

About Multi-Site Replication

Contents

- [Overview](#)
- [Multi-Site Replication Benefits](#)
- [About Active-Passive Topology](#)
- [About App-Layer Active-Active Topology](#)
- [About Failover and Switchover](#)
- [Infrastructure Requirements for Multi-Site Replication](#)
- [Multi-Site Replication Limitations](#)

This topic describes how the multi-site replication topology works in MySQL for VMware Tanzu and contains information to help users decide whether to use the Multi-Site Replication plan.

Overview

The multi-site replication topology in MySQL for VMware Tanzu is a disaster recovery solution that enables developers to provision a leader-follower instance across two foundations. This leader-follower service instance is comprised of two single node instances that are configured for replication. Operators can configure two foundations to be in the same data center or spread across multiple data centers or geographical regions.

The foundation types are:

- **Primary Foundation:** This foundation is deployed in your main data center. Generally, this data center receives the majority of app traffic. MySQL for VMware Tanzu assumes that the leader is deployed on this foundation when healthy.
- **Secondary Foundation:** This foundation is deployed in your failover data center. Generally, this data center receives less app traffic than the primary foundation, or no traffic at all. MySQL for VMware Tanzu assumes that the follower is deployed on this foundation unless a developer triggers a failover.

For more information about foundations, see [About Tanzu Foundations](#).

Note

The Multi-Site Replication plan type is configured separately from the leader-follower service plan.

For information for operators about enabling multi-site replication, see [Preparing for Multi-Site Replication](#).

For information for developers about using multi-site replication, see [Using MySQL for VMware Tanzu for Multi-Site Replication](#).

Multi-Site Replication Benefits

The Multi-Site Replication plan has the following benefits:

- **Resilience for Service Instances:** Developers can trigger a failover to maintain app uptime during a data center outage. For more information, see [Triggering Multi-Site Replication Failover and Switchover](#).
- **Data Center Upgrades with Zero Downtime:** Operators can upgrade data centers without taking databases offline by triggering a switchover first.
- **Support for Multiple Cloud Deployment Models:** Operators can configure multi-site replication with a single-cloud or hybrid-cloud deployment model. Both deployment models have the same end-user experience.
- **Support for Active-Passive and App-Layer Active-Active Disaster Recovery:** For more information, see [About Active-Passive Topology](#) and [About App-Layer Active-Active Topology](#) below.

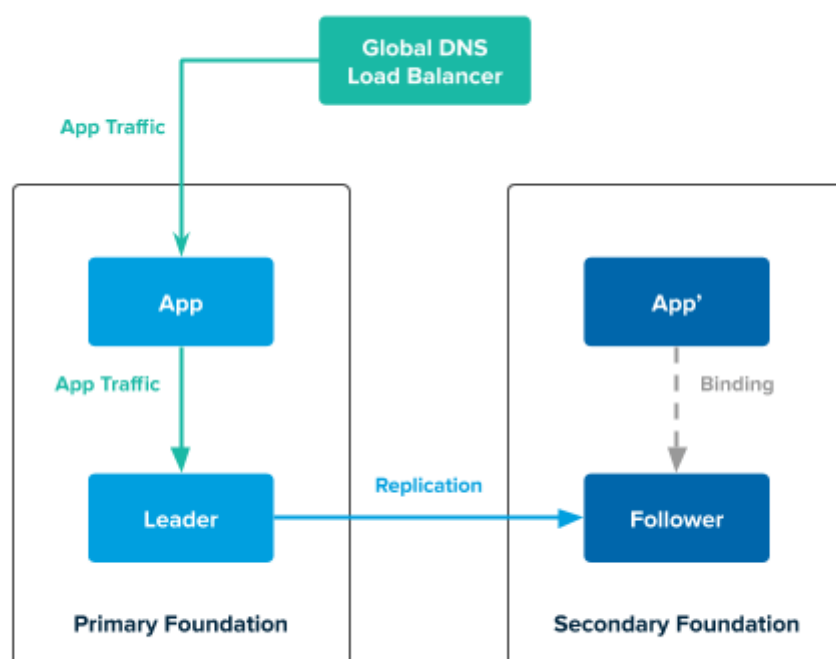
About Active-Passive Topology

In an active-passive topology, when all foundations and workloads are healthy, all app traffic is directed to the primary foundation. The secondary foundation receives no app traffic.

VMware recommends using this topology when your secondary foundation is scaled down, generally inactive, and does not receive significant app traffic.

For information about active-passive failover and switchover, see [About Failover and Switchover](#) below.

The following diagram describes the active-passive topology in a healthy state:



In the above diagram:

- The global DNS load balancer (GLB) directs traffic to the app in the primary foundation. This app issues transactions to the leader service instance.
- The follower service instance in the secondary foundation continuously replicates data from the leader service instance in the primary foundation.

