

Useful ADARUN Parameters

This chapter provides a reference for the ADARUN parameters useful to Adabas Cluster Services processing. For information on how to specify ADARUN parameters, read the *Adabas Operations* documentation provided with Adabas.

- CLOGMRG : Cluster Command Log Merge
- CLUCACHENAME : Cluster Cache Structure/Area Name
- CLUGROUPNAME : Cluster Group Name
- CLULOCKNAME : Cluster Lock Structure/Area Name
- CLUSTER : Adabas Nucleus Cluster Session Control
- DIRRATIO / ELEMENTRATIO : Cluster Cache Directory-to-Data Elements Ratio
- LRDP : Length of the Redo Pool
- MXCANCEL: Timeout Threshold for Canceled Peer Nucleus
- MXCANCELWARN : Timeout Threshold for Canceled Peer Nucleus Warning
- MXMSG: Timeout Threshold for Internucleus Command Processing
- MXMSGWARN : Timeout Threshold for Internucleus Command Processing Warning
- MXSTATUS : Member-Level XCF Status Monitoring Heartbeat Interval
- MXWTOR : Self-Termination Operator Query Interval
- NUCID : Adabas Cluster Nucleus ID
- PROGRAM : Program to Run

CLOGMRG : Cluster Command Log Merge

Parameter	Specify . . .	Possible Values	Default
CLOGMRG	whether the Adabas cluster is to run with automatic CLOG merge.	YES NO	NO

An automatic CLOG merge process provides a single cluster-wide command log containing, in chronological order, all Adabas commands executed by any of the cluster nuclei in the time period covered by the log.

CLOGMRG specifies whether the Adabas cluster is to run with the automatic CLOG merge:

Value	Action
YES	enables an automated process that copies dual or multiple command logs to sequential data sets and merges those data sets across a cluster into a single cluster-wide command log. This process automatically merges command logs when an ADARES CLCOPY is submitted from user exit 2 and executed. The cluster nuclei and the ADARES utility perform coordinated CLOG switches and CLCOPY functions with merging of CLOG records, much the same way the PLOG merge process works.
NO	(the default), each nucleus produces its own CLOG, independent of any other nucleus in the cluster; dual or multiple CLOGs can be manually copied using ADARES CLCOPY to a sequential data set; and the resulting sequential CLOGs can be manually merged across a cluster using the ADARES MERGE CLOG function.

CLOGMRG is a global (that is, a cluster-wide) parameter, which means that the setting of this parameter in the first cluster nucleus to become active is propagated to all nuclei that subsequently become active.

Parameter Dependencies

It is possible to specify both LOGGING=NO and CLOGMRG=YES. In a cluster environment, it is also possible that a CLOG will be written to even if LOGGING=NO. As long as CLOGMRG=NO, nothing will be written to the CLOG if LOGGING=NO. However, if LOGGING=NO and CLOGMRG=YES, control records necessary for ADARES CLCOPY will be written to the CLOG.

CLOGMRG is effective only in an Adabas cluster environment; that is, when CLUSTER is specified with a value other than NO and NUCID is specified with a nonzero value.

Dynamic Modification

The setting of the CLOGMRG parameter can be changed dynamically using the CLOGMRG command from the operator console, the ADADBS OPERCOM CLOGMRG function, or the Modify Parameter function of Adabas Online System.

Because CLOGMRG is a global parameter, the change request to one nucleus is automatically propagated to all nuclei in the cluster.

CLUCACHENAME : Cluster Cache Structure/Area Name

Parameter	Specify . . .	Possible Values	Default
CLUCACHENAME	the name of the cache structure/area of the Adabas cluster in which the current nucleus participates.	text	none

For Adabas Cluster Services, the CLUCACHENAME identifies the particular sysplex cache structure (defined in the CFRM policy) associated with a particular Adabas sysplex nucleus cluster. The same name must be set in each nucleus in the cluster. CLUCACHENAME must be unique for each cluster; that is, each database.

Adabas Parallel Services automatically creates a default name for each global cache area allocated by ADACOM. However, if you wish to control the name of the global cache area, you can set this parameter.

