

Adabas Cluster Services

Release Notes

Version 7.4.2

September 2009

This document applies to Adabas Cluster Services Version 7.4.2 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 2009. 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 Release Notes	1
2 Support for Previous Versions	3
3 Required Environment	5
Parallel Sysplex System	6
Software AG Product Environment	6
Sysplex System Prerequisites	7
4 Adabas Add-on Products Overview	9
5 Changes Since Adabas Cluster Services Version 7.2.2	11
Online Invert Support	12
Enforcement of Correct DSST Entries	12
Changed ADARUN Parameters	12
User Table Cleanup	13
New Parameter for ADARES PLCOPY NOPPT	14
Limited Support for Distributed Transactions	15
PLOG/CLOG Messages Changes	15
New ADAPLP Function	15
Triggers and Stored Procedures Facility Support	15
Support for System-Managed Processes	16
LOCAL Parameter Change	16
Cluster Nuclei	17
Online Recovery	17
New Formula for Lock Structure Size Estimation	17
6 Restricted Support for Adabas Features	19
7 Migration from Adaplex+	21
8 Switching Between Cluster and Noncluster Modes / PLOG Handling	23
Scenario 1	24
Scenario 2	24
Scenario 3	24
9 Applying Zaps to Cluster Service Components	27
10 Changes to Adabas Online System	29
Display Cluster Members	30
Nucleus File Status	31
Nucleus Status Flags	32
Cluster Usage	34
Maintain the User Table	40
Index	43

1 Release Notes

This documentation provides detailed information on Adabas Cluster Services version 7.4.2 and describes the updates and modifications of this version.

The Adabas Cluster Services Release Notes are organized in the following topics:

•	Support for Previous Versions
•	Required Environment
•	Adabas Add-on Products Overview
•	Changes Since Adabas Cluster Services Version 7.2.2
•	Restricted Support for Adabas Features
•	Migration from Adaplex+
•	Switching Between Cluster and Noncluster Modes / PLOG Handling
•	Applying Zaps to Cluster Service Components
•	Changes to Adabas Online System

2 Support for Previous Versions

Software AG plans to provide support for previous Adabas Cluster Services releases and SMs as follows:

Release	Date Released	End-of-Support Date
Version 7.2.1	March 2001	April 2002
Version 7.2.2	December 2001	May 2005

Support for a version level will be provided through the month specified in the End-of-Support Date column.

3 Required Environment

- Parallel Sysplex System 6
- Software AG Product Environment 6
- Sysplex System Prerequisites 7

This part of the documentation describes the environment settings required for running Adabas Cluster Services.

The chapter covers the following topics:

Parallel Sysplex System

One or more OS/390 or z/OS operating system images defined as members of a parallel sysplex are required

- connected with an IBM Sysplex Timer TM; and
- connected to a coupling facility; and
- running OS/390 version 2 (release 10), or z/OS (version 1, releases 1-4), or z/OS.e * (releases 3-4)

* Support for z/OS.e is currently restricted to client programs executing in batch, or under TSO or Com-plete.

Software AG Product Environment

Adabas Cluster Services version 7.4.2 requires Adabas version 7.4.2. Please check ServLine24 for cluster-related Adabas zaps when installing Cluster Services.

Adabas Cluster Services version 7.4.2 requires Entire Net-Work version 5.8 or above. The necessary components of Entire Net-Work are delivered with Adabas Cluster Services.

The following Entire Net-Work ZAPs must be applied to version 5.8:

ZAP	Solves a problem with...
WM58028	the interaction between Entire Net-Work and the Adabas 7.4 SVC.

If you install Adabas Online System (AOS), either the demo version delivered with Adabas 7.4 or the version 7.4 selectable unit is required. AOS is compiled under Natural version 3.1 and runs on that and all subsequent versions of Natural.

The full version 7.4 selectable unit is required for the triggers and stored procedures facility. The triggers and stored procedures facility is compiled under Natural 3.1 and runs on that and all subsequent versions of Natural.

Sysplex System Prerequisites

An operational IBM parallel sysplex environment is required. This includes the necessary hardware to support the configuration, the necessary links and channels for connectivity, and the power to provide the required performance during peak production periods.

Adabas Cluster Services uses

- a coupling facility cache structure
- a coupling facility lock structure
- multiple protection logs, command logs, and Work datasets on shared DASD

For testing purposes only, the hardware, operating system, and software environment may be totally different. You may simulate the sysplex cluster using a series of LPARs, including the coupling facility itself. However, under these circumstances, system performance cannot be accurately measured.