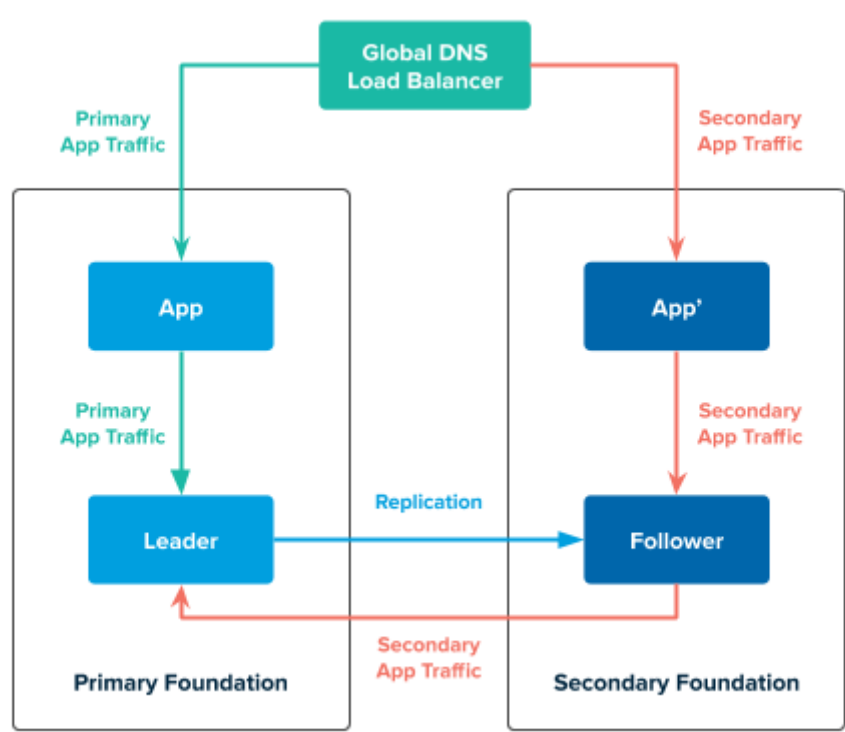
About App-Layer Active-Active Topology

In an app-layer active-active topology, when all foundations and workloads are healthy, app traffic is directed to both the primary and secondary foundation.

VMware recommends using this topology when both your primary and secondary foundations are scaled up and are expected to serve traffic.

For information about app-layer active-active failover and switchover, see [About Failover and Switchover](#) below.

The following diagram describes the app-layer active-active topology in a healthy state:



In the above diagram:

- The GLB directs traffic to the apps in the primary and secondary foundations. The apps in the primary and secondary foundation issue transactions to the leader service instance.

The app in the secondary foundation issues transactions to the leader as below:

1. Connects to the follower service instance
 2. Issues transactions to the follower service instance
 3. Forwards the transactions from the follower service instance to the leader
- The follower service instance in the secondary foundation continuously replicates data from the leader service instance in the primary foundation.

Note

In the Multi-Site Replication plan, the app-layer active-active topology does not enable multi-master replication. Therefore, the follower service instance is read-only. Apps can only write to the leader.

About Failover and Switchover

MySQL for VMware Tanzu prioritizes data consistency over availability. Therefore, MySQL for VMware Tanzu does not trigger failover or switchover automatically and developers must manually trigger a failover or switchover. For instructions about triggering failover or switchover, see [Triggering Multi-Site Replication Failover and Switchover](#).

The table below describes when you should trigger a failover or switchover:

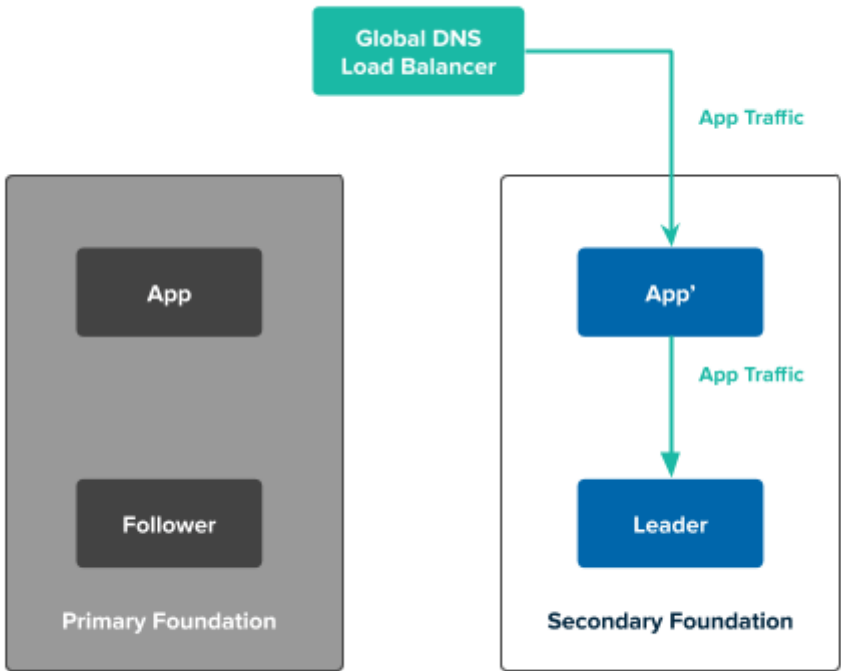
Failover if	Switchover if
<ul style="list-style-type: none">• The leader MySQL process has crashed or is unhealthy and is not automatically recovered by BOSH.• The leader VM is destroyed or unrecoverable.• The availability zone (AZ) for the leader VM experiences an unexpected outage.• The data center for the leader VM experiences an unexpected outage.	<ul style="list-style-type: none">• Both the leader and the follower instance are healthy.• You plan to do foundation or data center upgrades or maintenance on your primary site. For example, upgrading stemcells or data center hardware.• You plan to degrade performance on the primary site.

