kubernetes namespace that you created with HPE Ezmeral Runtime Enterprise with the HPE Ezmeral Runtime Analytics add-on.

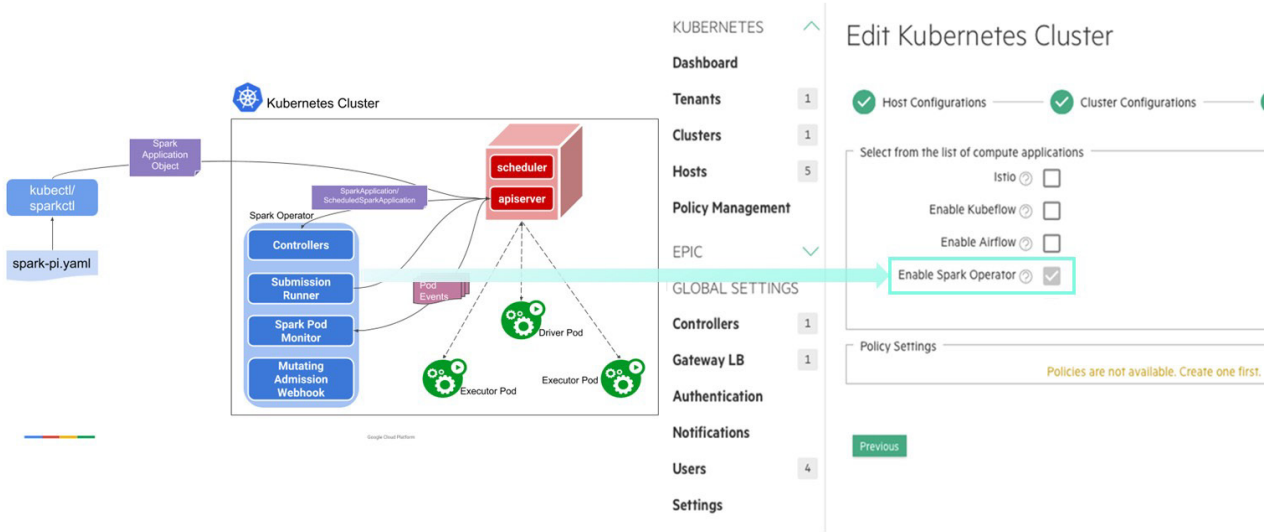

FIGURE 2. Enable the Spark Operator in your Kubernetes namespace with a simple check box selection

Using the Spark Operator batch types of Spark jobs can be **submitted** to the Spark Operator running at the Kubernetes cluster level.

There is also the option to interact with Spark through a Livy interface and Livy can be installed on-demand at the Kubernetes namespace/tenant level. Using Livy, a developer or multiple developers can have an interactive experience with Spark through this RESTful interface. As shown in Figure 3, multiple versions of Livy are supported in addition to many popular components and applications related to analytics workflows. It is important to note that only one version of Livy can be instantiated at a time in a single tenant.

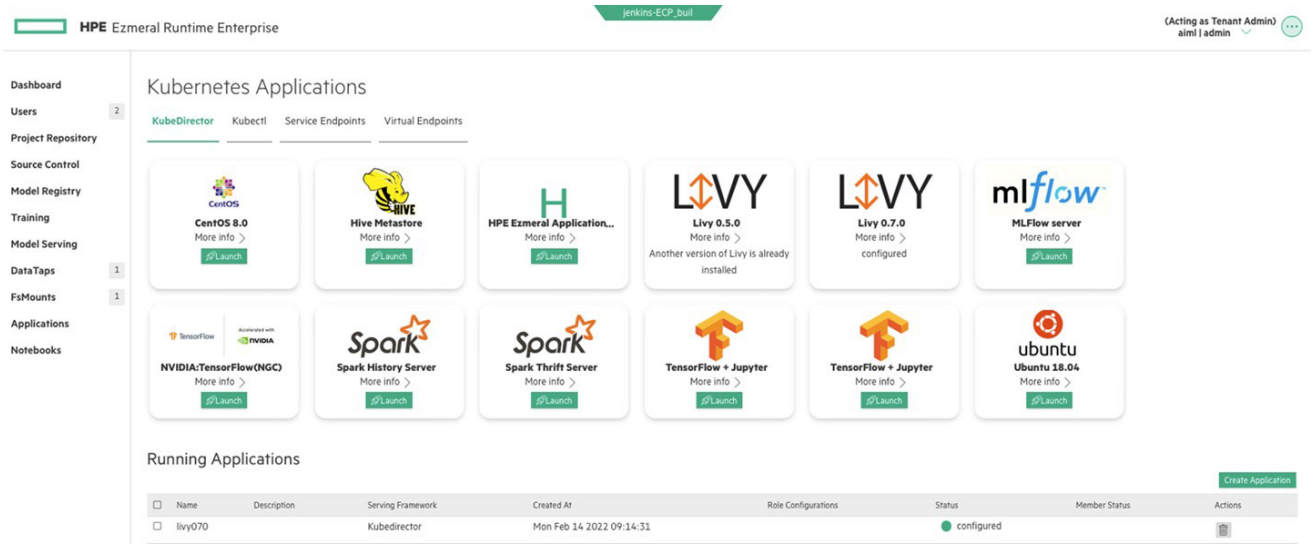

FIGURE 3. Support for Livy includes multiple versions

A Spark developer may also require access to a Spark Thrift Server, which enables interaction with external business intelligence systems such as Tableau and Power BI, which can then run SQL queries against Spark.

Hive Metastore can also be installed on-demand at the Kubernetes namespace/tenant level. The Hive Metastore stores the schema for the data that is being accessed by Spark jobs. This on-demand installation of the Hive Metastore at the Kubernetes namespace/tenant level gives the cluster and tenant administrator great flexibility.

Also included is the ability to launch a Spark History Server if a Spark developer wants to view Spark jobs that have or are being executed. The Spark History Server also includes metrics after Spark jobs are completed.

Both the History and Thrift Servers can be **launched** from the same Kubernetes namespace, managed as a **tenant**. It is simple enough that a developer can easily switch between developing their application and then self-service by adding the applications through the Kubernetes application window. Or this can all be done by a designated cluster or tenant administrator. It is up to the customer to decide what type of role-based access to give their team members.

Allowing a cluster or tenant administrator to have the choice to install or not to install all of these supporting components, as needed, provides great flexibility and delivers excellent resource efficiency.