Special Technical Requirements

Adabas Cluster Services requires a viable parallel sysplex environment if it is to provide 24x7 availability for an indefinite period of time. The knowledge and experience required to implement and maintain such an environment should be acquired from IBM.

4 Adabas Add-on Products Overview

Adabas Cluster Services version 7.4.2 supports add-on products as follows:

Adabas Product	Compatible Version Level
Adabas Caching Facility (ACF)	version 7.3 (initial release) or 7.4
Adabas Delta Save Facility (ADE)	version 7.4
Adabas Fastpath (AFP)	version 7.3 or above
Adabas Manager (AMA)	version 1.1.1 or above
Adabas Online System (AOS)	version 7.4
Adabas Review (REV)	version 4.3; upgrade Adabas Review to version 4.3 first; then upgrade Adabas to version 7.4.
Adabas SAF Security (AAF)	version 7.3 or above
Adabas Statistics Facility (ASF)	version 7.1
Adabas Transaction Manager (ATM)	version 7.4; Adabas Cluster Services does not yet fully support distributed transactions, but provides limited support through DTP=ET.
Adabas Vista (AVI)	version 7.3 or above

5

Changes Since Adabas Cluster Services Version 7.2.2

▪ Online Invert Support	12
▪ Enforcement of Correct DSST Entries	12
▪ Changed ADARUN Parameters	12
▪ User Table Cleanup	13
▪ New Parameter for ADARES PLCOPY NOPPT	14
▪ Limited Support for Distributed Transactions	15
▪ PLOG/CLOG Messages Changes	15
▪ New ADAPLP Function	15
▪ Triggers and Stored Procedures Facility Support	15
▪ Support for System-Managed Processes	16
▪ LOCAL Parameter Change	16
▪ Cluster Nuclei	17
▪ Online Recovery	17
▪ New Formula for Lock Structure Size Estimation	17

This section provides information on changes to Adabas Cluster Services since version 7.2.2.

The chapter covers the following topics:

Online Invert Support

Adabas Cluster Services 7.4 now supports the online invert function. Support of online reorder will be added in a subsequent release or version.

Enforcement of Correct DSST Entries

When deferred publishing of updates to the coupling facility cache structure is in effect (the default), Cluster Services 7.2.2 might produce incorrect DSST entries. Cluster Services 7.4.2 now contains logic for enforcing correct DSST entries.

Changed ADARUN Parameters

In order to better integrate with the new add-on product Adabas Parallel Services, which features multiple update nuclei running in the same operating system image and working on the same database, the ADARUN parameters used for invoking Cluster Services have been changed as follows:

Old Parameter Name	New Parameter Name
PLXID	NUCID
PLXCACHENAME	CLUCACHENAME
PLXGROUPNAME	CLUGROUPNAME
PLXLOCKNAME	CLULOCKNAME

Furthermore, there is the new parameter `CLUSTER` which specifies what kind of Adabas nucleus is to be run:

CLUSTER Parameter	Meaning
NO (default)	single, noncluster nucleus
LOCAL	Parallel Services nucleus
SYSPLEX	Cluster Services nucleus

For compatibility, Cluster Services 7.4 still supports the old `PLXID` and `PLXyyyyNAME` parameters. Software AG recommends that you change them and use the `CLUSTER`, `NUCID` and `CLUyyyyNAME` parameters instead.

The DD-names that the ADARES utility uses for the intermediate merge datasets in the automated PLOG and CLOG merge process have been changed:

Old DD-name	New DD-name
MERGIN(T)I	MERGIN1
MERGIN(T)O	MERGIN2

First, the DD-names have been made one character shorter (zap AU722069), so that they can be used as symbolic names for Parallel Services on VSE. Second, the implied reference to one intermediate dataset being for 'input' and the other being for 'output' has been removed, since ADARES determines by itself which one is to be read and which one to be written. There is no need to switch the DD-names between consecutive PLCOPY runs.

User Table Cleanup

The new operator command and ADADBS OPERCOM and AOS function CLUFREEUSER is available for situations where leftover User Table Elements in common storage need to be cleaned up. This function deletes such User Table Elements according to specified criteria.

By default, the cluster nucleus receiving this function deletes all User Table Elements assigned to itself which it does not know of (i.e., for which it has no User Queue Elements). If the GLOBAL option is specified, all User Table Elements are deleted that are assigned to nuclei that are no longer active, or that have no associated User Queue Elements in their assigned nuclei.

For more information about the User Table Cleanup, see the section [Maintain the User Table](#).

