

Adabas Parallel Services

Operations

Version 8.1.3

June 2008

Adabas Parallel Services

This document applies to Adabas Parallel Services Version 8.1.3 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © Software AG 2008. All rights reserved.

The name Software AG, webMethods and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA, Inc. Other company and product names mentioned herein may be trademarks of their respective owners.

Table of Contents

1 Operations	1
2 Conventions	6
3 Initialization	5
Sequence	6
ADACOM Initialization Process	6
Adabas Parallel Services Cluster Nucleus Process	7
PPT Processing	8
4 Termination	11
Normal Termination	12
Abnormal Termination	12
5 Using Alert and Timeout Settings	15
Using Cancel Alerts	16
Using Self-Termination Alerts	17
Using the Self-Termination Operator Query	18
Using Messaging Performance Statistics	19
6 Backout Processing	21
7 Restart/Recovery Processing	23
Offline Recovery (Session Autostart)	24
Online Recovery	24
Automatic Restart Management (ARM)	25
Archive Recovery	25
8 Utility Processing	27
ADADBS OPERCOM Commands	28
ADADBS REFRESHSTATS - RefreshStatistical Values	29
ADAICK PPTPRINT - Print/DumpParallel Participant Table	29
ADARAI - Adabas Recovery Aid	30
ADAREP - Checkpoint Information Extended	30
ADAPLP IPLOGPRI - Print Sequential Intermediate Data Sets	30
ADARES CLCOPY - Copy/Merge Nucleus ClusterCommand Logs	30
ADARES PLCOPY - Copy/Merge Nucleus ClusterProtection Logs	32
ADARES MERGE CLOG - Merge Nucleus ClusterCommand Logs	38
ADARES BACKOUT and REGENERATE— Uniquely Identifying Checkpoints	40
ADASAV Processing Change	41
ADASAV RESTPLOG -- Uniquely Identifying Checkpoints	41
9 Cluster Nucleus Session End Statistics	43
General Nucleus Information	44
Input/Output Statistics	44
Command Statistics	45
User Statistics	47
Efficiency Statistics	47
Global Cache Statistics (Cluster Nucleus Only)	48
Global Lock Statistics (Cluster Nucleus Only)	53

Data Set Activity Statistics	59
10 Performance and Tuning	61
ADARUN Parameter Settings	62
Allocating Work Data Set Space	62
Using Close (CL) Commands	62
Deferred Publishing	63
Tuning Buffer Flushes	64
Optimizing Global Cache and Lock Areas	64
Minimizing Communication with Global Areas	66
Optimizing Block Sizes	67
11 Adabas Online System Cluster Environment Screens	69
Display Cluster Members	70
Nucleus File Status	71
Nucleus Status Flags	72
Cluster Usage	74
Estimating Sizes for the Cache Structure in a Cluster Environment	81
Estimating Sizes for the Lock Structure in a Cluster Environment	84
Maintain the User Table	87
Index	89

1 Operations

This document provides information about initialization, termination, backout, restart, and recovery processes in an Adabas Parallel Services cluster environment. It tells you how to plan an outage and how to use the utility functions that are provided specifically for cluster environments. Finally, it provides a sample of the session-end statistics produced for a clustered nucleus and describes how to improve performance and tune Adabas Parallel Services.

This document is organized as follows:

•	<i>Initialization</i>	Describes topics related to Adabas Parallel Services initialization.
•	<i>Termination</i>	Describes topics related to Adabas Parallel Services termination.
•	<i>Using Alert and Timeout Settings</i>	Describes alert and timeout settings that can be specified to help you prevent or handle critical situations where system problems affect the availability of cluster members.
•	<i>Backout Processing</i>	Describes Adabas Parallel Services backout processing.
•	<i>Restart/Recovery Processing</i>	Describes Adabas Parallel Services restart and recovery processing.
•	<i>Utility Processing</i>	Describes Adabas Utilities processing pertinent to Adabas Parallel Services.
•	<i>Cluster Nucleus Session End Statistics</i>	Describes the statistics collected for a cluster nucleus session.
•	<i>Performance and Tuning</i>	Describes performance and tuning you can do for Adabas Parallel Services.
•	<i>Adabas Online System Cluster Environment Screens</i>	Describes the Adabas Online System version 7.4 screens that apply to the Adabas Parallel Services cluster environment.

2 Conventions

Notation *vrs* or *vr*: When used in this documentation, the notation *vrs* or *vr* stands for the relevant version, release, and system maintenance level numbers. For further information on product versions, see *version* in the *Glossary*.

3 Initialization

■ Sequence	6
■ ADACOM Initialization Process	6
■ Adabas Parallel Services Cluster Nucleus Process	7
■ PPT Processing	8

This chapter covers the following topics related to Adabas Parallel Services initialization:

Sequence

Due to the interdependence among ADACOM(s) and Adabas cluster nuclei, these programs have certain co-occurrence requirements.

- ADACOM must be running on the operating system before any nucleus in a cluster that it manages is started.
- A nucleus will not start if another nucleus with the same NUCID is already active under the same router (SVC or IDTNAME).

ADACOM Initialization Process

The ADACOM initialization process is recorded in `PLInnn` messages.

All `PLInnn` messages are printed to the console. `PLI001-049` messages are specific to a particular Router ID / DBID set and are written to the output data set/file specifically allocated for that set; starting with `PLI050`, the messages apply to ADACOM in general and are written to the `COMPRINT` data set.



Note: Prior to `PLI050`, the system default ADARUN parameter values for `MODE`, `DBID`, `DEVICE`, `SVC` or `IDTNAME`, and `AMODE` are displayed. These are not the values for the current ADACOM.

```
PLI050 00161  INITIALIZING ADACOM
PLI002 00161  INITIALIZING DBID=dbid SVC=svc
                ACQUIRING NEW PLXCB
                PLXCB IS LOCATED AT address
PLI063 00161  PROCESSING: ADACOM SVC=svc,DBID=dbid,NU=users
                INITIALIZATION COMPLETE
```

A new Adabas Parallel Services control block is acquired if none yet exists.

Adabas Parallel Services Cluster Nucleus Process

Each Adabas Parallel Services cluster nucleus serializes during initialization by means of an enqueue. This is done because each nucleus puts information into the CSA nucleus table.

```

PLX050 00006 ADACLU init DBID=00006 NUCID=06001
PLX003 00006 There are nucs/ADACOMs active
PLX006 00006 Max users for image 00000200
PLX006 00006 PLXCB is located at 1127E258
ADAN03 00006 Initializing NUCID=6001 INTNUCID=4
ADAX2I 00006 SS-2,memstate, id 06001 ind 00 ext 00000 03
ADAX2A 00006 TI-0, initialised, rc 00000004
PLX087 00006 Attempting to acquire dataspaces
PLX097 00006 Dataspaces acquired
PLX059 00006 Initialization of ADACLU complete
ADAL01 00006 2002-06-25 17:50:45 CLOG not active
ADAN03 00006 ADABAS coming up
ADAX31 00006 Opening Work data set for NUCID=6001
ADAX70 00006 Connected to lock structure
ADAX70 00006 Number of lock entries 39,608
ADAX70 00006 Max number of record elements 39,608
ADAX57 00006 Connected to cache structure
ADAX57 00006 Directory elements 60,493
ADAX57 00006 Data elements 15,123
ADAX57 00006 Data element size 1,024
ADAN19 00006 Buffer flush is asynchronous
ADAN01 00006 ADABAS Vv.r.s is active
ADAN01 00006 MODE = MULTI
ADAN01 00006 Running without recovery log

```

Access to the global lock area is established and acknowledged in the following message:

```
ADAX70 00006 Connected to lock structure
```

This message is followed by lock area statistics:

```

ADAX70 00006 Number of lock entries 32,768
ADAX70 00006 Max number of record elements 11,184

```

Access to the global cache area is established and acknowledged in the following message:

```
ADAX57 00006 Connected to cache structure
```

If the ADARUN parameter CLUCACHETYPE=V64 was specified, this message appears:

```
ADAX57 00006 Connecting to S64 cache at address
```

This message is followed by cache area statistics:

```
ADAX57 00006 Directory elements      11,490
ADAX57 00006 Data elements           2,872
ADAX57 00006 Data element size      1,024
ADAN19 00006 Buffer flush is asynchronous
ADAN01 00006 ADABAS Vv.r.s is active
ADAN01 00006 Mode = MULTI
ADAN01 00006 Running without recovery log
```

PPT Processing

The parallel participant table (PPT), which exists for both cluster and noncluster nuclei, is used to determine if any PLOGs still need to be copied from previous sessions. If the PPT indicates that PLOGs remain to be copied, the PLOG data sets are read and, if necessary, a user exit is invoked (user exit 2 for dual logs or user exit 12 for multiple logs).

This section covers the following topics:

- [First Cluster Nucleus Starts](#)
- [Subsequent Cluster Nucleus Starts](#)
- [Noncluster Nucleus Starts](#)
- [Different PLOG Detected](#)

First Cluster Nucleus Starts

The Adabas Parallel Services cluster nucleus that is the first to initialize checks all the PLOG entries from the previous session for all nuclei and marks any that are "still being written" to completed status. In this way, the user exit (user exit 2 for dual logs or user exit 12 for multiple logs) need not be called each time a cluster nucleus autostarts. The first nucleus then calls the user exit but waits only if the PLOGs that need to be copied are for its own NUCID.

Subsequent Cluster Nucleus Starts

A subsequent cluster nucleus checks only its own PLOGs and invokes the user exit (user exit 2 for dual logs or user exit 12 for multiple logs) if the PLOGs are still not copied/merged. It waits if the user exit instructs it to do so. If there is no user exit 2 or 12, it overwrites the PLOGs.

Noncluster Nucleus Starts

A noncluster nucleus checks whether the previous session was an Adabas Parallel Services cluster session and has a pending autorestart. If so, the noncluster nucleus is not allowed to start.

If PLOGs from a previous Adabas Parallel Services cluster session remain to be copied, ADARES invokes the merge or the PLCOPY as required. A noncluster nucleus always uses block 1 of the PPT and can only overwrite it when PLOGs from previous sessions have been processed to completion.

A user exit (user exit 2 for dual logs or user exit 12 for multiple logs) controls the copy/merge process. If there is no user exit 2 or 12, the PLOG and PPT entry are overwritten.

Different PLOG Detected

If PLOGRQ=FORCE and an uncopied PLOG is detected that does not match that specified in the last session, a parameter error occurs. If the PLOG has been copied, the PPT entry is overwritten and the nucleus starts.

4 Termination

■ Normal Termination	12
■ Abnormal Termination	12

This chapter covers the following topics related to Adabas Parallel Services termination:

Normal Termination

This section covers the following topics:

- [ADACOM](#)
- [Cluster Nuclei](#)

ADACOM

ADACOM must stay in operation in an active Adabas Parallel Services cluster environment. ADACOM will not terminate normally if any nucleus that it manages is still active.

When all nuclei that it manages have terminated, you can terminate ADACOM using an ADAEND command. See the section *Cluster Operator Commands*.

Cluster Nuclei

If the Adabas operator command ADAEND or HALT is issued, the nucleus stops with no pending autorestart. The other active nuclei in the cluster continue processing normally.

```
ADAN51 00006 2002-06-25 18:03:29 Operator type-in: ADAEND
ADAN42 00006 2002-06-25 18:03:29 Function accepted
ADAX2B 00006 TT-1, SMM terminating
PLX087 00006 Attempting to delete dataspaces
PLX092 00006 Dataspaces deleted
ADAM97 06001 This ASCB/Initiator will be terminated by MVS at EOJ
```

If the ADARUN parameter CLUCACHETYPE=V64 was specified, this message also appears:

```
ADAX5B 00006 Disconnecting from S64 cache
```

Abnormal Termination

This section covers the following topics:

- [ADACOM](#)

- Cluster Nuclei

ADACOM

If ADACOM terminates abnormally, a `PLInnn` error message is produced to explain the problem. All active Parallel Services nuclei in clusters managed by this instance of ADACOM will also abend. ADACOM is the owner of the cache and lock data spaces used by the cluster nuclei. If ADACOM goes away, the data spaces also go away, and the nuclei will abend when they attempt to continue accessing the data spaces.



