

Fault Tolerance

Introduction

JobScheduler can be operated for high availability and fault tolerance. A number of options are available:

- [Resilience](#) includes measures for operational robustness that are intended to cope with outages.
- [Redundancy](#) provides fail-over capabilities in case of outages:
 - A [Passive Cluster](#) including the set-up of a Primary JobScheduler and redundant Backup JobScheduler.
 - An [Active Cluster](#) that allows the execution of jobs on distributed server nodes.
 - The dynamic assignment of JobScheduler [Agents](#) on different server nodes.
- [Recovery Strategies](#) provide a set of measures to restore the scheduling service after an outage.

Resilience

- [Resilience](#) includes a number of measures for operational robustness:
 - *Master / Agent Reconciliation* allows continued execution of tasks in case of recoverable *Network Connection Loss*.
 - *Master Service Recovery* includes supported measures after a *Master Service Failure*.
 - *Database Service Recovery* includes the capability to recover in case of *Database Connection Loss*.
- Resilience mechanisms are available with the components without manual intervention.

Redundancy

Redundancy in a Passive Cluster

- A Primary JobScheduler and Backup JobScheduler are operated in a [Passive Cluster](#)
- Should the primary instance fail then the backup instance will take over the load from the primary instance.
- In event of failover, the processing of all jobs, job chains and orders will be continued from the status that they had with the primary instance.

Distributed Processing with an Active Cluster

- An [Active Cluster](#) consists of a number of JobScheduler instances on different server nodes.
- Jobs can be executed in any of the connected JobScheduler instances.
- Should a JobScheduler instance fail then the processing is continued by one of the connected JobScheduler instances.
- Actively clustered JobScheduler instances allow hot plug-in: i.e. an instance can be added to a cluster at any point in time and can be removed accordingly.

Dynamic Agent Assignment in a Master / Agent Cluster

- In a [Master / Agent Cluster](#) a number of [Agents](#) is operated on different server nodes ([Agent Cluster](#)).
- The Master can be configured to dynamically select the next available JobScheduler Agent.
- In case of failure of a JobScheduler Agent the Master JobScheduler will switch processing to the remaining Agents.
- FEATURE AVAILABILITY STARTING FROM RELEASE 1.8

Recovery Strategies

- [Recovery Strategies](#) provide a set of measures to restore the scheduling service after an outage.