HPE GreenLake

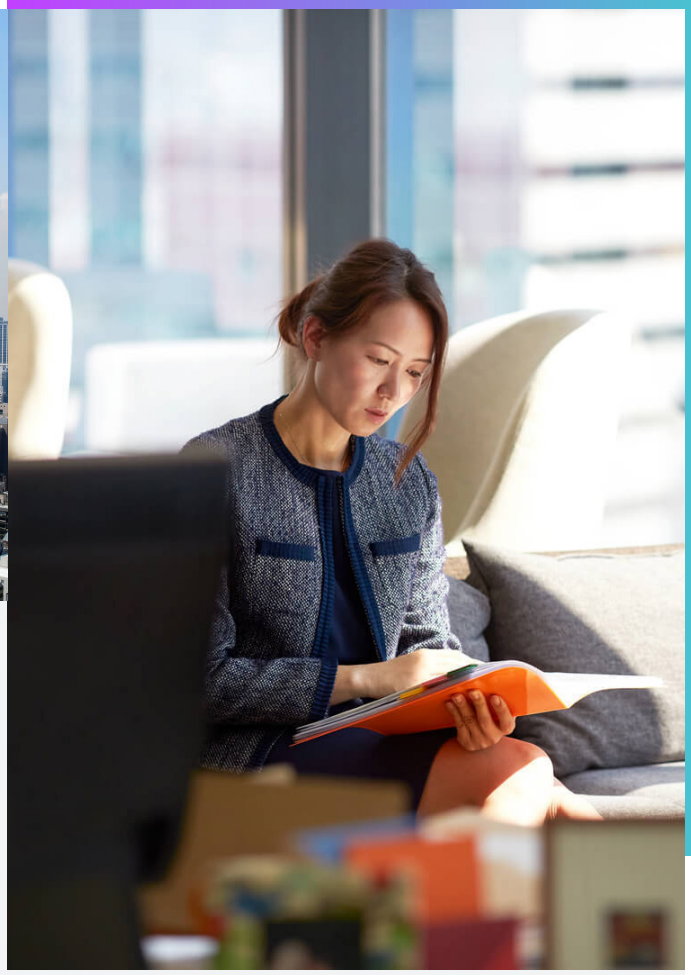

Multi-tenancy with HPE Ezmeral Data Fabric

Introduction

Organizations seek to share IT resources cost-effectively and securely among multiple data sets, user groups, and applications. Platforms that support this architecture are commonly known as multi-tenant technologies. Big Data platforms are increasingly expected to support multi-tenancy out-of-the-box, as non-native approaches are impractical. For example, maintaining separate clusters for separate tenants requires excessive administrative overhead and introduces significant processing latency. Also, partial solutions such as Apache Hadoop YARN or other resource managers only handle one definition of tenant (such as, the tasks/applications) and ignore the most-critical component of a multi-tenant environment—the data.

The key to multi-tenancy is the ability to ensure strict isolation of the distinct tenants while also allowing some level of sharing when necessary. The metaphor of an apartment building applies well, where some components allow isolation (residents in distinct, private units) and other components allow sharing (elevators, hallways, fitness center, pool, and such).

Use case overview

Multi-tenancy is useful and critical in a range of cases. Three common use cases are:

Enterprise data lake: Often, organizations start using HPE Ezmeral Data Fabric in a specific area, for example, data warehouse optimization for the marketing department. Soon, the customer service team also wants to run an app for churn prediction. Ultimately, questions arise as to who has access to which data and for what purposes, also taking into account regulatory issues (such as the Sarbanes–Oxley Act in banking or data protection legislation throughout the industry). With HPE Ezmeral Data Fabric, separate departments can share the same cluster, with some data shared across departments and some data kept private within each department. The multi-tenancy from HPE Ezmeral Data Fabric helps ensure the users and user groups can only access the data for which they are authorized.