Note: This is not true of Adabas Cluster Services, where the cache and lock structures exist until all connectors (Adabas nuclei) terminate.

Cluster Nuclei

When an Adabas Parallel Services cluster nucleus terminates abnormally, each surviving peer nucleus performs online recovery. Read [Restart/Recovery](#) for more information.

5

Using Alert and Timeout Settings

■ Using Cancel Alerts	16
■ Using Self-Termination Alerts	17
■ Using the Self-Termination Operator Query	18
■ Using Messaging Performance Statistics	19

In a cluster environment, the 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 issue intracluster commands to one another. If one cluster member does not respond to an intracluster command from another cluster member within a specified time period, the sending member cancels the unresponsive member.

Adabas Parallel Services provides alert and timeout settings that are designed to help you prevent or handle critical situations where system problems might cause the prolonged unavailability of one cluster member, thus endangering the ability of the cluster member (or in severe cases, the entire cluster) to provide service. These settings include:

- A cancel alert that generates an operator message can be invoked when a cluster member is unresponsive to an intracluster command for a specified period of time. If the cluster member does not respond before the message times out, it is canceled.
- Self-termination alerts that generate operator messages can be invoked when a canceled cluster member does not terminate as requested (in a specified period of time) and the other cluster members prepare to self-terminate.
- An operator query that prints an operator console message and requests a response from the operator. This can be invoked by cluster members that are preparing to self-terminate because a canceled peer member remains active.
- Messaging statistics provide information about the performance of message transmission events. These statistics can be used to determine the impact of messaging on system performance and to determine how to set the ADARUN parameters related to the alert and timeout settings.

This chapter covers the following topics:

Using Cancel Alerts

In cluster environments, the cluster members issue intracluster commands to one another to ensure that the database is kept physically and logically consistent. If one cluster member does not respond to an intracluster command from another cluster member within the time specified by the ADARUN `MXMSG` parameter, the sending member cancels the unresponsive member.

You can invoke a cancel alert before the unresponsive peer member is canceled. This alert generates an operator message that provides early warning information before the unresponsive cluster member is canceled.

The cancel alert is governed by the setting of the ADARUN `MXMSGWARN` parameter. This optional parameter specifies the number of seconds after which a cluster nucleus should generate an operator message warning about an outstanding intracluster response. If the cluster member does not respond within the time specified by ADARUN `MXMSGWARN`, message ADAX9C is issued. This warning message can be used to notify you sometime before the unresponsive cluster member is canceled.

As complements to the ADARUN `MXMSG` and `MXMSGWARN` parameters, two operator commands, `MXMSG` and `MXMSGWARN`, are provided that allow you to change the corresponding ADARUN settings dynamically, while the database is running. In addition, the `DPARM` operator command's output information includes the settings of the ADARUN parameters related to cluster alert and timeout enhancements.

For more information about the `MXMSG` and `MXMSGWARN` parameters, read *MXMSG: Timeout Threshold for Internucleus Command Processing* and *MXMSGWARN: Timeout Threshold for Internucleus Command Processing Warning*. For information about the `DPARM`, `MXMSG`, and `MXMSGWARN` operator commands, read *Cluster Operator Commands*.

Using Self-Termination Alerts

In cluster environments, if one cluster nucleus has issued a cancellation request for a second unresponsive cluster nucleus, but the canceled peer cluster nucleus does not terminate within the time specified by the ADARUN `MXCANCEL` parameter, the sending nucleus will either return response code 124, subcode 28 (if the intracluster communication occurred on behalf of an Adabas command) or terminate itself abnormally (if the intracluster communication occurred on behalf of an internal process that must not fail).

You can invoke a self-termination alert before a nucleus terminates itself because a canceled peer nucleus fails to terminate. This alert generates an operator message that provides early warning information regarding the pending self-termination.

Self-termination alerts are governed by the setting of the new ADARUN `MXCANCELWARN` parameter. This optional parameter specifies the number of seconds after which a requesting cluster nucleus should generate an operator message warning about the inability of a canceled peer nucleus to terminate quickly. If the canceled peer nucleus does not terminate within the time specified by ADARUN `MXCANCELWARN`, message ADAX9G is issued. This warning message can be used to notify you that the nucleus issuing the message is in danger of terminating itself.

As complements to the ADARUN `MXCANCEL` and `MXCANCELWARN` parameters, two operator commands, `MXCANCEL` and `MXCANCELWARN`, are provided that allow you to change the corresponding ADARUN settings dynamically, while the database is running. In addition, the `DPARM` operator command's output information has been enhanced to include the settings of ADARUN parameters related to Adabas Parallel Services alert and timeout enhancements.

For more information about the `MXCANCEL` and `MXCANCELWARN` parameters, read *MXCANCEL: Timeout Threshold for Canceled Peer Nucleus* and *MXCANCELWARN: Timeout Threshold for Canceled Peer Nucleus Warning*. For information about the `DPARM`, `MXCANCEL`, and `MXCANCELWARN` operator commands, read *Cluster Operator Commands*.

Using the Self-Termination Operator Query

You can invoke an operator query when a cluster member is in the process of self-terminating because a canceled peer nucleus fails to terminate. This gives you a chance to terminate the canceled cluster member manually, thus avoid the self-termination of the member that issued the ineffective cancel request.

This operator query prints a console message (message ADAX9J) explaining the situation and requesting instructions, waiting for a specified time for a response. The valid responses to message ADAX9J are:

- R (print the ADAX9J message again and continue to wait for resolution of this issue, but without setting a new wait period for the response)
- T (terminate the querying nucleus with message ADAX99 and userabend 79)
- W (continue to wait for another time period of length MXWTOR)

The amount of time the operator query waits for a response is governed by the setting of the ADARUN MXWTOR parameter. This optional parameter specifies the number of seconds the nucleus should wait for the operator response. If the operator does not respond in this time frame and if the canceled peer nucleus still has not terminated, the requesting nucleus issues message ADAX99 and terminates itself.

However, if the canceled cluster member terminates after all (whether due to operator intervention or another reason), the cluster nucleus that issued the operator query stays alive; it retracts the query and initiates an online recovery process.

As a complement to the ADARUN MXWTOR parameter, an operator command, MXWTOR, is provided that allows you to change the MXWTOR setting dynamically, while the database is running. In addition, the DPARM operator command's output information includes the settings of ADARUN parameters related to Adabas Cluster Services alert and timeout enhancements.

For more information about the MXWTOR parameter, read *MXWTOR : Self-Termination Operator Query Interval* . For information about the DPARM and MXWTOR operator commands, read *Cluster Operator Commands*.

Using Messaging Performance Statistics

Adabas Parallel Services messaging statistics provide information about the performance of message transmission events. These statistics can be used to determine the impact of messaging on system performance and to determine how to set the ADARUN `MXMSG` and `MXMSGWARN` parameters related to the other alert and timeout enhancements in Adabas Parallel Services.

The performance statistics are provided in the termination statistics of an Adabas nucleus as well as in response to the `DXMSG` operator command. The performance statistics are split into those that are subject to the ADARUN `MXMSG` parameter setting and those that are not; after each is reported separately in the output, a combined report is provided containing the summarization of the two for all messages.

For more information about the `DXMSG` operator command, read *Cluster Operator Commands*.

6

Backout Processing

Normal backout processing includes:

- BT command processing;
- backing out an update command that received a nonzero response code; and
- internal transaction backout due to, for example, a timeout.

Cluster nuclei perform normal Adabas backout processing. Each nucleus reads the protection data needed to back out updates from its own Work data set/file. It does not need to read protection data from the other Work data sets/files in the cluster.

7 Restart/Recovery Processing

■ Offline Recovery (Session Autorestart)	24
■ Online Recovery	24
■ Automatic Restart Management (ARM)	25
■ Archive Recovery	25

Restart/recovery occurs if a cluster nucleus fails. Restart/recovery uses the Work data sets/files of all nuclei to recover the database. The Work data sets/files are dynamically allocated from the data set names recorded in the PPT. Adabas Parallel Services version supports offline and online recovery.

This chapter covers the following topics:

Offline Recovery (Session Autorestart)

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

Offline recovery occurs if all active cluster nuclei in an Adabas Parallel Services cluster fail. Offline recovery relies only on information from the physical database and the Work data sets/files of each cluster nucleus. All information in the global cache and lock areas 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 data sets/files of all the failed nuclei.

The restarting nucleus retrieves the Work data set/file names from the PPT block for each terminated nucleus and opens these data sets/files using dynamic allocation. From that point, normal recovery processing occurs:

- the breakpoint on each Work data set/file is found;
- backward and forward repair is performed; and
- autobackout is performed.

While reading through the Work data sets/files, 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 global lock and cache areas 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 a 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` (read *ADARUN Parameter Usage in Cluster Environments*) 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 data sets 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 Adabas Parallel Services cluster environments. The restore/regenerate process is the same in both cluster and noncluster environments.

8

Utility Processing

■ ADADBS OPERCOM Commands	28
■ ADADBS REFRESHSTATS - RefreshStatistical Values	29
■ ADAICK PPTPRINT - Print/DumpParallel Participant Table	29
■ ADARAI - Adabas Recovery Aid	30
■ ADAREP - Checkpoint Information Extended	30
■ ADAPLP IPLOGPRI - Print Sequential Intermediate Data Sets	30
■ ADARES CLCOPY - Copy/Merge Nucleus ClusterCommand Logs	30
■ ADARES PLCOPY - Copy/Merge Nucleus ClusterProtection Logs	32
■ ADARES MERGE CLOG - Merge Nucleus ClusterCommand Logs	38
■ ADARES BACKOUT and REGENERATE— Uniquely Identifying Checkpoints	40
■ ADASAV Processing Change	41
■ ADASAV RESTPLOG -- Uniquely Identifying Checkpoints	41

Read your *Adabas Utilities* documentation for specific information about changes to utilities for use in an Adabas cluster environment.

This chapter covers the following topics:

ADADBS OPERCOM Commands

Changes have been made for ADADBS OPERCOM command processing in an Adabas Parallel Services cluster nucleus environment.

This section covers the following topics:

- [Global Commands](#)
- [Routing a Command to a Specific Nucleus](#)
- [Routing a Command to All Cluster Nuclei](#)

Global Commands

The following ADADBS OPERCOM commands have a GLOBAL option for making the request to all nuclei in an Adabas cluster:

- ADAEND
- CANCEL
- FEOFCL
- FEOFPL
- HALT

For example:

```
ADADBS OPERCOM ADAEND,GLOBAL
```

When GLOBAL is specified, the command is automatically propagated to all active cluster nuclei. When GLOBAL is *not* specified, a specific NUCID from the cluster must be specified and the command is sent to that NUCID.

Routing a Command to a Specific Nucleus

The NUCID option allows you to direct the OPERCOM commands to a particular nucleus in the cluster for execution.

The OPERCOM function's NUCID option is specified in a manner similar to the ADARUN function's NUCID parameter.

The following example sends the DSTAT command to the Adabas cluster nucleus designated with NUCID=3:

```
ADADBS OPERCOM DSTAT,NUCID=3
```

For inherently global commands, such as changing the setting of the TT parameter, the NUCID parameter is ignored.

Routing a Command to All Cluster Nuclei

When the NUCID option in the ADADBS OPERCOM function is not specified, the command is sent to all cluster nuclei and information is displayed for each nucleus in sequence.

ADADBS REFRESHSTATS - RefreshStatistical Values

The REFRESHSTATS function resets statistical values maintained by the Adabas nucleus for its current session. Parameters may be used to restrict the function to particular groups of statistical values.

