




Using DR to Ensure High Availability of GridGain Clusters

Konstantin Orlov

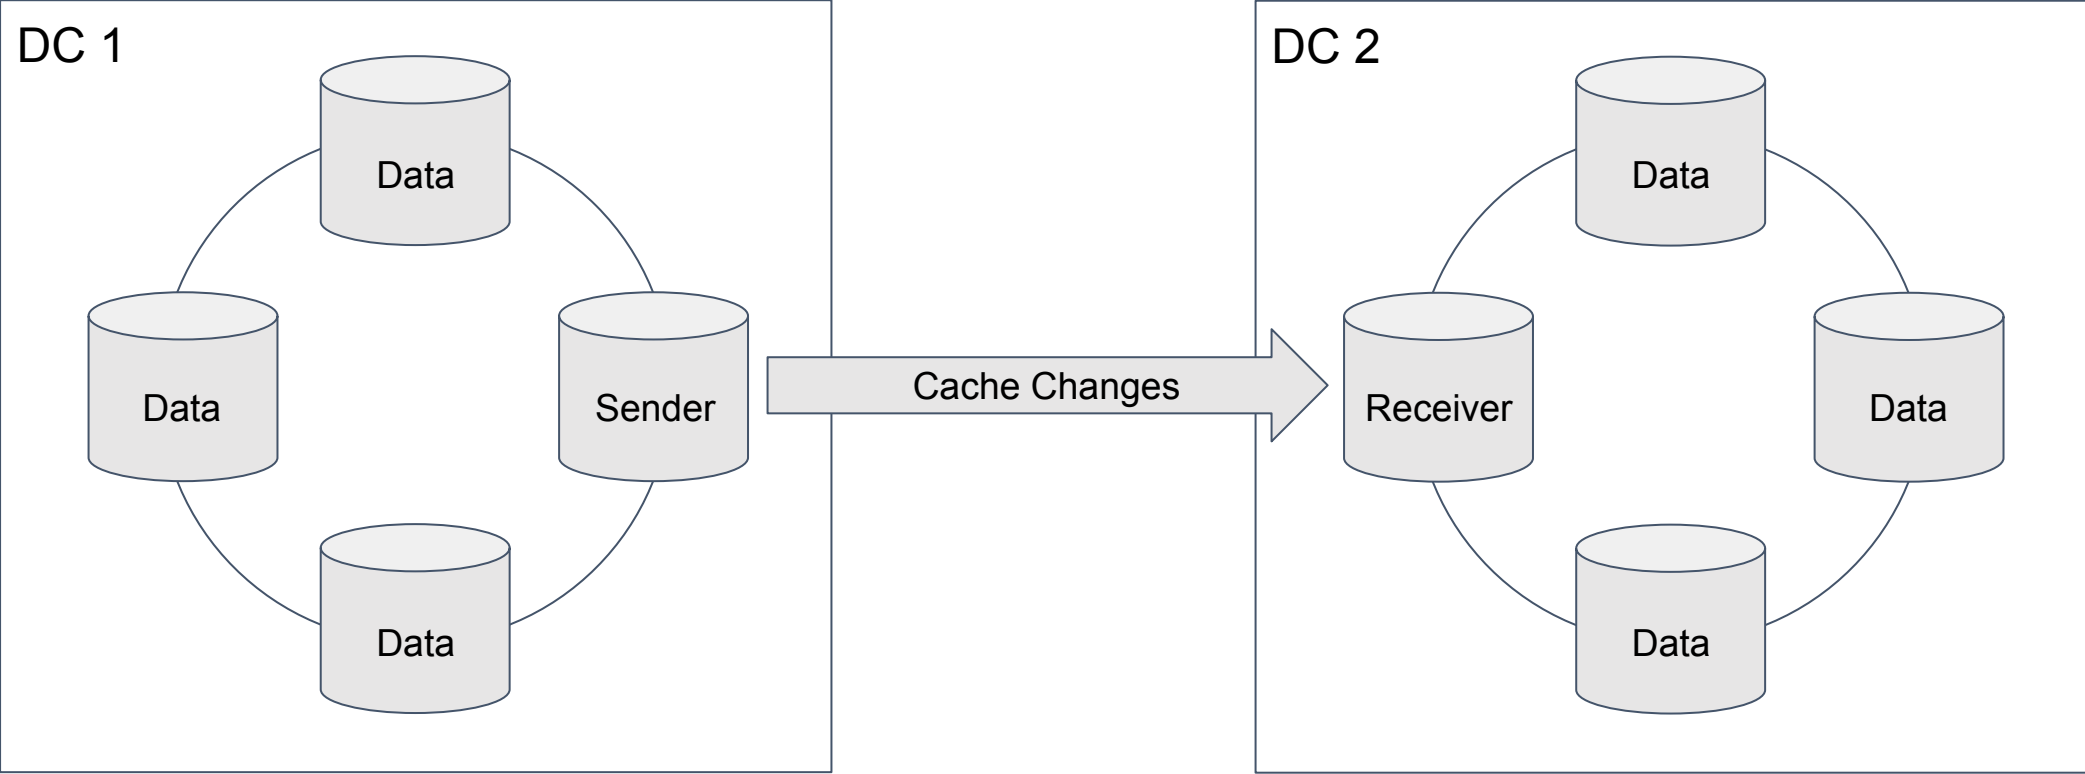