New Parameter for ADARES PLCOPY NOPPT

The new parameter `SBLKNUM` has been introduced for the ADARES PLCOPY NOPPT function. This function is for use in an emergency case when a last set of PLOGs must be copied and merged after the parallel participant table (PPT) in the Associator has been destroyed. In this case the ADARES utility does not know the last serial block number of the sequential PLOGs produced so far. By default, it starts numbering the blocks at 1 on the next sequential PLOG dataset. If this sequential PLOG is concatenated to its predecessor in a REGENERATE function, ADARES would report an error in the PLOG block number sequence. The `SBLKNUM` parameter can be used to specify the correct starting block number for the sequential PLOG produced by the PLCOPY NOPPT function.

There are two ways to copy and merge the PLOGs when the PPT has been destroyed:

1. Without using the `SBLKNUM` parameter.
 - Run ADARES PLCOPY NOPPT, specifying all direct-access PLOG datasets of all cluster nuclei as DD/PLOGnn datasets for the ADARES utility. ADARES copies and merges the relevant PLOGs and produces a sequential PLOG dataset starting with PLOG block number 1.
 - In case the database or individual files need to be restored and regenerated, run the regenerate in two (or more) job steps. The first REGENERATE should cover all sequential PLOGs produced by PLCOPY functions without the NOPPT parameter; these PLOGs may be concatenated into one sequential input. If the REGENERATE is performed with ET logic (REGENERATE database or with `CONTINUE` parameter), specify the `NOAUTOBACKOUT` parameter. The second REGENERATE should cover the sequential PLOG produced by the PLCOPY with NOPPT.
2. With using the `SBLKNUM` parameter:
 - Determine the highest PLOG block number written so far by the regular PLOG merge processes. This number can be determined from an up-to-date checkpoint listing, from an ADARAI LIST protocol, or from the job protocol of the most recent PLCOPY execution.
 - Run ADARES PLCOPY NOPPT, specifying all direct-access PLOG datasets of all cluster nuclei as DD/PLOGnn datasets for the ADARES utility. Specify the `SBLKNUM` parameter with value `n+1`, where `n` is the highest PLOG block number written so far, as determined in the previous step. ADARES copies and merges the relevant PLOGs and produces a sequential PLOG dataset starting with the specified PLOG block number.
 - In case the database or individual files need to be restored and regenerated, no special arrangements need to be made to run the REGENERATE in separate job steps.

Limited Support for Distributed Transactions

Adabas Cluster Services 7.4.2 can in a limited way participate in distributed Adabas transactions coordinated by the Adabas Transaction Manager, as follows:

- In each distributed transaction, a single cluster database is allowed to participate. Different cluster databases can participate in different distributed transactions.
- All nuclei of a cluster database participating in distributed transactions must be run with the new ADARUN parameter setting `DTP=ET`.
- Cluster databases can participate in distributed transactions only if all resource managers involved in these transactions are Adabas databases running with `DTP=RM`. `DTP=ET` cannot be supported when an external transaction framework is included in the transaction (CICS/RMI or RRMS).

Full support by cluster databases for distributed transactions will be provided in a subsequent release of Adabas Cluster Services.

PLOG/CLOG Messages Changes

The messages in the PLOG/CLOG merge protocol of the ADARES PLCOPY/CLCOPY functions have been revised to be more comprehensive and meaningful (zaps AU722055 and AU722069).

New ADAPLP Function

The new ADAPLP function IPLOGPRI prints out the contents of an intermediate dataset (MERGINI or MERGINO) in the PLOG merge process (zap AU722068). The dataset is specified in the JCL via DD-name/link name DD/PLOG. The output is similar to that of the SPLOGPRI function.

Triggers and Stored Procedures Facility Support

The triggers and stored procedures facility (SPT) is now fully supported in a sysplex cluster environment:

- The ADARUN `SPT` parameter is now global and must be set the same on all active nuclei.
- When a REFRESH is executed, the trigger table is passed to all active nuclei.
- Only the first nucleus to initialize reads the trigger file to create the trigger table; all subsequent nuclei obtain the trigger table from one of the already active nuclei.

If you choose to use triggers and stored procedures in a cluster environment, you must set the Error Action field in the SPT profile under Adabas Online System to "halt" or "reject", but *not* "ignore". Note that in "reject" mode, the way a command is handled across two nuclei can be different if the command has a trigger associated with it.

Support for System-Managed Processes

Adabas Cluster Services 7.4 supports certain processes that are managed by the operating system. These processes are supported for the cache and lock structures used by Adabas Cluster Services.

Two types of system-managed processes are supported: *rebuild processes* and *duplexing rebuild processes*. See the following IBM documentation for additional information:

- MVS Setting Up a Sysplex
- MVS Programming: Sysplex Services Guide
- MVS System Commands (the SETXCF operator command)

The minimum system requirements for support of system-managed processes is z/OS version 1, release 2 or above, and a coupling facility with CF level 12.