Software/Platform/Infrastructure as a service: Some organizations provide IT services—such as Big Data as a service—to internal or external customers. A basic requirement of such service providers is to isolate customers while achieving guaranteed service-level agreements (SLAs), be it in terms of availability or latency. In this scenario, one customer accessing another customer's data could be disastrous for the business. In addition, a service provider may require the flexibility to be able to run parts of the multi-tenancy deployment in a hybrid cloud setup, for example, to benefit from the elasticity of public clouds.

Data lifecycle management: Data often undergoes many transformations to suit the business needs of various departments and user groups. Multi-tenancy is a key enabler of data lifecycle management, where each group of users is necessarily responsible for different stages of the data. For example, stages may include raw ingested data, data scientist transformations, BI analyst-ready data, and archived data. Data in each stage can be stored as a separate tenant, to make sure only authorized users can access it. Also, with multi-tenancy, administrators can introduce constraints around how much compute and storage each stage is allowed to use.

HPE Ezmeral Data Fabric offers built-in multi-tenancy capabilities to support both isolation and sharing in a single cluster. Volumes are the foundation of multi-tenancy. They are logical partitions you create in the cluster upon which you can set specific policies. Volumes group together related directories, files, database tables, and streams as a cohesive unit to help simplify Big Data management. You can use volumes to enforce disk usage limits, set replication levels, establish ownership and accountability, and measure the costs of different projects or departments. A single cluster can have many volumes, up to hundreds of thousands. Volumes automatically grow across multiple nodes in the cluster as you insert more data in them, unlike static Linux® partitions.

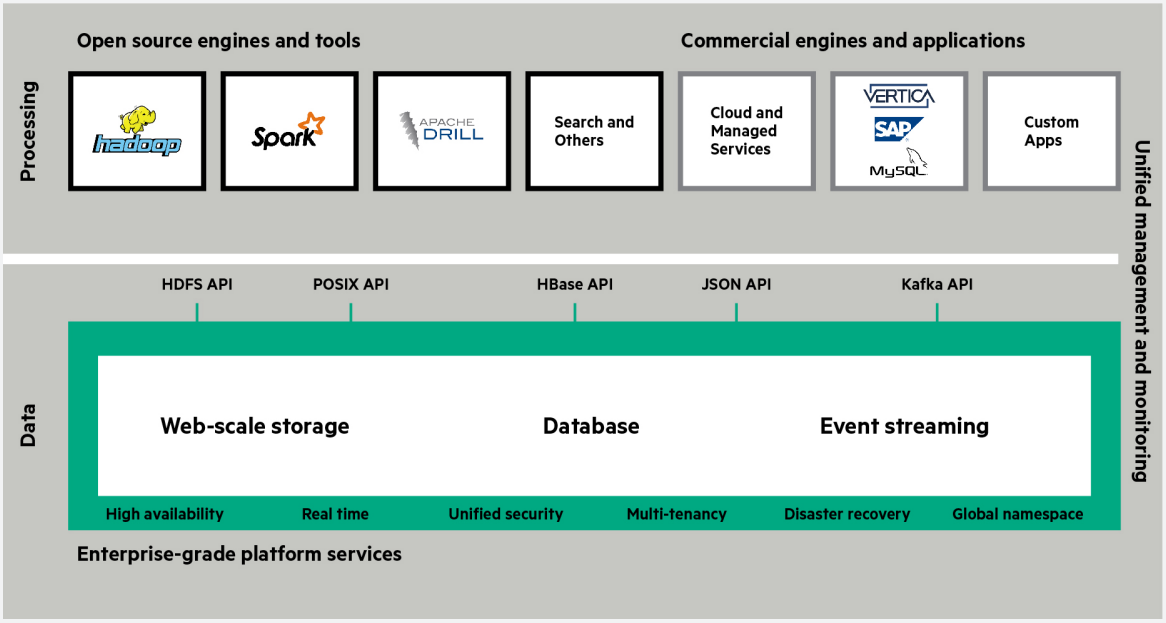

Figure 1. HPE Ezmeral Data Fabric

In a typical deployment, the private data for each user group or application is stored in a dedicated volume, so that it can be isolated and managed separately from the data of other user groups and applications. Access controls on the data ensure isolation and can be applied at the directory, file, database table, stream, and/or volume levels. Separate volumes can be created to allow information sharing, including sharable information from multiple user groups and applications.

