

## 2. Partition Manager & State Manager

### Partition Manager

Session locations are tracked by partitions. Partitions are evenly distributed across the hosting peers of a Service Space. The management of these partitions is performed by a Partition Manager, which:

- transfers partitions to other peers;
- repopulates lost partitions; and
- updates version and location of partitions.

### Partition Rebalancing

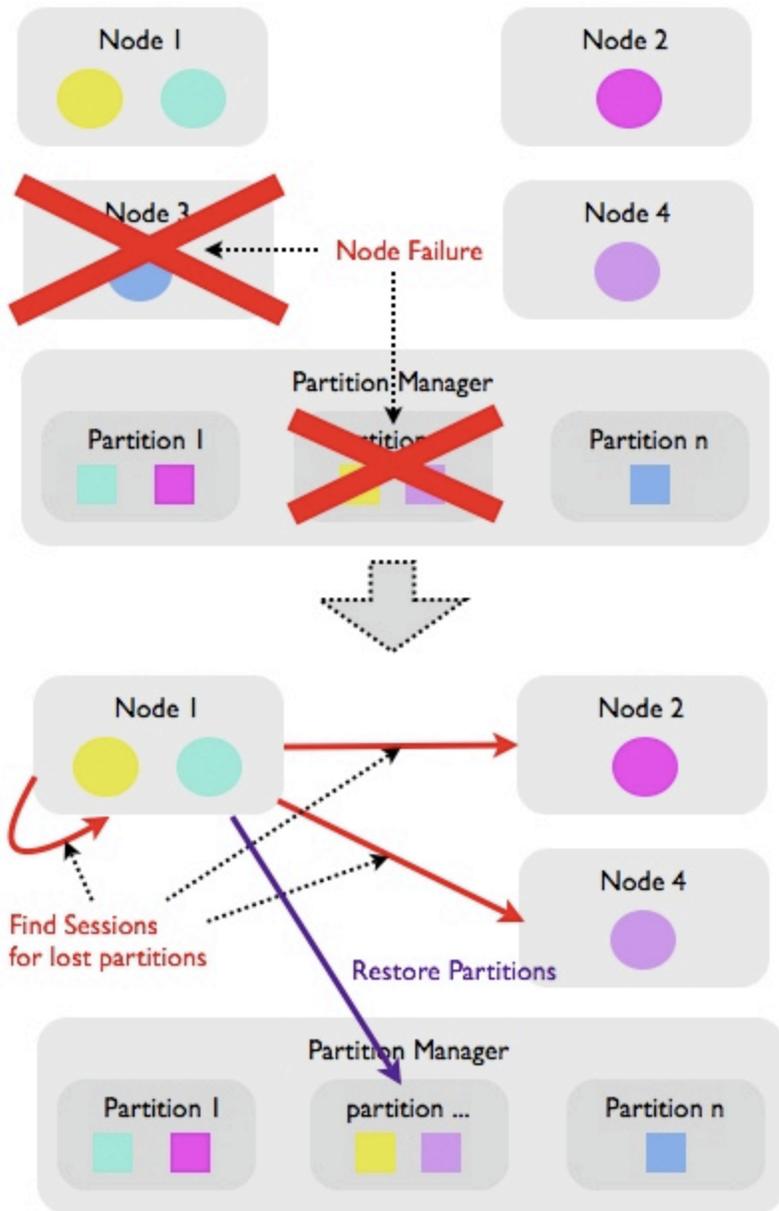
Each time that a cluster membership change is detected, e.g. when a peer joins a Service Space, a partition rebalancing singleton service is responsible for computing a new distribution plan for the partitions so that partitions are evenly distributed.

For instance, the rebalancing of 12 partitions works as follows:

- When the first hosting peer, Node 1, of a given Service Space starts:
  - all the partitions are allocated to Node 1.
- When a second hosting peer, Node 2, starts:
  - 6 partitions owned by Node 1 are transferred to Node 2 by Node 1.
- When a third hosting peer, Node 3, starts:
  - 2 partitions owned by Node 1 are transferred to Node 3 by Node 1; and
  - 2 partitions owned by Node 2 are transferred to Node 3 by Node 2.

### Partition Repopulation

When a peer hosting some partitions fail, lost partitions are repopulated by the peers becoming the new owner of these partitions:



Example:

1. Node 3, hosting some partitions, dies.
2. The node hosting the partition rebalancing singleton service queries all the nodes in the cluster for the partitions they are owning.
3. The partition rebalancing service identifies that:
  - some indexes have disappeared and need to be recreated; and
  - the disappeared partitions are to be managed by Node 1.
4. The partition rebalancing service provides to all the nodes a new partition distribution plan.
5. Node 1, which is instructed to acquire some lost partitions, queries the cluster for all the sessions belonging to the lost partitions.
6. Node 1 repopulates the lost partitions based on the responses of the above query.

## State Manager

State Manager is responsible for updating the session locations tracked by partitions when sessions are created, evacuated or destroyed:

- When a session is created, State Manager adds a new location entry to the relevant partition for this session;
- When a session is evacuated, i.e. when a session is migrated from one node to another because the source node is shutting down, State Manager updates the location of this session to reflect the new session location; and
- When a session is destroyed, State Manager removes the location entry for this session.