

Adabas Parallel Services

Reference

Version 8.4.1

October 2019

This document applies to Adabas Parallel Services Version 8.4.1 and all subsequent releases.

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

Copyright © 2019 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

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

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third-Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Use, reproduction, transfer, publication or disclosure is prohibited except as specifically provided for in your License Agreement with Software AG.

Document ID: ASM-AASMREF-841-20210909

Table of Contents

Preface	v
1 About this Documentation	1
Document Conventions	2
Online Information and Support	2
Data Protection	3
2 Initialization Parameters	5
3 ADACOM Initialization Parameters	7
DBID - Database Identification	9
FORCE - Force Restart after Abnormal Termination	9
LOC - Specify PLXCB Location	10
NU - Number of User Table Elements on the OS Image	10
SVC - Interregion Communication Number	11
4 About ADARUN Parameters for Cluster Nuclei	13
Global ADARUN Parameters	14
ADARUN Parameter Usage in Cluster Environments	15
Specifying ADARUN Parameters for Cluster Nuclei	25
5 Useful ADARUN Parameters for Adabas Parallel Services	27
CLOGMRG Parameter: Cluster Command Log Merge	28
CLUCACHEEXTRA Parameter: Adjust Shared Cache Vector Size	29
CLUCACHESIZE Parameter: Size of the Global Cache Area	29
CLUCACHETYPE Parameter: Global Cache Area Storage Type Option	31
CLUCACHEUNCHANGED Parameter: Global Cache Unchanged Block Option	33
CLUGROUPNAME Parameter: Cluster Group Name	34
CLULOCKSIZE Parameter: Size of the Global Lock Area	34
CLUPUBLPROT Parameter: Publishing of Protection Data	35
CLUSTER Parameter: Adabas Nucleus Cluster Session Control	36
CLUWORK1CACHE Parameter: Number of WORK Blocks in Global Cache	36
DIRRATIO / ELEMENTRATIO Parameters: Cluster Cache Directory-to-Data Elements Ratio	38
LDEUQP Parameter: Length of Unique Descriptor Pool	39
LRDP Parameter: Length of the Redo Pool	40
MLWTO Parameter: Multi-Line Console Message Control	41
MXCANCEL Parameter: Timeout Threshold for Canceled Peer Nucleus	41
MXCANCELWARN Parameter: Timeout Threshold for Canceled Peer Nucleus Warning	42
MXMSG Parameter: Timeout Threshold for Internucleus Command Processing	43
MXMSGWARN Parameter: Timeout Threshold for Internucleus Command Processing Warning	44
MXWTOR Parameter: Self-Termination Operator Query Interval	45
NUCID Parameter: Adabas Cluster Nucleus ID	46
PROGRAM Parameter: Program to Run	47

SMFDETAIL Parameter: Selecting Detail SMF Record Sections	48
UPDATECONTROL Parameter: Control Scheduling of Update Commands During Buffer Flush	50
V64BIT Parameter: 64-Bit Virtual Storage Option	51
6 Cluster Operator Commands	53
ADACOM Operator Commands	54
Adabas Cluster Nucleus Operator Commands	64
Index	87

Preface

This documentation provides reference information for Adabas initialization parameters, cluster operator commands as well as the Adabas Online System screens pertinent to Adabas Parallel Services.

This document is organized as follows:

- Initialization Parameters* Describes the initialization parameters needed for an Adabas Parallel Services environment.
- Cluster Operator Commands* Describes operator commands for ADACOM and for an Adabas Parallel Services cluster nucleus.

1 About this Documentation

▪ Document Conventions	2
▪ Online Information and Support	2
▪ Data Protection	3

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Monospace font	Identifies service names and locations in the format <i>folder.subfolder.service</i> , APIs, Java classes, methods, properties.
<i>Italic</i>	Identifies: Variables for which you must supply values specific to your own situation or environment. New terms the first time they occur in the text. References to other documentation sources.
Monospace font	Identifies: Text you must type in. Messages displayed by the system. Program code.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.
	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the symbol.
[]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information and Support

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at <https://documentation.softwareag.com>.

Software AG Empower Product Support Website

If you do not yet have an account for Empower, send an email to empower@softwareag.com with your name, company, and company email address and request an account.

Once you have an account, you can open Support Incidents online via the eService section of Empower at <https://empower.softwareag.com/>.

You can find product information on the Software AG Empower Product Support website at <https://empower.softwareag.com>.

To submit feature/enhancement requests, get information about product availability, and download products, go to [Products](#).

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the [Knowledge Center](#).

If you have any questions, you can find a local or toll-free number for your country in our Global Support Contact Directory at https://empower.softwareag.com/public_directory.aspx and give us a call.

Software AG Tech Community

You can find documentation and other technical information on the Software AG Tech Community website at <https://techcommunity.softwareag.com>. You can:

- Access product documentation, if you have Tech Community credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.
- Access articles, code samples, demos, and tutorials.
- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.
- Link to external websites that discuss open standards and web technology.

Data Protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR). Where applicable, appropriate steps are documented in the respective administration documentation.

2 Initialization Parameters

This documentation provides detailed information on the initialization parameters needed for an Adabas Parallel Services environment.

ADARUN parameters are used to customize the Adabas environment. For Adabas Parallel Services environments, two types of initialization parameters must be specified, based on the setting of the ADARUN `PROGRAM` parameter:

- To run an Adabas nucleus, the name must be ADANUC.
- To set up an ADACOM task, the name must be ADACOM (no other ADARUN parameters are recognized for the ADACOM program run).

The chapter is organized in the following parts:

<i>ADACOM Initialization Parameters</i>	Describes the ADACOM initialization parameters required by Adabas Parallel Services.
<i>About ADARUN Parameters for Cluster Nuclei</i>	Describes the Adabas nucleus (ADANUC) initialization parameters required by Adabas Parallel Services.
<i>Useful ADARUN Parameters</i>	Provides a reference for all of the ADARUN parameters useful to Adabas Parallel Services processing.

3 ADACOM Initialization Parameters

- DBID - Database Identification 9
- FORCE - Force Restart after Abnormal Termination 9
- LOC - Specify PLXCB Location 10
- NU - Number of User Table Elements on the OS Image 10
- SVC - Interregion Communication Number 11

An unlimited number of Adabas Parallel Services nucleus clusters, each with up to 31 members sharing a common database, can be defined for an operating system image.

The ADACOM initialization task (ADARUN PROGRAM=ADACOM) must be run in order to set up the environment, and it must be maintained in order to monitor and control the nuclei of one or more Adabas Parallel Services clusters.

ADACOM initialization parameters specify the Router ID / DBID combinations (sets) that the ADACOM is to manage.

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

Other ADACOM parameters are discussed in the following sections.

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

ADACOM prints global messages that apply to all Router ID / DBID sets to two output data sets or files. One is identified by the DD or link name *Pssdddd*, where *ss* is the last two digits of the SVC number and *dddd* is the DBID. The *Pssdddd* data set or file is also used by Adabas Cluster Services. The second output data set or file is identified by the DD or link name *Dssdddd* and is used only by Adabas Parallel Services. ADACOM uses the *Dssdddd* data set or file to record which cache and lock spaces it allocates and releases for the associated cluster database, as the Adabas Parallel Services nuclei start and end.

On z/OS systems, ADACOM automatically allocates these data sets in the spool with SYSOUT=*, if they are not explicitly specified.

For a sample job to run ADACOM, read the z/OS section *Create a Startup Procedure for ADACOM* in the *Adabas Parallel Services Installation Guide*.

DBID - Database Identification

Parameter	Specify...	Minimum	Maximum	Default
DBID	the database to be used.	1	65000	none

The DBID parameter designates the ID of an Adabas Parallel Services cluster's external physical database; that is, the database number that the user programs address to send commands to the single physical database of an Adabas Parallel Services nucleus cluster.

This number must be unique among all Adabas database IDs, NUCIDs, Natural buffer pool IDs, etc.

FORCE - Force Restart after Abnormal Termination

Parameter	Specify...	Possible Values	Default
FORCE	whether ADACOM forces a restart after an abnormal termination.	YES NO	NO

If Adabas Parallel Services believes there are still active nuclei, ADACOMs, or users on an image after a termination, a NU parameter value that is different from the NU value already in effect will not be recognized on restart. If you are certain that Adabas Parallel Services is wrong in its belief, you can use FORCE=YES to force a clean restart.

Note, however, that if a cluster nucleus or an ADACOM with the same SVC number or IDTNAME is active on the image where FORCE=YES is used, it will crash.



Notes:

1. FORCE=YES works only when the NU parameter value is being changed and has the effect of resetting the environment.
2. The nucleus ADARUN parameter FORCE is neither seen nor used by the SVCCLU. The ADACOM parameter FORCE may only be used to change the NU that is valid for the cluster.
3. If you use the ADACOM FORCE parameter, remember to remove it from ADACOM after you have reset the environment.

LOC - Specify PLXCB Location

Parameter	Specify...	Possible Values	Default
<u>L</u> OC	where to place the PLXCB structures that maintain information about nuclei and users.	CSA <u>D</u> SP	CSA

To provide relief from 31-bit ECSA storage constraints you may direct the PLXCB be placed in a dataspace. This is available only under z/OS.



Notes:

1. The dataspace is mapped into 31-bit private storage in the ADACOM and Adabas nucleus address spaces and requires virtual storage space there (less than 1 MB for every 10,000 configured users, as per the NU parameter).
2. A mapped dataspace can be captured in SYSUDUMP.
3. The dataspace is a common storage dataspace (i.e., allocated with SCOPE=COMMON). The number of dataspaces of this kind in one operating system image is limited by the MAXCAD parameter in SYS1.MACLIB(IEASYSxx), which can range from 10 through 250 and has a default value of 50.