HDFS-based data platforms do not support volumes and thus require administrative trade-offs. Policies defined at the cluster level are too coarse-grained and conflict with the notion of multi-tenancy. Policies can be defined at the file or directory level, but that requires significant administrative overhead and adds high risk for error, especially with regard to gaps in access controls. As a workaround, organizations using HDFS-based data platforms create separate physical clusters for each tenant, which add architectural complexity thus higher risk of other errors.

Multi-tenancy in HPE Ezmeral Data Fabric also has a significant total cost of ownership (TCO) advantage. Organizations can leverage a single cluster for multiple use cases rather than maintaining a large number of isolated clusters. This reduces overall administrative overhead as well as increases efficiency due to shared hardware allocation.

Data placement control

HPE Ezmeral Data Fabric supports volume topology, also known as data placement control, which lets you restrict a volume to a specific set of nodes in the cluster. This feature lets you take advantage of a cluster of heterogeneous hardware, by placing specific data sets on hardware servers most appropriate for the usage of the data set. For example, data placement control can be used for a multi-temperature data strategy in which operational data is stored on fast nodes with SSDs and archived data on slower, disk-dense nodes. This feature also helps with resource management to provide SLAs and resource availability for specific user groups by letting them run only their workloads on their own dedicated hardware.

Job placement control

HPE Ezmeral Data Fabric provides the ability to run jobs on specific nodes, which is handled either by applying labels to the queues to which jobs are submitted or to the jobs themselves. A queue label specifies which set of nodes will run the jobs in that given queue and a job label specifies which set of nodes will run the given job. This capability lets administrators determine SLAs for specific applications and create separation between different applications or business units. This feature also allows administrators to designate a subset of the nodes for low-priority jobs (such as experiments) or jobs that require access to external systems through the corporate firewall.

Security and access control

HPE Ezmeral Data Fabric provides cryptographically secure authentication and encryption. Organizations that have a Kerberos infrastructure can leverage it for authentication, while organizations that do not have a Kerberos infrastructure can leverage the native ticket-based, user name/password authentication scheme that provides the same security without the complexity associated with deploying and managing Kerberos.

HPE Ezmeral Data Fabric offers fine-grained access control with access control expressions (ACEs) for user and role-based access control (RBAC) on volumes, files, directories, database tables, documents, column families, columns/elements, and streams. ACEs at the volume level (whole-volume ACEs) provide a safeguard mechanism against user errors around access controls. For a user to have access to a specific item, such as a file or directory, the user must have permission for that item, as well as on the volume in which it resides.

For example, if you have a volume for customer X, you set a whole-volume ACE on it, so only customer X can access the data in the volume. If a user gives permission on specific items in the volume to users outside of customer X, the whole-volume ACE will still block the access. No matter what permissions are set on the files and directories in that volume, the whole-volume ACE allows only customer X can access the data.

Administration and reporting

From an administrative perspective, HPE Ezmeral Data Fabric allows organizations to define and enforce storage, CPU, and memory quotas. Flexible policies can be set at the volume, user, and group levels to match security and/or regulatory requirements. Monitoring from HPE Ezmeral Data Fabric provides real-time usage information, so administrators can easily understand where resources are used.

For service providers who need to provide accurate usage and billing information, HPE Ezmeral Data Fabric offers reporting on resource usage on over 60 different metrics, available via the control system browser-based user interface. It is also available through the command-line interface and REST API for integration with third-party systems.

Conclusion

Enterprise architectures that include data from many sources often require multi-tenancy capabilities to provide data isolation as well as data sharing. HPE Ezmeral Data Fabric offers multi-tenancy support out-of-the-box to manage distinct data sets, user groups, and applications in the same cluster. This capacity is achieved through volumes, enabling isolation for both compute and storage, and includes security and reporting.

Learn more at

hpe.com/info/data-fabric

Hewlett Packard Enterprise

a50001593ENW, Rev. 2