System-managed rebuild dynamically rebuilds a coupling facility structure in the same or another coupling facility, while an active application (in this case, Adabas Cluster Services) may be using the structure. It can be used to reduce planned outage time for applications using the coupling facility.

System-managed duplexing rebuild dynamically duplicates a coupling facility structure into another coupling facility, while an application may be using the structure. It can be used to reduce unplanned outage time for applications using the coupling facility.

The options of the SETXCF operator command initiating these system-managed processes are SETXCF START,REBUILD and SETXCF START,REBUILD,DUPLEX, respectively.

LOCAL Parameter Change

The meaning of the LOCAL parameter has been adapted to that of a non-cluster database. It specifies whether or not the database can be accessed via Net-Work from other systems. If LOCAL=NO, the cluster database can be accessed from all systems to which a Net-Work connection exists. If LOCAL=YES, all cluster nuclei must run on the same system image and the database can be accessed only from that system.



Note: For a cluster database, the `LOCAL` parameter applies to the common DBID target established by the cluster nuclei. In contrast, the `FORCE` parameter (force overwrite of ID Table Entry) applies to the `NUCID` parameter of each individual nucleus.

Cluster Nuclei

Cluster nuclei now perform an automatic check during session start to verify that resource locks via Global Resource Serialization (GRS) are effective against one another. If GRS has been incorrectly configured such that an ENQ performed by a starting nucleus has no effect against one of the already active nuclei, the starting nucleus will terminate with parm-error 106.

Online Recovery

At the beginning of online recovery after a peer nucleus failure, the nuclei participating in the recovery process will now in most cases print out the NUCID of the nucleus performing the autorestart.

New Formula for Lock Structure Size Estimation

The minimum lock structure size can be roughly estimated as

$$(NU*2 + NH + LDEUQP/16 + MAXFILES*4 + 50) * 240 + 500,000 \text{ bytes}$$

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 term in parenthesis $(NU*2 + NH + LDEUQP/16 + MAXFILES*4 + 50)$ is the minimum number of lock records that the cluster nuclei expect to have available.

Further help is available at IBM's Coupling Facility Structure Sizer website at <http://www-1.ibm.com/servers/eserver/zseries/cfsizer>:

1. Read the text at the 'DISCLAIMER' link.
2. Click on 'OEM LOCK' link in the left panel.
3. Fill in the template:
 - Set the 'Number of Lock Entries' to a power of 2 (32,768 or higher).
 - Set the 'Number of Systems Sharing the Lock Structure' to the number of system images where cluster nuclei are to run.
 - Set the 'AAI' field to the number in parenthesis above, that is:

$NU*2 + NH + LDEUQP/16 + MAXFILES*4 + 50$

- Check the 'OEM Lock Structure' box.
4. Click on the 'Submit' button.
 5. The CFSIZER tool then proposes a lock structure size.



Note: This link and procedure were valid in December 2002. Software AG cannot guarantee the future availability of this website, nor the accuracy of its results.

Note that the provided formula and procedure yield a minimum lock structure size. If it turns out that the lock structure of an Adabas cluster incurs a lot of 'false contention' (i.e., the appearance of resource conflicts which do not really exist), then it may be advisable to increase the lock structure in order to reduce false contention and thus improve performance.

6 Restricted Support for Adabas Features

This section describes facilities internal to Adabas that are not supported for sysplex cluster nuclei running under Adabas Cluster Services version 7.4.2. The facilities are supported normally for noncluster nuclei; however, no Cluster Services functionality is available to them.

For an Adabas cluster nucleus running in an parallel sysplex environment, the following features are not available and cannot be specified:

- `MODE=SINGLE`
- sequential protection log (DDSIBA)
- synchronous buffer flush (LFIOP=0)

The following features are not currently supported for an Adabas cluster nucleus running in an parallel sysplex environment under Adabas Cluster Services version 7.4.2, but will be supported in subsequent versions of the product:

- `READONLY=YES` (receives PARM ERROR 71 if attempted).
- `UTIONLY=YES` can be specified for a sysplex cluster nucleus; if you start cluster nuclei with conflicting settings of `UTIONLY`, the system will change them to conform to the setting of the first active nucleus. Currently, however, the `UTIONLY` setting cannot be changed using an ADADBS OPERCOM or Adabas Online System function. Once the cluster is started, the only way to change the `UTIONLY` setting is to bring down the whole cluster and restart it with a different setting.
- advance file lock (ALOCKF).
- two-phase commit processing (`DTP=RM`), but limited support (`DTP=ET`) for distributed transactions is available. See section *Limited Support for Distributed Transactions* for more information.
- online reorder.