NU - Number of User Table Elements on the OS Image

Parameter	Specify...	Minimum	Maximum	Default
<u>N</u> U	the number of users that can be active in parallel on cluster nuclei in the image.	20	16 777 215	200

In the extended CSA (ECSA), the `SVCCLU` maintains a user table with entries (UTEs) containing information about every active user in the cluster nuclei on the operating system image. A UTE is assigned when a user issues an `OP` command or (if the user did not issue an `OP` command) at the first Adabas command. The UTE is released when the user issues a `CL` command or when a timeout occurs.

The ADACOM `NU` parameter specifies the number of concurrent users (UTEs) allowed for all the nuclei of a cluster. The first ADACOM started governs the value for `NU`: different values set for subsequent nuclei or ADACOMs are ignored.

To manually change the `NU` value, you must stop all cluster nuclei and ADACOM tasks in the image, modify the `NU` value for all the nucleus and ADACOM jobs, and then restart.

When the `NU` parameter is set to zero, any space allocated for the user table is freed, the Adabas Parallel Services control block is freed, and the ADACOM task terminates. However, if ADACOM believes that there are active nuclei, ADACOMs, or users, and `FORCE=YES` is not used, no action is taken.

For example, the following parameter would run the Adabas Parallel Services cluster nuclei with 500 elements in the user table:

```
NU=500
```

SVC - Interregion Communication Number

Parameter	Specify...	Possible Values	Default
<code>SVC</code>	the supervisor call number to be used for the Adabas Parallel Services nucleus cluster session.	z/OS: 200-255 z/VSE: see text	z/OS: 249 z/VSE: 45

The `SVC` is used to perform various Adabas internal functions; the number is used to communicate between the users and the database.

The `SVC` number is specified as an integer and must correspond to the number used for the Adabas `SVC` (ADASVC).

- For z/OS, valid `SVC` values are 200-255.
- For z/VSE, 45 is the recommended value but any free `SVC` value can be used. Read the *Adabas Installation Guide* for information about finding free z/VSE Adabas `SVC` values.

For example, the following parameter will execute an Adabas Parallel Services nucleus cluster session on a z/OS image using ADASVC 202.

```
SVC=202
```


4 About ADARUN Parameters for Cluster Nuclei

- Global ADARUN Parameters 14
- ADARUN Parameter Usage in Cluster Environments 15
- Specifying ADARUN Parameters for Cluster Nuclei 25

Software AG recommends that you use the default settings (or your existing values) of the Adabas ADARUN parameters for each Adabas nucleus in an Adabas Parallel Services cluster, and then tune the values after analyzing the performance of the node or cluster. Read *Performance and Tuning* in the *Adabas Parallel Services Operations* for information about expected differences.

Session statistics can be used to determine the best settings for each parameter. The statistics are printed automatically at the end of a session, but can also be displayed using nucleus or ADACOM operator commands during the session.

For your convenience, ADARUN parameters that are most significant to Adabas Parallel Services usage are described in *Useful ADARUN Parameters*, elsewhere in this guide.

This chapter covers the following topics:

- *Global ADARUN Parameters*
- *ADARUN Parameter Usage in Cluster Environments*
- *Specifying ADARUN Parameters for Cluster Nuclei*
- *Useful ADARUN Parameters for Adabas Parallel Services*

Global ADARUN Parameters

ADARUN parameters that must be the same for all nuclei in the cluster are called *global*.

Some global parameters are set at nucleus startup and cannot be changed during the ensuing session; other global parameters can be changed during a session:

- When you set a value for a global parameter that cannot be changed after the first nucleus in an Adabas cluster has started, intracuster messages are used to communicate the ADARUN parameter settings of the first nucleus to all subsequent cluster nuclei. Each subsequent nucleus receives this information during initialization and determines whether its global nonchangeable parameters are equal to those of the first nucleus.

If they are not equal, the nucleus fails with a parameter error. The nonequal global changeable parameters are reset to the value retrieved from the intracuster messages and a corresponding message is printed.

- If you change the value of a global parameter that can be changed during a running session, the nucleus on which you make the change acquires a "parameter change lock", makes the changes in its local parameter area, and communicates the changes to the other cluster nuclei using intracuster messages.

All other nuclei in the cluster receive the intracuster messages containing the global parameters that have changed, change the parameters in their local parameter area, and send an "acknowledge" message.

ADARUN Parameter Usage in Cluster Environments

A cluster nucleus makes use of:

- *global* parameters, whose values are enforced by Adabas Parallel Services to be equal for all nuclei in a cluster. Some of these parameters can be modified (globally modifiable) during a session using an operator command or the Adabas Online System (NISNHQ, NONDES, and AOSLOG are only modifiable using AOS); others are fixed (globally fixed) and cannot be modified while the nuclei in the cluster are running.
- *local* parameters, which can be different for each nucleus. Some of these parameters are modifiable (locally modifiable) using an operator command or the Adabas Online System; others are fixed (locally fixed) and cannot be modified.

A few Adabas ADARUN parameters are not available to a cluster nuclei (No).

In the following table, the "N/A" column indicates which ADARUN parameters are not available to a cluster nucleus, the "LF" column indicates whether the parameter is a locally fixed parameter, the "LM" column indicates whether the parameter is a locally modifiable parameter, the "GF" column indicates whether the parameter is a globally fixed parameter, and the "GM" column indicates whether the parameter is a globally modifiable parameter.



Note: The parameters for which links are provided have particular use in Adabas Parallel Services environments and are described in [Useful ADARUN Parameters for Adabas Parallel Services](#), elsewhere in this guide. A description of every ADARUN parameter (including those not described in the Adabas Parallel Services documentation) can be found in the *Adabas Operations Manual*.

Parameter	Usage	N/A	LF	LM	GF	GM
AOSLOG	Log to DDPRINT commands issued by AOS or ADADBS OPERCOM that modify the active nucleus			LM		
AREXCLUDE	Exclude file(s) from autorestart		LF			
ARMNAME	Name used to activate ARM		LF			
ARNWORKBUFFERS	Number of work buffers allocated during autorestart processing.		LF			
ASSOCACHE	Controller caching control for the Associator component		LF			
ASSOSPACEWARN	Associator space notifications					GM
ASYTVS	Asynchronous buffer flush based on volser			LM		
AUTOCQENV	Type of user affected by the AUTOCQTIME parameter setting		LF			

Parameter	Usage	N/A	LF	LM	GF	GM
AUTOQCQTIME	Maximum wait time for commands in the command queue during online recovery		LF			
AUTOINCASSOSIZE	Associator automatic database increase size					GM
AUTOINCASSOTHRESHOLD	Associator space threshold size					GM
AUTOINCASSOTOTAL	Associator space maximum size					GM
AUTOINCDATASIZE	Data Storage automatic database increase size					GM
AUTOINCDATATHRESHOLD	Data Storage space threshold size					GM
AUTOINCDATATOTAL	Data Storage space maximum size					GM
CACHE	Load ADACSH (Adabas Caching Facility)		LF			
CACTIVATE	RABN range activation (Adabas Caching Facility)			LM		
CASSODSP	Associator RABNs cached for data space (Adabas Caching Facility)			LM		
CASSOEXT	Associator RABNs cached for extended memory (Adabas Caching Facility)			LM		
CASSOG64	Cache an Associator RABN or RABN range backed by 2G large pages in the virtual 64-bit storage cache.			LM		
CASSOHSP	Associator RABNs cached for hiperspace (Adabas Caching Facility)			LM		
CASSOL64	Cache an Associator RABN or RABN range backed by 1M large pages in the virtual 64-bit storage cache.			LM		
CASSOV64	Associator RABNs cached for virtual 64 (Adabas Caching Facility)			LM		
CASSOMAXS	Associator cache space unit size (Adabas Caching Facility)			LM		
CBUFNO	Read buffer count for concurrent I/O (Adabas Caching Facility)			LM		
CCTIMEOUT	Cache space area inactivity time (Adabas Caching Facility)			LM		
CDATADSP	Data Storage RABNs cached for data space (Adabas Caching Facility)			LM		
CDATAEXT	Data Storage RABNs cached for extended memory (Adabas Caching Facility)			LM		
CDATAG64	Cache a Data Storage RABN or RABN range backed by 2G large pages in the virtual 64-bit storage cache.			LM		
CDATAHSP	Data Storage RABNs cached for hiperspace (Adabas Caching Facility)			LM		

Parameter	Usage	N/A	LF	LM	GF	GM
CDATAL64	Cache a Data Storage RABN or RABN range backed by 1M large pages in the virtual 64-bit storage cache.			LM		
CDATAV64	Data Storage RABNs cached for virtual 64 (Adabas Caching Facility)			LM		
CDATAMAXS	Data Storage cache space unit size (Adabas Caching Facility)			LM		
CDEMAND	Lowest acceptable Adabas buffer efficiency level (Adabas Caching Facility)			LM		
CDISPSTAT	RABN range statistic display location (Adabas Caching Facility)			LM		
CDXnn	Collation descriptor user exit(s)		LF			
CEXCLUDE	Excluded command types (Adabas Caching Facility)			LM		
CFILE	File or file range to be cached (Adabas Caching Facility)			LM		
CLOGBMAX	Maximum size of a log buffer		LF			
CLOGDEV	Multiple command log device		LF			
CLOGLAYOUT	Define command log format				GF	
CLOGMAX	Maximum size of all logged buffers allowed for an Adabas command.		LF			
CLOGMRG	Automatic command log merge control in a cluster environment					GM
CLOGSIZE	Multiple command log size (blocks)		LF			
CLUCACHEEXTRA	The additional vector size, in blocks, required to allow Adabas Cluster Services and Adabas Parallel Services to track the blocks held by Adabas Caching Facility in the different cluster nuclei.		LF			
CLUCACHENAME	Cluster cache structure name (Adabas Cluster Services only)				GF	
CLUCACHESIZE	Cluster cache area size (Adabas Parallel Services only)				GF	
CLUCACHETYPE	Cluster cache area storage type (Adabas Parallel Services only)				GF	
CLUCACHEUNCHANGED	Shared cache unchanged block control (Adabas Parallel Services only)		LF			
CLUGROUPNAME	Cluster group name (Adabas Cluster Services only)				GF	
CLULOCKNAME	Cluster lock structure name (Adabas Cluster Services only)				GF	
CLULOCKSIZE	Cluster lock area size (Adabas Parallel Services only)				GF	