Note:

Software AG recommends that you allow the system to build the default name.

For Adabas Parallel Services, the CLUCACHENAME identifies a particular global cache area associated with all cluster nuclei that have the same SVC/DBID set as the current nucleus. The same name must be set in each cluster nucleus that has the same SVC/DBID set. CLUCACHENAME must be unique for each SVC/DBID set.

In all cases, the name can be 1-16 characters long. Valid characters are numeric, uppercase alphabetic, the special characters \$, @, and #. The name must begin with an alphabetic character (A-Z). IBM names begin with SYS, UNDESIG, or the characters A through I (these characters are not reserved).

Note:

Additionally, the special character underscore (_) is allowed in Adabus Cluster Services; that is, sysplex environments.

Example

In the following example, the sysplex cache structure name is defined as "SAG1SMP#CACHE1".

```
ADARUN PROG=ADANUC, CLUCACHENAME=SAG1SMP#CACHE1
```

CLUGROUPNAME : Cluster Group Name

Parameter	Specify . . .	Possible Values	Default
<u>CLUGROUPNAME</u>	the name of the XCF messaging group in Adabus Cluster Services.	text	none

The CLUGROUPNAME parameter provides:

- the XCF group name to be used by the sysplex cluster nuclei for their communication using XCF.
- the name to be used by the cluster nuclei for their communication.

All nuclei in the same cluster must have the same CLUGROUPNAME. Nuclei belonging to different clusters must have different CLUGROUPNAMEs.

In addition, the CLUGROUPNAME must be different from the XCF group names used by other subsystems that communicate using XCF.

The name can be 1-8 characters long. Valid characters are numeric, uppercase alphabetic, and the special characters \$, @, and #. The name must begin with an alphabetic character (A-Z). IBM names begin with SYS, UNDESIG, or the characters A through I (these characters are not reserved).

Example

The following examples specifies the name of the XCF messaging group as "SAG1CLU".

```
ADARUN  PROG=ADANUC, CLUGROUPNAME=SAG1CLU
```

CLULOCKNAME : Cluster Lock Structure/Area Name

Parameter	Specify . . .	Possible Values	Default
CLULOCKNAME	the name of the lock structure defined in the coupling facility's CFRM policy or the global lock area in which the current nucleus participates	text	none

For Adabas Cluster Services, CLULOCKNAM identifies a particular sysplex lock structure (defined in the CFRM policy) associated with a particular Adabas nucleus cluster. The same name must be set in each nucleus in the cluster. CLULOCKNAME must be unique for each cluster; that is, each database.

Adabas Parallel Services automatically creates a default name for each global lock area allocated by ADACOM. However, if you wish to control the name of the global lock area, you can set this parameter.

Note:

Software AG recommends that you allow the system to build the default name.

For Adabas Parallel Services, the CLULOCKNAME identifies a particular global lock area associated with all Adabas Parallel Services cluster nuclei that have the same SVC/DBID set as the current nucleus. The same name must be set in each cluster nucleus that has the same SVC/DBID set. CLULOCKNAME must be unique for each SVC/DBID set.

In all cases, the name can be 1-16 characters long. Valid characters are numeric, uppercase alphabetic, the special characters \$, @, and #. The name must begin with an alphabetic character (A-Z). IBM names begin with SYS, UNDESIG, or the characters A through I (these characters are not reserved).

Note:

Additionally, the special character underscore (_) is allowed in Adabus Cluster Services; that is, sysplex environments.