May 13



Part 1. Concepts and Terms



What Is The Data Center Replication

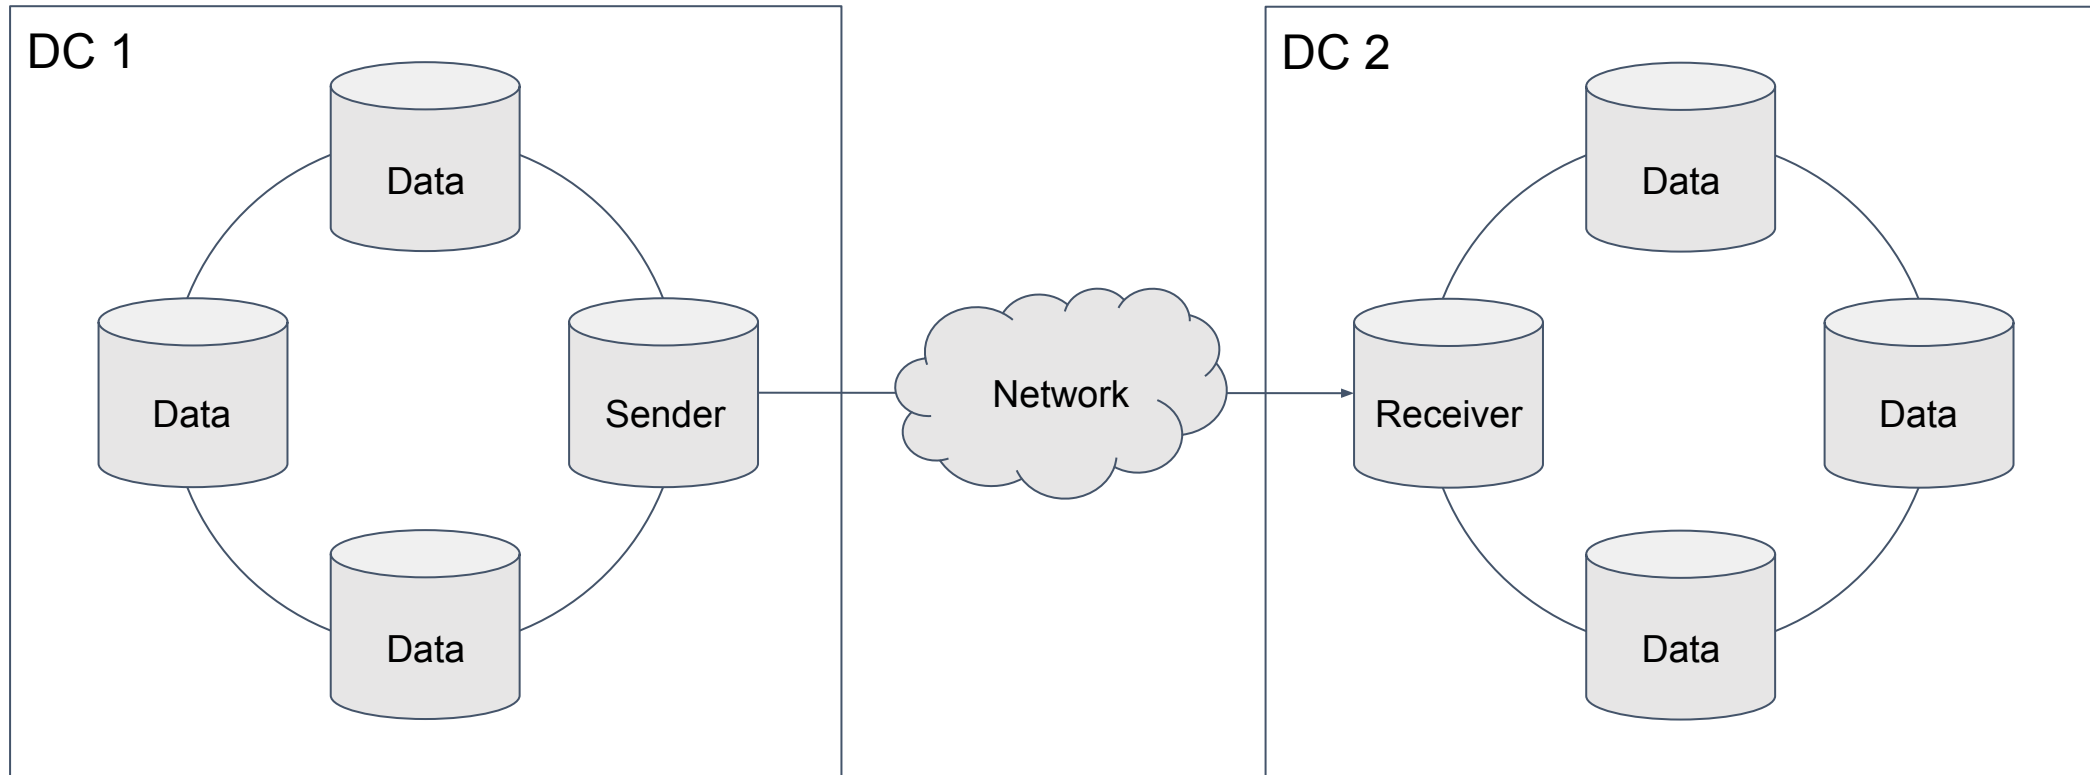