For information about multi-site replication topologies in a healthy state, see [About Active-Passive Topology](#) and [About App-Layer Active-Active Topology](#) above.

About Failover

In both the active-passive and app-layer active-active topologies, if a developer triggers a failover, all app traffic is directed to the leader in the secondary foundation and the primary foundation receives no app traffic.

The following diagram describes what happens when you trigger a failover for a multi-site replication topology:



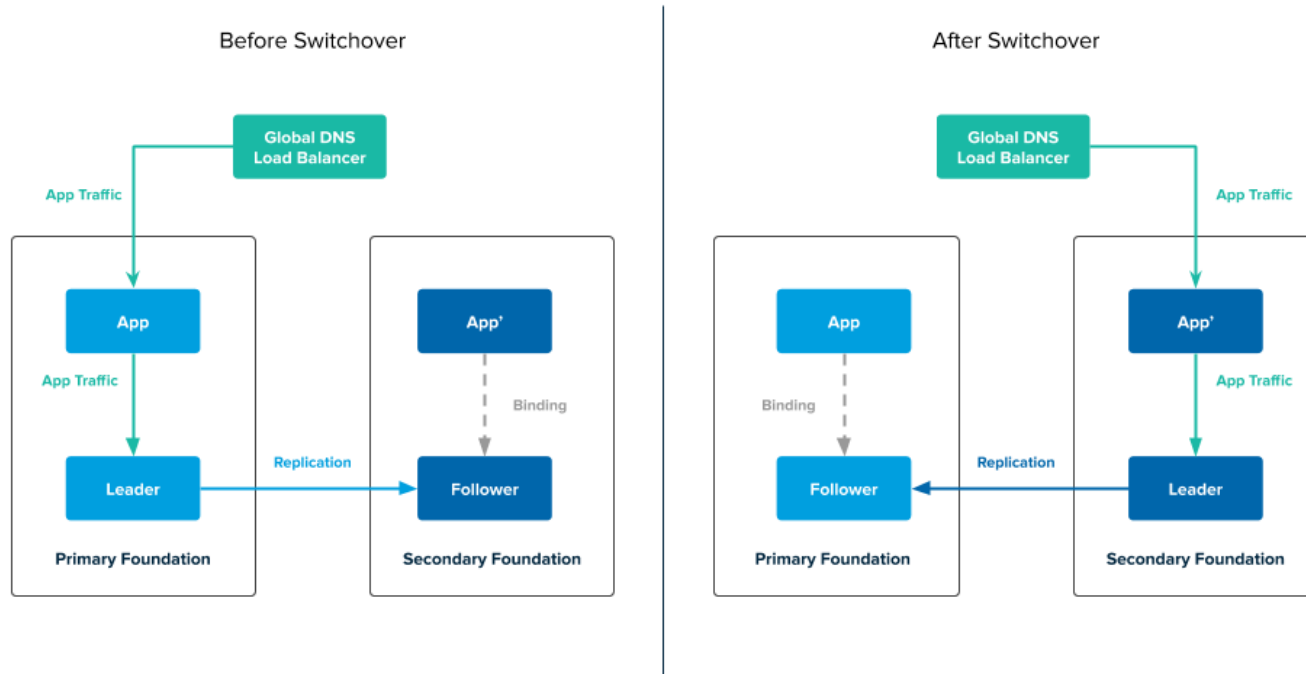
In the above diagram:

- The GLB directs traffic to the app in the secondary foundation. This app issues transactions to the leader service instance in the secondary foundation.
- The leader service instance in the secondary foundation does not replicate data to another service instance.