Example

The following example set the name of the lock structure or global lock area as "SAG1SMP#LOCK1".

```
ADARUN  PROG=ADANUC, CLULOCKNAME=SAG1SMP#LOCK1
```

CLUSTER : Adabas Nucleus Cluster Session Control

Parameter	Specify . . .	Possible Values	Default
<u>CLUSTER</u>	whether the Adabas nucleus will participate in a cluster of nuclei working on the same Adabas database, and, if so, the type of cluster (Adabas Parallel Services or Adabas Cluster Services) in which it will participate.	NO SYSPLEX LOCAL	NO

The CLUSTER parameter specifies whether the Adabas nucleus will participate in a cluster of nuclei working on the same Adabas database, and if so, which type of cluster:

- If CLUSTER=NO (the default), the nucleus will not participate in an Adabas cluster. If CLUSTER=NO, the NUCID parameter must be set (or default) to 0.
- If CLUSTER=LOCAL, the nucleus will participate in a cluster where all nuclei run on the same operating system image. The add-on product Adabas Parallel Services is required to run in this mode.
- If CLUSTER=SYSPLEX, the nucleus will participate in a cluster where the nuclei may run on different systems within a parallel sysplex. This is supported only for z/OS systems. The add-on product Adabas Cluster Services is required to run in this mode.

If CLUSTER=LOCAL or SYSPLEX, the NUCID parameter must be set to a nonzero value that identifies the individual cluster nucleus. Also, the parameter settings MODE=SINGLE, READONLY=YES, LFIOP=0 and DTP=TM, and the use of sequential protection log datasets (DD-name/link name DDSIBA) are not allowed for cluster nuclei.

Example

The following example includes the nucleus in a cluster run with Adabas Cluster Services:

```
ADARUN DBID=53, CLUSTER=SYSPLEX, NUCID=5301
```

DIRRATIO / ELEMENTRATIO : Cluster Cache Directory-to-Data Elements Ratio

The DIRRATIO and ELEMENTRATIO parameters together define the ratio of directory elements to data elements in the Adabas Cluster Services coupling facility cache structure or the Adabas Parallel Services global cache area.

This ratio influences the allocation of the cache structure/area when the first cluster nucleus starts. It has no influence on the cache structure/area when a subsequent nucleus starts.

Parameter	Specify . . .	Minimum	Maximum	Default
<u>DIRRATIO</u>	the directory portion of the directory-to-data element ratio	1	32767	4
<u>ELEMENTRATIO</u>	the element portion of the directory-to-data element ratio	1	32767	1

The DIRRATIO / ELEMENTRATIO ratio is optionally used in Adabas cluster environments to tune the cache structure/area for better performance.

Directory elements are used to keep track of the cluster members that have individual blocks in their local caches (buffer pool and Adabas Caching Facility caches), so that the blocks can be invalidated should other members modify them.

Data elements are used to hold the contents of Associator and Data blocks from the database.

If the number of directory elements is insufficient, existing directory elements are reused and the blocks associated with those directory elements are invalidated and can then no longer be reused. When blocks become invalid even though they have not been modified, they must be reread and registered again the next time they are referenced and validated.

It is generally better to reassign storage for data elements to keep more Associator and Data blocks in the cache structure/area than to define too many directory elements in the cache. More data elements than necessary can be used to keep additional blocks to improve the cache efficiency.

The number of directory elements need not be greater than the sum of the sizes of all buffer pools plus the amount of space represented by data elements (number of data elements times data element size), divided by the smallest Associator block size.

You can use the Adabas Online System cache structure size calculator described in *Estimating Sizes for the Cache Structure in a Cluster Environment* to determine appropriate global cache area sizes and DIRRATIO/ELEMENTRATIO parameter values.

When connecting to the cache structure/area during startup, the ADAX57 message reports the number of directory and data elements allocated and the size of a data element.

LRDP : Length of the Redo Pool

Parameter	Specify . . .	Minimum	Maximum	Default
<u>LRDP</u>	the size (in bytes) of the redo pool.	0; or 80000	2147483647	LFIO value

The LRDP parameter specifies the size of the redo pool in an Adabas Cluster Services or Adabas Parallel Services nucleus. It has an effect only when CLUSTER=SYSPLEX or LOCAL has also been specified. Its default is the size of the flush I/O pool (LFIO parameter) of the cluster nucleus.

- If LRDP is set to zero, the cluster nucleus performs immediate publishing: Updated database blocks are always written to the global cache at the time each update takes place.