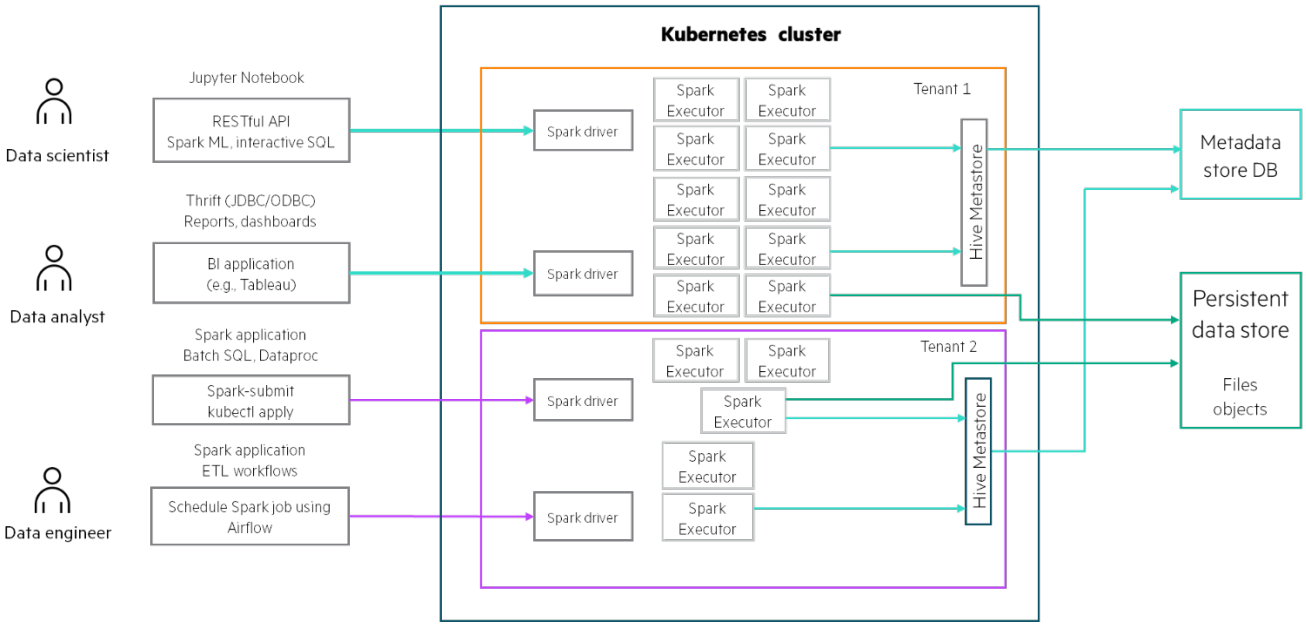

FIGURE 4. HPE Ezmeral Runtime Analytics enables ultimate flexibility for Spark execution

WIZARD-DRIVEN UI EXPERIENCE CREATING SPARK APPLICATIONS

If a DevOps engineer or developer needs to build their customer Apache Spark-based application, they can utilize this wizard-driven web UI.

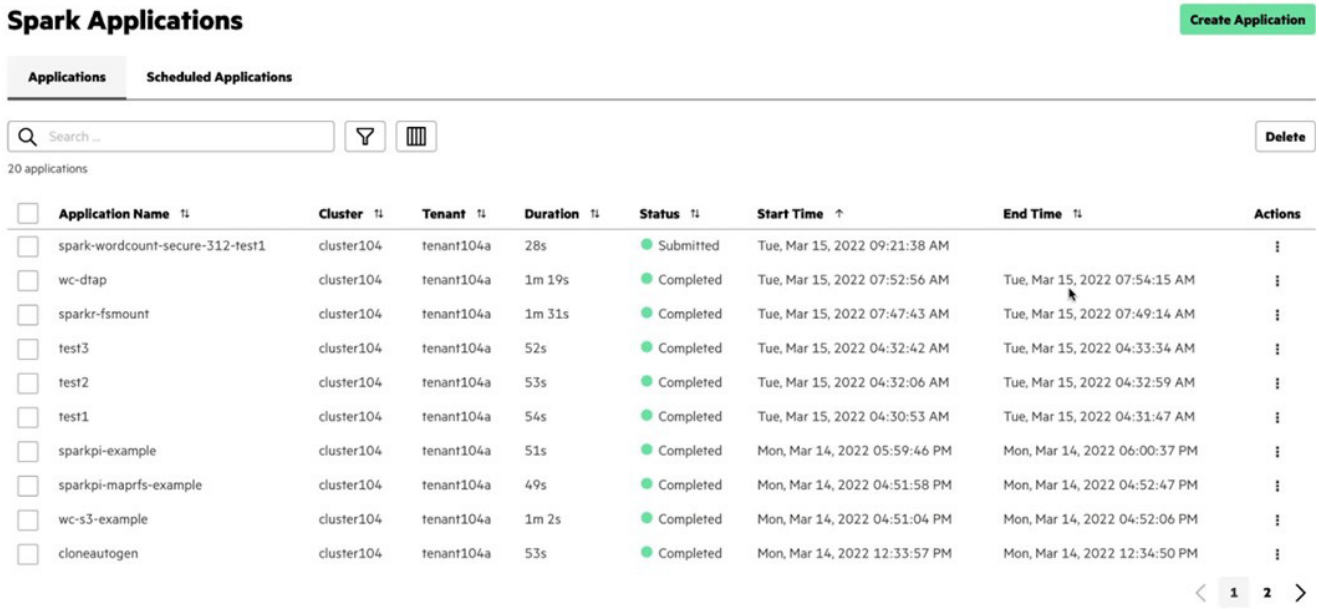

FIGURE 5. Monitor and create Spark applications from the HPE Ezmeral Runtime Enterprise web UI

An Apache Spark developer can use this step-by-step, graphical, application-building workflow to simplify the process of creating a custom, containerized Spark application in a very structured and intuitive way.

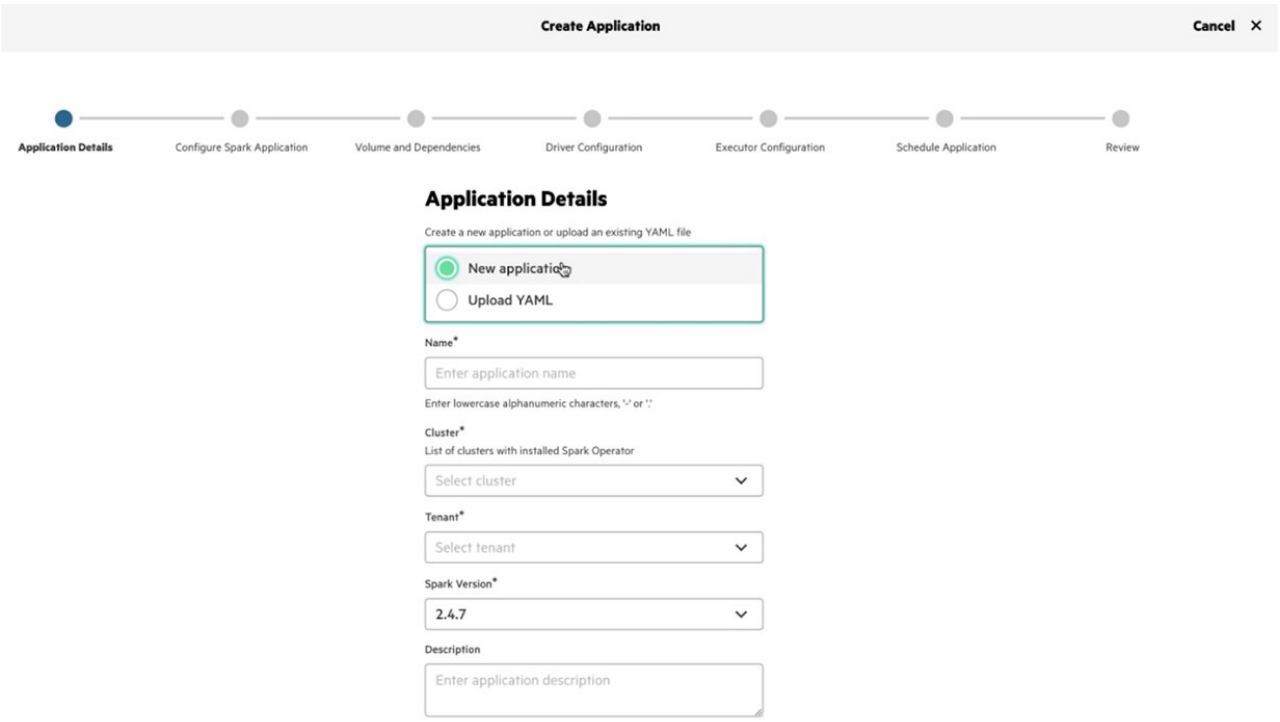

FIGURE 6. HPE Ezmeral Runtime Enterprise web UI-based Spark application development workflow

As Figure 6 indicates, a developer can create an entirely new application or upload a YAML file that defines the application. Many details required to launch and manage an application within a Kubernetes cluster are provided in this wizard-driven workflow.

