

Gap Recovery

The gap recovery is the weakest recovery class and it ensures neither input nor output preservation. Therefore, a system failure results in missing results for all incoming elements for all incoming elements before Odysseus can process them again (at-most-once). Gap recovery addresses the needs of applications that operate solely on the most recent information (e.g., sensor-based environment monitoring) where dropping old data is tolerable for reduced recovery time and run-time overhead. Because elements may be aggregated, missing elements may also result in incorrect output values. (Reference: J.H. Hwang et al., High-Availability Algorithms for Distributed Stream Processing, 2005)

In Odysseus, all relevant actions like the installation of (1) a source connection, (2) a sink connection or (3) a query are logged in a system log, if the recovery feature is installed. That system log is stored in the HOME folder of Odysseus (.odysseus). After a crash of Odysseus and a restart, all source connections, sink connections and queries will be reinstalled. Queries that were running, will be started again.

Requirements: There are no requirements. The gap recovery is always active, if the recovery feature is installed and if no other recovery technique is chosen. But there is one additional feature, if the gap recovery is explicitly set for a query: An operator, that checks when the results are again the same as is would be without any error, will be inserted after a crash. The trust meta attribute signals then when elements may be different to a run without any error (e.g. different aggregation result).

Example

```
#PARSER PQL
#RECOVERYCONFIGURATION GapRecovery
[Definition of queries to apply gap recovery]
```