- If LRDP is set to (or defaults to) a nonzero value, the cluster nucleus performs deferred publishing: The writing of updated database blocks to the global cache is deferred until later. However, updates are always published in the global cache before the transactions they belong to are committed (by ET commands). The redo pool is used to maintain descriptions of all unpublished updates, in case they need to be redone due to concurrent updates to the same blocks by other nuclei in the cluster.

By combining multiple updates to the same block into a single write-to-cache operation, deferred publishing tends to result in fewer read and write operations from and to the global cache than immediate publishing. Depending on the type of application workload, this may have a large positive effect on performance. This is especially the case for Adabas Cluster Services, where frequent communication with the cache in the coupling facility may cause a high overhead.

Recommendation: For Adabas Cluster Services, specify a nonzero LRDP parameter (or use the default).

Different nuclei in the same cluster can have different settings of LRDP. It is also possible, although not recommended, to run one nucleus with LRDP=0 and another one with LRDP>0.

If the redo pool specification is too small and the pool runs full, the nucleus will write additional updated blocks to the global cache before the surrounding transactions end. Such additional cache writes may hurt system performance.

The nucleus reports on the usage (high watermark) of the redo pool in a shutdown statistic and in the response to the DRES operator command or the corresponding AOS or ADADBS OPERCOM functions.

MXCANCEL: Timeout Threshold for Canceled Peer Nucleus

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXCANCEL</u>	the timeout threshold, in seconds, in which the processing of an intracluster cancellation request must occur.	15	2147483647	MXMSG parameter value

The optional MXCANCEL parameter specifies the maximum number of seconds for which one cluster nucleus (the sending nucleus) waits for the termination of another nucleus (the peer nucleus) in the cluster that it has canceled for not responding to an internal, intracluster command. If a value for MXCANCEL is not specified, the value specified for the ADARUN MXMSG parameter is used. Values from zero through 14 are not valid MXCANCEL values.

If the sending nucleus does not receive notification that the canceled nucleus has terminated within the MXCANCEL timeout period, it does one of the following things:

- It returns response 124, subcode 28 (no reply) to the user on whose behalf the intracluster command was issued.
- It terminates itself abnormally if the communication occurred on behalf of an internal process that must not fail.

When a cluster nucleus fails, the other nuclei in the cluster can recover from the failure only once they know for sure that the first nucleus has disconnected from the shared structures and is no longer writing to the database. When an unresponsive cluster nucleus has been canceled, it is important that it terminate quickly, so that it does not hold up the recovery process performed by the other nuclei in the cluster.

Note:

Do not specify an MPMDUMP JCL statement for a cluster nucleus. An MPMDUMP is written before the failing nucleus disconnects from the shared structures and may significantly delay the start of the online recovery process performed by the other nuclei in the cluster. Instead, on z/OS systems, consider specifying a //SVCDUMP DD DUMMY statement, which usually produces a very quick dump prior to disconnecting. On other systems, and as a backup method on z/OS, specify a regular system dump statement (e.g., SYSUDUMP on z/OS) for a dump that is produced after disconnecting.

The value for MXCANCEL can be changed dynamically (for the time the Adabas nucleus is running) using the MXCANCEL operator command.

MXCANCELWARN : Timeout Threshold for Canceled Peer Nucleus Warning

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXCANCELWARN</u>	the timeout threshold, in seconds, for an intracluster cancellation request to occur before an early warning about the failure of the cancellation request is issued.	0 (no warning) or 4	MXCANCEL – 1	MXCANCEL / 4

The optional MXCANCELWARN parameter specifies the number of seconds after an intracluster cancellation request is made when the cluster nucleus requesting the cancellation should issue warnings about the inability of the canceled peer nucleus to terminate. If no value is specified for the MXCANCELWARN parameter, the value used is calculated as the value of the ADARUN MXCANCEL parameter divided by four. When dividing by four to determine the default value, the quotient is rounded up to the next integer value.

The value of MXCANCELWARN must be less than the value of MXCANCEL. If MXCANCELWARN is set to a value greater than or equal to the value specified for MXCANCEL, the value is rejected. If MXCANCELWARN is set to a value of zero (0), no warnings are issued about canceled cluster nuclei that are slow to terminate.