In Adabas cluster environments, you must specify the specific nucleus (NUCID) for which statistical values are to be refreshed. If NUCID is not specified, statistical values will be refreshed for all active nuclei in the cluster.

ADAICK PTPRINT - Print/DumpParallel Participant Table

The PTPRINT function has been added to the Adabas ADAICK utility to support an Adabas cluster environment. It is used to dump/print the parallel participant table (PPT) for the Adabas cluster.

Each of the 32 blocks (RABNs) allocated for the PPT represents a single nucleus in the cluster and comprises:

- a single header of fixed length; and
- multiple entries of variable length.

Note that in the dump/print, "PPH" is the tag for the PPT header and "PPE" is the tag for the PPT entries.

ADARAI - Adabas Recovery Aid

Adabas cluster products support the Adabas Recovery Aid (ADARAI).

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 ADARAI LIST function supports Adabas version 7 and above RLOGs; Adabas version 6 RLOGs are not supported.

ADAREP - Checkpoint Information Extended

Given that each cluster nucleus has its own PLOG data sets, checkpoints are no longer identified only by their name, PLOG number, and PLOG block number, but also by the ID of the nucleus that writes the checkpoint.

Several new parameters have been introduced for utilities that need to identify checkpoints on the PLOG.

ADAPLP IPLOGPRI - Print Sequential Intermediate Data Sets

The IPLOGPRI function is used to print the sequential intermediate data sets created from the PLOG merge process. Input to ADAPLP IPLOGPRI must be a MERGIN1/MERGIN2 data set created by the ADARES utility and specified in the JCL with DD name/link name DD/PLOG.

ADARES CLCOPY - Copy/Merge Nucleus ClusterCommand Logs

Sample JCL has been added for allocating the intermediate data sets/files MERGIN1 and MERGIN2 required for automated CLOG copy/merge processing in nucleus cluster environments.



Notes:

1. When intermediate data sets/files are used for both CLCOPY and PLCOPY, the data set names must be unique so that they are not overwritten.

- The data set BLKSIZE used must be greater than or equal to the largest CLOG BLKSIZE plus eight. The LRECL must be set to the BLKSIZE minus four.

- [z/OS](#)
- [z/VSE](#)

z/OS

```
//ALLOC JOB
/*
/* Example to allocate the ADARES CLCOPY intermediate data sets
/*
//CM1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//OUTFIL DD DISP=(NEW,CATLG),SPACE=(CYL,(3,10)),UNIT=SYSDA,
// VOL=SER=volser,DCB=(RECFM=VB,LRECL=23472,BLKSIZE=23476),
// DSN=EXAMPLE.CLOG.MERGIN1
//INPFIL DD *
/*
//SYSIN DD *
REPRO INFILE(INPFIL) -
OUTFILE(OUTFIL)
/*
/*
//CM2 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//OUTFIL DD DISP=(NEW,CATLG),SPACE=(CYL,(3,10)),UNIT=SYSDA,
// VOL=SER=volser,DCB=(RECFM=VB,LRECL=23472,BLKSIZE=23476),
// DSN=EXAMPLE.CLOG.MERGIN2
//INPFIL DD *
/*
//SYSIN DD *
REPRO INFILE(INPFIL) -
OUTFILE(OUTFIL)
/*
```

z/VSE

```
//JOB ALLOC
/* Example to allocate the ADARES CLCOPY intermediate data sets
//EXEC PGM=IDCAMS
DEFINE CLUSTER (NAME (MERGIN1.CLCOPY) -
VOLUME(xxxxxx) CYLINDERS(3,10) RECSZ(23472) -
/*
//EXEC PGM=IDCAMS
DEFINE CLUSTER (NAME (MERGIN2.CLCOPY) -
VOLUME(xxxxxx) CYLINDERS(3,10) RECSZ(23472) -
/*
/&
```

ADARES PLCOPY - Copy/Merge Nucleus ClusterProtection Logs

In an Adabas nucleus cluster environment, the protection logs (and optionally, the command logs: see the ADARES MERGE CLOG function and the ADARUN CLOGMRG parameter) of all individual nuclei in the cluster must be merged into single log files in chronological order for the cluster database shared by all the nuclei as a whole. The chronological order is determined by timestamps on all individual nucleus log records.

Protection logs are automatically merged when an ADARES PLCOPY is executed. In an Adabas cluster environment, the PLCOPY process accesses the parallel participant table (PPT) to determine which protection logs to copy and opens the appropriate data sets/files using dynamic allocation. PLCOPY copies/merges as much data as possible; if a nucleus is still writing to a protection log data set, PLCOPY 'partially' merges the data set.

The merge begins with the lowest timestamp from all protection logs being merged and ends with the lowest of the ending timestamps from all data sets/files. Records beyond this point are written to an 'intermediate' data set, which must be supplied as input to the subsequent merge. A cross-check ensures that the correct intermediate data set has been supplied.

ADARES expects that at least one of the protection logs being merged is at 'completed' status. If this is not the case, ADARES reports that there is no data to be copied.

A sample user exit (USEREX2P for dual logs or UEX12 for multiple logs) is provided to illustrate the necessary change for the intermediate data set.

- [z/OS Sample Jobs](#)
- [z/VSE Sample Jobs](#)

z/OS Sample Jobs

This section describes the sample jobs provided for z/OS systems.

ALLOC Job

The following sample JCL illustrates the allocation of the intermediate data sets MERGIN1 and MERGIN2 which are required for automated PLOG copy/merge processing in nucleus cluster environments.



Notes:

1. When intermediate data sets are used for both CLCOPY and PLCOPY, the data set names must be unique so that they are not overwritten.
2. The data set BLKSIZE used must be greater than or equal to the largest PLOG BLKSIZE plus eight. The LRECL must be set to the BLKSIZE minus four.

```
//ALLOC JOB
/*
/* Example to allocate the ADARES PLCOPY intermediate data sets
/*
//CM1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//OUTFIL DD DISP=(NEW,CATLG),SPACE=(CYL,(3,10)),UNIT=SYSDA,
// VOL=SER=volser,DCB=(RECFM=VB,LRECL=23472,BLKSIZE=23476),
// DSN=EXAMPLE.PLOG.MERGIN1
//INPFIL DD *
/*
//SYSIN DD *
REPRO INFILE(INPFIL) -
OUTFILE(OUTFIL)
/*
/*
//CM2 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//OUTFIL DD DISP=(NEW,CATLG),SPACE=(CYL,(3,10)),UNIT=SYSDA,
// VOL=SER=volser,DCB=(RECFM=VB,LRECL=23472,BLKSIZE=23476),
// DSN=EXAMPLE.PLOG.MERGIN2
//INPFIL DD *
/*
//SYSIN DD *
REPRO INFILE(INPFIL) -
OUTFILE(OUTFIL)
/*
```

ADARESPM Job

A sample job ADARESPM is provided on the JOBS data set to illustrate the manual execution of the PLCOPY merge function. Two intermediate data sets must be supplied. ADARES analyzes the data sets to determine which is to be used as input and which for output. Specific cross-checks determine whether the correct intermediate data set has been supplied; if not, ADARES will not continue. Continuing without the correct input can result in lost updates and inconsistencies if the output is used for REGENERATE or BACKOUT functions.

Once DDPLOGRn statements have been supplied on the session startup JCL, you do not need to supply them again for ADARES as these are opened using dynamic allocation. If the DD statements are supplied, they are ignored.

The following sample JCL illustrates the ADARES PLCOPY merge function:

```
//ADARESPM JOB
//*
//* ADARES : COPY/MERGE DUAL/MULTIPLE PROTECTION LOG
//* TWO COPIES OF OUTPUT ARE TO BE CREATED
//* FOR USE WITH AN ADABAS NUCLEUS CLUSTER
//*
//RES EXEC PGM=ADARUN
//STEPLIB DD DISP=SHR,DSN=ADABAS.Vvrs.LOAD
//*
//DDASSOR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.ASSOR1
//DDDATAR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.DATAR1
//MERGIN1 DD DISP=SHR,DSN=EXAMPLE.PLOG.MERGIN1
//MERGIN2 DD DISP=SHR,DSN=EXAMPLE.PLOG.MERGIN2
//DDSIAUS1 DD DSN=EXAMPLE.DByyyyy.PLOG1(+1),
// VOL=SER=ADAxXX,UNIT=TAPE,DISP=(NEW,CATLG)
//DDSIAUS2 DD DSN=EXAMPLE.DByyyyy.PLOG2(+1),
// VOL=SER=ADAxXX,UNIT=TAPE,DISP=(NEW,CATLG)
//DDDRUCK DD SYSOUT=X
//DDPRINT DD SYSOUT=X
//SYSUDUMP DD SYSOUT=X
//DDCARD DD *
ADARUN PROG=ADARES,MODE=MULTI,SVC=svc,DEVICE=3380,DBID=yyyyyy
/*
//DDKARTE DD *
ADARES PLCOPY TWOCOPIES
/*
```

ADARES PLCOPY NOPPT—Ignore PPT

NOPPT is for emergency use when the PPT has been overwritten or is otherwise unavailable. It specifies that the PLOG data sets of all cluster nuclei are being supplied with DD names DDPLOGnn in the JCL.



Caution: Use this parameter cautiously since it ignores the PPT and all control-type information typically coming from the PPT.

When you use this parameter, you must supply

- the correct intermediate data set; and
- the correct input protection logs from all nuclei with DD names DDPLOG01-nn.

The optional parameter SBLKNUM can be used to specify the starting block number for the sequential merge output.



Caution: Without the PPT, ADARES cannot perform any extensive validations on the input data sets

ADARESIP Job

The following sample JCL illustrates the ADARES PLCOPY NOPPT merge function:

```
//ADARESIP JOB
//*
//* ADARES : COPY/MERGE DUAL/MULTIPLE PROTECTION LOGS FROM ALL
//* NUCLEI IN AN ADABAS NUCLEUS CLUSTER
//* PPT IS TO BE IGNORED
//* THIS IS ONLY FOR EMERGENCY USE WHEN THE PPT HAS BEEN
//* OVER-WRITTEN - USE CAUTION WHEN SUBMITTING
//*
//RES EXEC PGM=ADARUN
//STEPLIB DD DISP=SHR,DSN=ADABAS.Vvrs.LOAD <=== ADABAS LOAD
//*
//DDASSOR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.ASSOR1 <=== ASSO
//DDATAR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.DATAR1 <=== DATA
//DDPLOG01 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.PLOGR1.NUC1 <=== PLOG1 NUC1
//DDPLOG02 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.PLOGR2.NUC1 <=== PLOG2 NUC1
//DDPLOG03 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.PLOGR1.NUC2 <=== PLOG1 NUC2
//DDPLOG04 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.PLOGR2.NUC2 <=== PLOG2 NUC2
//DDPLOG05 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.PLOGR1.NUC3 <=== PLOG1 NUC3
//DDPLOG06 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.PLOGR2.NUC3 <=== PLOG2 NUC3
//MERGIN1 DD DISP=SHR,DSN=EXAMPLE.PLOG.MERGIN1 <=== INTERMEDIATE 1
//MERGIN2 DD DISP=SHR,DSN=EXAMPLE.PLOG.MERGIN2 <=== INTERMEDIATE 2
//DDSIASU1 DD DSN=EXAMPLE.DByyyyy.PLOG1(+1), <=== PLOG COPY
// VOL=SER=ADAxxx,UNIT=TAPE,DISP=(NEW,CATLG)
//DDDRUCK DD SYSOUT=X
//DDPRINT DD SYSOUT=X
//SYSUDUMP DD SYSOUT=X
//DDCARD DD *
ADARUN PROG=ADARES,MODE=MULTI,SVC=svc,DEVICE=3380,DBID=yyyyy
/*
//DDKARTE DD *
ADARES PLCOPY NOPPT
/*
//
```

z/VSE Sample Jobs

This section describes the sample jobs provided for z/VSE systems.