When Should You Use Replication

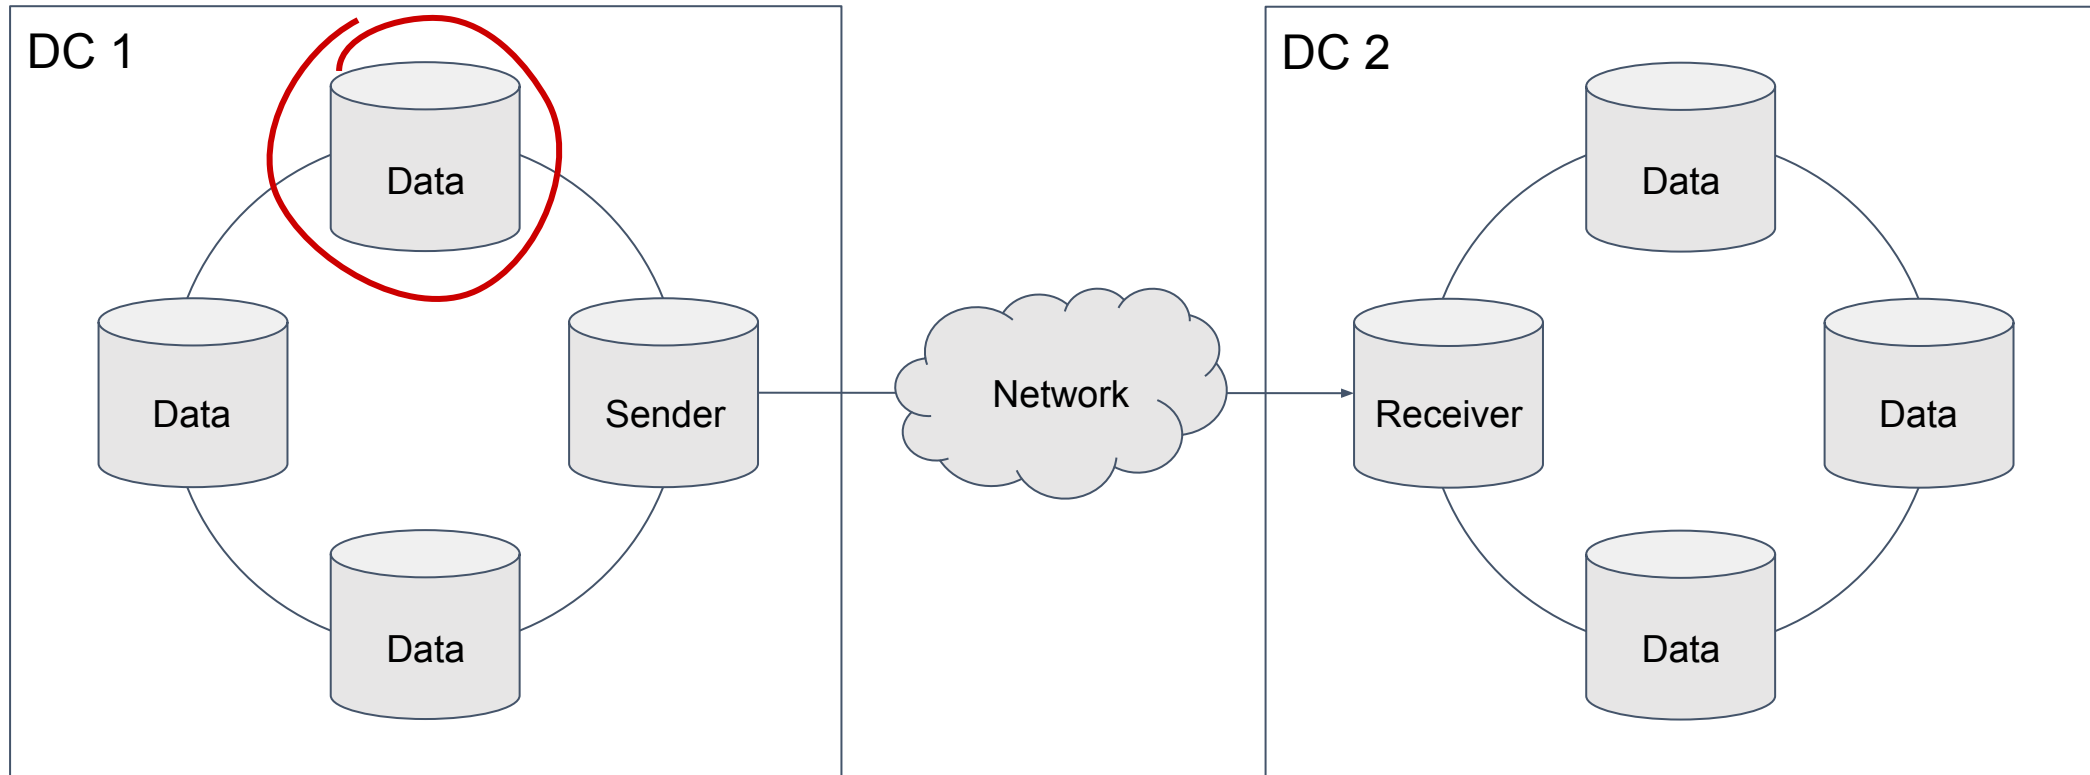


- When you need an online replica for recovery
- When you need a read-only replica for OLAP
- When you need an active replica for load balancing