HPE EZMERAL ML OPS

An additional license may be purchased to add-on HPE Ezmeral ML Ops to the HPE Ezmeral Runtime Enterprise platform. This enables a great many tools for use in the creation of end-to-end ML pipelines from data ingestion, data transformation, data analytics to ML model training, deployment, and serving.

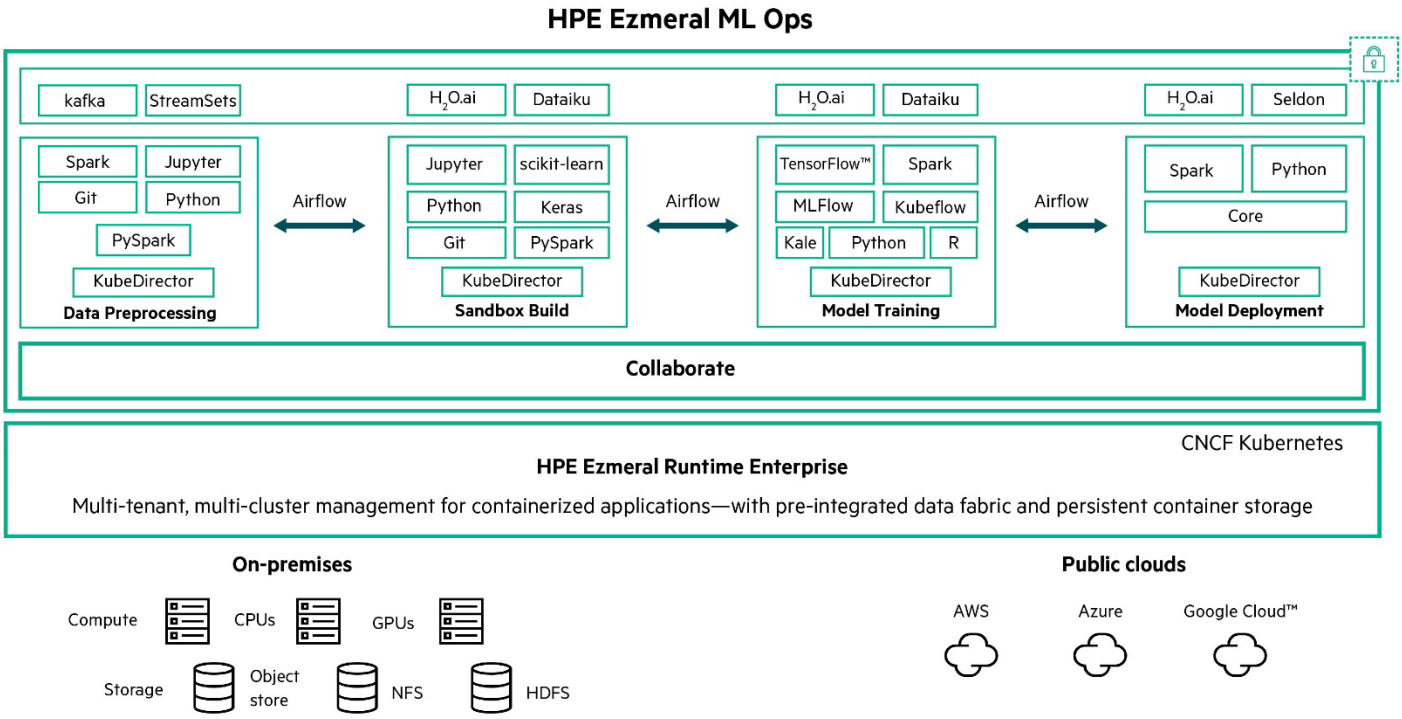

FIGURE 7. HPE Ezmeral ML Ops

As Figure 7 illustrates, there are many tools, data sources, and compute resources that must be architected to work together. This is the key value proposition of not only the HPE Ezmeral Runtime Enterprise software in general but also of HPE Ezmeral ML Ops—tight integration of all the resources a data scientist and data engineer needs. Let's explore some highlights of the key features of HPE Ezmeral ML Ops.

Supporting an ML pipeline

To enable business processes, it all starts and ends with the data that flows into and out of a given business. Take for example a business that lends money to customers. That business may need to: identify new customers, validate customer data, detect and prevent fraud, approve loans, deny loans, and so on.

All of these business processes can be automated and improved by leveraging ML with tightly integrated data management. HPE Ezmeral ML Ops provides an end-to-end ML pipeline with integrated data management. As indicated in Figure 7, this pipeline flows from data preprocessing to sandbox building, then to model training, and finally to model deployment. The pipeline is not static, and in fact, recurses back onto itself as the data changes.

How the software is deployed

HPE Ezmeral Runtime Enterprise is first deployed on a set of compute resources we call hosts. These could be physical hosts in a local data center or virtual machines (VMs). HPE Ezmeral Runtime Enterprise is installed as separate pieces: one or more hosts are configured as the controller and one or more hosts are configured as the network gateway. Then, a user can add as many hosts as needed, at any time to function as workers in clusters. Figure 8 gives a simple example of the infrastructure required to install and run the software itself.

