



This is the goal state that we are trying to maintain in the HA Deployment.

RackHD - We can only have one set of RackHD core services running at one time so we want to keep all of these services running with minimal downtime. We must monitor the state of these services and if one fails, we must failover the services to another node. Current goal is to have services down no more than 2 minutes on failover. These services must be reachable at the same IP addresses across failover so that we do not have to reconfigure any northbound or southbound interfaces. There is no service registry in play at this time.

Mongo and RabbitMQ – These services maintain state for RackHD core services. It is critical to have no loss of state. Therefore, we use the native clustering functionality for these services so that at least one server with current state is active at all times. Our goal is to make sure that these clusters stay fully functional. Therefore, if one or more nodes in the cluster fail, we must start up the service on another server. These services must be reachable at the same IP addresses across failover so that we do not have to reconfigure RackHD.

Reference HA Deployment