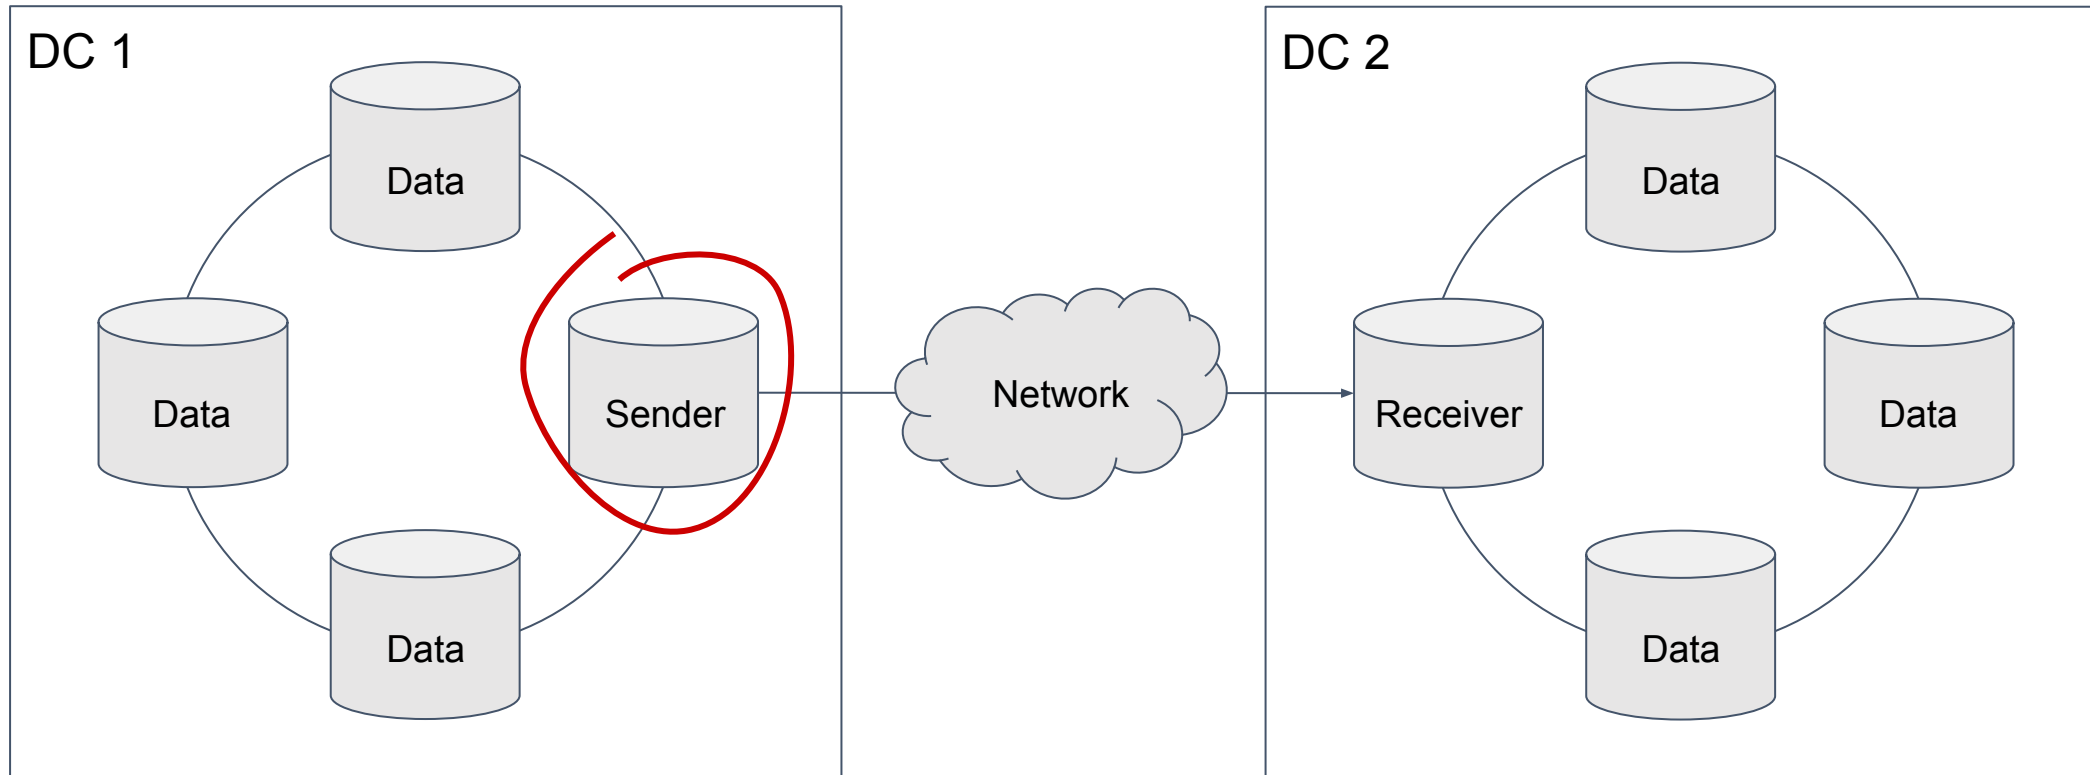
Data Center Replication Parts



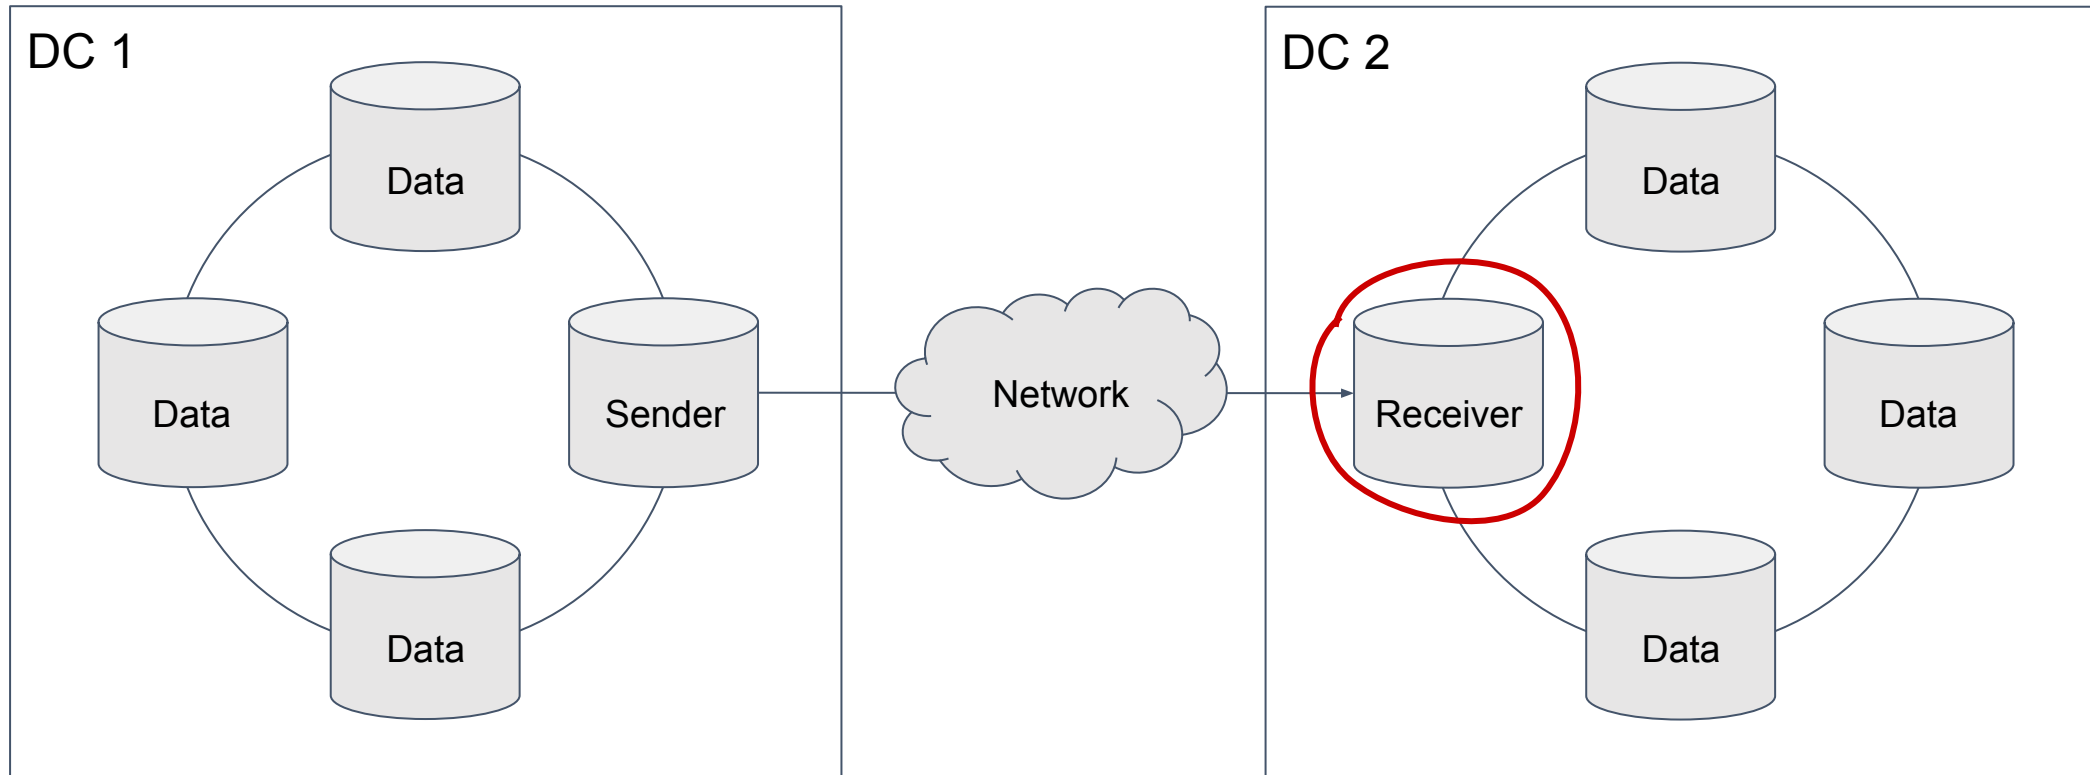
Data Center Replication Parts



Data Center Replication Parts



Data Center Replication Parts



Part 2. Demo



Scenario 1. Preconditions



- Replication mode is active-passive
- Replication is started
- Cache is empty

Scenario 1. Configuration Key Points



- All nodes have a GridGainConfiguration with proper dataCenterId configured
- At least one node from the active topology has a proper DrSenderConfiguration
- At least one node from the passive topology has a proper DrReceiverConfiguration
- Replicated cache has a proper GridGainCacheConfiguration

Scenario 2. Preconditions



- Replication mode is active-passive
- Replication is started
- Cache is NOT empty when replication starts up

Scenario 2. Replication limitations



- State transfer is required if replications starts on non-empty cache