If the target nucleus does not terminate within the time period set by MXCANCELWARN, the nucleus that issued the cancellation request writes message ADAX9G to the operator console, identifying the target nucleus by its nucleus ID, job name, and system name. This message indicates that the nucleus writing the message may be in danger of terminating itself if the canceled peer nucleus does not terminate within the timeout period defined by the MXCANCEL parameter. You can use the ADAX9G message to trigger an alert or take other appropriate action before the MXCANCEL timeout period expires.

If the peer nucleus targeted for cancellation terminates after the ADAX9G message is written, the nucleus that issued the cancellation request writes message ADAX9K to the operator console, indicating that the cancellation occurred and the requesting nucleus is no longer in danger of terminating itself as a result of this particular incident.

The value for MXCANCELWARN can be changed dynamically (for the time the Adabas nucleus is running) using the MXCANCELWARN operator command.

Note:

If the value of the MXCANCEL parameter is dynamically changed using the MXCANCEL operator command while the nucleus is running, but its new value is less than or equal to the current MXCANCELWARN setting, the value of MXCANCELWARN is automatically set to zero. You can then use the MXCANCELWARN operator command to dynamically set the MXCANCELWARN parameter to a new value that is less than the new MXCANCEL value.

MXMSG: Timeout Threshold for Internucleus Command Processing

Parameter	Specify . . .	Minimum	Maximum	Default
MXMSG	the timeout threshold, in seconds, in which the processing of an intracenter command must occur.	15	32767	300

The optional MXMSG parameter specifies the maximum number of seconds in which one cluster nucleus waits for the response of another nucleus in the cluster to an internal, intracenter command. If a value for MXMSG is not specified, a value of 300 seconds (5 minutes) is used.

If the receiving nucleus does not respond to the intracenter command within the time period set by the MXMSG parameter, the requesting cluster nucleus issues a cancellation request for the unresponsive nucleus. The cancellation is requested to preserve the ability of all of the other active nuclei in the cluster to process user commands.

Adabas nuclei working on the same cluster database must collaborate to keep the database physically and logically consistent while processing user commands. To do this, they intermittently submit internal, intracenter commands to one another. If a nucleus in the cluster does not respond to these intracenter requests in a timely manner, the other nuclei in the cluster may quickly or gradually become unable to continue processing user commands.

Problems that may prevent a cluster nucleus from responding to intracenter commands fall into two broad categories:

- Problems have occurred in the cluster nucleus, such as a CPU loop or a resource deadlock. In such cases, the nucleus experiencing the problem must be terminated to resolve the problem and the sooner this happens, the sooner the other cluster nuclei can recover from the failure and get back to normal processing.
- Problems have occurred outside the cluster nucleus, such as CPU starvation or a problem in the host system. In such cases, the cause of the delayed response might possibly be removed by system measures unrelated to the cluster database, and letting the other cluster nuclei wait some time for a resolution might prevent an unnecessary abnormal termination of a cluster member.

The value of MXMSG should strike a balance between these two categories, so that faulty cluster nuclei that cannot possibly stay active are canceled reasonably quickly while cluster nuclei that are victims of temporary system problems are not unnecessarily terminated. For assistance in setting the value of the MXMSG parameter, review the statistics provided on the messaging performance reports produced at

termination or by the DXMSG operator command.

The value for MXMSG can be changed dynamically (for the time the Adabas nucleus is running) using the MXMSG operator command.

MXMSGWARN : Timeout Threshold for Internucleus Command Processing Warning

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXMSGWARN</u>	the timeout threshold, in seconds, for a response to an intracluster command to occur before an early warning about the failure of the intracluster request is issued	0 (no warning) or 4	MXMSG – 1	MXMSG / 4

The optional MXMSGWARN parameter specifies the number of seconds after an intracluster request is made when the cluster nucleus issuing the intracluster request should issue warnings about the failure of the target cluster nucleus to respond. If no value is specified for the MXMSGWARN parameter, the value used is calculated as the value of MXMSG divided by four. When dividing by four to determine the default value, the quotient is rounded up to the next integer value.

