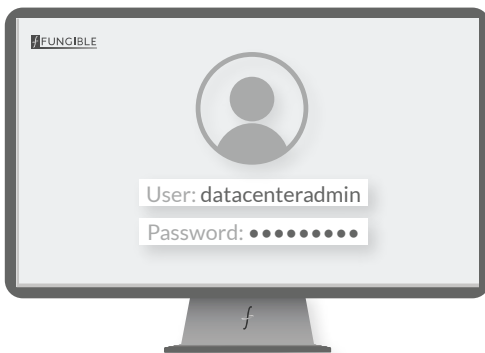
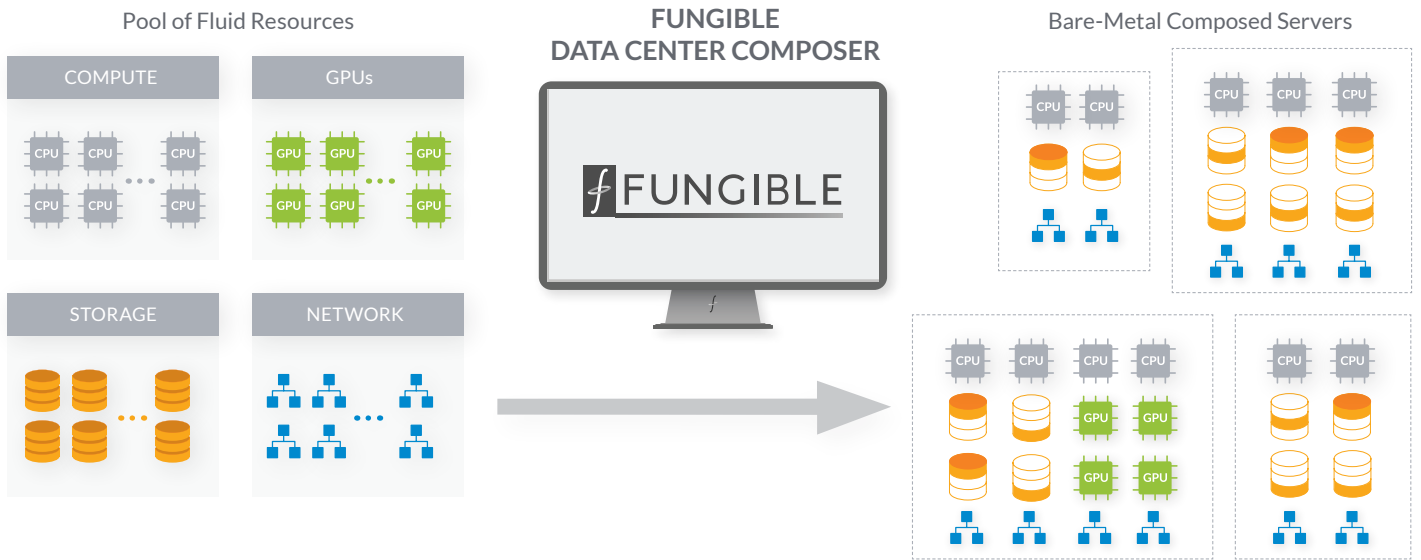


THE FUNGIBLE DATA CENTER – HOW IT WORKS

The Fungible Data Center (FDC) comprises pools of compute, storage, and networking hardware managed by the Fungible Data Center Composer software.



The Fungible Data Center Composer supports role-based access control, presenting two levels of administrative control to FDC users.




Data Center Administrator (DCA) – A role that manages the overall Fungible Data Center. The DCA partitions the compute, storage, and network resources and allocates them to different Logical Partitions (LPs).



Logical Partition Administrator (LPA) – A role that supports the tenant for which the partition is allocated to. The LPA uses the resources allocated by the DCA to compose bare-metal servers based on the requirements of application workloads. The LPA is also responsible for deploying these application workloads within the LP.

A Logical Partition (LP) is an abstraction representing one tenant.



The DCA is responsible for supporting multiple tenants, ensuring there is complete isolation amongst the tenants.



DCAs can see the aggregated utilization metrics of a logical partition but not how resources are being used inside a logical partition.

Once resources are allocated to the LP, the LPA is responsible for managing the resources and management privileges for that LP. LPAs are empowered to be self-sufficient through self-service management portals and APIs.

WORKLOAD PROVISIONING

Within the Logical Partition, the role of a Logical Partition Admin (LPA) includes provisioning an application workload to its tenant. LPAs can deploy application workloads within the Logical Partitions with the server, storage, and networking resources allocated to them by the DCA.

Logical Partition
DEVOPS
Admin-tenant1

An LP can have multiple environments. A specific environment would contain a set of resources needed to run a particular application workload.

ENVIRONMENT A Workload Splunk
Subfunction: Search head
Subfunction: Indexer

ENVIRONMENT B Workload Cloudera

ENVIRONMENT C Workload Kubernetes Containers

For example, in a workload like Splunk, there are two distinct sub-functions: search heads and indexer nodes that require different server configurations, hence, different composition profiles.

ENVIRONMENT A

SEARCH HEAD SERVER PROFILE

COMPUTE SPECIFICATION:

- CPU = 24 cores
- Memory = 128 GB
- Model = SuperMicro SuperServer 1029U-TR4

STORAGE SPECIFICATION:

- Boot disk = 20 GB
- Data disk = 6 TB
- Cache disk = 2 TB

NETWORK SPECIFICATION:

- NIC 1 = vNET A
- NIC 2 = vNET B

INDEXER SERVER PROFILE

COMPUTE SPECIFICATION:

- CPU = 16 cores
- Memory = 64 GB

STORAGE SPECIFICATION:

- Boot disk = 20 GB
- Data disk = 4 TB

NETWORK SPECIFICATION:

- NIC 1 = vNET C

ONE-CLICK WORKLOAD DEPLOYMENT

To further simplify workload deployment, LPAs can create environments from scratch or instantiate environments from existing templates. The FDC allows for templates to be created to represent a workload environment and to be published into an application marketplace. Such templates allow for fast, one-click deployment of applications.

APPLICATION MARKETPLACE

Template A

Template B

Template X