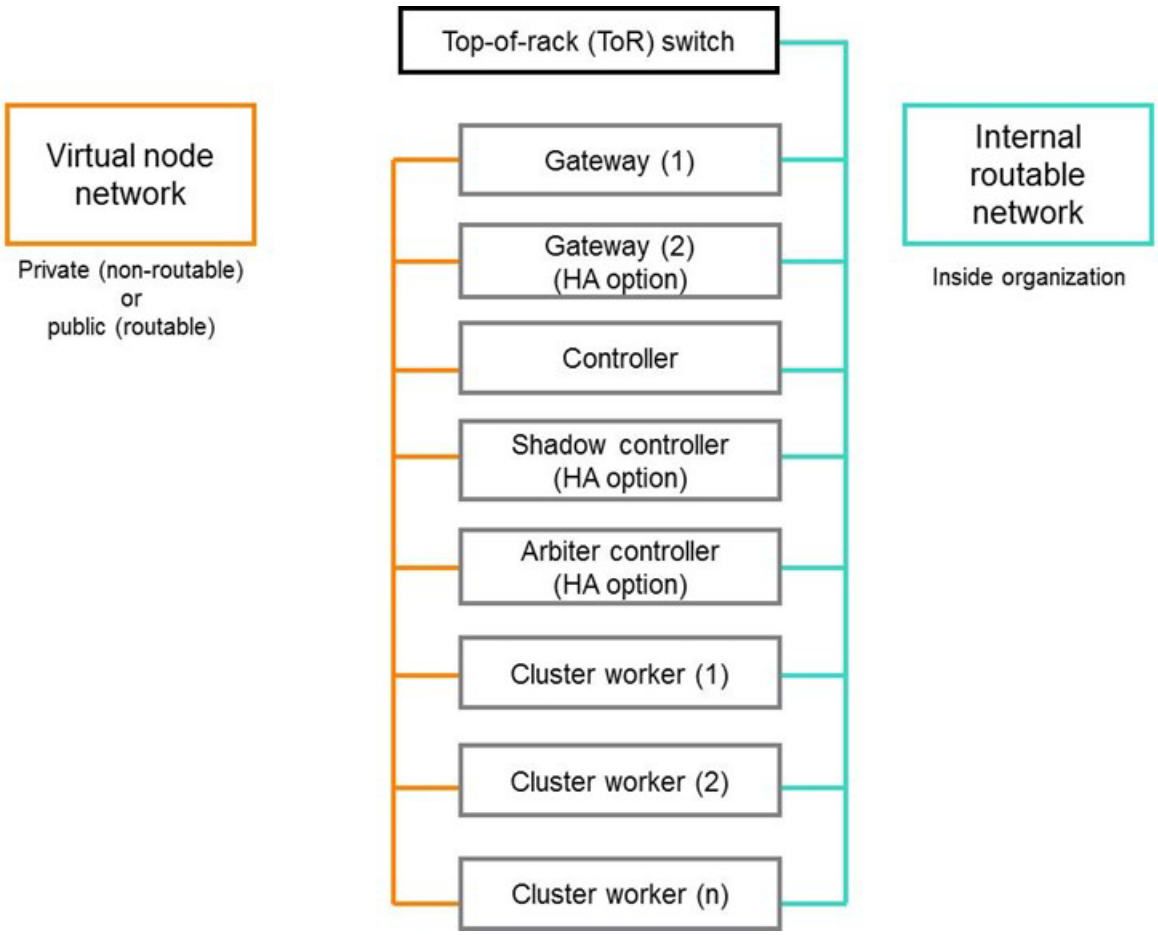

FIGURE 8. The dual network configuration of HPE Ezmeral Runtime Enterprise

Once installed onto these hosts, you will see that there are two networks configured.

- **One network for the controller, worker, and gateway hosts:** This network must be routable and part of the organization. If the HPE Ezmeral Runtime Enterprise HA feature is enabled, then both the primary controller and shadow controller must be in the same subnet of this network.
- **The second network is for the container network (Docker/K8s containers):** This can be either public (routable) or private (non-routable) and is configured and managed by HPE Ezmeral Runtime Enterprise. For a detailed description, see the gateway hosts and load balancing links to HPE Ezmeral Runtime Enterprise online documentation. Also, during deployment, you could consult the "Network Requirements" section.

CLUSTER MANAGEMENT

Web UI and RESTful API access

The web UI provides a feature-rich GUI for easy point-and-click management of your clusters. In addition, there is full support of a RESTful API. This provides the flexibility to use the API or UI to create clusters, deploy applications, manage ML projects, and more.

Secure, authenticated access to the web UI

The method used to launch and log in to the web interface will vary slightly depending on the authentication configuration that the user sets up during the initial deployment of HPE Ezmeral Runtime Enterprise, for example, using SSL or not.

Once you have installed the software, you can then access the web UI from the Fully Qualified Domain Name (FQDN) or IP address of the gateway. For example: http://xhost.xthing.xdivision.hpecorp.net/bdswebui/login/

If you have configured the controller and gateway for high availability, the FQDN will point you to the primary controller for HPE Ezmeral Runtime Enterprise. If there is a failure of the primary controller, the same FQDN will automatically point you to the other controller, so you never lose connection to the platform.

When you enter the IP address or FQDN into a web browser, a login screen similar to Figure 9 will then be displayed for you to log in.

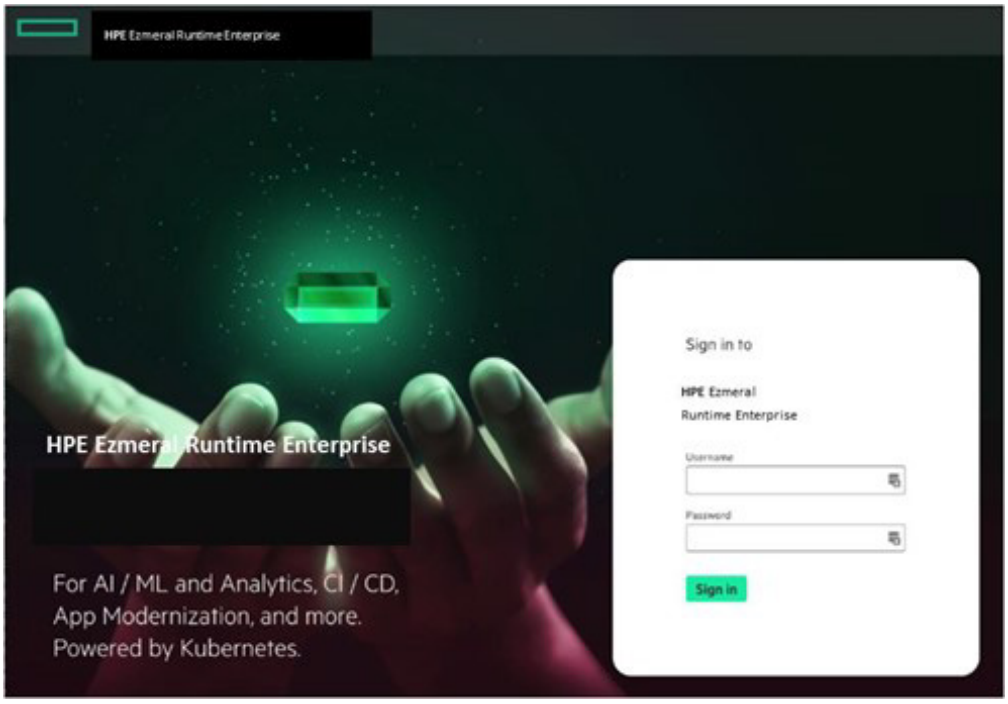

FIGURE 9. Platform-authenticated web UI login screen

Role-based access to tenants/namespaces and clusters

When you install the platform, you have the option to include details for connection to your organization's AD or LDAP user authentication services. Also, provided is local user authentication that is managed by HPE Ezmeral Runtime Enterprise itself. Each user that is added to the platform is assigned a specific role that controls what they can or cannot do. For example, you can add a user to the platform as a:

- **Tenant member:** This type of user role allows the user to be assigned to specific Kubernetes namespaces also called tenants. These users have no access to the cluster creation or deletion functions of the site or management of hosts. Within the tenant, they are assigned a member role; they can launch the new application and use applications already running, as well as access to shared data is available for file management.
- **Tenant administrator:** This role not only inherits all the member role access but also allows the user to create dataTaps and FSMounts (storage used by the applications in the tenant) and adds additional users to the tenant. This user has no cluster or container platform management access.
- **Site administrator:** Any user assigned to this role can create and delete clusters and tenants, as well as manage the configuration of HPE Ezmeral Runtime Enterprise itself.