ALLOC Job

The following sample JCL is provided for allocating the intermediate data sets MERGIN1 and MERGIN2 which are required for automated PLOG copy/merge processing in nucleus cluster environments.



Notes:

1. When intermediate data sets are used for both CLCOPY and PLCOPY, the data set names must be unique so that they are not overwritten.
2. The data set block size used must be greater than or equal to the largest PLOG block size plus eight. The record length must be set to the block size minus four.

```
// JOB MRGFILE
// ASSGN SYS004,SYSIPT
// ASSGN SYS005,DISK,VOL=VSE209,SHR
// DLBL UOUT,'QA.DB1010.MERGIN1',0
// EXTENT SYS005,VSE209,1,0,10575,300
// EXEC PGM=OBJMAINT
./ CARD DLM=$$
./ COPY
$$
/*
// ASSGN SYS004,SYSIPT
// ASSGN SYS005,DISK,VOL=VSE209,SHR
// DLBL UOUT,'QA.DB1010.MERGIN2',0
// EXTENT SYS005,VSE209,1,0,11625,300
// EXEC PGM=OBJMAINT
./ CARD DLM=$$
./ COPY
$$
/*
/&
```

ADARESPM Job

A sample job ADARESPM is provided to illustrate the manual execution of the PLCOPY merge function. Two intermediate data sets must be supplied. ADARES analyzes the data sets to determine which is to be used as input and which for output. Specific cross-checks determine whether the correct intermediate data set has been supplied; if not, ADARES will not continue. Continuing without the correct input can result in lost updates and inconsistencies if the output is used for REGENERATE or BACKOUT functions.

Once PLOGRn statements have been supplied on the session startup JCL, you do not need to supply them again for ADARES as these are opened using dynamic allocation. If the PLOGRn statements are supplied, they are ignored.

The following sample JCL illustrates the ADARES PLCOPY merge function:

```
// JOB ADARESPM
// *
// * ADARES : COPY/MERGE DUAL/MULTIPLE PROTECTION LOG
// * TWO COPIES OF OUTPUT ARE TO BE CREATED
// * FOR USE WITH AN ADABAS NUCLEUS CLUSTER
// *
// LIBDEF PHASE,SEARCH=(SAGLIB.ASM742,SAGLIB.ADA742)
// DLBL ASSOR1,'EXAMPLE.DByyyyy.ASSOR1'
// DLBL DATAR1,'EXAMPLE.DByyyyy.DATAR1'
// DLBL MERGIN1,'MERGIN1.PLCOPY'
// DLBL MERGIN2,'MERGIN2.PLCOPY'
// TLBL SIAUS1,'EXAMPLE.DByyyyy.PLOG1'
// TLBL SIAUS2,'EXAMPLE.DByyyyy.PLOG2'
// ASSGN SYS021,TAPE
// EXEC ADARUN,SIZE=ADARUN
PROG=ADARES,MODE=MULTI,SVC=svc,DEVICE=3380,DBID=yyy
YY
/*
ADARES PLCOPY TWOCOPIES
/*
/&
```

ADARES PLCOPY NOPPT—Ignore PPT

NOPPT is for emergency use when the PPT has been overwritten or is otherwise unavailable. It specifies that the PPT is to be ignored and that the PLOG data sets of all cluster nuclei are being supplied with link names PLOGnn in the JCL.



Caution: Use this parameter cautiously since it ignores the PPT and all control-type information typically coming from the PPT.

When you use this parameter, you must supply

- the correct intermediate data set; and
- the correct input protection logs from all nuclei with link names PLOG01-nn.



Caution: Without the PPT, ADARES cannot perform any extensive validations on the input data sets.

ADARESIP Job

The following sample job ADARESIP is used with the ADARES PLCOPY NOPPT function:

```
// JOB ADARESIP
// *
// * ADARES : COPY/MERGE DUAL/MULTIPLE PROTECTION LOG
// * FOR USE WITH AN ADABAS NUCLEUS CLUSTER
// * CAUTION: NOPPT EXECUTION!
// LIBDEF PHASE,SEARCH=(SAGLIB.ASM742,SAGLIB.ADA742)
// DLBL ASSOR1,'EXAMPLE.DByyyyy.ASSOR1'
// DLBL DATAR1,'EXAMPLE.DByyyyy.DATAR1'
// DLBL PLOG01,'EXAMPLE.DByyyyy.PLOGR1.NUC1'
// DLBL PLOG02,'EXAMPLE.DByyyyy.PLOGR2.NUC1'
// DLBL PLOG03,'EXAMPLE.DByyyyy.PLOGR1.NUC2'
// DLBL PLOG04,'EXAMPLE.DByyyyy.PLOGR2.NUC2'
// DLBL PLOG05,'EXAMPLE.DByyyyy.PLOGR1.NUC3'
// DLBL PLOG06,'EXAMPLE.DByyyyy.PLOGR2.NUC3'
// DLBL MERGIN1,'MERGIN1.PLCOPY'
// DLBL MERGIN2,'MERGIN2.PLCOPY'
// TLBL SIAUS1,'EXAMPLE.DByyyyy.PLOG1'
// ASSGN SYS021,TAPE
// EXEC ADARUN,SIZE=ADARUN
PROG=ADARES,MODE=MULTI,SVC=svc,DEVICE=3380,DBID=yyyyy
/*
ADARES PLCOPY NOPPT
/*
/&
```

ADARES MERGE CLOG - Merge Nucleus ClusterCommand Logs

In an Adabas cluster environment, command logs (CLOGs) from the cluster nuclei may be manually merged using the ADARES MERGE CLOG NUMLOG=nn function.

The NUMLOG parameter is required: it specifies the number of command log data sets/files to be included in the merge process. The maximum number is 32.

Sequential data sets/files are expected as input to the MERGE CLOG function; therefore, the ADARES CLCOPY function must be executed prior to the ADARES MERGE function.

The timestamp contained in the CLOGLAYOUT=5 format of the CLOG is required for the proper merging of command logs records.

- [z/OS](#)

■ z/VSE

z/OS

The following sample job ADAARESCM for z/OS systems (see the JOBS data set) illustrates the execution of the ADARES MERGE CLOG function:

```
//ADARESCM JOB
//*
//* ADARES : MERGE SEQUENTIAL COMMAND LOGS
//* FOR USE WITH AN ADABAS NUCLEUS CLUSTER
//*
//RES EXEC PGM=ADARUN
//STEPLIB DD DISP=SHR,DSN=ADABAS.Vvrs.LOAD <=== ADABAS LOAD
//*
//DDASSOR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.ASSOR1 <=== ASSO
//DDDATAR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.DATAR1 <=== DATA
//DDWORKR1 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.WORKR1 <=== WORK
//DDCLOG01 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.CLOGR1.NUC1 <=== CLOG1 NUC1
//DDCLOG02 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.CLOGR1.NUC2 <=== CLOG1 NUC2
//DDCLOG03 DD DISP=SHR,DSN=EXAMPLE.DByyyyy.CLOGR2.NUC3 <=== CLOG2 NUC3
//DDSIAUS1 DD DSN=EXAMPLE.DByyyyy.CLOGM, <=== OUTPUT OF
// VOL=SER=ADAxxx,UNIT=TAPE,DISP=(NEW,CATLG) CLOG MERGE
//DDDRUCK DD SYSOUT=X
//DDPRINT DD SYSOUT=X
//SYSUDUMP DD SYSOUT=X
//DDCARD DD *
ADARUN PROG=ADARES,MODE=MULTI,SVC=svc,DEVICE=3380,DBID=yyyyyy
/*
//DDKARTE DD *
ADARES MERGE CLOG,NUMLOG=3
/*
//
```

z/VSE

The following sample job for z/VSE systems (see the JOBS data set) illustrates the execution of the ADARES MERGE CLOG function:

```
//JOB ADARESCM
//*
//* ADARES : MERGE SEQUENTIAL COMMAND LOGS
//* FOR USE WITH AN ADABAS NUCLEUS CLUSTER
//*
// LIBDEF PHASE,SEARCH=(SAGLIB.ASM742,SAGLIB.ADA742)
// DLBL ASSOR1,'EXAMPLE.DByyyyy.ASSOR1'
// DLBL DATAR1,'EXAMPLE.DByyyyy.DATAR1'
// DLBL WORKR1,'EXAMPLE.DByyyyy.WORKR1'
// DLBL ASSOR1,'EXAMPLE.DByyyyy.ASSOR1'
```

```
// DLBL CLOG01, 'EXAMPLE.DByyyyy.CLOGR1.NUC1 '  
// DLBL CLOG02, 'EXAMPLE.DByyyyy.CLOGR1.NUC2 '  
// DLBL CLOG03, 'EXAMPLE.DByyyyy.CLOGR2.NUC3 '  
// TLBL SIAUS1, 'EXAMPLE.DByyyyy.CLOGM '  
// ASSGN SYS021,TAPE  
// EXEC ADARUN,SIZE=ADARUN  
ADARUN PROG=ADARES,MODE=MULTI,SVC=svc,DEVICE=3380,DBID=YYYYY  
/*  
ADARES MERGE CLOG,NUMLOG=3  
/*  
/&
```

ADARES BACKOUT and REGENERATE— Uniquely Identifying Checkpoints

The protection log merge process usually changes the numbering of the PLOG blocks. The PLOG block number of a checkpoint on the original PLOG data set will not necessarily be the same as the block number after the merge. To uniquely identify the checkpoint in this situation, it is necessary to also specify the `NUCID` for all ADARES functions that can specify a `TOBLK` / `FROMBLK` parameter; that is, `BACKOUT` and `REGENERATE`.

The merge process ensures that there is at most one checkpoint per block. It records the (old) block number prior to the merge and the `NUCID` that wrote the checkpoint. When you then specify the block number and `NUCID` as reported in `ADAREP`, ADARES is able to uniquely identify the block.



Note: In an Adabas nucleus cluster environment, `ADAREP` includes the `NUCID` when printing all checkpoint information.

The additional parameters that are required in an Adabas nucleus cluster environment are `NUCID`, `TONUCID`, and `FROMNUCID`.

If the `NUCID` is the same for the starting and ending checkpoint, only the `NUCID` needs to be specified.



Notes:

1. An `ADAREP CPEXLIST` function can be used to determine the original block number and `NUCID` that wrote the checkpoint. This is the block number prior to the merge and the one that ADARES `REGENERATE` and `BACKOUT` expects.
2. `BACKOUT DPLOG` and `BACKOUT MPLOG` are not allowed for a cluster database. The PLOG must be merged before the backout can be performed.

ADASAV Processing Change

Sample JCL is located in the ADASAVRW member of the JOBS data set.

For the following ADASAV functions, the Work data sets/files of all cluster (or noncluster) nuclei for the database that may have been active at the time of the ABEND must be reset:

- RESTONL (database)
- RESTONL GCB
- RESTORE (database)
- RESTORE GCB

This can be done either:

- manually (e.g., by using ADAFRM WORKRESET FROMRABN=1,SIZE=1B); or
- by specifying the Work data sets/files with DD names/link names DD/WORKRn (n=1-9) or DD/WORKnn (nn=10-32) in the JCL for the RESTONL/RESTORE function.

Otherwise, the nuclei that did not have their Work data sets/files reset will give parm-error 42 when started.

The DD/PLOGRn and DD/CLOGRn data sets are not reset in the restore process. They must be either copied/merged by ADARES PLCOPY/CLCOPY or reset by ADAFRM.

ADASAV RESTPLOG -- Uniquely Identifying Checkpoints

After the protection log merge process, the block number will not necessarily be the same. To uniquely identify the checkpoint in this situation, it is necessary to also specify the NUCID parameter for the ADASAV RESTPLOG function when specifying the SYN1 or SYN4 parameter.



Note: An ADAREP CPEXLIST function can be used to determine the original block number and NUCID that wrote the checkpoint. This is the block number prior to the merge and the one that ADASAV RESTPLOG expects.