Parameter	Usage	N/A	LF	LM	GF	GM
CLUPUBLPROT	Cluster publishing of protection data					GM
CLUSTER	Adabas cluster session control				GF	
CLUWORK1CACHE	Cluster number of WORK blocks in global cache		LF			
CMADDR	Starting address of the GETMAIN common memory pool above the 16M line		LF			
CMAXCSPS	Storage area count for ADACSH (Adabas Caching Facility)			LM		
CMDQMODE	Command queue memory pool location		LF			
CMFIX	GETMAIN common memory pool fixed location indicator above the 16M line		LF			
CMLADDR	Starting address of the GETMAIN common memory pool below the 16M line		LF			
CMLFIX	GETMAIN common memory pool fixed location indicator below the 16M line		LF			
CMLSCOPE	Control for access to the GETMAIN common memory pool below the 16M line		LF			
CMLSIZE	Indicator for GETMAIN performance in common memory pool below the 16M line		LF			
CMSCOPE	Control for access to the GETMAIN common memory pool above the 16M line		LF			
CMSIZE	Indicator for GETMAIN performance in common memory pool above the 16M line		LF			
CRETRY	Cache space retry interval (Adabas Caching Facility)			LM		
CSTORAGE	RABN caching activation (Adabas Caching Facility)		LF			
CT	Command time limit (seconds)					GM
CWORKSTORAGE	Cache space type (Adabas Caching Facility)			LM		
CWORK2FAC	Percentage of Work part 2 cached (Adabas Caching Facility)			LM		
CWOFK3FAC	Percentage of Work part 3 cached (Adabas Caching Facility)			LM		
CXFILE	Excluded files (Adabas Caching Facility)			LM		
DATACACHE	Controller caching control for the Data Storage component		LF			
DATASPACEWARN	Data Storage space notifications					GM
DBID	Database ID (physical)				GF	
DEVICE	Device type of the first ASSO extent				GF	
DIRRATIO	Ratio of directory entries to data elements in a cluster cache structure/area					GM

Parameter	Usage	N/A	LF	LM	GF	GM
DSF	Delta Save Facility control				GF	
DSFEX1	Delta Save Facility user exit		LF			
DTP ¹	Distributed transaction processing control				GF	
DUALCLD	Dual command log device		LF			
DUALCLS	Dual command log size (blocks)		LF			
DUALPLD	Dual protection log device		LF			
DUALPLS	Dual protection log size (blocks)		LF			
ELEMENTRATIO	Ratio of directory entries to data elements in a cluster cache structure/area					GM
EXCPVR	EXCP or EXCPVR indicator when APF-authorized		LF			
FASTPATH	Adabas Fastpath control				GF	
FMXIO	Limit parallel I/O operations by LFIOP flush processing			LM		
FORCE	Overwrite IDTE		LF			
GROUPS	Indicator for BS2000 interprocess communication limited to users with the same logon ID		LF			
HEX nn	Hyperdescriptor exit(s)		LF			
IDTNAME	Name for alternate ID table		LF			
IDTPSUP	IDT name suppression indicator		LF			
IGNDIB	Ignore DIB entry		LF			
IGNDTP	Ignore distributed transaction processing area (Work part 4)		LF			
INFOBUFFERSIZE	Size of information buffer pool		LF			
INTAUTO	The time interval (in seconds) between autorestart progress messages.		LF			
INTNAS	Interval between nucleus statistic checkpoints (SYNS 60)					GM
LARGEPAGE	Large page use indicator		LF			
LBP	Length of buffer pool		LF			
LCP	Length of security pool		LF			
LDEUQP	Length of unique (UQ) descriptor pool		LF			
LDTP	Length of distributed transaction processing area (Work part 4)		LF			
LFIOP ²	Length of asynchronous flush pool		LF			
LFP	Length of internal format buffer pool		LF			
LI	Length of ISN list table (TBI)		LF			
LNKGNAME	Link globals table name		LF			
LOCAL ³	Nucleus (cluster) unreachable to the network				GF	

Parameter	Usage	N/A	LF	LM	GF	GM
LOGABDX	Log ABDs		LF			
LOGCB	Log control block			LM		
LOGCLEX	Log command log extension (CLEX)		LF			
LOGFB	Log format buffer			LM		
LOGGING	Logging of Adabas commands			LM		
LOGIB	Log ISN buffer			LM		
LOGIO	Log I/O activity			LM		
LOGMB	Log multifetch buffers		LF			
LOGRB	Log record buffer			LM		
LOGSB	Log search buffer			LM		
LOGSIZE	Maximum command log size		LF			
LOGUX	Log user exit B data			LM		
LOGVB	Log value buffer			LM		
LOGVOLIO	Extended I/O list log indicator			LM		
LOGWARN	PLOG/CLOG status check frequency			LM		
LP	Length of data protection area (Work part 1)		LF			
LQ	Length of sequential command table		LF			
LRDP	Length of the redo pool in cluster environments.		LF			
LRPL	Size of Adabas or Event Replicator replication pools		LF			
LS	Length of sort area			LM		
LTZ	The size of the time zone pool used by a user session.		LF			
LU	Length of intermediate user buffer					GM
LWKP2	Length of ISN list processing area (Work part 2)		LF			
LWP	Length of Adabas work pool		LF			
MLWTO	Whether multi-line messages that are normally written to the operator console as a series of single-line write-to-operator (WTO) messages should instead appear as multi-line WTO messages.		LF			
MODE ²	Mode of operation				GF	
MSGBUF	Size of the message buffer		LF			
MSGCONSL	Case of messages on message console		LF			
MSGDRUCK	Case of messages in DD/DRUCK data set		LF			
MSGPRINT	Case of messages sent to DD/PRINT data set		LF			

Parameter	Usage	N/A	LF	LM	GF	GM
MXCANCEL	Timeout threshold for a cancel request between cluster nuclei(Adabas Cluster Services and Adabas Parallel Services)			LM		
MXCANCELWARN	Timeout threshold for a cancel request warning between cluster nuclei (Adabas Cluster Services and Adabas Parallel Services)			LM		
MXMSG	Maximum message reply time between cluster nuclei (Adabas Cluster Services and Adabas Parallel Services)			LM		
MXMSGWARN	Timeout threshold for a message reply warning between cluster nuclei (Adabas Cluster Services and Adabas Parallel Services)			LM		
MXSTATUS	XCF status monitoring heartbeat interval (Adabas Cluster Services)			LM		
MXTNA	Maximum inactivity time limit override for a user					GM
MXTSX	Maximum Sx execution time limit override for a user					GM
MXTT	Maximum transaction time limit override for a user					GM
MXWTOR	Self-termination query operator response interval (Adabas Cluster Services and Adabas Parallel Services)			LM		
NAB	Number of attached buffers		LF			
NC	Number of command queue elements		LF			
NCLOG	Number of command logs		LF			
NH	Number of hold queue elements		LF			
NISNHQ	Number of ISNs in hold queue for user					GM
NONDES	Non-descriptor searches					GM
NPLOG	Number of protection logs		LF			
NPLOGBUFFERS	Number of PLOG protection I/O buffers		LF			
NQCID	Number of active command IDs per user					GM
NSISN	Number of ISNs per ISN table element			LM		
NT	Number of threads		LF			
NU	Number of user queue elements		LF			
NUCID	Cluster nucleus ID		LF			
NWORK1BUFFERS	Number of Work part 1 protection I/O buffers		LF			
OPENRQ	Open command required				GF	
PAMREQS	Maximum number of BS2000 PAM blocks that can be transferred in a single PAM macro request		LF			

Parameter	Usage	N/A	LF	LM	GF	GM
PGFIX	I/O control block page request indicator		LF			
PLOGDEV	Multiple protection log device		LF			
PLOGRQ	Protection log required				GF	
PLOGSIZE	Multiple protection log size (blocks)		LF			
PREFETCH ⁴	Prefetch/multifetch feature control (see note below)	N/A				
PREFICMD ⁴	Include command from prefetch/multifetch (see note below)	N/A				
PREFIFIL ⁴	Include file from prefetch/multifetch (see note below)	N/A				
PREFNREC ⁴	Multifetch record count (see note below)	N/A				
PREFSBL ⁴	Prefetch single buffer length (see note below)	N/A				
PREFSTDD ⁴	Job statement label for location for multifetch statistics	N/A				
PREFTBL ⁴	Prefetch total buffer length (see note below)	N/A				
PREFXCMD ⁴	Exclude command from prefetch/multifetch (see note below)	N/A				
PREFXFIL ⁴	Exclude file from prefetch/multifetch (see note below)	N/A				
PROGRAM	The program to be run		LF			
QBLKSIZE	Sequential data set block size (optimized by ADAIOR)		LF			
READONLY ²	Read-only session control			LM		
REFSTPRT	Whether statistics should be printed after they are refreshed		LF			
REPLICATION	Replication indicator				GF	
REVFILTER	Adabas Review record filtering indicator		LF			
REVIEW	Adabas Review control			LM		
REVLOGBMAX	Maximum number of bytes of a logged buffer for Adabas Review		LF			
REVLOGMAX	Maximum size of all logged buffers for an Adabas Review command		LF			
RIAFTERUPDATE	Control RI command behavior				GF	
RPLCONNECTCOUNT	Number of connection attempts for replication		LF			
RPLCONNECTINTERVAL	Connection attempt interval for replication		LF			
RPLPARMS	Location of replication definitions		LF			
RPLSORT	Replication transaction data sorting indicator				GF	
RPWARNINCREMENT	Replication pool usage warning message interval		LF			
RPWARNINTERVAL	Replication pool usage warning message suppression interval		LF			