STORAGE OPTIONS FOR CLUSTERS

HPE Ezmeral Data Fabric File & Object Store is discussed in greater detail later in this document, and it is important to understand that ultimately it is the underlying storage manager. This brief description explains how storage is organized for use by applications running in the tenants on the Kubernetes clusters.

- **Tenant or project storage:** Tenant storage is an optional storage location that is shared by all nodes within a given tenant. The platform administrator configures tenant storage during installation and can change it at any time thereafter. Tenant storage can be configured to use either an HPE Ezmeral Data Fabric File & Object Store installation (configured on the host storage) or a remote HDFS or NFS system. Alternatively, you can create a tenant without dedicated storage. See the online documentation for details of tenant/project storage.
- **Node storage:** Node storage (referred to as ephemeral storage in Kubernetes clusters) is built from the local storage in each host and is used for the disk volumes that back the local storage for each virtual node. Using self-encrypting drives (SEDs) will ensure that any data written to node storage is encrypted on write and decrypted on read by the OS. A tenant can optionally be assigned a quota for how much storage the nodes in that tenant can consume. See the "Node Storage" section for details of node storage.
- **dataTaps:** dataTaps expand access to shared data by specifying a named path to a specified storage resource. Applications running within clusters that can use the HDFS file system protocols can then access paths within the resource using that name. This allows you to run jobs using your existing data systems without the need to make time-consuming copies or transfers of your data.
  See the "About dataTaps" section for details of dataTap.
- **FSMounts:** The file system mount feature allows the automatic addition of NFS v3 or v4 volumes or mounts to a given tenant. This allows applications within the tenant to directly access NFS shares as if they were local directories. This eliminates the need to manually copy common files to individual virtual nodes.

Monitoring and alerting

There are many useful dashboards included in HPE Ezmeral Runtime Enterprise to help you manage and monitor your Kubernetes clusters. These dashboards include detailed monitoring at the cluster level. Figure 10 is an example Kubernetes dashboard showing various utilization metrics of the underlying resources, including usage, load, and services status. Also included is policy violation reporting on the dashboard to alert administrators when any Kubernetes objects break the policy rules that were set up.

Kubernetes dashboard

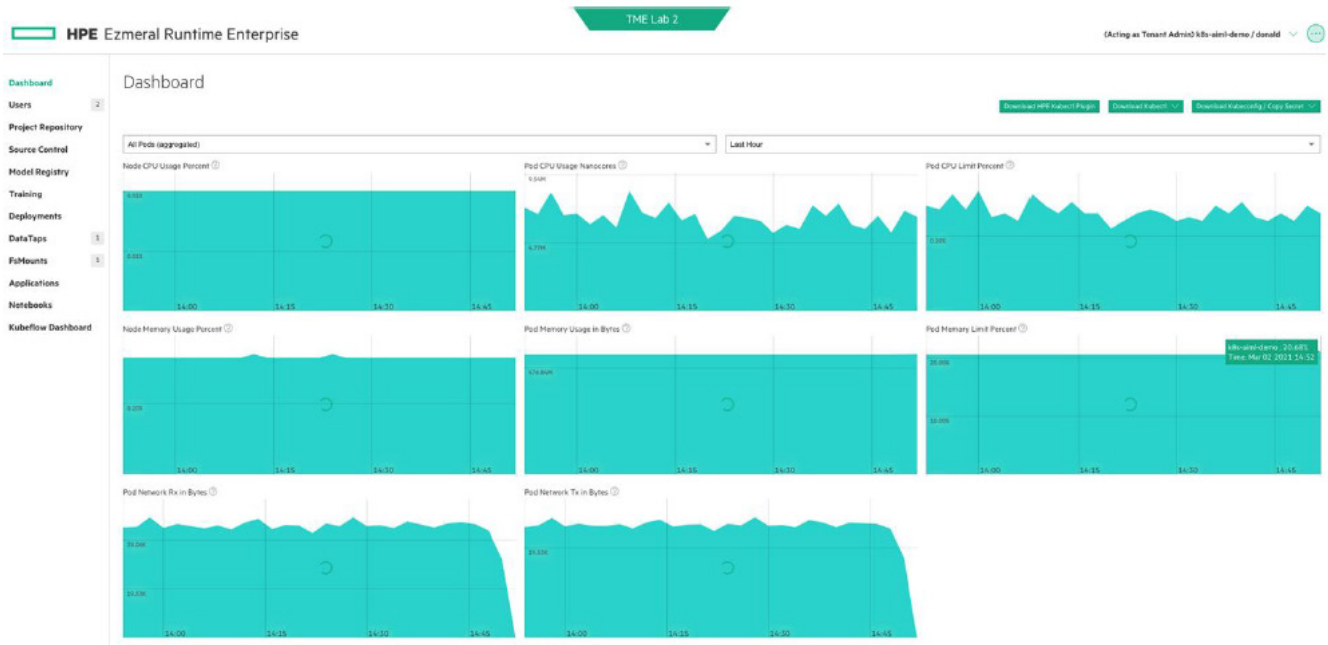

FIGURE 10. Example Kubernetes dashboard

DATA ORCHESTRATION WITH HPE EZMERAL DATA FABRIC FILE & OBJECT STORE

HPE Ezmeral Data Fabric File & Object Store complements HPE Ezmeral Runtime Enterprise by providing large-scale distributed data storage, motion, and management. An external deployment of HPE Ezmeral Data Fabric can be registered as the tenant storage provider for HPE Ezmeral Runtime Enterprise. HPE Ezmeral Runtime Enterprise is concerned primarily with managing computational resources in multiple computational clusters, whereas the data fabric manages data in multiple distributed storage clusters.

The data stored in HPE Ezmeral Data Fabric File & Object Store can scale in several ways, in size, the number of objects, the number of applications supported on any cluster, and terms of geographical distribution. This scalability is not only possible partly because of the advanced technology used to implement the data fabric but also because key aspects of managing the data are exposed to users so that they can help the fabric manage data correctly. As far as possible, the data fabric is designed so that key concerns can be separated. Developers for different applications and even administrators can work efficiently when they need to access common data.

In addition, the data fabric allows access to data using many access methods, allowing legacy applications and modern containerized applications to have shared access to common data. This means that virtualized workloads can generate data that legacy Hadoop programs analyze and then Kubernetes-based ML systems can use to build AI models, all without needing more than one platform and without requiring extraneous copies of data.

This ubiquitous access is enhanced when data fabric allows local and remote access to data. The fabric also supports fabric-level reliable replication of data. This means that an application running in a central facility can access raw data stored in edge clusters in the field using a pathname that doesn't depend on where the application is running. Remote access is particularly useful if only a small part of a large data set is needed. If the entire data set is to be processed intensively, the data fabric can reliably mirror remote data to a local cluster as well. For more detailed information about HPE Ezmeral Data Fabric File & Object Store, visit hpe.com/us/en/software/ezmeral-data-fabric.html.

MODERNIZING STATEFUL APPLICATIONS WITH KUBEDIRECTOR

HPE Ezmeral Runtime Enterprise provides the central control plane to manage your compute resources for your containerized workloads, along with HPE Ezmeral Data Fabric File & Object Store that manages storage for those applications. Getting a stateful, data-driven application into a containerized solution in the first place is dramatically simplified using KubeDirector.