9

Cluster Nucleus Session End Statistics

■ General Nucleus Information	44
■ Input/Output Statistics	44
■ Command Statistics	45
■ User Statistics	47
■ Efficiency Statistics	47
■ Global Cache Statistics (Cluster Nucleus Only)	48
■ Global Lock Statistics (Cluster Nucleus Only)	53
■ Data Set Activity Statistics	59

In addition to the end-of-session statistics printed by every Adabas nucleus, the statistics for a cluster nucleus also include global cache and lock statistics.

If you are running the selectable unit Adabas Online System (AOS), all of the statistics shown in the following sample output are displayed.

If you are running only the demo version of AOS delivered with Adabas, the statistics displayed are limited as follows:

Section	Displays statistics only for...
Global Cache Statistics	totals, DS, and NI
Global Lock Statistics	buffer flush, hold ISN, new data RABN, and global update command sync locks

This chapter covers the following topics:

General Nucleus Information

```
The  A d a b a s  nucleus session

Started  2001-02-13 22:58  and  ended  2001-02-13 23:05

Duration      00000:06:59  hours
Wait-time     00000:02:26  hours
Cpu-time      00000:00:53  hours
```

Input/Output Statistics

- [I/O Counts \(Including Initialization\)](#)
- [Log Reads and Buffer Efficiency](#)
- [Distribution of ASSO/DATA I/Os by Volser Number \(Excluding Initialization\)](#)

I/O Counts (Including Initialization)

READS	WRITES	

ASSO	4710	6913
DATA	1750	2853
WORK	3	7251
PLOG	0	0
CLOG	0	0

Total	6463	17017

Log Reads and Buffer Efficiency

Log. reads	173,393
Buffer eff.	26.8

Distribution of ASSO/DATA I/Os by Volser Number (Excluding Initialization)

Vol-ser	High RABN	Count
-----	-----	-----
WRKMO1 (ASSO	8082)	11599
WRKMO1 (DATA	5990)	4603
-----	-----	-----
TOTAL		16202

Command Statistics

- [Count of Calls Executed and Threads Used](#)
- [Distribution of Commands by Source](#)
- [Distribution of Commands by Thread](#)
- [Distribution of Commands by File](#)
- [Distribution of Commands by Type](#)

Count of Calls Executed and Threads Used

A d a b a s	executed	10,249 calls
	in	8 threads

Distribution of Commands by Source

Source	Number
-----	-----
Remote commands	0
Local commands	10,102
Internal commands	144
Operator commands	3

Distribution of Commands by Thread

Thread	Number

1	2,657
2	1,803
3	1,401
4	1,300
5	1,193
6	977
7	917
8	1

Total	10,249

Distribution of Commands by File

File	Number

0	4,281
30	5,968

Total	10,249

Distribution of Commands by Type

Cmd-type	Number

A1/4	1,968
CL	44
ET	4,040
N1/2	2,000
OP	43
UC	7
REST	2,147

Total	10,249

User Statistics

```

There were          43  users  participating
Most calls (        303)    initiated by user  USADFMB2
Most I/O-s (        331)    initiated by user  USADFMB2
Most thr.-time (00:00:08) was used by user  USADFMB1

```

Efficiency Statistics

```

46 Formats had to be translated

      0 Formats had to be overwritten

      0 Autorestarts were done

      0 Throw-backs due to  ISN  problem

      0 Throw-backs due to space problem

    143 Bufferflushes  were done

```

- [Buffer Flush Information](#)
- [Actual High-water Marks for Major Pools \(Except the Buffer Pool\)](#)

Buffer Flush Information

```

Flush phases                212
Blocks flushed              28,503
Flush I/Os

Flush requests:
  Return immediately        52,658
  Return after logical flush      0
  Return after entire  flush     15

```

Actual High-water Marks for Major Pools (Except the Buffer Pool)

AREA	ADARUN	PARM	HIGH-WATER-MARK
AB -POOL	NAB=	2000	51712 (0 %)
CQ -POOL	NC =	96000	3840 (4 %)
DUQ -POOL	LDE=	5000	0 (0 %)
FI -POOL	LFP=	20000	6560 (32 %)
HQ -POOL	NH =	16856	588 (3 %)
SC -POOL	LCP=	10000	0 (0 %)
TBI -POOL	LI =	10000	0 (0 %)
TBS -POOL	LQ =	100000	0 (0 %)
UQ -POOL	NU =	500	8844 (6 %)
UQF -POOL	NU =	500	1512 (3 %)
WORK-POOL	LWP=	800000	114296 (14 %)
XID -POOL	XID=		

Global Cache Statistics (Cluster Nucleus Only)

Cast-out dir :	188
Synchronous :	188
Asynchronous :	0
Unlock cast-out:	212
Synchronous :	132
Asynchronous :	80
Directory reads:	3
Synchronous :	0
Asynchronous :	3

- Totals
- Address Converter (AC)
- Data Storage (DS)
- Data Storage Space Table (DSST)
- File Control Block (FCB)
- Normal Index (NI)
- Upper Index (UI)

■ File Statistics for Files with More than 25% of the Total Cache Statistics

Totals

Reads	:	15,006
Synchronous	:	15,006
Asynchronous	:	0
In cache	:	6,245
Not in cache	:	8,761
Area full	:	0
Writes	:	66,726
Synchronous	:	66,726
Asynchronous	:	0
Written	:	66,726
Not written	:	0
Area full	:	0
Validates	:	327,623
Block invalid	:	0
Cast-out reads	:	28,503
Synchronous	:	28,503
Asynchronous	:	0
Deletes	:	0
Timeouts	:	0

Address Converter (AC)

Reads	:	8
Synchronous	:	8
Asynchronous	:	0
In cache	:	0
Not in cache	:	8
Area full	:	0
Writes	:	2,004
Synchronous	:	2,004
Asynchronous	:	0
Written	:	2,004
Not written	:	0
Area full	:	0
Validates	:	5,983
Block invalid	:	0

Cast-out reads :	72
Synchronous :	72
Asynchronous :	0
Deletes :	0
Timeouts :	0

Data Storage (DS)

Reads :	2,775
Synchronous :	2,775
Asynchronous :	0
In cache :	26
Not in cache :	2,749
Area full :	0
Writes :	4,972
Synchronous :	4,972
Asynchronous :	0
Written :	4,972
Not written :	0
Area full :	0
Validates :	9,965
Block invalid :	0
Cast-out reads :	2,921
Synchronous :	2,921
Asynchronous :	0
Deletes :	0
Timeouts :	0

Data Storage Space Table (DSST)

Reads :	2
Synchronous :	2
Asynchronous :	0
In cache :	0
Not in cache :	2
Area full :	0
Writes :	2,004
Synchronous :	2,004
Asynchronous :	0

Written	:	2,004
Not written	:	0
Area full	:	0
Validates	:	4,490
Block invalid	:	0
Cast-out reads	:	69
Synchronous	:	69
Asynchronous	:	0
Deletes	:	0
Timeouts	:	0

File Control Block (FCB)

Reads	:	5
Synchronous	:	5
Asynchronous	:	0
In cache	:	0
Not in cache	:	5
Area full	:	0
Writes	:	4,970
Synchronous	:	4,970
Asynchronous	:	0
Written	:	4,970
Not written	:	0
Area full	:	0
Validates	:	56,029
Block invalid	:	0
Cast-out reads	:	119
Synchronous	:	119
Asynchronous	:	0
Deletes	:	0
Timeouts	:	0

Normal Index (NI)

Reads	:	12,057
Synchronous	:	12,057
Asynchronous	:	0
In cache	:	6,219
Not in cache	:	5,838
Area full	:	0
Writes	:	44,096
Synchronous	:	44,096
Asynchronous	:	0
Written	:	44,096
Not written	:	0
Area full	:	0
Validates	:	25,685
Block invalid	:	0
Cast-out reads	:	22,973
Synchronous	:	22,973
Asynchronous	:	0
Deletes	:	0
Timeouts	:	0

Upper Index (UI)

Reads	:	159
Synchronous	:	159
Asynchronous	:	0
In cache	:	0
Not in cache	:	159
Area full	:	0
Writes	:	8,680
Synchronous	:	8,680
Asynchronous	:	0
Written	:	8,680
Not written	:	0
Area full	:	0
Validates	:	225,471
Block invalid	:	0
Cast-out reads	:	2,349

Synchronous	:	2,349
Asynchronous	:	0
Deletes	:	0
Timeouts	:	0

File Statistics for Files with More than 25% of the Total Cache Statistics

File	30:	
Reads	:	14,998
Writes	:	64,710
Validates	:	323,105

Global Lock Statistics (Cluster Nucleus Only)

- General Control Block (GCB) Lock
- Security Lock
- File Space Table (FST) Lock
- File Lock Table Lock
- Online Save Lock
- Buffer Flush Lock
- Global ET Sync Lock
- Recovery Lock
- Hold ISN Locks
- Unique Descriptor Locks
- ETID Locks
- New Data RABN Locks
- Checkpoint Lock
- ET Data Lock
- Global Update Command Sync Lock
- Parameter Lock

General Control Block (GCB) Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0

Synchronous	:	0
Asynchronous	:	0

Security Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

File Space Table (FST) Lock

Obtains - Conditional	:	1
Granted	:	1
Rejected	:	0
Unconditional	:	1
Synchronous	:	2
Asynchronous	:	0
Releases - Issued	:	2
Synchronous	:	2
Asynchronous	:	0

File Lock Table Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	5
Synchronous	:	5
Asynchronous	:	0
Releases - Issued	:	5
Synchronous	:	5
Asynchronous	:	0

Online Save Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

Buffer Flush Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	152
Synchronous	:	152
Asynchronous	:	0
Releases - Issued	:	152
Synchronous	:	152
Asynchronous	:	0

Global ET Sync Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

Recovery Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

Hold ISN Locks

Obtains - Conditional	:	3972
Granted	:	3972
Rejected	:	0
Unconditional	:	0
Synchronous	:	3972
Asynchronous	:	0
Releases - Issued	:	3972
Synchronous	:	3972
Asynchronous	:	0

Unique Descriptor Locks

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

ETID Locks

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

New Data RABN Locks

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	1000
Synchronous	:	1000
Asynchronous	:	0
Releases - Issued	:	1000
Synchronous	:	1000
Asynchronous	:	0

Checkpoint Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	4
Synchronous	:	4
Asynchronous	:	0
Releases - Issued	:	4
Synchronous	:	4
Asynchronous	:	0

ET Data Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

Global Update Command Sync Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	143
Synchronous	:	143
Asynchronous	:	0
Releases - Issued	:	143
Synchronous	:	143
Asynchronous	:	0

Parameter Lock

Obtains - Conditional	:	0
Granted	:	0
Rejected	:	0
Unconditional	:	0
Synchronous	:	0
Asynchronous	:	0
Releases - Issued	:	0
Synchronous	:	0
Asynchronous	:	0

Data Set Activity Statistics

```
.....  
ADAI03 DDWORKR1          3 READS      7251 WRITES  
ADAI03 DDDATAR1       1750 READS      2853 WRITES  
ADAI03 DDASSOR1       4710 READS      6913 WRITES
```

10

Performance and Tuning

■ ADARUN Parameter Settings	62
■ Allocating Work Data Set Space	62
■ Using Close (CL) Commands	62
■ Deferred Publishing	63
■ Tuning Buffer Flushes	64
■ Optimizing Global Cache and Lock Areas	64
■ Minimizing Communication with Global Areas	66
■ Optimizing Block Sizes	67

This chapter covers the following topics:

ADARUN Parameter Settings

Software AG recommends that you use your existing Adabas ADARUN parameters (or the default values) for each nucleus in an Adabas cluster, and then tune the values after analyzing the performance of the cluster.

Session statistics can be used to determine the best settings for each parameter. The statistics can be displayed using operator commands during the session; the statistics are also printed automatically at the end of a session or in response to an ADADBS REFRESHSTATS command.