The value of MXMSGWARN must be less than the value of MXMSG. If MXMSGWARN is set to a value greater than or equal to the value specified for MXMSG, the value is rejected. If MXMSGWARN is set to zero, no warnings are issued about cluster nuclei that are slow to respond.

If the target nucleus does not respond within the time period set by MXMSGWARN, the nucleus that issued the request writes message ADAX9C to the operator console, identifying the target nucleus by its nucleus ID, job name, and system name. This message indicates that the target nucleus may be in danger of being canceled by the nucleus that issued the request if it does not respond within the timeout period defined by the MXMSG parameter. You can use the ADAX9C message to trigger an alert or take other appropriate action before the MXMSG timeout period expires.

If the target nucleus responds to the intracluster command after the ADAX9C message is written, the nucleus that issued the intracluster-command writes message ADAX9D to the operator console, indicating that a response was given and the target nucleus is no longer in danger of being canceled as a result of this particular incident.

The value for MXMSGWARN can be changed dynamically (for the time the Adabas nucleus is running) using the MXMSGWARN operator command.

Note:

If the value of the MXMSG parameter is dynamically changed using the MXMSG operator command while the nucleus is running, but its new value is less than or equal to the current MXMSGWARN setting, the value of MXMSGWARN is automatically set to zero. You can then use the MXMSGWARN operator command to dynamically set the MXMSGWARN parameter to a new value that is less than the new MXMSG value.

MXSTATUS : Member-Level XCF Status Monitoring Heartbeat Interval

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXSTATUS</u>	the interval, in seconds, at which the member-level XCF status monitoring heartbeat should occur.	0 (no member-level status monitoring) or 15	21474836	15

The optional MXSTATUS parameter specifies the number of seconds between member-level XCF status monitoring heartbeats. If no value is specified for the MXSTATUS parameter, a value of "15" is used.

Note:

Note: The MXSTATUS parameter is only used by Adabas Cluster Services and not by Adabas Parallel Services. Adabas Parallel Services does not use XCF and ignores this parameter when it is set.

If MXSTATUS is set to zero, member-level status monitoring will not occur. However, system-level status monitoring always occurs.

The value for MXSTATUS can be changed dynamically (for the time the Adabas nucleus is running) using the MXSTATUS operator command.

XCF status monitoring provides a second method by which Adabas Cluster Services can warn you that a cluster nucleus might be unable to respond in a timely way to intracluster commands. The first method is, of course, via the normal intracluster communication that occurs between cluster members. If a nucleus has heartbeat exceptions (as determined by XCF status monitoring), it most likely will be unable to process and respond to an intracluster command; if a nucleus is slow to respond to an intracluster command, it might or might not have a heartbeat monitor exception (a nucleus may appear to be active to XCF but be unable to respond to an intracluster command). If the ADARUN MXMSGWARN parameter for a cluster nucleus is nonzero, it produces warning messages (ADAX9B or ADAX9C) when intracluster communication with other nuclei in the cluster is too slow; likewise, when XCF status monitoring determines that a nucleus is missing its heartbeat updates, it produces warning messages (ADAX22 and ADAX04). You can use an automated mechanism set up at installation to raise an alert or take other appropriate action based on the existence of these messages, as they identify existing or potential problems in the cluster.

MXWTOR : Self-Termination Operator Query Interval

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXWTOR</u>	the timeout interval, in seconds, in which an operator response to a self-termination operator query must occur.	0 (no warning) or 15	64800	0

The optional MXWTOR parameter specifies the maximum number of seconds that a cluster nucleus should wait for an operator response to a self-terminating operator query. If no value is specified for the MXWTOR parameter, a value of zero is used. Values from 1 through 14 are not valid MXWTOR values.

A cluster nucleus issues a self-terminating operator query to the console via message ADAX9J if a canceled peer nucleus remains active after the time period defined by the MXCANCEL parameter. The MXWTOR parameter specifies how long the nucleus should wait for a response to this query.