KubeDirector uses the standard Kubernetes (K8s) facilities of custom resources and API extensions to implement stateful scale-out application clusters. This approach enables transparent integration with K8s user/resource management and existing K8s clients and tools.

In broad terms, KubeDirector is a Kubernetes custom controller.

The KubeDirector custom controller is itself deployed into K8s and watches for custom resources of a given type to be created or modified within some K8s namespaces. In such an event, KubeDirector uses K8s APIs to create or update the resources and configuration of a cluster to bring it under the specifications defined in that custom resource.

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

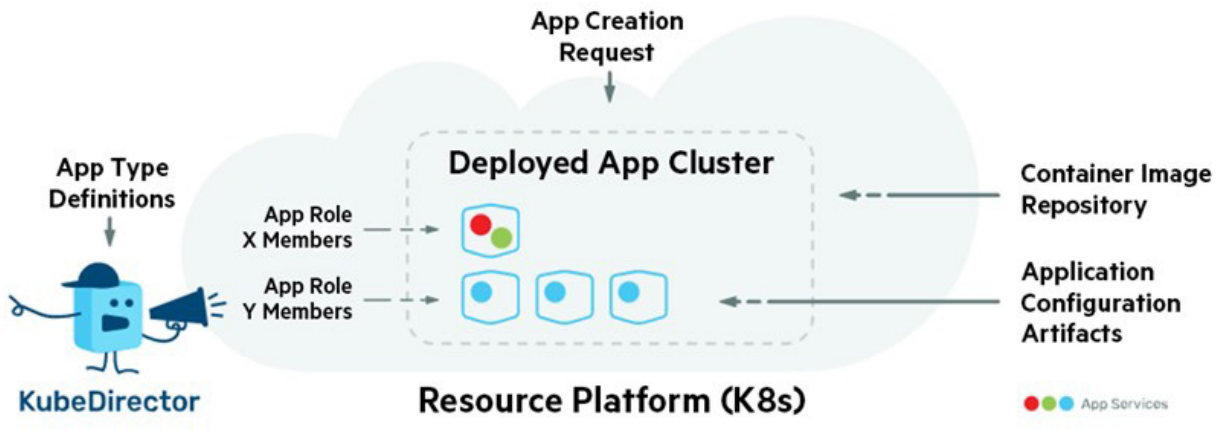

FIGURE 11. Easily deploy stateful apps on Kubernetes

KubeDirector offers these features that simplify importing an application into your Kubernetes cluster:

- Preconfigured operator to deploy stateful apps
- Register and deploy stateful apps through UI and API
- Built-in web terminal for each tenant with appropriate access controls
- Integrated Helm 3 clients in web terminal
- A shared browsable file system for uploading and deploying compute workloads/apps
- Ability to browse and edit YAML files for easy application deployment
- Automatically mapping internal ports to gateways and exposing endpoints as services on UI and through API
- Significantly improved productivity
- Simplifies deployment of stateful applications

From the web UI perspective, you get an application store experience providing a list of application tiles to choose from. KubeDirector manages the launch and monitoring of the applications for you.

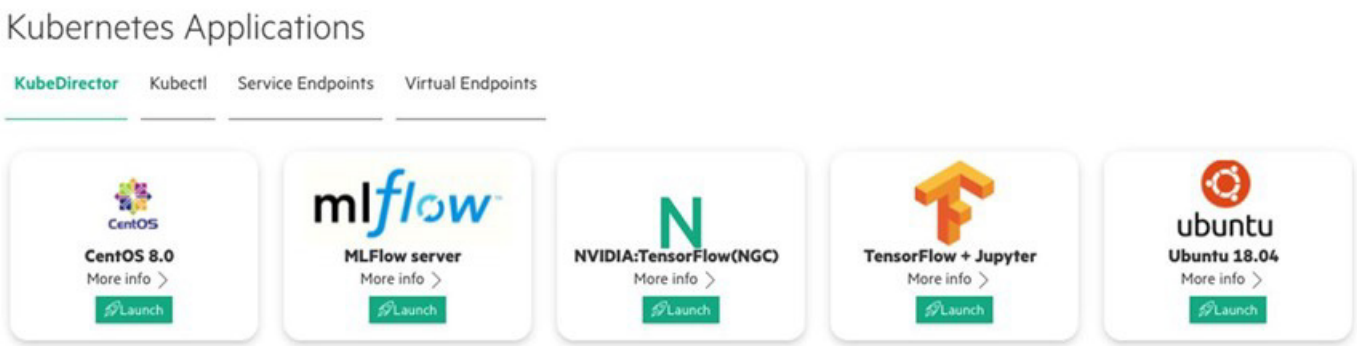

FIGURE 12. Application store

SECURITY AND ACCESS CONTROL

HPE Ezmeral Runtime Enterprise implements security throughout the platform. This includes several areas to consider:

- User authentication
- Role-based access
- Network security
  - Open Policy Agent security for any Kubernetes object security
- Kubernetes native security features
- Data fabric security
- Support for SPIFFE

User authentication

When HPE Ezmeral Runtime Enterprise is initially deployed by the site administrator, the installer has the option of integrating the organization's AD and LDAP configuration. Users can then be added and authenticated against AD/LDAP before being granted access to clusters and tenant resources. In addition, a local authentication option is built into the platform to allow users to be authenticated solely within HPE Ezmeral Runtime Enterprise.

Role-based access

Once an authenticated user is added to HPE Ezmeral Runtime Enterprise, the site administrator then makes decisions on several levels about which resources a given user can have access to, based on their role.

NETWORK SECURITY

As described in the "Cluster Management" section of this document, the first level of network security comes with two network interface designs and the use of a gateway. Second, the platform can be deployed using SSL. The end-user access to services in the containers (such as SSH or web applications) is routed through a gateway host that runs the HAProxy service.

All other traffic, including access to remote HDFS or other enterprise systems such as AD, MIT KDC (Kerberos provider), SSO (identity providers), and Certificate Authority (CA), is performed via the host network interface masquerading, as opposed to the gateway host port proxying. See also the "Network Planning" section of the online documentation for details.

Open Policy Agent (OPA) security policy management for Kubernetes objects

HPE Ezmeral Runtime Enterprise provides out-of-the-box security by providing built-in policies. Whenever a Kubernetes object, such as a POD violates a policy, the violation is reported on the dashboard.

Kubernetes native security features

A detailed discussion of the full range of networking security features built into Kubernetes is beyond the scope of this document. Since HPE Ezmeral Runtime Enterprise uses standard CNCF-certified Kubernetes, all security options included within Kubernetes clusters are supported in any Kubernetes cluster deployed using the container platform.

Data fabric security-secure by default

Security semantics are applied automatically as data is being stored and retrieved from the platform. HPE Ezmeral Data Fabric File & Object Store supports all four pillars of security-authentication, authorization, auditing, and encryption—as described in the HPE Ezmeral Data Fabric File & Object Store online documentation's "Policy-Based Security" section.

SPIFFE and SPIRE zero trust-authenticating software services across hybrid platforms

In addition to the methods discussed here around user authentication, network security, Kubernetes container security, and data security, there is another level of security available to users of HPE Ezmeral Runtime Enterprise—it is SPIRE zero trust. This security implements the CNCF-certified SPIFFE. This is a set of open-source standards for software identity.