For parameters that allot processing resources to the cluster nuclei (such as NU, NH, LP, etc.), Software AG recommends that you set them large enough that each individual cluster nucleus could handle the entire load on the database if the other nuclei were to terminate abnormally.

Allocating Work Data Set Space

Each Adabas cluster nucleus requires its own Work data set to hold its temporary data.

The individual sizes of the different Work parts (1, 2, and 3) as specified by ADARUN parameters such as LP and LWKP2 can be different among the nuclei; however, the overall size of each Work data set must be the same, because the total Work size is stored in the Adabas general control block (GCB). Software AG recommends that you use the same LP and LWKP2 values on each nucleus active for the same database.

For each nucleus, you need to specify shared access to DD/WORKR1. During an offline or online recovery, a nucleus may access the Work data sets belonging to other nuclei in the cluster.

Using Close (CL) Commands

Users are assigned to a nucleus for their entire sessions and should therefore issue Adabas close (CL) commands as appropriate. The close command ends the user's session, making the user eligible for reassignment to another nucleus when the user again issues an Adabas open (OP) command. This allows Adabas Parallel Services to rebalance the workload over the participating nuclei.

Deferred Publishing

Publication of updated blocks to the global cache area is usually deferred until just before the end of the associated transaction. Multiple updates to a block may produce only a single write of the block to the cache rather than a cache write for each update.

The greater the number of database update in parallel transactions, the greater the expected improvement in performance.



Note: Deferred publishing creates an asymmetry between users on the update nucleus, who see uncommitted updates (unless they read with hold), and users on other cluster nuclei, who may or may not see uncommitted updates.

- [Redo Pool](#)
- [ADARUN Parameter LRDP](#)

Redo Pool

Since the write of updated blocks to the cache may fail due to conflicting updates to the same blocks by other nuclei in the cluster, every cluster nucleus must be capable of redoing the updates it has not yet written to the cache. The nucleus maintains information about these updates in the "redo pool".

ADARUN Parameter LRDP

The size of the redo pool is specified by the new ADARUN parameter `LRDP`. The `LRDP` parameter is effective only in a cluster nucleus, that is, when a nonzero `NUCID` is specified.

If `LRDP` is not specified, the nucleus takes as default the value of the `LFIOF` parameter. If `LRDP` is explicitly set to zero, the nucleus writes each update immediately to the cache.

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 a peer nucleus with `LRDP>0`.



Note: If one nucleus runs with `LRDP=0` and a peer nucleus runs with `LRDP>0` and the different cluster nuclei concurrently update the same Data Storage blocks, incorrect DSST entries may be produced. These are reported by `ADADCK`. Such errors are harmless and do not affect the results of the application programs.

The nucleus reports on the use (high watermark) of the redo pool in a shutdown statistic and in the response to the `DRES` command from the operator console or from `ADADBS OPERCOM`.

Tuning Buffer Flushes

When the update load on the database is so high that the buffer flush becomes the bottleneck, you can improve performance by reducing the duration of buffer flushes.

Instead of starting one I/O per volume, a buffer flush can initially start a predetermined number of I/Os on each volume and then start a new one when another I/O on the same volume finishes. This occurs independently on each volume.

The ADARUN parameters `LFIO` and `FMXIO` (see the *Adabas Operations* documentation for details) can be used to control buffer flushes. The `LFIO` parameter enables asynchronous buffer flush operation and sets the I/O pool size. The `FMXIO` parameter sets the limit on the number of I/O operations that can be started in parallel by `LFIO` flush processing.

- [Effect of ASYTVS Parameter Setting](#)
- [Dynamically Modifying the FMXIO Parameter Setting](#)

Effect of ASYTVS Parameter Setting

The meaning of the `FMXIO` parameter is affected by the setting of the `ASYTVS` parameter:

When `ASYTVS=YES` (buffer flushes occur by volume), `FMXIO` specifies the number of I/Os to be started in parallel *on each volume*. The minimum and default number is 1; the maximum number is 16. If you specify a number greater than 16, it is reduced to 16 without returning a message.

When `ASYTVS=NO` (buffer flushes occur in ascending RABN sequence without regard to the distribution of the blocks over volumes), the minimum, default, and maximum values continue to be 1, 60, and 100, respectively.

Dynamically Modifying the FMXIO Parameter Setting

The setting of `FMXIO` can be modified dynamically using the `FMXIO=nn` command from the operator console or the Modify Parameter function of Adabas Online System.

Optimizing Global Cache and Lock Areas

As a user, you must allocate and define sizes that are appropriate to your application needs for the global cache and lock areas.

This section provides guidelines for determining optimal sizes for these areas based on current experience.



Note: There may be sites for which these guidelines are not appropriate.

- [Global Cache Area Size](#)
- [Global Lock Area Size](#)

Global Cache Area Size

The global cache area must be large enough to retain:

- *directory elements* for all blocks that reside in all the buffer pools; and
- enough data elements to keep the changed blocks between buffer flushes (cast-outs).

Directory elements are used to keep track of the cluster members that have a particular block in their buffer pools so that the block can be invalidated should any member modify it.

If the number of directory elements is insufficient, Adabas Parallel Services reuses existing directory elements and invalidates the blocks associated with those directory elements, because they can no longer be tracked. These blocks must then be reread from the database and registered again the next time they are referenced and validated, even though they did not change.

It is generally better to reassign storage for data elements to keep more ASSO and DATA blocks in the global cache area than to define too many directory elements in the global cache area. More data elements than necessary can be used to keep additional blocks to improve the local buffer efficiency.

The number of directory elements need not be greater than the sum of the sizes of all buffer pools divided by the smallest block size in use for ASSO and DATA.

When connecting to the global cache area during startup, the ADAX57 message reports the number of directory elements and data elements. The ADARUN parameters `DIRRATIO` and `ELEMENTRATIO` determine the ratio between the number of directory and data elements.

Global Lock Area Size

All nuclei in a database cluster share the global lock area.

Locks are held for a variety of entities, for example unique descriptor values. These lock types tend to occur with very different frequencies. The amount of lock activity during a session for each lock type is displayed in the shutdown statistics.

It is often the case that ISN locks show the greatest activity. The sum of high-water marks for NH yields an upper limit for the number of ISN locks that were held concurrently during the session.

The global lock manager uses a hash table to allocate and find a specific lock entry.

When the global lock manager receives a lock request (for example, to put an ISN of a file into hold status), it allocates a specific lock entry unless another member of the cluster has already

made a conflicting allocation. A conflicting allocation produces lock contention because another member holds the same lock. Depending on its type, the lock request is then rejected or remains pending, waiting for the associated resource to become available.

The minimum lock structure size can be roughly estimated as:

```
(NU*3 + NH + LDEUQP/16 + MAXFILES*4 + 50) * 240 + 500,000 bytes
```

where `MAXFILES` is the maximum number of files in the database (set in `ADADEF` or `ADAORD`) and `NU`, `NH`, and `LDEUQP` are the `ADARUN` parameters of the cluster nuclei. The formula in parentheses $(NU*3 + NH + LDEUQP/16 + MAXFILES*4 + 50)$ is used to calculate the minimum number of lock records that the cluster nuclei expect to have available.

Minimizing Communication with Global Areas

Most of the additional processing required for Adabas Parallel Services environments compared to a single Adabas nucleus involves communication with the global areas.

For this reason, optimizing the performance of an Adabas Parallel Services environment means minimizing the need for communication with the global areas. It is also important to keep the time required for each communication as short as possible.

- [Avoiding the Hold Option](#)
- [Reducing Direct Interaction with the Global Cache Area](#)

Avoiding the Hold Option

Lock requests usually depend on application requirements. Under data-sharing, the hold option is more expensive and access with the hold option should be avoided unless records will in fact be updated or must be protected from concurrent updates.

Reducing Direct Interaction with the Global Cache Area

Cache area requests occur when blocks:

- that are referenced do not exist in the local buffer pool;
- exist in the local buffer pool but have become invalid due to concurrent updates by other cluster members or from directory reuse; or
- are updated.

The first and second situation require registering and (re)reading the blocks from the global cache area. This is more expensive than just validating blocks.

The first situation is related to the buffer efficiency in a noncluster environment. In a cluster environment, buffer efficiency represents the combined effect of the local buffer pool and the global cache area. In order to reduce the interaction with the global cache, the local buffer pool (LBP) should not be decreased from what would be used in a noncluster nucleus. A large `LBP` parameter and the usage of forward index compression are recommended to improve the buffer efficiency in the local buffer pool.

Optimizing Block Sizes

Although earlier versions of Adabas often worked well with large block sizes, the buffer pool manager and forward index compression features introduced with Adabas version 7 make smaller block sizes more attractive, especially in data-sharing mode.

Use the following guidelines when selecting an optimal block size for ASSO and DATA:



Note: Only general recommendations can be given.

1. Avoid 4-byte RABNs

If the database is not extremely large, avoid 4-byte RABNs as this increases the number of AC blocks by 33%. When growth considerations are taken into account, this may require larger block sizes or limit reductions in block size. The same holds true for the maximum compressed record length.

2. Use forward index compression

Forward index compression can significantly reduce the number of index blocks in a database. Apply forward index compression to all frequently accessed files (or to all files, regardless of their frequency of use). Choose the ASSO block size that is as small as possible but large enough to keep the number of index levels down to 3 or 4.

3. Minimize frequently updated descriptors

When files are updated frequently, the number of blocks that are modified and need to be written to the global cache area often depends on the number of descriptors that have been defined and modified during update processing. Support for additional keys whose descriptor values are subject to frequent modifications becomes even more expensive in a data-sharing environment.

11

Adabas Online System Cluster Environment Screens

■ Display Cluster Members	70
■ Nucleus File Status	71
■ Nucleus Status Flags	72
■ Cluster Usage	74
■ Estimating Sizes for the Cache Structure in a Cluster Environment	81
■ Estimating Sizes for the Lock Structure in a Cluster Environment	84
■ Maintain the User Table	87

This chapter describes the Adabas Online System screens that apply to a cluster environment.

Display Cluster Members

From the Session Monitoring menu, a new function *Display cluster members* (option A) produces the following screen:

```

16:21:45                ***** A D A B A S      BASIC SERVICES *****                2002-07-19
DBID 105                -   Display Cluster Members   -                                PACA002

Total number of nuclei in the cluster: 4

I Sel I Nuc ID I System ID I Jobname   I Status    I Available Services          I
-----
I  _   I  1      I DAEMVS      I ADANUC01 I Active    I All        I
I  _   I  2      I DAEMVS      I ADANUC02 I Inactive  I Lock       I
I  _   I  3      I DDZMVS      I ADANUC03 I Active    I All        I
I  _   I  4      I DDZMVS      I ADANUC04 I Active    I All        I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I
I      I          I              I           I           I           I

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit                Refresh                Menu

```

The screen includes a list of nuclei participating in the cluster and information about the current status of each nucleus.

- ▶ To select a nucleus for additional processing

- Type "S" in the Sel column opposite that nucleus.

- ▶ To display additional information about a nucleus

- Type "D" in the Sel column opposite that nucleus.

For an Adabas cluster nucleus that has a nonzero nucleus ID, its entry in the parallel participant table (PPT) is displayed in a screen similar to the following:

```

16:21:45          ***** A D A B A S  BASIC  SERVICES  *****          2002-07-19
DBID 105          -  Display PPT Entry  -          PACA002

Nuc ID. . .    3 Active Nucleus

Name            Status                                Data Set Name
-----
WORK1
PLOGR1  Ready to be copied/merged  SAG.ADABAS.DB105.PLOGR1
PLOGR2  Being written by nucleus   SAG.ADABAS.DB105.PLOGR2

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help          Exit      Refresh          Menu
  
```

Nucleus File Status

From the Resource Utilization menu, the *Nucleus file status* (option N) has been added and is the equivalent of the DNFV operator command.

```

16:03:17          ***** A D A B A S  BASIC  SERVICES  *****          2002-05-29
DBID 1955          -  Nucleus File Status  -          PACUN02
NucID 1021

          Locking
File      NucID  Access count  Update count  State
-----
24                0              0  Access
25                0              0  Access, Update
  
```

Last page
PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help Repos Exit Refresh - + Menu

In an Adabas cluster environment, the file may be locked for exclusive use by another cluster nucleus. If this is the case and the file is in the nucleus file status table, the Locking NucID column for the file shows the ID of the nucleus that has exclusive control.