Parameter	Usage	N/A	LF	LM	GF	GM
RPWARNMESSAGELIMIT	Replication pool usage warning message limit before suppression		LF			
RPWARNPERCENT	Replication pool usage threshold		LF			
SECUID	Security system user ID requirement level					GM
SMF	Adabas SMF recording control		LF			
SMF89	Type 89 SMF record control		LF			
SMFDETAIL	SMF record detail section control			LM		
SMFINTERVAL	SMF interval record control			LM		
SMFRECNO	SMF user-defined record number control			LM		
SMFSUBSYS	IBM or user-defined SMF subsystem control			LM		
SMGT	Error handling (PIN) facility control		LF			
SORTCACHE	Controller caching control for the Adabas sort area component		LF			
SPT	Adabas triggers and stored procedures control				GF	
SRLOG	Spanned record logging control		LF			
SUBMPSZ	Common memory pool size		LF			
SVC	SVC number		LF			
SWITCHNR	BS2000 job switch number		LF			
TAPEREL	End-of-file processing for tapes or cartridges		LF			
TARGETID	Unique Entire Net-Work target ID		LF			
TASKCTGY	BS2000 Adabas task category		LF			
TCPIP	TCP/IP access control			LM		
TCPURL	TCP/IP universal resource locator (URL)			LM		
TEMPCACHE	Controller caching for the Adabas temp area component		LF			
TFLUSH ²	Synchronous buffer flush time			LM		
TLSCMD	Time limit for S1, S2, and S4 complex searches (seconds)					GM
TMDRQ ³	Number of entries in Adabas Transaction Manager internal request queue		LF			
TMETDATA ³	Databases storing Adabas Transaction Manager ET data		LF			
TMGTT ³	Time limit in which an Adabas Transaction Manager global transaction can be opened without being prepared		LF			
TMLOG ³	Logging option for Adabas Transaction Manager			LM		
TMMSGSEV ³	Severity threshold for suppression of Adabas Transaction manager warning messages		LF			

Parameter	Usage	N/A	LF	LM	GF	GM
TMRESTART ³	Adabas Transaction Manager problematic transaction restart handling		LF			
TMSYNCMGR ³	Indicator for Adabas Transaction Manager interaction with external transaction coordinator		LF			
TMTCIDPREF ³	Adabas Transaction Manager prefix for dynamically allocated client IDs		LF			
TNAA	Non-activity time limit (access-only users)					GM
TNAE	Non-activity time limit (ET logic users)					GM
TNAX	Non-activity time limit (exclusive update users)					GM
TT	Transaction time limit					GM
UEX nn	User exits: 1, 3, 4, 5, 8		LF			
	User exits: 2, 12				GF	
	User exits: 6, 9 (for utilities)	N/A				
UEXSMF	SMF user exit module name		LF			
UTIONLY	Utilities-only session				GF	
V64BIT	Virtual storage above the 2G bar usage indicator		LF			
VISTA	Adabas Vista control				GF	
WORKCACHE	Controller caching for the Adabas work area component		LF			



Notes:

1. Adabas Cluster Services and Adabas Parallel Services 8.2 do not support DTP=TM, but it does support DTP=RM.
2. Adabas Cluster Services and Adabas Parallel Services 8.2 do not support LFIOP=0, MODE=SINGLE, READONLY=YES, or TFLUSH.
3. Adabas Cluster Services and Adabas Parallel Services 8.2 do not support the LOCAL=YES and all TM*parameters
4. The PREF xxx parameters are used with application programs (PROGRAM=USER) making Adabas calls. They have no effect when specified for an Adabas nucleus.

Specifying ADARUN Parameters for Cluster Nuclei

When specifying ADARUN session parameters for Adabas Parallel Services cluster nuclei:

- ensure that the correct program to be executed is specified (PROG=ADANUC); and
- determine which setting is applicable for the SVC parameter for the session.

The CLOGMRG, CLUSTER, CLUCACHESIZE, CLUCACHETYPE, CLUCACHEUNCHANGED, CLULOCKSIZE, DIRRATIO / ELEMENTRATIO, LRDP, and NUCID parameters are used by the Adabas Parallel Services cluster nucleus and its environment.

If protection logs or command logs are used in a cluster environment, they must be dual or multiple logs and all nuclei must use them. All cluster nuclei must have the same PLOGRQ setting.

The remaining Adabas cluster nucleus parameters are the same as those of a standard Adabas nucleus. For more information, read the *Adabas Operations Manual*.

5

Useful ADARUN Parameters for Adabas Parallel Services

▪ CLOGMRG Parameter: Cluster Command Log Merge	28
▪ CLUCACHEEXTRA Parameter: Adjust Shared Cache Vector Size	29
▪ CLUCACHESIZE Parameter: Size of the Global Cache Area	29
▪ CLUCACHETYPE Parameter: Global Cache Area Storage Type Option	31
▪ CLUCACHEUNCHANGED Parameter: Global Cache Unchanged Block Option	33
▪ CLUGROUPNAME Parameter: Cluster Group Name	34
▪ CLULOCKSIZE Parameter: Size of the Global Lock Area	34
▪ CLUPUBLPROT Parameter: Publishing of Protection Data	35
▪ CLUSTER Parameter: Adabas Nucleus Cluster Session Control	36
▪ CLUWORK1CACHE Parameter: Number of WORK Blocks in Global Cache	36
▪ DIRRATIO / ELEMENTRATIO Parameters: Cluster Cache Directory-to-Data Elements Ratio	38
▪ LDEUQP Parameter: Length of Unique Descriptor Pool	39
▪ LRDP Parameter: Length of the Redo Pool	40
▪ MLWTO Parameter: Multi-Line Console Message Control	41
▪ MXCANCEL Parameter: Timeout Threshold for Canceled Peer Nucleus	41
▪ MXCANCELWARN Parameter: Timeout Threshold for Canceled Peer Nucleus Warning	42
▪ MXMSG Parameter: Timeout Threshold for Internucleus Command Processing	43
▪ MXMSGWARN Parameter: Timeout Threshold for Internucleus Command Processing Warning	44
▪ MXWTOR Parameter: Self-Termination Operator Query Interval	45
▪ NUCID Parameter: Adabas Cluster Nucleus ID	46
▪ PROGRAM Parameter: Program to Run	47
▪ SMFDETAIL Parameter: Selecting Detail SMF Record Sections	48
▪ UPDATECONTROL Parameter: Control Scheduling of Update Commands During Buffer Flush	50
▪ V64BIT Parameter: 64-Bit Virtual Storage Option	51

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

CLOGMRG Parameter: Cluster Command Log Merge

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

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

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

Value	Action
YES	enables an automated process that copies dual or multiple command logs to sequential data sets and merges those data sets across a cluster into a single cluster-wide command log. This process automatically merges command logs when an ADARES CLCOPY is submitted from user exit 2/12 and executed. The cluster nuclei and the ADARES utility perform coordinated CLOG switches and CLCOPY functions with merging of CLOG records, much the same way the PLOG merge process works. For the automated CLOG MERGE to synchronize correctly, ADARES must be invoked from the exit.
NO	(the default), each nucleus produces its own CLOG, independent of any other nucleus in the cluster; dual or multiple CLOGs can be manually copied using ADARES CLCOPY to a sequential data set; and the resulting sequential CLOGs can be manually merged across a cluster using the ADARES MERGE CLOG function.



Note: If CLOGMRG=NO is specified, at the end of an ADARES CLCOPY job no additional invocations of the ADARES CLCOPY will occur. The additional invocations only occur in a CLOGMRG=YES environment.

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

Parameter Dependencies

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

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

Dynamic Modification

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

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

CLUCACHEEXTRA Parameter: Adjust Shared Cache Vector Size

Parameter	Specify . . .	Possible Values	Default
CLUCACHEEXTRA	the additional vector size, in blocks, required to allow Adabas Cluster Services and Adabas Parallel Services to track the blocks held by Adabas Caching Facility in the different cluster nuclei.	0 - 2147483647	1000

This parameter is local and unchangeable. It has effect only if both a cluster product (Adabas Cluster Services or Adabas Parallel Services) and Adabas Caching Facility are active.

If Adabas Cluster Services or Adabas Parallel Services are running on the same system as the Adabas Caching Facility, use the CLUCACHEEXTRA parameter to adjust the vector size of the shared cache. This ensures that Adabas Cluster Services and Adabas Parallel Services can track all of the blocks held privately by the Adabas Caching Facility.

Set the value to the maximum number of blocks expected to be cached by the Adabas Caching Facility. This number can be derived from the CASSOMAXS, CDATAMAXS, and CMAXCSPS parameter settings of the Adabas Caching Facility and the ASSO and DATA block sizes of the database.

CLUCACHESIZE Parameter: Size of the Global Cache Area

Parameter	Specify . . .	Minimum	Maximum	Default
CLUCACHESIZE	the amount of storage to allocate for the global cache area that services the Adabas Parallel Services cluster in which the current nucleus participates.	128K	depends on the type of cache (CLUCACHETYPE parameter)	none

The `CLUCACHESIZE` parameter specifies the amount of storage to be allocated for the global cache area which all nuclei in the Adabas Parallel Services cluster use to share Associator and Data blocks from the database. It has an effect only if `CLUSTER=LOCAL` is also specified.

