

Cluster Tasks

Adabas Parallel Services ensures data integrity and enables intercommunication among the Adabas nuclei in each cluster.

All nuclei in a cluster share the same database (ASSO and DATA data sets, DBID), global cache and lock areas, and global nucleus and user table.

ADARAI maintains a recovery log (RLOG) for each database; all nuclei in the cluster support a database write to the same RLOG and concurrent updates to the RLOG are controlled by a lock.

The Adabas Parallel Services characteristics of a particular DBID are fully transparent to the Adabas link routines.

Two types of tasks together build an Adabas Parallel Services cluster:

- a control task (ADACOM);
- one or more nuclei up to a maximum of 31.

This chapter covers the following topics:

- ADACOM
 - Cluster of Adabas Nuclei
 - Router Cluster Component SVCCLU
-

ADACOM

Note:

Router ID is the SVC number.

The ADACOM task must be run to set up the Adabas Parallel Services cluster environment before any cluster nucleus is started. A common storage area is allocated for maintaining information about active users and the nuclei to which they are assigned. Any changes are sent to ADACOM and to the cluster nuclei.

ADACOM also holds the cache and lock data spaces that the cluster nuclei use for sharing information. The data spaces can exist only as long as their owner, ADACOM, is active. For this reason, ADACOM cannot terminate if it is holding data spaces for an active Adabas Parallel Services cluster. Note that with Adabas Cluster Services, ADACOM does not have to be active all of the time.

Important:

Canceling ADACOM causes all Adabas Parallel Services nuclei it manages to fail.

One ADACOM can manage multiple cluster databases on different routers. ADACOM parameters specify the Router ID / DBID combinations (sets) that the ADACOM is to manage:

- The Router ID identifies the SVC number on z/OS or z/VSE systems. The Router ID value must be the same within a cluster; however, the same Router ID may be used for different clusters.
- The DBID identifies the external physical database shared by a particular cluster of nuclei and known to the application.

You can dynamically add or terminate Router ID / DBID combinations during an Adabas Parallel Services session. ADACOM itself, however, cannot be terminated as long as any of the Router ID / DBID sets it manages is still active.

Although a single ADACOM job can run all Router ID / DBID sets in a cluster environment, it is possible to run multiple ADACOM tasks simultaneously with the same, mixed, or completely different Router ID / DBID sets. An ADACOM subtask is attached to each Router ID / DBID set for each ADACOM task in which it occurs.

A COMPRINT data set is used to hold global messages that apply to all Router ID / DBID sets defined to an ADACOM task.

In addition, on z/OS platforms, Adabas Parallel Services dynamically allocates a spool data set for command output to each Router ID / DBID set. The format of its DD name is *Pssdddd* where *ss* is the last two digits of the SVC number and *dddd* is the DBID.

Alternatively, for each Router ID / DBID combination, you can set up an output data set in the startup JCL by specifying the DD name/link name in the format *Pssdddd* as just explained. The logical record length (LRECL) must be set to 80. The command output for the corresponding Router ID / DBID combination, whether defined in the startup JCL or dynamically allocated, is written to that data set. You might want to preset the spool output modules to control the data set's link name or the message class, which otherwise defaults to "X".

When you dynamically terminate a Router ID / DBID combination, the related spool output module is retained. On z/VSE systems, this information is written to SYSLST.

The first ADACOM started governs the size of the global user table (NU parameter): different values set for subsequent ADACOMs or for nuclei are ignored. If NU is set to zero, ADACOM frees the environment. If you do not set a value for NU, it defaults to 200. ADACOM parameters are discussed in the section *Initialization Parameters*.

After initialization, ADACOM functions as a command manager to monitor and control cluster nuclei running under Router ID / DBID sets it manages and to dynamically create or terminate Router ID / DBID sets. ADACOM commands are available to:

- exercise some control over the assignment of users to nuclei
- display the number of commands and user assignments for each active nucleus
- dynamically create or terminate Router ID / DBID combinations (sets).

Cluster of Adabas Nuclei

The ADARUN parameter CLUSTER=LOCAL is used to identify a nucleus as a member of an Adabas Parallel Services cluster. The NUCID parameter identifies the Adabas Parallel Services cluster nucleus for communication via the Adabas router. The NUCID must be a unique number in the range 1-65000. A

non-cluster nucleus has the parameter CLUSTER=NO (the default value); its NUCID is zero.

The NUCID is the only ID table entry for an Adabas Parallel Services nucleus.

Additionally, each nucleus has its own:

- command queue
- local buffer pool (LBP) and manager;
- Work data set, which must have the same device type and physical size for each nucleus in the cluster;
- dual or multiple protection log (PLOG) data sets (optional)
- dual or multiple command log (CLOG) data sets (optional)
- workpool
- format pools, etc.

Dual or multiple protection logs are optional but if supplied, each nucleus must have its own and all must use the same type. The same is true for command logs.

The nucleus handles all updates and protection logging. Updates are handled in a timely manner to ensure that the other cluster nuclei access the most recent data.

The ADADBS OPERCOM functions provide the NUCID parameter to route the command to a specific nucleus in the Adabas Parallel Services cluster. If NUCID is not specified, the command goes either to the nucleus assigned to ADADBS or AOS, or to all nuclei, depending on the inherent logic of the command.

Router Cluster Component SVCCLU

The Adabas router, (ADASVC on z/OS and z/VSE systems) always contains the SVCCLU component. The SVCCLU component includes functionality for Adabas cluster environments. The resulting SVC routes commands to the cluster nuclei in an Adabas Parallel Services cluster. To make routing decisions, it uses the nucleus and user tables in local common storage, which are updated based on nucleus or ADACOM information.

There is no predefined limit to the number of Adabas Parallel Services clusters, each servicing a separate database, that can run under a single router (ID table).

SVCCLU is invoked by the router as an "exit". It:

- maintains a table of active users and the nucleus to which each is currently assigned;
- assigns new users to the least loaded active cluster nucleus; and
- routes user commands according to the established assignment to cluster nuclei.