The Access count / Update count fields display the number of access or update users, respectively, that refer to the specified file in their user queue elements (UQEs). These users either have specified the file in an OP command with R-option or are using the file in an as yet incomplete transaction.

A State field indicates when the file is used for access only or for access and update. The State field indicates to what extent a nucleus can use a file on its own. If the requested use exceeds the given state, the nucleus must first communicate with the other nuclei in the cluster in order to upgrade the state.

Nucleus Status Flags

From the Resource Utilization menu, a second screen has been added to the *System status* (option S), which displays I/O counts for the ASSO, DATA, WORK, and PLOG data sets; remote and local call distribution; and other current session status information.

18:50:16	***** A D A B A S BASIC SERVICES *****				2002-05-30
DBID 1955	- System Status -				PACUS02
NucID: 1022					
	Physical				
	Reads	Writes	Call Distribution		
	-----	-----	-----		
ASSO	370	67	Remote Logical		0
DATA	3	18	Remote Physical		0
WORK	2	104	Local Logical		860
PLOG		67	Local Physical		0
Logical Reads		349	Logical Reads (binary)	0000015D	
Buffer Efficiency		0.9	No. of HQEs active		0
			No. of UQEs in User Queue ..		2
Format Translations ..		51	No. of CQEs waiting in CQ ..		0
Format Overwrites		0			
			Total intern. Autorestarts .		0
Throw Backs for ISN ..		0	No. of PLOG switches		0
Throw Backs for Space.		0	No. of Bufferflushes		18
page 1 of 2					

```
PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit      Refresh                +      Menu
```

Press PF8 to display an additional screen that indicates if one or more of the following are in progress:

- Online database save running;
- ADAEND in progress;
- Online file save running;
- READONLY/UTIONLY transition;
- READONLY status;
- Update processing suspended;
- ET-sync in progress;
- UTIONLY status; and
- Exclusive-DB-control utility running.

Otherwise, "Adabas operation normal" is displayed.

```
16:47:41          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID 1955          -   System Status   -          PACUS02
NucID: 1021
```

```

                Nucleus Status Flags
                -----
                Adabas operation normal

```

page 2 of 2

```
PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit      Refresh                +      Menu
```

Cluster Usage

From the Resource Utilization menu, *Cluster usage* (option X) displays nucleus cluster statistics that are equivalent of those displayed using the DXCACHE, DXLOCK, and DXFILE operator commands.

The equivalent direct command is:

```
DISPLAY CLUSTERSTATUS
```

```
16:10:31          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                                - Cluster Usage -                                PACUX02

                                Code      Service
                                ----      -
                                C        Cache statistics
                                F        File statistics
                                L        Lock statistics
                                ?        Help
                                .        Exit
                                ----      -

Code ..... _
File Number .. 0
Database ID .. 1955      (WIS1955)                      NucID .. 1021

Command ==>
PF1----- PF2----- PF3----- PF4----- PF6----- PF10----- PF11----- PF12-----
Help              Exit              Fuse      Flist      Menu
```

This section covers the following topics:

- [Cache Statistics](#)
- [File Statistics](#)

- Lock Statistics

Cache Statistics

Choosing *cache statistics* (option C) from the Cluster Usage menu displays the following menu:

```

16:14:23          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                                - Cache Statistics -                                PACUX12

      Code      Service
      ----      -
      K      Cast-out / Directory
      P      Publishing requests
      X      Individual cache blocks
      .      Exit
      ?      Help
      ----      -

Code .....
Database ID .. 1955      (WIS1955)                      NucID .. 1021

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help              Exit      Refresh              Menu

```

The rest of this section describes each of the options on this screen.

- Cast-out / Directory
- Publishing Requests
- All Cache Blocks

Cast-out / Directory

Choosing *cast-out / directory* (option K) from the Cache Statistics menu display the following:

```
16:14:23          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955                      - Cast-out / Directory -                      PACUX12
NucID 1021

      Cast-out Directory Reads          Directory Reads
      -----
Total .....          28          Total .....          5
  Sync .....          1          Sync .....          1
  Async ....          27          Async ....          4

Unlock Cast-out Calls
-----
Total .....          28
  Sync .....          1
  Async ....          27

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help                Exit      Refresh                Detail      Menu
```

Counters have a multiplier column with the following values:

Value	The total shown is in ...
blank	(factor of 1)
K	kilo (factor of 1,000)
M	mega (factor of 1,000,000)
G	giga (factor of 1,000,000,000)

If a number has a multiplier shown, it has been divided by the multiplier, showing the significant digits to 9 places with no decimal point.

Press PF9 to see the entire value. This value is the exact count up to 20 digits in length.

Publishing Requests

Choosing *publishing requests* (option P) from the Cache Statistics menu display the following:

```

16:26:21          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - Publishing Requests -          PACUX12
NucID 1021

          Publishing Request Category
          -----
          Update sync .....          34

          BT or CL or ET ....          162

          Redo threshold ....          2

          Full bufferpool ...          0

          All blocks .....          84

          Specific RABN .....          0

          File DS blocks ....          4

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help       Exit      Refresh      Detail      Menu

```

All Cache Blocks

Choosing *all cache blocks* (option X) from the Cache Statistics menu display the following:

```

16:27:05          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - All Cache Blocks -          PACUX12
NucID 1021

  Reads                                     Writes
  -----                                     -----
  Total .....          167          Total .....          38,176
    Sync .....          24          Sync .....          15,148
    Async .....         143          Async .....          23,028

    In cache .....          49          Written .....          38,176
    Not in cache ..         118          Not written .....          0
    Struc. full ...          0          Struc. full .....          0

  Cast-out Reads                             Other
  -----                                     -----
  Total .....          212          Validates .....          187,677
    Sync .....          212          Invalid .....          43
    Async .....          0          Deletes .....          0

```

Timeouts						0	
Redo processes						0	
PF1-----	PF2-----	PF3-----	PF4-----	PF7-----	PF8-----	PF9-----	PF12-----
Help	Repos	Exit	Refresh	PrevBlk	NxtBlk	Detail	Menu

Use PF7 and PF8 to scroll through the cache blocks; use PF2 to reposition.

Statistics are displayed for the following:

- All cache blocks
- Address converter (AC) cache blocks
- Data Storage (DS) cache blocks
- Data Storage space table (DSST) cache blocks
- File control block (FCB) cache blocks
- Normal index (NI) cache blocks
- Upper index (UI) cache blocks

Press PF9 from the above screen to display the following detail screen:

16:27:05		***** A D A B A S BASIC SERVICES *****		2002-05-29	
DBID 1955		- All Cache Blocks -		PACUX12	
NucID 1021					
Reads		Writes			
-----		-----			
Total		167	Total	38,176	
Sync		24	Sync	15,148	
Async		143	Async	23,028	
In cache..		49	Written ...	38,176	
Not in ...		118	Not writ ..	0	
Stru.full.		0	Stru.full .	0	
Cast-out Reads		Other			
-----		-----			
Total		212	Validates ...	187,677	
Sync		212	Invalid ...	43	
Async		0	Deletes	0	
			Timeouts ..	0	
			Redo procs ..	0	
Press Enter to continue					

File Statistics

Choosing *file statistics* (option F) from the Cluster Usage menu for file 25 displays the following menu:

```

16:37:02          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - File 25 Statistics -          PACUX22
NucID 1021

  Reads                                     Writes
  -----                                     -----
Total ..... 67                               Total ..... 20,157
  Sync ..... 0                               Sync ..... 7,583
  Async ..... 67                             Async ..... 12,574

  In cache ..... 0                           Written ..... 20,157
  Not in cache .. 67                         Not written ..... 0
  Struc. full ... 0                           Struc. full ..... 0

Cast-out Reads                             Other
-----                                     -----
Total ..... 78                               Validates ..... 79,248
  Sync ..... 78                               Invalid ..... 0
  Async ..... 0                               Deletes ..... 0
                                              Timeouts ..... 0
                                              Redo processes ..... 0

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help      Repos      Exit      Refresh                                Detail      Menu

```

Lock Statistics

Choosing *lock statistics* (option L) from the Cluster Usage menu displays the following menu:

```

16:38:16          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                                - Lock Statistics -          PACUX32

Code  Service                                     Code  Service
-----                                     -----
A    Buffer flush lock                             I    Global update command sync lock
B    Checkpoint lock                               J    Hold ISN lock
C    DSF lock                                       K    New-Data-RABN lock
D    ETID lock                                     L    Online save lock
E    File-lock-table lock                           M    Parameter lock
F    FST lock                                       N    Recovery lock
G    GCB lock                                       O    RLOG lock
H    Global ET sync lock                           P    Security lock
.    Exit                                           Q    Spats lock
?    Help                                           R    Unique descriptor lock
-----                                     -----

```

```
Code ..... _
Database ID .. 1955      (WIS1955)          NucID .. 1021

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help              Exit      Refresh                      Menu
```

Each of the options on the Lock Statistics menu displays statistics for a particular lock. For each lock, the screen displays obtain and release information about the various types of that lock that are currently in use by a cluster nucleus:

- The system may obtain locks conditionally or unconditionally, synchronously or asynchronously. A conditional request for a lock may be granted or rejected.
- Releases may be performed synchronously or asynchronously.

Hold ISN Lock

Choosing *hold ISN lock* (option J) from the Lock Statistics menu displays the following:

```
16:38:16          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955                  - Hold ISN Lock -                  PACUX32
NucID 1021

      Obtains                      Releases
      -----                      -----
Conditional ....          16,017      Issued .....          16,017
      Granted .....          16,017      Sync .....          15,971
      Rejected ....           0          Async .....           46
Unconditional ..           0

Sync .....              158
Async .....            15,859

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help      Repos    Exit      Refresh                      PrevLok  NxtLok  Menu
```

Use PF7 and PF8 to scroll through the locks; use PF2 to reposition.

Estimating Sizes for the Cache Structure in a Cluster Environment

These instructions describe how to use the structure size calculator to estimate the size for the cache structure in a cluster environment. The cache structure should be made large enough to provide sufficient space for:

- Tracking all blocks kept in the buffer pools of all connected cluster nuclei (directory elements) and
- Keeping all changed blocks until they are written to the database (data elements).

The assignment of total cache space into directory and data elements is done via the `DIRRATIO` and `ELEMENTRATIO` ADARUN parameters.

The actual cache structure size value is dependent on coupling facility internals and may vary across different coupling facility levels. If the estimates by this structure size calculator are too far off the real allocations reported by cluster nuclei, you can tweak the directory element size value to bring the calculator more in line with the actual coupling facility being used.

▶ To access and use the space calculator for the cache structure in a cluster environment

- 1 Select option **S** on the **Basic Services Main Menu** in AOS.

The **Space Calculation** menu appears.

- 2 Select option **C** on the **Space Calculation** menu.

