

Components and Processes

The Adabas Cluster Services components and processes ensure intercommunicability and data integrity among the operating system images and the associated Adabas nuclei in each cluster in a sysplex environment.

The chapter covers the following topics:

- Cluster of Adabas Nuclei
 - Adabas Sysplex Cluster
 - Single Physical Database Per Cluster
 - ADACOM
 - ADASVC Component SVCCLU
 - Cache Structure
 - Lock Structure
 - XCF Group
 - Messaging Services
-

Cluster of Adabas Nuclei

The ADARUN parameter CLUSTER=SYSPLEX identifies a nucleus as a member of an Adabas cluster in a sysplex environment. The NUCID parameter is used to number the Adabas cluster nuclei 1-65000. A NUCID of zero (0) identifies a noncluster nucleus. The NUCID is the only ID table entry for a nucleus.

Each nucleus has its own

- 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;
- (optional) direct-access protection log (PLOG) data sets; and
- (optional) command log (CLOG) data sets.

Protection logs are optional but if supplied, each nucleus must have its own. The same is true for command logs.

Adabas Sysplex Cluster

Each Adabas sysplex cluster is defined by its

- database ID;
- sysplex group name;
- sysplex cache and lock structure names.

In a sysplex environment, the Adabas nuclei in a cluster can be distributed over multiple operating system images that are synchronized by an IBM Sysplex Timer™.

Although more than one Adabas cluster nucleus may be run under an operating system image, several nuclei typically run under different operating system images to minimize the impact on the cluster when an operating system image must be stopped for maintenance or upgrades. The cluster database would continue to be serviced by the cluster nuclei that are active under different operating system images. For more information, read *Planning an Outage*, in *Adabas Cluster Services Operations*.

Single Physical Database Per Cluster

Using Adabas Cluster Services, up to 32 Adabas nuclei running under multiple operating system images access a single physical database simultaneously. A "single physical database" is one set of Associator and Data Storage data sets identified by a single database ID number (DBID).

An unlimited number of Adabas clusters can operate in the same sysplex system under the same or different SVCs; that is, an unlimited number of separate databases can be processed, each with its own Adabas cluster of up to 32 nuclei.

Each database has an ID table header (IDTH) prefix element.

Applications see only one database target; that is, no interface changes are required. Applications still communicate with their intended databases and communicate with the Adabas cluster of nuclei without modification.

ADACOM

The ADACOM initialization module or a cluster nucleus, whichever is the first to start, sets up the Adabas sysplex cluster environment on an operating system image. See the section *Getting Started*. An ECSA area is allocated for maintaining information about active users and the nuclei to which they are assigned. Using the ADACLU module and Entire Net-Work, any changes are sent to all peer ADACOMs or cluster nuclei.

The ADACOM initialization task:

- must be run at least once per IPL on any operating system image that has cluster users but no cluster nuclei to set up the environment.
- can optionally be used to monitor and control the nuclei of one or more sysplex clusters.

Parameters specify the SVC/DBID combinations (sets) that the ADACOM is to manage. The DBID identifies the external physical database shared by a particular cluster of nuclei and known to the application. The same SVC may be used for different clusters.

The first nucleus or ADACOM started governs the value for NU: different values set for subsequent nuclei or ADACOMs are ignored. The NU parameter in each ADACOM and each cluster nucleus should be large enough for all possible parallel local users and locally assigned remote users, so that the user table in common storage and the user queue in each cluster nucleus is big enough to continue providing service in case only one nucleus remains active in the cluster. If NU is set to zero, ADACOM frees the environment and exits. If no value is set for NU, it defaults to 200.

After initialization, ADACOM can either be quiesced or maintained in operation as an internucleus command manager to monitor and control cluster nuclei running under specified SVC/DBID sets and to dynamically create or terminate SVC/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 on the local image only or on every image known to the local image; and
- dynamically create or terminate SVC/DBID combinations (sets).

Software AG recommends that you execute and maintain the ADACOM module in each operating system image that participates in the sysplex so that the internucleus commands are available.

For operating system images that have users but no nuclei, Software AG recommends that you keep ADACOM in operation so that it is available not only as a command manager but also to reset the environment if Entire Net-Work goes down for any reason and comes back up.

Although a single ADACOM job can run all SVC/DBID sets in a sysplex environment, it is possible to run multiple ADACOM tasks simultaneously with the same, mixed, or completely different SVC/DBID sets. Two subtasks are attached for each SVC/DBID pair in each ADACOM task in which it occurs.

The COMPRINT data set prints global messages that apply to all SVC/DBID sets defined to an ADACOM task. In addition, two SYSOUT data sets are dynamically allocated for command output to each SVC/DBID set. The format of these file names are "*Pssdddd*" and "*Dssdddd*" where *ss* is the last two digits of the SVC and *dddd* is the DBID.