SPIRE provides a universal identity control plane that reduces reliance on secrets or network-based security controls by leveraging strongly attested cryptographic identity to authenticate services across platforms using:

- **Authentication:** SPIRE's authentication mechanisms can be compared to user authentication. Implementing SPIRE, therefore, extends authentication to the applications, processes, and services running on HPE Ezmeral Runtime Enterprise.
- **Attestation-proof of identity:** SPIRE uses multiple factors to verify the identity of a service running on your platforms such as the location the service is running on, the OS version required for that service, and verification of the signature issued to the service.
- **Secretless authentication:** Using this strongly attested identity verification to confirm the authenticity of the services, the need for managing secrets is eliminated. This vastly simplifies application and service deployment and management.

As described earlier, SPIRE is used to implement SPIFFE on HPE Ezmeral Runtime Enterprise. SPIFFE is composed of five key parts:

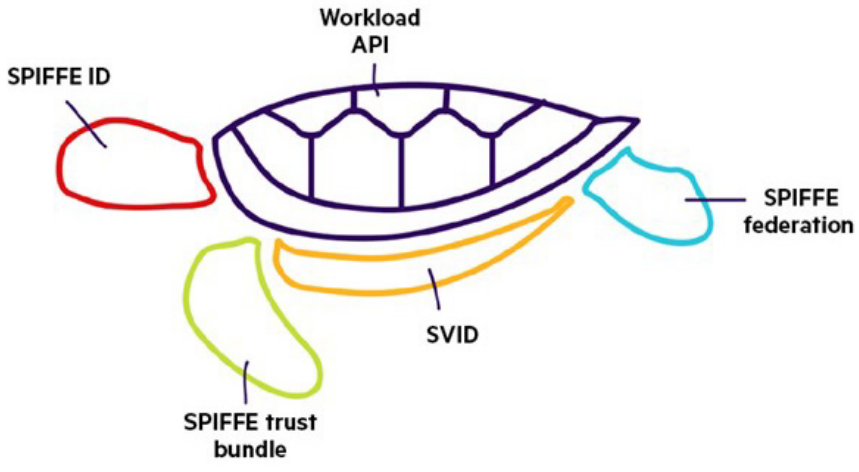

FIGURE 13. SPIFFE's five key parts

- SPIFFE ID: How a software service's name (or identity) is represented
- SVID: SPIFFE Verifiable Identity Document-a cryptographically verifiable document used to prove a service's identity to a peer
- Workload API: A simple node-local API that services use to obtain their identities without the need for authentication
- SPIFFE trust bundle: A format for representing a collection of public keys in use by a given SPIFFE issuing authority
- SPIFFE federation: A mechanism for sharing SPIFFE trust bundles

Using SPIFFE and SPIRE increases the security of the solutions being deployed on HPE Ezmeral Runtime Enterprise and further enables you to modernize and transform your applications so that they can leverage a hybrid- and multi-cloud infrastructure—securely.

All of this and more are detailed in the e-book: Solving the Bottom Turtle.

HIGH AVAILABILITY

HPE Ezmeral Runtime Enterprise provides a centralized management platform containing many pieces that work together to modernize your application and the data that drives them. Each of these pieces is protected from a single point of failure, as you would expect from an enterprise-class solution. In this section, we provide an overview of how the HPE Ezmeral software is highly available along with links to the online documentation for more details.

Levels of HA

The various types of HA that we will focus on include:

- Platform-level HA: Protection from controller host failure
- Kubernetes cluster HA: Protection within a cluster from master node failure
- Network/gateway host HA: Protection from gateway host failure
- Data fabric HA: Protection of data within a data fabric cluster

Platform-level HA

Platform-level HA protects against the failure of any one of the three hosts being used to provide this protection. The warning message will, therefore, appear if either the shadow controller or arbiter host fails, even if the controller host itself is functioning properly.

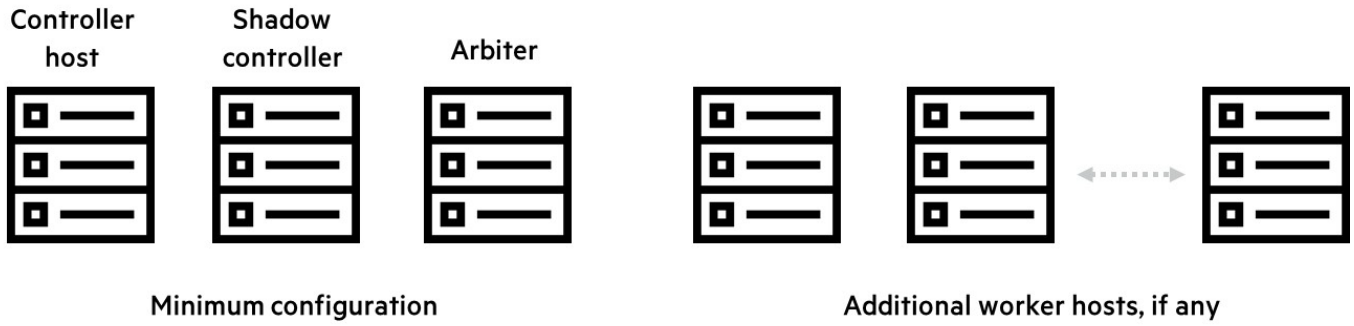

FIGURE 14. Platform HA

For more details on platform HA, see the HPE Ezmeral Runtime Enterprise online documentation.

Kubernetes cluster HA

When creating a Kubernetes cluster, it is possible to designate that one or more hosts be used as the master node. HPE Ezmeral Runtime Enterprise provides a simple interface that allows a user to simply select one or more available hosts as the master nodes.