The location and maximum size of the global cache area depends on the `CLUCACHETYPE` parameter. If `CLUCACHETYPE=DSP`, the global cache is kept in a dataspace shared between the Adabas Parallel Services nuclei. Otherwise, the cache is kept in shared 64-bit virtual storage (z/OS systems only).

`CLUCACHESIZE` must be specified for Adabas Parallel Services nuclei; there is no default.

Each Adabas Parallel Services nucleus has its own buffer pool sized by the ADARUN parameter `LBP`. Optionally, it may also have one or more local cache spaces administered by the Adabas Caching Facility (refer to the description of the ADARUN parameter `CACHE`). In addition, Adabas Parallel Services uses a global cache area to ensure that current images of Associator and Data blocks from the database are available to all Adabas Parallel Services cluster nuclei.

The global cache area must be large enough to retain:

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

The ADARUN parameters `DIRRATIO` and `ELEMENTRATIO` determine the ratio between the number of directory and data elements.

`CLUCACHESIZE` may be specified in bytes, in kilobytes followed by a "K", in megabytes followed by an "M", or in gigabytes followed by a "G":

- the minimum size is 131072 bytes (or 128K);
- for `CLUCACHETYPE=DSP`, the maximum size is 2147483147 bytes (or 2097152K; or 2048M; or 2G); and
- the maximum allowed numeral is 2147483147; for `CLUCACHETYPE=V64`, use the K, M or G suffix (for example, 10G or 10240M or 10485760K for a global cache of 10 gigabytes size).

The value is rounded up to the nearest four-kilobyte boundary. If `CLUCACHETYPE=L64`, the value is rounded up to the nearest one-megabyte boundary. If `CLUCACHETYPE=G64`, the value is rounded up to the nearest two-gigabyte boundary.



Note: If "G64" is specified for a cache space, we recommend that you specify a `CLUCACHESIZE` value that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. In this case, specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused. If "G64" is specified and 2G large pages are not available for allocation, the specified size will be rounded up to the nearest multiple of one megabyte, and the storage will not be backed by large pages. A console message will report on this case.

Appropriate sizes for the global cache area depend very much on the application load in the cluster. For more details about these calculations, read *Global Cache Area Size* in the *Adabas Parallel Services Operations Guide*. You can use the Adabas Online System cache structure size calculator described in *Estimating Sizes for the Cache Structure in a Cluster Environment* in the *Adabas Parallel Services Operations Guide* to determine appropriate global cache area sizes.

Example

In the following example, the global cache area allocated is 500 M.

```
ADARUN PROG=ADANUC,CLUCACHESIZE=500M
```

CLUCACHETYPE Parameter: Global Cache Area Storage Type Option

Parameter	Specify . . .	Valid Values	Default
CLUCACHETYPE	the virtual storage type for the global cache area	DSP G64 L64 V64	DSP

The CLUCACHETYPE parameter specifies the type of virtual storage to be used for the global cache of an Adabas Parallel Services cluster. It has an effect only if CLUSTER=LOCAL is also specified. Valid values are "DSP" (available on all supported operating systems), "V64" (available on z/OS and z/VSE systems), "G64" (available only on z/OS systems) and "L64" (available only on z/OS systems).

The size of the global cache is specified in the [CLUCACHESIZE parameter](#).

CLUCACHETYPE Value	Description
DSP	This is the default value. It indicates that a shared dataspace will be used for the global cache. In this case, the maximum cache size is two gigabytes (2G).
G64	<p>If "G64" is specified (on z/OS 2.1 or later systems running on IBM zEC12 machines), the global cache space shared by the cluster nuclei will reside in 64-bit virtual storage that is backed by page-fixed 2G large pages (provided the large page pool has been configured to a sufficient size and is available in the system). If the system cannot satisfy the allocation request with storage backed by 2G large pages, storage backed by pageable 4K pages will be allocated instead.</p> <p>Note: Allocation requests for storage to be backed by 2G large pages will be rounded up to the nearest multiple of two gigabytes. Any difference between the specified size and the allocated size will remain unused. To avoid leaving a potentially large amount of allocated storage unused when specifying CLUCACHETYPE=G64, we recommend that you supply a multiple of two gigabytes for the CLUCACHESIZE parameter. If "G64" is specified and 2G large pages are not available for allocation, the specified size will be rounded up to the nearest multiple of one megabyte, and the storage will not be backed by large pages. A console message will report on this case.</p>

CLUCACHETYPE Value	Description
L64	If "L64" is specified (on z/OS 1.9 or later systems running on IBM z10 machines), the global cache space shared by the cluster nuclei will reside in 64-bit virtual storage that is backed by page-fixed one-megabyte (1M) large pages (provided the large page pool has been configured to a sufficient size and is available in the system). If the system cannot satisfy the allocation request with storage backed by 1M large pages, storage backed by pageable four-kilobyte (4K) pages will be allocated instead.
V64	If "V64" is specified (on z/OS or z/VSE systems), a storage area in 64-bit virtual storage will be used for the global cache. Users in z/OS environments should note that the cache can be made much larger than 2G, but is still subject to the MEMLIMIT value in effect for the ADACOM job for the Adabas Parallel Services database.

**Notes:**

1. If one ADACOM job is used to hold the shared cache and lock spaces for multiple Adabas Parallel Services clusters, L64 and G4 cache spaces of different clusters are slightly less isolated from one another than DSP or V64 cache spaces, creating a marginally higher risk that a cache-related error in one cluster might induce an error in a different cluster. This very small risk can be avoided altogether by using different ADACOM jobs for different Adabas Parallel Services clusters that use L64 or G64 cache spaces, but this is not mandatory.
2. Virtual 64-bit storage backed by page-fixed 1M large pages can only be used on IBM z10 mainframes running z/OS Version 1.9 or above and for which IBM large page support has been enabled via IBM APAR OA20902. You must also allocate the size of the large page pool (use the LFAREA parameter in the IEASYSxx member of SYS1.PARMLIB). The LFAREA parameter allows you to specify the amount of real storage to be used for page-fixed 1M large pages; this parameter cannot be changed dynamically and, if it is not set, page frames will remain allocated at a size of 4K.
3. Virtual 64-bit storage backed by page-fixed 2G large pages can only be used on IBM zEC12 mainframes running z/OS version 2.1 or above or an earlier z/OS version for which 2G large page support has been enabled via IBM APAR. You must also allocate the size of the large page pool (use the LFAREA parameter in the IEASYSxx member of SYS1.PARMLIB). The LFAREA parameter is extended to allow you to specify the amount of real storage to be used for page-fixed 1M and 2G large pages; this parameter cannot be changed dynamically and, if it is not set, page frames will remain allocated at a size of 4K.

Example

In the following example, the global cache area will be allocated in shared 64-bit virtual storage:

```
ADARUN PROG=ADANUC,CLUCACHETYPE=V64
```

CLUCACHEUNCHANGED Parameter: Global Cache Unchanged Block Option

Parameter	Specify . . .	Possible Values	Default
CLUCACHEUNCHANGED	whether unchanged blocks are written to global cache during Adabas Parallel Services processing.	YES NO	NO

The CLUCACHEUNCHANGED parameter specifies whether an Adabas Parallel Services nucleus should write unchanged blocks to the global cache. It has an effect only if CLUSTER=LOCAL is also specified. Valid values are "YES" and "NO".

- The default value is "NO", which indicates that the nucleus will write only changed (updated) Associator and Data blocks from the database to the global cache.
- If "YES" is specified, the nucleus will write changed and unchanged Associator and Data blocks to the global cache.

Recommendation

If a lot of virtual storage can be made available for local and global caching by the Adabas Parallel Services cluster (where local caching is controlled by the ADARUN LBP parameter and possibly the Adabas Caching Facility add-on product), assigning the bulk of the available storage to the global cache (ADARUN CLUCACHESIZE parameter), which is shared by all Adabas Parallel Services nuclei for the database, provides for more economic use of the storage than assigning it to the local caches of the individual cluster nuclei (ADARUN LBP parameter and the Adabas Caching Facility ADARUN CASSOMAXS and CDATAMAXS parameters), which are private to each nucleus. However, each cluster nucleus requires sufficient local buffer pool space (LBP parameter) to hold the entire working set of Associator and Data blocks needed for parallel command processing.

In cases where the global cache size (CLUCACHESIZE parameter) is set to a large value, CLUCACHEUNCHANGED should be set to "YES", so that the global cache achieves a better hit rate (efficiency).

Example

In the following example, the CLUCACHEUNCHANGED parameter is set to that both changed and unchanged Associator and Data blocks are written to the global cache.

```
ADARUN PROG=ADANUC, CLUCACHEUNCHANGED=YES
```

CLUGROUPNAME Parameter: Cluster Group Name

Parameter	Specify . . .	Possible Values	Default
CLUGROUPNAME	the name of the XCF messaging group in Adabas Cluster Services.	text	none
	the name of the messaging group in Adabas Parallel Services.	text	the router name

The CLUGROUPNAME parameter need not be specified for Adabas Parallel Services except when Adabas Parallel Services is used to recover an Adabas Cluster Services database after a failure that resulted in a remnant DIB entry. In this case, the CLUGROUPNAME parameter for Adabas Parallel Services must be set to the same value as it was for the failed Adabas Cluster Services nuclei.

For Adabas Cluster Services, the CLUGROUPNAME parameter provides the XCF group name to be used by the cluster nuclei for their communication using XCF.

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

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

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

Example

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

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

CLULOCKSIZE Parameter: Size of the Global Lock Area

Parameter	Specify . . .	Minimum	Maximum	Default
CLULOCKSIZE	the amount of storage allocated for the global lock area	128K	2G	none

The CLULOCKSIZE parameter specifies the amount of storage to be allocated for the global lock area which all nuclei in the Adabas Parallel Services cluster use to synchronize their processing. It has an effect only if CLUSTER=LOCAL is also specified.