ADASVC Component SVCCLU

Prelinked to the version 8.2 Adabas SVC is the cluster SVC component SVCCLU that routes commands to local and, using Entire Net-Work, to remote nuclei. To make routing decisions, it uses local ECSA tables, which are updated based on nucleus or ADACOM information. When a user is moved to another system by CICSplex, for example, the ADASVC containing the SVCCLU on that system knows where to route the user's command.

An unlimited number of Adabas clusters each servicing a separate database can run under a single version 8.2 ADASVC.

Cache Structure

Each Adabas cluster in a sysplex environment uses a cache structure to hold ASSO and DATA blocks that have been updated during the session. Caching only the updated blocks improves the efficiency of information transfer and reduces the overall size of the required cache area. The cache structure synchronizes the nuclei, users, and operating system images and ensures data integrity.

The sysplex cache structure resides in the coupling facility in accordance with IBM's definition of coupling facility cache operations; it is thus accessible to all nuclei. The name and size of the cache structure is defined in the coupling facility resource management (CFRM) policy. The cache structure name for an Adabas cluster in a sysplex environment is specified for each cluster nucleus using the ADARUN parameter CLUCACHENAME. The same name must be provided to each nucleus in the cluster, and the same structure name may not be used outside the cluster.

Every nucleus in an Adabas cluster has a local buffer pool and manager. The buffer pool manager (BPM) oversees all nucleus requests for database blocks.

For each block in its local buffer pool, the BPM

- registers its interest in the block with the coupling facility;
- checks the coupling facility for the status of the registered block to ensure that its nucleus always has the most current copy of the block; and
- writes changed blocks to the coupling facility cache structure.

Whenever an updated block is written to the cache structure, the coupling facility notifies all nuclei that have "registered" an interest in that particular block that their local copy has become invalid.

Buffer Flush

From the cache structure, modified blocks are "cast out" to the flush I/O pool (FIOP) buffer before they are written to disk. The FIOP buffer is sized using the ADARUN LFIOP parameter, and the frequency of buffer flushing depends on the limit set. Until such blocks are written to the database, the cache structure holds more current information than the database.

The coupling facility handles the deletion of blocks in the cache structure when it becomes necessary to overwrite them.

Lock Structure

Each Adabas cluster in a sysplex environment uses a lock structure to manage the setting, status, and release of various locks imposed during multiple update nucleus processing. The lock structure synchronizes the nuclei, users, operating system images, and transaction processing to ensure data integrity.

The sysplex lock structure resides in the sysplex coupling facility in accordance with IBM's definition of coupling facility lock operations; it is thus accessible to all nuclei. The name and size of the lock structure is defined in the sysplex coupling facility resource management (CFRM) policy. The lock structure name for an Adabas cluster in a sysplex environment is specified for each cluster nucleus using the ADARUN parameter CLULOCKNAME. The same name must be provided to each nucleus in the cluster, and the same structure name may not be used outside the cluster.

Adabas Cluster Services uses the lock structure name when initializing its lock manager. The lock manager

- connects to and disconnects from the lock structure;

- obtains (conditional or unconditional), releases, and alters ownership (shared or exclusive) of resource locks; and
- reads recovery information about a failed lock.

Lock manager calls may be asynchronous, meaning that the nucleus may continue processing in other threads before a call has completed.

If a lock request is conditional, it is rejected if the lock is not free; if the lock request is unconditional, the nucleus thread waits until the requested lock is free before continuing.

In general, locks are used to prevent two cluster nuclei from using the same resource at the same time. Such resources include

- data records, which are protected by hold queue element (HQE) locks;
- unique descriptor values, which are protected by unique descriptor element (UQDE) locks;
- end-transaction IDs, which are protected by the ETID lock; and
- various other single-instance resources.

XCF Group

Each Adabas cluster in a sysplex environment uses a cross-system coupling facility (XCF) group to control communication among the cluster nuclei. XCF automatically creates an XCF group for a cluster when the first-starting cluster nucleus attempts to join the group.

When the sysplex COUPLExx data sets are first formatted, the maximum number of XCF groups to be used for the entire sysplex environment is defined. However, XCF group names are *not* predefined in the CFRM policy.

A name for a cluster's XCF group must be provided to each cluster nucleus using the ADARUN parameter CLUGROUPNAME. The same name must be provided to each nucleus in the cluster and the name may not be used outside the cluster.

Messaging Services

The messaging services component is essential to communication among the active nuclei of an Adabas sysplex cluster. For example, messages need to be sent between nuclei when:

- the general control block (GCB) has been modified
- a global parameter has been modified
- a buffer flush occurred
- a PLOG is switching

Adabas cluster messaging service statistics are produced during normal nucleus termination and in response to the DXMSG operator command.