One way that the operator may reply to the ADAX9J message is to make sure that the canceled cluster nucleus mentioned in the message actually terminates; once this happens, the ADAX9J message is retracted automatically, an online recovery process is performed, and the nucleus that issued the ADAX9J message remains active.

If the operator does not respond to the ADAX9J message within the time period set by the MXWTOR parameter and the canceled peer nucleus has not terminated in the meantime, the nucleus that requested the operator response and printed message ADAX9J terminates itself with message ADAX99.

If MXWTOR is set to zero (the default), the operator query does not occur (no ADAX9J messages are written) and no operator intervention is expected. Instead, the nucleus terminates itself immediately with the ADAX99 message and user abend code 79.

Some installations may want to specify MXWTOR value long enough to allow the operator a good opportunity to bring down the canceled peer nucleus. Other installations may prefer not to get such operator queries at all.

The value for MXWTOR can be changed dynamically (for the time the Adabas nucleus is running) using the MXWTOR operator command.

NUCID : Adabas Cluster Nucleus ID

Note:

The NUCID parameter replaces the earlier parameters PLXID and SMPID.

Parameter	Specify . . .	Possible Values	Default
NUCID	the ID for a nucleus in an Adabas cluster.	0-65000	0

NUCID identifies an Adabas cluster nucleus. The cluster may span operating system images under Adabas Cluster Services, or it may be confined to a single operating system image under Adabas Parallel Services. A value greater than 0 (the default) identifies this nucleus as an Adabas cluster nucleus.

Value	Meaning
0	The default setting. This is not an Adabas cluster nucleus.
1-65000	This is an Adabas cluster nucleus. This value must be unique among all target IDs declared to the Adabas router (ADASVC): Adabas databases (DBID), cluster nuclei (NUCID), Natural buffer pools and other MPM servers, Entire Net-Work, and internal targets used by Adabas Cluster Services and Adabas Parallel Services (restricted range 65,001-65,479). Nucleus IDs for the same cluster need not be specified contiguously.

Guidelines

- Each nucleus being run must have its own Work data set.
- All nuclei in a cluster must be run against the same database.
- Software AG strongly recommends that you run utilities or user applications against a nucleus in the local z/OS image. Although a remote open nucleus may be used, performance degradation due to network overhead is likely.

Example

The following example starts an Adabas cluster session with the specified cluster nuclei:

```
ADARUN  PROG=ADANUC,NUCID=1, ...
```

PROGRAM : Program to Run

Parameter	Specify:	Possible Values	Default
<u>P</u> ROGRAM	the program to be executed.	see table below	USER

This parameter specifies what to execute. The possible values are described in the following table:

Specify:	To start:
ADACOM	an ADACOM task (used in Adabas Cluster Services and Adabas Parallel Services environments) For more information, refer to your Adabas Cluster Services and Adabas Parallel Services documentation.
ADANUC	an Adabas nucleus For more information about executing an Adabas nucleus, read <i>Adabas Session Execution</i>, in <i>Adabas Operations</i> documentation.
ADAREV	an Adabas Review hub. Specify this in conjunction with the ADARUN REVIEW parameter. For more information, refer to your Adabas Review documentation.
NETWRK	an Entire Net-Work node For more information, refer to your Entire Net-Work documentation.
RENTUSER	a user program to be run using a reentrant Adabas batch/TSO link routine. For more information, refer to your Adabas TP monitor installation documentation.
USER	a user program to be run using a non-reentrant Adabas batch/TSO link routine. For more information, read <i>Linking Applications to Adabas</i>, in <i>Adabas Operations</i> documentation.
<i>utility-name</i>	an Adabas utility Specify an Adabas utility for <i>utility-name</i>. For more information, refer to <i>Adabas Utilities</i> documentation.

Examples

The following example starts an Adabas nucleus.

```
ADARUN PROGRAM=ADANUC
```

The following example starts an Adabas Review hub.

```
ADARUN PROGRAM=ADAREV, REVIEW=202
```