CLULOCKSIZE must be specified for Adabas Parallel Services nuclei; there is no default.

CLULOCKSIZE may be specified in bytes, in kilobytes followed by a "K", in megabytes followed by an "M", or in gigabytes followed by a "G".

- The minimum size is 131072 bytes (or 128K)
- The maximum size is 2147483147 bytes (or 2097152K, 2048M, or 2G).

The value is rounded up to the nearest 4-kilobyte boundary.

Appropriate sizes for the global lock area depend on the expected amount of parallel command processing in the cluster. For more details about these calculations, read *Global Cache Area Size* in the *Adabas Parallel Services Operations Guide*. You can use the Adabas Online System lock structure size calculator described in *Estimating Sizes for the Lock Structure in a Cluster Environment* in the *Adabas Parallel Services Operations Guide* to determine appropriate global lock area sizes.

CLUPUBLPROT Parameter: Publishing of Protection Data

Parameter	Specify . . .	Possible Values	Default
CLUPUBLPROT	whether or not cluster nuclei make their protection data available within the cluster before they make the related updates available to the other nuclei.	YES NO	NO

The optional CLUPUBLPROT parameter specifies whether or not cluster nuclei make their protection data available within the cluster before they make the related updates available to the other nuclei. Valid values are YES and NO. The default value is NO. CLUPUBLPROT is a global, modifiable parameter - that is, all active nuclei in the cluster run with the same setting and this setting can be changed at runtime. When a starting nucleus joins a cluster, it takes over the current setting of the nuclei that are already active.

If CLUPUBLPROT is specified as YES, each nucleus in the cluster writes updated ASSO and DATA blocks to the global cache only after it has first written the protection data for the latest updates to these blocks either to its WORK dataset or to the global cache. With this setting, when one nucleus in the cluster performs a buffer flush, the other nuclei need not be able to participate in the execution of the flush - the nuclei are more independent of one another.

If CLUPUBLPROT is not specified or is specified as NO, the nuclei in the cluster do not write protection data to the global cache. They continue to write protection data to their respective WORK datasets when this is necessary to provide for the recovery of the database after any failure of a nucleus. With this setting, when one nucleus in the cluster performs a buffer flush, all other nuclei must also take a few actions required for the execution of the flush (such as writing out their latest protection data to their WORK datasets).

Setting CLUPUBLPROT=YES is performance-sensitive for update-intensive workloads. The related CLUWORK1CACHE parameter should be configured suitably for all nuclei in the cluster before CLUPUBLPROT is set to YES before or during a nucleus session.

CLUSTER Parameter: Adabas Nucleus Cluster Session Control

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

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

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

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

Example

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

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

CLUWORK1CACHE Parameter: Number of WORK Blocks in Global Cache

Parameter	Specify . . .	Possible Values	Default
CLUWORK1CACHE	the maximum number of WORK blocks that the cluster nucleus should keep in the global cache at a time. If the specified value for CLUWORK1CACHE is larger than NWORK1BUFFERS+1 or LP / 20, the smaller of these two values is taken as the effective CLUWORK1CACHE value.	0 - 32767	0

Parameter	Specify . . .	Possible Values	Default
	If the effective CLUWORK1CACHE value is equivalent to more than 1% of the total size of the data elements in the global cache, it is reduced to 1% of that size, in order to limit the portion of the cache that may be used for protection data.		

The optional `CLUWORK1CACHE` parameter specifies how a cluster nucleus should make its latest protection data available within the cluster, if the `CLUPUBLPROT` parameter is set to `YES`. Valid values are integers in the range 0 through 32,767. The default value is 0. `CLUWORK1CACHE` is a local, unmodifiable parameter — that is, different nuclei in the same cluster may run with different `CLUWORK1CACHE` values and the value for a nucleus cannot be changed while that nucleus is running.

If the specified value for `CLUWORK1CACHE` is larger than `NWORK1BUFFERS+1` or `LP / 20`, the smaller of these two values is taken as the effective `CLUWORK1CACHE` value. If the effective `CLUWORK1CACHE` value is equivalent to more than 1% of the total size of the data elements in the global cache, it is reduced to 1% of that size, in order to limit the portion of the cache that may be used for protection data. Any reductions of the `CLUWORK1CACHE` value to these boundaries are reported in the `ADAX5E` and `ADAX5F` messages.

`CLUWORK1CACHE` specifies the maximum number of `WORK` blocks that the cluster nucleus should keep in the global cache at a time. It also defines the size of an extension of the area on the `WORK` dataset where the nucleus stores protection data. Another nucleus in the cluster may copy protection data to this extension area from the global cache if this nucleus is slow or unable to participate in a buffer flush.

Apart from defining the size of the protection data extension area on `WORK`, the `CLUWORK1CACHE` parameter has an effect only when the related `CLUPUBLPROT` parameter is set to `YES`.

If `CLUWORK1CACHE` is specified as or defaults to zero and `CLUPUBLPROT=YES`, the nucleus always writes protection data directly to its `WORK` dataset before it writes updated `ASSO` and `DATA` blocks to the global cache, if the protection data relates to updates to those blocks. The nucleus does not write protection data to the global cache.

If `CLUWORK1CACHE` is specified greater than zero and `CLUPUBLPROT=YES`, the nucleus writes protection data to the global cache that is not yet present on `WORK` and relates to updates to be written to the cache. At most the number of `WORK` blocks specified by `CLUWORK1CACHE` will be stored in the cache. If further protection data needs to be made available within the cluster before updates can be written to the cache, older protection data is written to `WORK` first to make space for new `WORK` blocks in the cache.

When `CLUPUBLPROT=YES`, the setting of the `CLUWORK1CACHE` parameter is performance-sensitive for update-intensive workloads. Setting `CLUWORK1CACHE=0` may lead to a significant number of additional `WORK` writes. Setting the parameter to a large value may lead to a significant

number of additional cache writes. Any number in-between may lead to some additional WORK writes and some additional cache writes. Trials with different settings may be necessary to find the right balance that is best for the overall performance of the cluster for a given workload.

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

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

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

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

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

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

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

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

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

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

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

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

LDEUQP Parameter: Length of Unique Descriptor Pool

Parameter	Specify . . .	Minimum	Maximum	Default
LDEUQP	the size of the unique descriptor pool.	5000	2147483647	5000

The unique (UQ) descriptor pool comprises elements that are the length of a unique descriptor value plus 14 bytes. The UQ descriptor pool should be large enough to hold all unique descriptor values that are being changed by *all* transactions at any point in time.

Adabas will make an entry in this pool every time a unique descriptor value is deleted from or inserted into a file's index. A full update of a unique descriptor value produces two entries: one for the deletion of the old value and one for the insertion of the new value. The entry for an insertion is normally removed as soon as the index update has been performed. The entry for a deletion is removed when the current transaction ends. However, Adabas Cluster Services and Adabas Parallel Services nuclei always keep entries for insertions until the end of the transaction too.

Specific Product Recommendations

- Consider adjusting the value of this parameter when the SLOG file is defined for use on the Event Replicator Server, causing replication data to be written and read from the SLOG file.
- If the Event Replicator Server is set to support updates by multiple concurrent users to Adabas targets (when the NPADACALLS initialization parameter is set to any value greater than "1"), consider adjusting the value of this parameter in the target Adabas nucleus to ensure the target nucleus can handle updates from multiple concurrent users.

Example

For example, assuming that the average length of unique descriptor values is about 16 bytes, the following Adabas session can process a maximum of about 1000 unique descriptors at one time.

```
ADARUN PROG=ADANUC ,LDEUQP=30000
```

LRDP Parameter: Length of the Redo Pool

Parameter	Specify . . .	Minimum	Maximum	Default
LRDP	the size (in bytes) of the redo pool.	0; or 80000	2147483647	<p>For Adabas Parallel Services 8.2 or later, the default is 0 (zero).</p> <p>For Adabas Cluster Services 8.2 or later, the default is either the value of the ADARUN LFIOP parameter or 10 megabytes (10M), whichever is smaller.</p> <p>For earlier releases of Adabas Parallel Services and Adabas Cluster Services, the default is the value of the ADARUN LFIOP parameter.</p>

The LRDP parameter specifies the size of the redo pool in an Adabas Cluster Services or Adabas Parallel Services nucleus. It has an effect only when CLUSTER=SYSPLEX or LOCAL has also been specified. For Adabas Parallel Services 8.2 or later, its default is zero (0). For Adabas Cluster Services 8.2 or later, its default is either 10 megabytes (10M) or the value of the ADARUN LFIOP parameter, whichever is smaller. For earlier releases of Adabas Parallel Services and Adabas Cluster Services, the default is the value of the ADARUN LFIOP parameter.

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

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

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

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

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

redo pool, review the redo high-water mark in the output of the DRES operator command and in the cluster nucleus shutdown statistics. This information is also available in corresponding Adabas Online System (AOS) displays and in the output of the ADADBS OPERCOM function.

MLWTO Parameter: Multi-Line Console Message Control



Note: At this time, this parameter only affects the appearance of replication-related and cluster-related messages (messages with prefixes ADAF*, ADAX*, PLI*, and PLX*).

Parameter	Specify . . .	Possible Values	Default
MLWTO	Whether multi-line messages that are normally written to the operator console as a series of single-line write-to-operator (WTO) messages should instead appear as multi-line WTO messages.	YES NO	NO

To activate multi-line console messages, set MLWTO to YES in the nucleus startup job. When MLWTO is set to NO (or is allowed to default), no multi-line console messages will appear.

Multi-line console messages are only supported in z/OS and z/VSE environments. Specifying MLWTO=YES in BS2000 environments has no effect on the console messages.