Enhanced error recovery is supported; however, option changes are effective only for the local nucleus.

TCP/IP direct links are supported; however, the IP address/port is tied to an individual nucleus.

7 Migration from Adaplex+

All nuclei of Adabas Adaplex+ must be migrated to Adabas Cluster Services version 7.4 at one time. The Adabas ADACNV utility must be used to convert Adabas internal structures and all PLOGs must be reformatted before the Adabas Cluster Services version 7.4 nucleus is started.

It is not possible for a nucleus running Adabas Adaplex+ to participate in an Adabas Cluster Services version 7.4 environment. And the reverse: it is not possible for an Adabas Cluster Services version 7.4 nucleus to run in an Adabas Adaplex+ environment.

8 Switching Between Cluster and Noncluster Modes / PLOG

Handling

■ Scenario 1	24
■ Scenario 2	24
■ Scenario 3	24

Switching from cluster to noncluster mode (or vice versa) is possible only after normal termination. A starting nucleus checks in the PPT whether the previous session ended abnormally with a pending autorestart. If this is the case and the previous nucleus ran in the same mode as the starting nucleus, the session autorestart logic will be executed. If the previous nucleus ran in a different mode than the starting nucleus, the session start will terminate with an error.

The following sections illustrate a few scenarios where a cluster nucleus starts after the normal termination of a noncluster nucleus. PLOGRQ is not set to FORCE. These scenarios apply to two PLOGs as well as up to eight PLOGs.

Scenario 1

The previous session was noncluster mode, there are remaining PLOGs to be copied, there is no UEX2/12 in use, and the PLOG datasets are different from what was used in the previous session. The results of this scenario are as follows:

- The information in the PPT entry of the noncluster nucleus remains, and the new entry of the cluster nucleus is written.
- Initialization continues.

Scenario 2

The previous session was noncluster mode, there are remaining PLOGs to be copied, there is no UEX2/12 in use, and the PLOG datasets are the same as what was used in the previous session of a noncluster nucleus. The results of this scenario are as follows:

- A warning that the PLOG is being overwritten will occur and the PLOG flag in the previously used PPT block will be reset or the PPT entry will be overwritten (whichever is appropriate).
- Initialization continues.

Scenario 3

The previous session was noncluster mode, there are remaining PLOGs to be copied, UEX2/12 is in use, and the PLOG datasets are different from what was used in the previous session of a noncluster nucleus. The results of this scenario are as follows:

- UEX2/12 is called to submit a PLCOPY job that will copy and merge the PLOGs. The information in the PPT entry of the noncluster nucleus remains, and the new entry of the cluster nucleus is written.

- Initialization continues.



Note: If ADARES detects that there is data to be copied both from a cluster nuclei and from a noncluster nucleus (different PLOGs), it will copy the oldest data first.

9 Applying Zaps to Cluster Service Components

Usually, zaps for components of Cluster Services (Adabas nuclei, the SVCCLU component of the Adabas SVC, the ADACOM task) can be applied and made active one component at a time. That is, individual components can be shut down, have the zap applied, and be brought up again without ever shutting down the entire cluster. This is the default method of applying zaps, which is in effect if the zap description does not explicitly state otherwise.

In some cases, it may be possible that applying and activating a zap one component at a time would introduce erroneous behavior in the components that have not yet been zapped. If this is the case for a zap, it will be clearly indicated in the zap description, and instructions will be given for how to apply and activate the zap properly.

10 Changes to Adabas Online System

▪ Display Cluster Members	30
▪ Nucleus File Status	31
▪ Nucleus Status Flags	32
▪ Cluster Usage	34
▪ Maintain the User Table	40

This section describes changes to Adabas Online System version 7.4 that apply to the Adabas cluster environment.

The chapter covers the following topics:

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    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                               DataSet 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          0          Access
25          0          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 datasets; 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
```

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
    
```

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
-----
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
-----
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.

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	is a character string or hexadecimal byte string as follows:	
	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' '.
	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. 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. If UID not specified, UTES may be deleted regardless of their user IDs.	
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
 - support for earlier releases, 3
- Adabas features
 - restricted support for, 19
- Adaplex+
 - migrating from, 21

C

- Cluster
 - display usage statistics
 - using Basic Services, 34

M

- Migration
 - from Adaplex+, 21

N

- Nucleus file
 - display status
 - using Basic Services, 31

S

- Session
 - display status
 - using Basic Services, 32
- System
 - display status
 - using Basic Services, 32

T

- Triggers and stored procedures
 - restricted support for, 15