- Cache in remote DC should be empty since state transfer does not transfers removes

Scenario 3. Preconditions



- Replication mode is active-active
- Replication is started on both sides
- Cache is empty on both sides

Scenario 3. Configuration Key Points



- All nodes have a GridGainConfiguration with proper dataCenterId configured
- At least one node from both topologies has a proper DrSenderConfiguration
- At least one node from both topologies has a proper DrReceiverConfiguration
- Cache configuration has a proper GridGainCacheConfiguration with conflict resolver configured

Scenario 4. Preconditions



- Replication mode is active (doesn't matter)
- Replication is started
- Cache is empty on both sides

Scenario 4. Conclusion



GridGain has no standard mechanism to handle failover.

Possible approaches to handle it:

- Reconnection to another topology
- Keeping connections to several topologies and switching between them
- Mixing both approaches

Part 3. Known issues



Occasionally data nodes fail with OOM



CacheDrSenderConfiguration#setBatchSendFrequency

CacheDrSenderConfiguration#setBatchSendSize

CacheDrSenderConfiguration#setMaxBatches

CacheDrSenderConfiguration#setStateTransferThreadsCount

Occasionally senders fail with OOM



DrSenderConfiguration#setMaxQueueSize

DrSenderConfiguration#setFullStateTransferBufferSize

Replication stops when FST is running



CacheDrSenderConfiguration#setStateTransferThrottle
CacheDrSenderConfiguration#setStateTransferThrottleBytes
CacheDrSenderConfiguration#setBatchSendFrequency
CacheDrSenderConfiguration#setBatchSendSize
DrSenderConfiguration#setMaxQueueSize

Links



DR documentation

gridgain.com/docs/latest/administrators-guide/data-center-replication/introduction

GridGain Webinars

gridgain.com/resources/webinars

Demo

github.com/GridGain-Demos/gridgain-data-center-replication-demo