Example:

```
ADARUN PROG=ADANUC,MLWTO=YES
```

MXCANCEL Parameter: Timeout Threshold for Canceled Peer Nucleus

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

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

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

- It returns response code 124 (ADARSP124), subcode 28 (no reply) to the user on whose behalf the intracluster command was issued.

- It terminates itself abnormally if the communication occurred on behalf of an internal process that must not fail.

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



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

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

MXCANCELWARN Parameter: Timeout Threshold for Canceled Peer Nucleus Warning

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

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

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

If the target nucleus does not terminate within the time period set by MXCANCELWARN, the nucleus that issued the cancellation request writes message ADAX9G to the operator console, identifying

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

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

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



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

MXMSG Parameter: Timeout Threshold for Internucleus Command Processing

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

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

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

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

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

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

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

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

MXMSGWARN Parameter: Timeout Threshold for Internucleus Command Processing Warning

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

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

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

If the target nucleus does not respond within the time period set by `MXMSGWARN`, plus any extra time allowed for some types of requests (such as ET synchronization), the nucleus that issued the request writes message `ADAX9C` to the operator console, identifying the target nucleus by its nucleus ID, job name, and system name. This message indicates that the target nucleus may be in danger of being canceled by the nucleus that issued the request if it does not respond within the

timeout period defined by the `MXMSG` parameter. You can use the `ADAX9C` message to trigger an alert or take other appropriate action before the `MXMSG` timeout period expires.

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

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



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

MXWTOR Parameter: Self-Termination Operator Query Interval

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

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

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

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

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

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

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

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

NUCID Parameter: Adabas Cluster Nucleus ID

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

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

Value	Meaning
0	The default setting. This is not an Adabas cluster nucleus.
1-65000	This is an Adabas cluster nucleus. This value must be unique among all target IDs declared to the Adabas router (ADASVC): Adabas databases (DBID), cluster nuclei (NUCID), Natural buffer pools and other MPM servers, Entire Net-Work, and internal targets used by Adabas Cluster Services and Adabas Parallel Services (restricted range 65001-65479). Nucleus IDs for the same cluster need not be specified contiguously.

Guidelines

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

Example

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

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

PROGRAM Parameter: Program to Run

Parameter	Specify:	Possible Values	Default
PROGRAM	the program to be executed.	see table below	USER

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

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

Examples

The following example specifies that an Adabas nucleus is running.

```
ADARUN PROGRAM=ADANUC
```

The following example specifies that an Adabas Review hub is running.

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

The following example specifies that an Entire Net-Work node is running.

SMFDETAIL Parameter: Selecting Detail SMF Record Sections

Parameter	Specify . . .	Valid Values	Default
SMFDETAIL	the type of detail sections in SMF Interval and Termination records (subtypes 2 and 3) that should be included in Adabas SMF records. All specifications (regardless of how many) for this parameter should be enclosed in parentheses.	NONE ALL comma-delimited list of detail section names	NONE

Adabas SMF records can contain a variable set of detail sections in Interval and Termination records (subtypes 2 and 3). This parameter allows you to select the detail section types in Interval and Termination records that should be included in the Adabas SMF records. The value or values you specify for this parameter must be enclosed in parentheses. If you specify NONE or ALL, they should be specified alone. However, you can specify one or more of the other detail section names (CMD, CSHB, CSHF, CSHG, CSHP, FILE, IOOD, LOCK, MSGB, MSGC, MSGH, PARM, STG, THRD, ZIIP or USER) in one SMFDETAIL parameter, separating each value with a comma.



Note: The CSHB, CSHF, CSHG, LOCK, MSGB, MSGC, and MSGH detail section names apply only to cluster environments (Adabas Cluster Services or Adabas Parallel Services must be installed). In addition, the CSHP detail section name only applies in cluster environments with Adabas Parallel Services installed.

The following table describes the meaning of the possible detail section names that can be used in the SMFDETAIL parameter:

Detail Section Name	Description
ALL	Generate all possible detail sections in the Adabas SMF records. If this value is specified, no others can be specified on the SMFDETAIL parameter.
CMD	Generate Adabas command activity detail sections in the Adabas SMF records.
CSHB ¹	Generate global cache activity by block detail sections in the Adabas SMF records.
CSHF ¹	Generate global cache activity by Adabas file number detail sections in the Adabas SMF records.
CSHG ¹	Generate global cache activity detail sections in the Adabas SMF records.
CSHP ²	Generate Adabas Parallel Services cache activity detail sections in the Adabas SMF records.
FILE	Generate Adabas file activity detail sections in the Adabas SMF records.
IOOD	Generate I/O activity by DD statement detail sections in the Adabas SMF records.
LOCK ¹	Generate global lock activity detail sections in the Adabas SMF records.

Detail Section Name	Description
MSGB ¹	Generate internucleus messaging control block activity detail sections in the Adabas SMF records.
MSGC ¹	Generate internucleus messaging count detail sections in the Adabas SMF records.
MSGH ¹	Generate internucleus messaging service time histogram detail sections in the Adabas SMF records.
NONE	Generate no detail sections in the Adabas SMF records. If this value is specified, no others can be specified on the SMFDETAIL parameter.
PARAM	Generate ADARUN parameter value detail sections in the Adabas SMF records.
SESS	Generate Adabas session statistics sections in the Adabas SMF records.
STG	Generate Adabas storage pool detail sections in the Adabas SMF records.
THRD	Generate thread activity detail sections in the Adabas SMF records.
ZIIP	Generate zIIP statistics detail sections in the Adabas SMF records.
USER	Generate user-defined details sections in the Adabas SMF records. If USER is specified, a value for the UEXSMF parameter should also be specified to identify the user exit to be invoked to generate the user-defined detail section.
¹ Available only in cluster environments (Adabas Cluster Services or Adabas Parallel Services must be installed).	
² Available only in cluster environments with Adabas Parallel Services installed.	

This parameter can be altered by operator commands while the nucleus is running.

Example

In the following example, detail sections for Adabas command, file, and thread activity are generated and written in the Adabas SMF records.

```
ADARUN PROG=ADANUC,SMFDETAIL=(CMD,FILE,THRD)
```

In the following example, all detail sections are generated and written in the Adabas SMF records.

```
ADARUN PROG=ADANUC,SMFDETAIL=(ALL)
```

UPDATECONTROL Parameter: Control Scheduling of Update Commands During Buffer Flush

Parameter	Specify . . .	Possible Values	Default
UPDATECONTROL	whether or not to shortly delay the scheduling of update commands at the end of buffer flushes.	DELAY NODELAY	DELAY

The `UPDATECONTROL` parameter controls how the nucleus schedules update commands at the end of a buffer flush. Its setting influences the techniques that Adabas uses to recover the database in a session autorestart, in the case that this nucleus session terminates abnormally.

When `UPDATECONTROL` is set to `DELAY`, at the end of every buffer flush the nucleus delays the start of new update commands until all active update commands have finished. In a cluster, this happens in all nuclei in the cluster at the same time. Normal update command scheduling resumes once all previously active update commands have ended. In the case that the nucleus terminates abnormally, the following session autorestart recovers the database by making use of the presence of a point in time before the failure where no update processing was in progress.

When `UPDATECONTROL` is set to `NODELAY`, the nucleus does not delay the start of new update commands at the end of buffer flushes. The buffer flushes do not impact normal update command scheduling. In the case that the nucleus terminates abnormally, the following session autorestart recovers the database even if update commands were in progress continually before the failure. Omitting the suspension of update commands after buffer flushes avoids the related short delays in update processing and, in a cluster, eliminates a category of periodic processes where all nuclei in the cluster must collaborate.

The default value of `UPDATECONTROL` is `DELAY`. The `UPDATECONTROL` setting cannot be changed at runtime. In a cluster, `UPDATECONTROL` must be set to the same value in all nuclei in the cluster. If the `UPDATECONTROL` setting is to be changed, all nuclei in the cluster must be shut down before the first one can be restarted with the new setting.

If `UPDATECONTROL` is set to `NODELAY`, the `INDEXUPDATE` parameter must be set to `ADVANCED`.



Note: If an Adabas nucleus running with `UPDATECONTROL=NODELAY` terminates abnormally, the `UPDATECONTROL` setting in the restarted nucleus does not matter for the autorestart.

V64BIT Parameter: 64-Bit Virtual Storage Option



Note: This parameter is applicable only to z/OS and z/VSE systems.

Parameter	Specify . . .	Possible Values	Default
V64bit	whether or not the Adabas nucleus should use virtual storage above the two gigabyte bar (64-bit virtual storage). In order for the nucleus to use 64-bit virtual storage, the operating system must also support 64-bit virtual storage.	YES NO	NO

This parameter indicates whether or not the Adabas nucleus should use 64-bit virtual storage. The default is NO.

Currently, the following storage areas will be allocated above the two gigabyte bar in an Adabas nucleus:

- The flush I/O pool (see ADARUN parameter LFIOP)
- The PLOG protection I/O buffers (see ADARUN parameter NPLOGBUFFERS)
- The Work part 1 protection I/O buffers (see ADARUN parameter NWORK1BUFFERS)
- The cast-out directory buffer, in cluster environments (automatically derived from the attributes of the global cache space for Adabas Cluster or Adabas Parallel Services).

The use of 64-bit virtual storage will help alleviate virtual storage constraints below the two gigabyte bar in installations that use large amounts of virtual storage in Adabas nucleus sessions (for example, nucleus sessions with ADARUN parameter settings of LBP=1200M and LFIOP=200M).

Value Meaning

YES The Adabas nucleus will use 64-bit virtual storage (storage above the two gigabyte bar).

NO The default setting. The Adabas nucleus will not use 64-bit virtual storage.