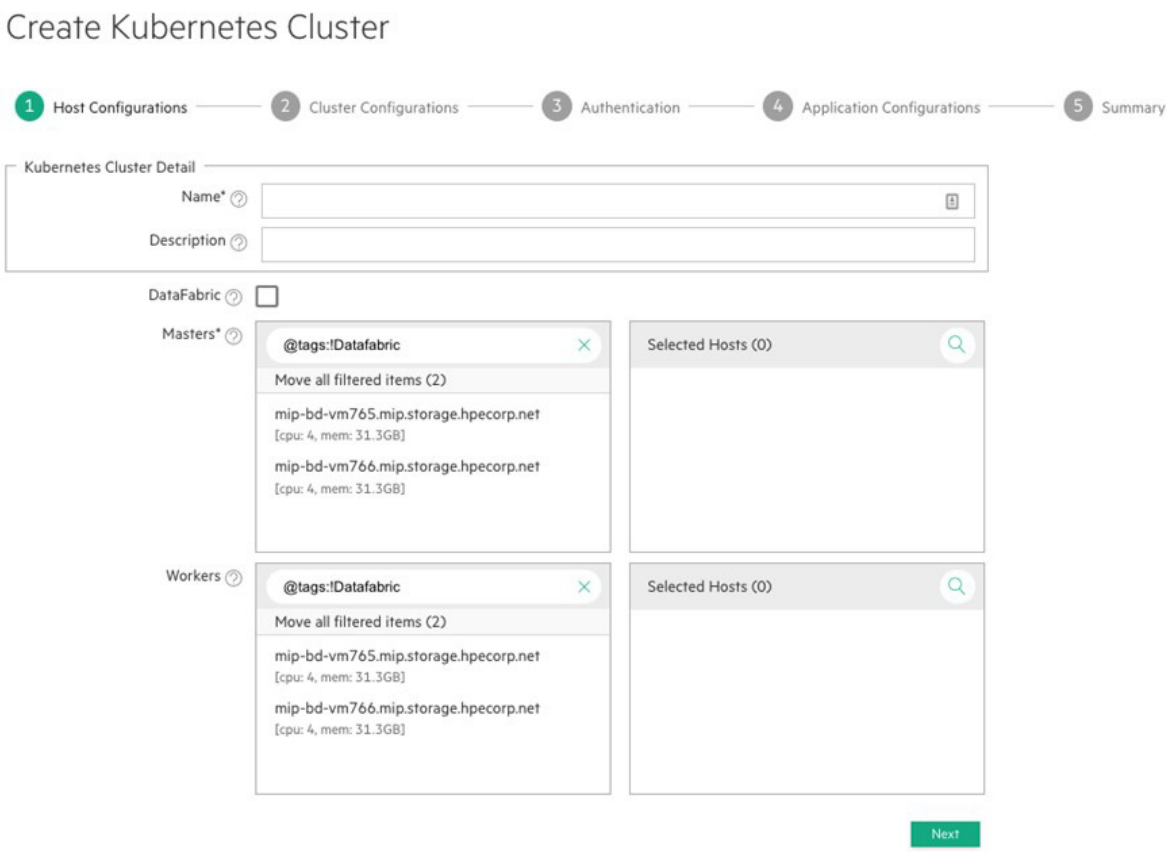

FIGURE 15. A user-friendly method for creating Kubernetes clusters

Once multiple masters have been selected, the Kubernetes HA master feature is automatically set up. The inner workings of Kubernetes HA are outside the scope of this document but can be reviewed at Kubernetes.io.

Network/gateway host HA

You can add redundancy for gateway hosts by mapping multiple gateway host IP addresses to a single hostname, as described in the "Gateway Hosts" section of this document. When this is done, then 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 basis. This ensures that there is no single point of failure for the gateway.

Data fabric HA

HPE Ezmeral Data Fabric File & Object Store provides HA management and data processing services for automatic continuity throughout the data fabric cluster. MapReduce services such as the Resource Manager, management services such as the ZooKeeper, and data access services such as NFS provide continuous service during any system failure.

There are several important cluster management components within a data fabric platform responsible for this automatic continuity: ZooKeeper, Warden, Container Location Database, and Control System. These and more are described in the online documentation.

Data fabric HA-specific features

In addition to being a HA platform itself, a user has many enterprise features to choose from when organizing their data. Starting with the concept of a volume itself, as is described in detail in the online documentation.

Volumes

Volumes are logical structures in the data fabric file system that hold files, directories, tables, and streams. Volumes are useful to administrators for applying policies to a large collection of items at once instead of an item at a time.

Volumes can enforce disk consumption limits (especially good for multi-tenancy purposes), and the volumes can also be useful for specifying replication levels (4x or 2x instead of 3x). Volumes also manage ownership (permissions), accountability (resource usage), and more.

Volumes snapshots

A snapshot is a consistent point-in-time copy of a volume that does not change over time. Snapshots are fast since they only capture metadata and consume very little space. They can run on a schedule or run on-demand as needed. They can be useful for recovering from human error or application error. Snapshots are also useful for ML model training since the data does not change over time. If the data changes when the model changes, it makes training difficult. See the "Managing Snapshots" section in the online documentation for more information.

Mirror volumes

A mirror volume is a read-only copy of a normal volume. Mirror volumes can be local (in the same cluster) or remote (in a different cluster). Local mirror volumes can offer additional read capacity for applications that need read-only access to data. Mirror volumes are updated incrementally, using the 8K blocks in the data fabric for atomic change propagation. Plus, block compression minimizes WAN bandwidth consumption. Mirror volumes can be promoted to full read/write if the primary volume becomes unavailable. Learn more in the "Mirror Volumes" section of the online documentation.

Node and volume topologies

Node topologies describe where each of the data nodes is inside the cluster. If you think about a large, multirack cluster, then you might think in terms of where each data node lives in a rack. Why configure the data fabric to tell which data nodes are in which racks? Because it helps when a whole rack fails or falls off the network due to the failure of a shared resource. The data fabric uses node topologies to spread data replicas so that a rack failure doesn't take down all the copies of the data. Learn more in the "Setting Up Node Topology" section of the online documentation.

SUMMARY OF KEY BENEFITS

The many benefits of the HPE Ezmeral portfolio help us understand that it truly enables, simplifies, and accelerates your journey to application modernization. In summary, we have discussed:

- **Greater choice:**
  Hybrid architecture provides the foundation for next-gen lakehouse analytics, delivering greater freedom and choice to data engineers, analysts, and scientists.
- **Greater flexibility:**
  A unified platform for orchestration of cloud-native and non-cloud-native applications on-premises, in any cloud, and at the edge.
- **Improved productivity:**
  Containerized applications deliver faster results and higher throughput on bare metal.
- **Reduced risk:**
  Extensive policy-based privilege management and control let you easily define and tailor access, trust levels, and privileges for people, teams, and data spaces.
- **No lock-in:**
  Using 100% open-source components and open APIs, you can pick up your application and data and move it without needing to refactor.
- **Improved ROI:**
  The utilization of hardware resources is improved by sharing a common data infrastructure across teams and workloads and provides a cloud-like experience for non-cloud-native monolithic applications, increasing the return on hardware investment.
- **How HPE Ezmeral Runtime Enterprise makes Kubernetes easy:**
  With a central control plane that offers both a point-and-click secure web UI experience and a REST API interface, you can deploy and manage your compute and storage resources anywhere—on-premises, in hybrid- / multi-cloud environments, or at the edge.
- **Data orchestration with HPE Ezmeral Data Fabric File & Object Store:**
  The product offers enterprise-grade data management with a global namespace architecture designed for distributed, large-scale secure data-intensive workloads.

- **Modernizing stateful applications with KubeDirector**
  Once your Kubernetes infrastructure is deployed, you can use the KubeDirector custom operator for transforming your previously on-premises siloed, data-driven applications into a modern, containerized infrastructure.
- **Security**
  HPE Ezmeral software offers security at the platform, network, user, and even at the zero-trust software services level. Your entire IT infrastructure can be modernized and deployed anywhere with the highest levels of security.
- **High availability**
  HPE Ezmeral software is built not just for enterprise-scale security and management but also for enterprise-grade high availability. The platform and every software component are architected to remain available in case of unexpected hardware, software, network, or application failure.

RESOURCES

HPE Ezmeral Runtime Enterprise

Request a demo

KubeDirector

HPE Ezmeral software

HPE Ezmeral Data Fabric File & Object Store

HPE Ezmeral Runtime Enterprise Online Documentation

HPE Ezmeral Data Fabric File & Object Store Online Documentation

HPE Developer—HPE Ezmeral Runtime Enterprise

HPE Developer—HPE Ezmeral Data Fabric

Learning portal

Customer references

LEARN MORE AT

HPE Ezmeral Runtime Enterprise

a50003809ENW, Rev. 2

Hewlett Packard Enterprise