

Restart/Recovery Processing

Restart/recovery occurs if a cluster nucleus fails. Restart/recovery uses the Work datasets of all nuclei to recover the database. The Work datasets are dynamically allocated from the dataset names recorded in the PPT. Adabas Cluster Services 7.4 supports offline and online recovery.

This chapter covers the following topics:

- Offline Recovery (Session Autorestart)
 - Online Recovery
 - Automatic Restart Management (ARM)
 - Archive Recovery
-

Offline Recovery (Session Autorestart)

- If a cluster nucleus session terminates, start one of the cluster nuclei to invoke autorestart.
- If a noncluster nucleus session terminates, restart the noncluster nucleus to invoke autorestart.

Offline recovery occurs if all active cluster nuclei in an Adabas sysplex cluster fail. Offline recovery relies only on information from the physical database and the Work datasets of each cluster nucleus. All information in the coupling facility is lost.

The first cluster nucleus to restart repairs any physical inconsistencies in the database and backs out all incomplete commands and transactions. The restarted nucleus obtains recovery information from blocks in the common database and from the Work datasets of all the failed nuclei.

The restarting nucleus retrieves the Work dataset names from the PPT block for each terminated nucleus and opens these datasets using dynamic allocation. From that point, normal recovery processing occurs:

- the breakpoint on each Work dataset is found;
- backward and forward repair is performed; and
- autobackout is performed.

While reading through the Work datasets, the restarting nucleus on the fly merges the protection records by their timestamps into chronological sequence.

Online Recovery

When one or more cluster nuclei have failed while one or more other nuclei in the same cluster remain active, online recovery processing is performed by collaboration of all surviving nuclei.

All surviving cluster nuclei quiesce their operations and reinitialize their working storage. Command processing is quiesced and the internal status variables, tables, and pools are repaired.

The peer nuclei compete for the recovery lock: when one of the nuclei obtains it, it invokes offline recovery processing. It repairs any physical inconsistencies in the database and backs out all incomplete command and transactions. Open transactions executed by the surviving nuclei are backed out as well. All information in the lock and cache structures is discarded.

Once this recovery processing has completed, normal processing resumes.

Users are affected by online recovery as follows:

- users assigned to failed nuclei lose their commands, transactions, sequential processes, and search results. They may receive response codes 9, 21, 148, or 251, depending on the status of their session at the time of the failure.
- users assigned to surviving nuclei may or may not lose their commands/transactions, depending on whether they managed to complete them in the quiesce phase. They retain their sequential processes and search results, but they may experience an increased response time. Users that do lose their commands/transactions will subsequently receive response code 9 and might possibly get response code 21 as well.

Automatic Restart Management (ARM)

Automatic restart management (ARM) is an OS/390 and z/OS facility that can be used to automatically restart a nucleus when it ABENDs. Automatic restart is suppressed when the ABEND is intentional; for example, when it results from a parameter error.

ARM can be used for Adabas nuclei in both cluster and noncluster environments.

The ADARUN parameter `ARMNAME` is used to identify the element in the ARM 'policy' that is to be activated. Each element specifies when, where, and how often an automatic restart is to be attempted.

If an ARM policy has not been defined, the `ARMNAME` parameter has no effect.

Archive Recovery

Archive recovery occurs if the container datasets of the database are damaged or restart/recovery is not effective.

Archive recovery

- restores the database; and
- regenerates the updates from the protection logs.

The protection logs to be regenerated are the output of the ADARES PLCOPY protection log copy and merge process that occurs in sysplex cluster environments. The restore/regenerate process is the same in both cluster and noncluster environments.