The **Cache Structure Calculator** panel appears.

```

13:01:16          ***** A D A B A S  BASIC  SERVICES *****          2007-10-02

DBID 1955          -  Cache Structure Calculator  -                      PSPC002

Smallest block size in DB ..... 4092

Largest block size in DB ..... 27990

Buffer pool size (LBP) ..... 104857600____

Size proper for caching blocks .. 104800000____

Max nuclei in cluster ..... 3

Directory element size ..... 400

```

```

Cache structure size (in KB) .... 256000____

For minimum calculation, leave cache structure size field empty.

Modify values, press Enter to provide estimates below.


Cache CFRM SIZE/INITSIZE ..... 256000      ( 250.0      MB)

ADARUN DIRRATIO ..... 62

ADARUN ELEMENTRATIO ..... 49

Cache directory elements ..... 128597

Cache data elements ..... 101633

Cache data element size ..... 2048


PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----

Help                Exit          Lock                Menu

```

3 Specify values for the following fields on the **Cache Structure Calculator** panel.

Field	Description	Default
Smallest block size	Specify a value between 1024 and 32768 bytes.	The smallest block size of the current Adabas Online System database ID.
Largest block size	Specify a value between 1024 and 32768 bytes. If the value of the Smallest block size field exceeds this value, then the Smallest block size value is swapped in.	The largest block size of the current Adabas Online System database ID.
Buffer pool size	Specify a value between 80,000 and 999,999,999,999 bytes.	The value of the LBP parameter setting for the current Adabas Online System database ID.
Size proper for caching blocks	Specify a value between 100000 - 999,999,999,999 bytes.	The LBP parameter of the current Adabas Online System database ID, rounded down to nearest 100000. "Size proper" means that this does not include the overhead in the cache structure required for administering these blocks. Thus, this value specifies how much space should be available in the cache

Field	Description	Default
		structure for keeping changed blocks between buffer flushes and for buffering blocks so that the cluster nuclei do not have to read them from the database.
Max Nuclei in cluster	Specify a value between 2 and 32.	3
Directory element size	Specify a value between 100 and 999 bytes. This value specifies how much space (including the overhead for the access paths) will be used in the cache structure by each directory element.	400
Cache Structure size	Specify a blank for the minimum calculation or specify a value between 100 and 999,999,999 (KB). Although this value is given as an output field, you may want to propose a cache structure size, to see how to allocate the cache space (directory and data elements).	blank

- 4 Press Enter after all values in the previous step are specified.

The following output fields on the screen are filled.

Field	Description
Cache CFRM SIZE/INITSIZE	The recommended cache structure SIZE or INITSIZE specification in the coupling facility resource management policy.
ADARUN DIRRATIO	The recommended ADARUN DIRRATIO parameter settings for the cluster nuclei.
ADARUN ELEMENTRATIO	The recommended ADARUN ELEMENTRATIO parameter settings for the cluster nuclei.
Cache directory	The estimated directory and data element counts resulting from the SIZE/INITSIZE and DIRRATIO settings.
Cache data elements	The estimated directory and data element counts resulting from the SIZE/INITSIZE and ELEMENTRATIO settings.
Cache data element size	This (accurate) value depends only on the largest Asso/Data/Work block size in the database.

Estimating Sizes for the Lock Structure in a Cluster Environment

These instructions describe how to use the structure size calculator to estimate the size for the lock structure in a cluster environment. The lock structure should be made large enough to provide sufficient space for:

- Keeping the lock record elements for all locks held at the same time, and
- Avoiding too much false contention between locks on different resources.

The number of lock table entries and record elements are shown in the results for comparison with the related cluster nucleus message (ADAX70) and to aid your own calculations.

The actual lock structure size value is dependent on coupling facility internals and may vary across different coupling facility levels. If the estimates by this structure size calculator are too far off the real allocations reported by cluster nuclei, you can tweak the lock record element size value to bring the calculator more in line with the actual coupling facility being used.

► To access and use the space calculator for the lock structure in a cluster environment

- 1 Select option **S** on the **Basic Services Main Menu** in AOS.

The **Space Calculation** menu appears.

- 2 Select option **C** on the **Space Calculation** menu.

The **Cache Structure Calculator** panel appears.

```

13:01:16          ***** A D A B A S  BASIC  SERVICES *****          2007-10-02

DBID 1955          -  Cache Structure Calculator  -                      PSPC002

Smallest block size in DB ..... 4092

Largest block size in DB ..... 27990

Buffer pool size (LBP) ..... 104857600____

Size proper for caching blocks .. 104800000____

Max nuclei in cluster ..... 3

Directory element size ..... 400

Cache structure size (in KB) .... 256000____

```


For minimum calculation, leave cache structure size field empty.

Modify values, press Enter to provide estimates below.

Cache CFRM SIZE/INITSIZE 256000 (250.0 MB)

ADARUN DIRRATIO 62

ADARUN ELEMENTRATIO 49

Cache directory elements 128597

Cache data elements 101633

Cache data element size 2048

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----

Help

Exit

Lock

Menu

3 Press PF4 to access the **Lock Structure Calculator** panel:

13:42:29 ***** A D A B A S BASIC SERVICES ***** 2007-08-20

DBID 1955 - Lock Structure Calculator - PSPL002

Max files in database (MAXFILES) 400

Max number of parallel users (NU) 200_____

Number of hold queue elements (NH) 40000

Unique descriptor pool size (LDEUQP) .. 50000

Lock record element size 260

Lock structure size (in KB)

For minimum calculation, leave lock structure size field empty.

Modify values, press Enter to provide estimates below.

Lock CFRM SIZE/INITSIZE 13232 (12.9 MB)

Number of lock table entries 131072

Number of lock record elements 46157 Required min .. 45175

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----

Help Exit Cache Menu

- 4 Specify values for the following fields on the **Lock Structure Calculator** panel.

Field	Description	Default
Max files in database	Specify a value between 3 and 5000. This is the same as the MAXFILES parameter of the ADADEF and ADAORD utilities.	The MAXFILES parameter setting of the current AOS database ID.
Max number of parallel users	Specify a value between 20 and 16,777,215.	The NU parameter setting of the current AOS database ID.
Number of hold queue elements	Specify a value between 20 and 16,777,215. .	The NH parameter setting of the current AOS database ID.
Unique descriptor pool size	Specify a value between 1 and 999,999,999.	The LDEUQP parameter setting of the current AOS database ID.
Lock record element size	Specify a value between 100 and 999. This parameter specifies how much space (including the overhead for the access paths) will be used by each lock record element in the lock structure.	260
Lock structure size	Specify a blank for the minimum calculation or specify a value between 100 and 999,999,999 (KB). Although this value is given as an output field, you may want to propose a lock structure size,	blank

Field	Description	Default
	to see how to allocate the lock table entries and lock table elements.	

- 5 Press Enter after all values in the previous step are specified.

The following output fields on the screen are filled.

Field	Description
Lock CFRM SIZE/INITSIZE	The recommended lock structure SIZE or INITSIZE specification in the coupling facility resource management policy.
Number of lock table entries	The calculated count of lock table entries resulting from the SIZE/INITSIZE setting.
Number of lock record elements	The estimated count of lock record elements resulting from the SIZE/INITSIZE setting. You must actually start a cluster nucleus with the specified parameters to see how many lock record elements it gets from the lock structure. The number on the right side is the minimum number of lock record elements required to be available by the starting cluster nuclei.

Maintain the User Table



Note: This option is available in Adabas nucleus cluster environments only.

A new function has been added to the Session Opercoms menu to support the CLUFREEUSER command. When option V (maintain user table) is selected, the following screen is displayed:

```

16:59:29          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                  - User Table Maintenance -                             PACIV02

                  Code      Service
                  ----      -
                  C          Begin CLUFREEUSER process
                  ?          Help
                  .          Exit
                  ----      -

Code ..... _
TNA ..... 0_____
UID ..... _____
Force ..... _
Global ..... _

Database ID .. 1955  (WIS1955)          NucID .. 1022

```

Command ==>

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help Exit Menu

The CLUFREEUSER command is only valid in cluster environments. It can be issued against the local nucleus only or, with the Global option, against all active and inactive nuclei in the cluster.

The command is used to delete leftover user table elements (UTES) in common storage that are no longer associated with user queue elements (UQEs) in a nucleus where

TNA	is a decimal number specifying the timeout value in seconds. UTes that are not used during the time specified may be deleted if other conditions are fulfilled. If TNA is not specified, UTes may be deleted without regard to their recent use.	
UID	ccccccc	where the argument is 1-8 letters, digits, or embedded '-' signs without surrounding apostrophes.
	'ccccccc'	where the argument is 1-8 characters with surrounding apostrophes.
	X'xxxxxxxxxxxxxxxx'	where the argument is an even number of 2-16 hexadecimal digits enclosed by X' '.
	<p>A character string must be enclosed in apostrophes if it contains characters other than letter, digits, or embedded '-' signs. If a specified character string is less than 8 characters long, it is implicitly padded with blanks. If a specified hexadecimal string is shorter than 16 hexadecimal digits, it is implicitly padded with binary zeros.</p> <p>If the last 8 bytes of a user's 28-byte communication ID match a specific user ID or user ID prefix, that user's UTE may be deleted if other conditions are fulfilled.</p> <p>If UID not specified, UTes may be deleted regardless of their user IDs.</p>	
FORCE	Delete leftover UTes even if the users are due a response code 9, subcode 20. If FORCE is not specified, such UTes are not deleted. Before using the FORCE parameter, ensure that the users owning the UTes to be deleted will not expect any of their transactions to remain open. Specify FORCE on this screen by marking the Force field with any character.	
GLOBAL	Delete leftover UTes throughout the Adabas cluster if they are no longer associated with UQEs and are eligible according to the other specified parameters. Additionally and subject to the other rules, delete leftover UTes if their assigned nuclei have terminated since their last use. If GLOBAL is not specified, only UTes assigned to the local nucleus and used since the nucleus start are eligible for deletion. Specify GLOBAL on this screen by marking the Global field with any character.	

Index

A

- Adabas Parallel Services
 - starting, 5
 - stopping, 11
- ADACOM
 - abnormal termination, 13
 - initialization process, 6
 - normal termination, 12
 - starting, 6
- ADADBS utility
 - REFRESHSTATS function, 29
- ADAICK utility
 - PPTPRINT function, 29
- ADAPLP utility
 - IPLOGPRI function, 30
- ADARES utility
 - CLCOPY function, 30
 - MERGE CLOG function, 38
 - PLCOPY function, 32
- ADARUN parameters
 - determining correct settings, 62
- alert settings
 - cancel alerts, 16
 - messaging performance statistics, 19
 - overview, 15
 - self-termination alerts, 17
 - self-termination query alerts, 18
- ARM
 - explanation, 25
- Autorestart
 - session, 24

C

- cache and lock size, estimating, 66
- cancel alerts, 16
- CLnn command
 - close a nucleus to new users, 62
- Cluster
 - display usage statistics
 - using Basic Services, 74
 - maintain user table, 87
- Command log
 - merging multiple, 30, 38
- COMPRINT
 - messages written to, 6

I

- Initialization
 - ADACOM, 6
 - nucleus, 7
 - sequence, 6

M

- messaging performance statistics, 19

N

- NUCID
 - ADADBS utility OPERCOM parameter
 - for command routing, 29
- Nucleus
 - abnormal termination, 13
 - initialization process, 7
 - normal termination, 12
 - routing utility-issued operator commands, 29
 - serialization during initialization, 7
 - starting, 6
- Nucleus file
 - display status
 - using Basic Services, 71

O

- operator commands
 - self-termination query alerts, 18

P

- Parallel participant table
 - print/dump
 - using utility, 29
- Parameters
 - ADADBS OPERCOM
 - NUCID, 29
- Performance considerations, 61
- Protection log
 - merging multiple, 32

S

- self-termination alerts, 17
- self-termination operator query, 18

- Serialization
 - of nucleus, 7
- Session
 - display status
 - using Basic Services, 72
 - reset statistical values for
 - using utility, 29
- Space allocation, 62
- System
 - display status
 - using Basic Services, 72

T

- Termination
 - abnormal, 12
 - ADACOM, 12
 - normal, 12
 - nucleus, 12
- timeout settings
 - cancel alerts, 16
 - messaging performance statistics, 19
 - overview, 15
 - self-termination alerts, 17
 - self-termination query alerts, 18

U

- User table
 - maintain, 87
- Utilities
 - ADADBS OPERCOM
 - NUCID parameter, 29
- utilities
 - cluster environment, 28

W

- WORK data set allocation, 62