About Active-Passive Switchover

In an active-passive topology, if a developer triggers a switchover, all app traffic is directed to the leader in the secondary foundation. The primary foundation receives no app traffic.

The following diagram describes what happens when you trigger a switchover in an active-passive topology:



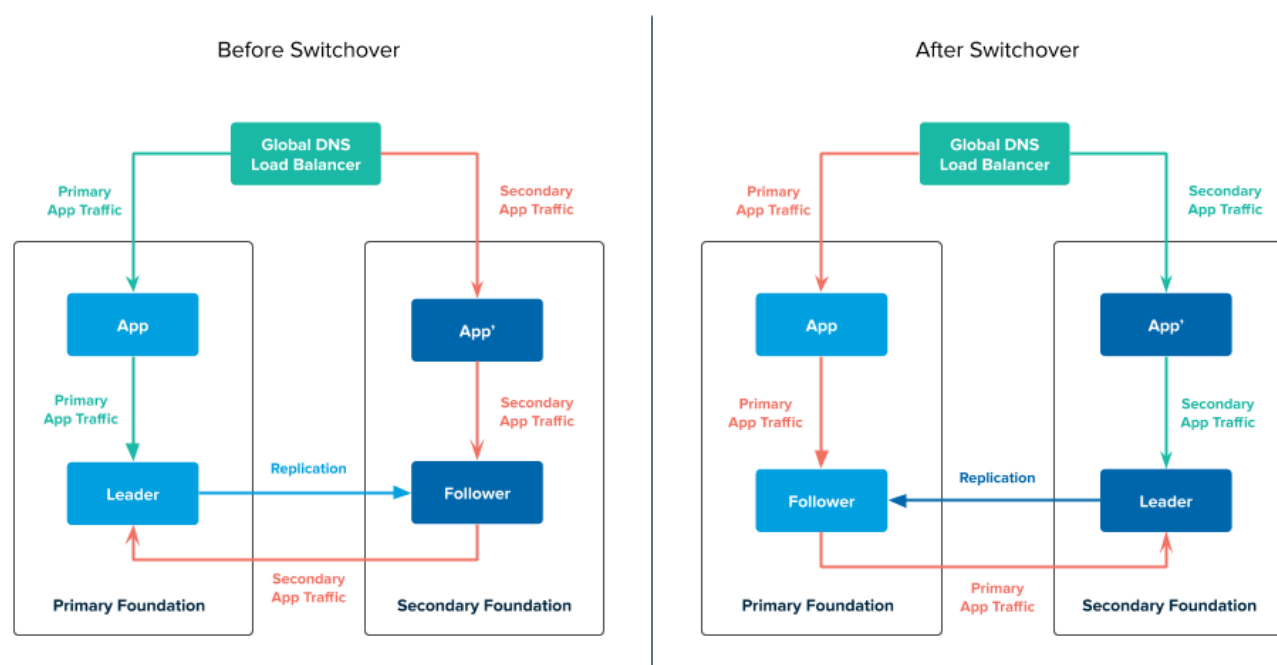
In the above diagram:

- The GLB directs traffic to the app in the secondary foundation. This app issues transactions to the leader service instance in the secondary foundation.
- The leader service instance in the secondary foundation replicates data to the follower service instance in the primary foundation.

About App-Layer Active-Active Switchover

If a developer triggers a switchover, app traffic is still directed to both the primary and secondary foundation. However, the leader service instance is in secondary foundation.

The following diagram describes what happens when you trigger a switchover in an active-active topology:



In the above diagram:

- The GLB directs traffic to the apps in the primary and secondary foundations. The apps in the primary and secondary foundation issue transactions to the leader service instance in the secondary foundation.

The app in the primary foundation issues transactions to the leader as below:

1. Connects to the follower service instance
 2. Issues transactions to the follower service instance
 3. Forwards the transactions from the follower service instance to the leader
- The follower service instance in the primary foundation continuously replicates data from the leader service instance in the secondary foundation

Infrastructure Requirements for Multi-Site Replication

Before using the Multi-Site Replication plan, consider the following requirements.

Capacity Planning

When calculating IaaS usage, you must take into account that each Multi-Site Replication instance requires two VMs. Therefore, the resources used for the Multi-Site Replication plan are twice the amount of the single node plan.

For more information, see [Setting Limits for On-Demand Service Instances](#) and [Persistent Disk Usage](#).

Networking Requirements

Operators should consider that Multi-Site Replication instances that are deployed in data centers that are geographically farther apart experience higher latency.

For information about standard networking rules, see [Required Networking Rules for Multi-Site Replication](#).

Multi-Site Replication Limitations

The Multi-Site Replication plan has the following limitation:

- **Synchronous replication is not supported:** For Multi-Site Replication plans, any data that is written to the leader is asynchronously replicated to the follower.
-