Users in z/OS environments should note that virtual storage above the two gigabyte bar (64-bit virtual storage) is shown in dumps to SVCDUMP or SYSMDUMP but not in dumps to MPMDUMP, SYSABEND, or SYSUDUMP. Therefore, we recommend that you use the SVCDUMP in the JCL of an Adabas nucleus running with ADARUN V64BIT=YES to ensure that 64-bit virtual storage is shown in any SVC dumps that are produced. If, for some reason, you do not want to produce SVC dumps, we recommend that you use SYSMDUMP to ensure 64-bit virtual storage is shown in a dump of the Adabas nucleus address space.



Note: The replication pool will be allocated in 64-bit virtual storage when an Adabas nucleus runs with ADARUN parameters REPLICATION=YES and V64BIT=YES.

Example

The following example indicates that the Adabas nucleus will make use of virtual storage above the two-gigabyte bar and should use 1M large pages.

```
ADARUN PROG=ADANUC,LARGEPAGE=YES,V64BIT=YES
```

6 Cluster Operator Commands

- ADACOM Operator Commands 54
- Adabas Cluster Nucleus Operator Commands 64

This chapter describes operator commands for ADACOM and for an Adabas Parallel Services cluster nucleus.

ADACOM Operator Commands

Special ADACOM operator commands exist to display and control the multiprocessing environment. These commands, which are similar to regular Adabas operator commands, are issued to the local ADACOM initialization job.

This section is organized in the following topics:

- [Entering z/OS System Commands](#)
- [Entering z/VSE System Commands](#)

Entering z/OS System Commands

This section describes the format for entering ADACOM operator commands on z/OS systems. It contains the following topics:

- [ADACOM Environment](#)
- [ADACOM Command Syntax and Processing](#)
- [Controlling Subtasks](#)
- [Controlling PLXCBs](#)
- [DIM - Display Image](#)
- [DN - Display Active Nuclei](#)
- [DUMP - Dump Storage Areas](#)
- [SN - Set Nucleus Status](#)

ADACOM Environment

The main ADACOM task controls the initial processing and routing of operator commands. It also attaches and detaches a pair of subtasks (load modules ADACOT and ADADSP) for each specific SVC/DBID combination. After initialization ADACOT processes SVC/DBID-specific operator commands. ADADSP allocates, owns and deallocates dataspace.

The SVC/DBID-specific subtasks may be attached as specified in ADACOM Initialization Parameters DDKARTE statements or by operator commands after ADACOM has initialized.



Note: ADACOM writes global ADACOM messages into the output data set with the DD name COMPRINT. It writes messages pertinent to an individual SVC/DBID combination into the output data set with the DD name 'P*ss*dddd', where *ss* is the last two digits of the SVC number and *dddd* is the database ID. Similarly, dataspace-related messages are written to DD name 'D*ss*dddd'

ADACOM Command Syntax and Processing

When running the Adabas Cluster and Parallel Services routine ADACOM, the operator commands described in this section can be specified in following input formats:

```

MODIFY  jobname,command
F      jobname,command

```

where

jobname	name of the ADACOM job or started task
command	one of the operator commands described in this section

Every operator command is first processed by the ADACOM main task and is echoed with the message

```
PLI060 SVC=svc DBID=dbid Operator Command:xxx
```

Certain SVC/DBID-specific commands must identify which SVC/DBID is being addressed:

```

MODIFY  jobname,SVC=svc,DBID=dbid,command
F      jobname,SVC=svc,DBID=dbid,command

```

After initialization, any command issued is directed to the last SVC/DBID pair encountered in the input. Whenever a new default SVC/DBID is established by an operator command this message will appear:

```
PLI052 Commands will go to SVC=svc DBID=dbid
```

Commands to be serviced by the ADACOT SVC/DBID-specific subtask show this message:

```
PLI062 Command queued
```

Message PLI062 is issued by the ADACOM main task. The subtask will respond to the command.

If you enter an operator command for an SVC/DBID and a previous command to the same SVC/DBID has not yet completed, you will receive message PLI062 but only one of the commands will be processed. If the first command has not yet been accepted by the subtask the second command will take precedence and the first command will be ignored. If processing has begun on the first command the second command will be ignored.

A command can be given only if the SVC/DBID subtasks are already active (has been specified before, and not terminated). Otherwise, specify the SVC/DBID pair without a command to activate the SVC/DBID (as shown above), and then issue the command separately (as shown below). If, when starting an SVC/DBID pair, a PLXCB is not found one will be allocated using default values. Specific values may also be given:

```
MODIFY jobname,SVC=svc,DBID=dbid [,CMDMGR=NO][,NU= maxusers] [,FORCE=YES] [,LOC=DSP] ,command
F      jobname,SVC=svc,DBID=dbid [,CMDMGR=NO][,NU= maxusers] [,FORCE=YES] [,LOC=DSP] ,command
```

When you change or add a SVC/DBID, the new SVC/DBID becomes the default for all commands issued until the SVC/DBID is changed.

Controlling Subtasks

The ADACOM subtasks used to allocate and monitor the PLXCB and subordinate data structures terminate after processing if CMDMGR=NO has been specified for that SVC/DBID pair. If CMDMGR=YES (the default setting) has been specified for the SVC/DBID pair, the ADACOM subtasks stay active and accepts operator commands for displaying and controlling the multiprocessing environment.

Any SVC/DBID pair specifying the PLXCB to be placed in a dataspace requires ADACOM to be running with active SVC/DBID subtasks.

Parallel Services nuclei (ASM) require ADACOM to be running with active SVC/DBID subtasks in order to use dataspace for Cache and Lock.

SVC/DBID subtasks may be terminated by sending an ADAEND command to the specific SVC/DBID:

```
MODIFY jobname,SVC=svc,DBID=dbid,command,ADAEND
F      jobname,SVC=svc,DBID=dbid,command,ADAEND
```

 **Caution:** You must specify the SVC and DBID when terminating an SVC/DBID combination. Otherwise, the ADAEND command terminates the entire ADACOM job. The currently set default does not apply when using ADAEND.

The command will be rejected if there are active nuclei using the PLXCB.

When the subtasks terminate all dataspace are deleted. The PLXCB remains unless it was allocated in a dataspace.

Controlling PLXCBs

A PLXCB and associated structures (PLXNuc, PLXMap, PLXUser) must exist for each SVC/DBID combination with active nuclei. These are created as specified by DDKARTE initialization parameters or operator commands as described above.

PLXCB parameters such as NU and LOC cannot be changed after the PLXCB is allocated. Any existing PLXCB must be deleted and a new one allocated. This is done automatically when ADACOM is started if the change is initiated by DDKARTE input statements as long as there are no active nuclei using the PLXCB.

When the subtasks aren't running and no nuclei are using an existing PLXCB, an operator command directed to the SVC/DBID specifying or defaulting NU, LOC, FORCE or CMDMGR parameters will delete any existing PLXCB and allocate a new one using the specified parameters.

You can delete an existing PLXCB without allocating a new one as long as no nuclei are using the existing PLXCB by specifying NU=0 on either a DDKARTE statement or an operator command. Using the operator command requires the subtasks not be active.

The FORCE=YES operand will override the requirement that no active nuclei are using the existing PLXCB, as may be the case after certain abnormal terminations. Note that if one or more active nuclei are using the PLXCB they will ABEND after the existing PLXCB is deleted.

LOC=DSP requires the SVC/DBID subtasks be active.

DIM - Display Image

```
DIM [ image-name ]
```

This command is serviced by an SVC/DBID subtask.

The DIM command displays information about each active cluster nucleus on the specified (or all) operating system image. Since with Adabas Parallel Services (in contrast to Adabas Cluster Services) all cluster nuclei run on the same system image, its output for DIM is equivalent to that for the DN command. See [DN - Display Nuclei](#) for sample output.

Specifying the image name is optional.

DN - Display Active Nuclei

This command is serviced by an SVC/DBID subtask.

Use DN to display the number of commands processed and the number of currently active users for each active nucleus.

Sample Output

```
PLI060 SVC=svc DBID=dbid OPERATOR COMMAND:DN
PLI004 image-name NUCID UP LO RO -#USERS- -#CMNDS- LURA= RULA=
      jobname      00001 Y Y N 00000152 00000001
PLI004 image-name NUCID UP LO RO -#USERS- -#CMNDS- LURA= RULA=
      jobname      00002 Y Y N 00000089 00000000
```

The display uses the following indicators:

<i>image-name</i>	the name of the image
<i>jobname</i>	the job name of a cluster nucleus active on the local image
NUCID	unique cluster nucleus identifier between 1 and 65000
UP	whether (Y or N) the specified nucleus is available for normal processing
LO	whether the specified nucleus is open (Y); or closed (N) for new users
RO	not applicable to Parallel Services
#USERS	the number of users currently assigned to the specified nucleus
#CMNDS	the number of commands currently in progress in the specified nucleus
LURA=	not applicable to Parallel Services
RULA=	not applicable to Parallel Services

DUMP - Dump Storage Areas

This command is serviced by an SVC/DBID subtask.

The DUMP command displays internal information about various storage areas that Adabas Cluster Services and Adabas Parallel Services use for controlling the communication with the cluster nuclei. Use this command as directed by Software AG Support. Its output might be helpful in the diagnosis of communication-related problems with cluster databases.

```
DUMP stg-area
```

where *stg-area* can be any of the following storage areas: CLUDSP, IDTE, IDTH, IDTHPREFIX, PLXCB, PLXMAP, PLXNUC, PLXUSER, or SVC.

SN - Set Nucleus Status

This command is serviced by an SVC/DBID subtask.

```
SN {RMTALL | image-name | nucleus-id} {OP | CL}
SN {LCLALL | nucleus-id} {OP | CL} {LCL | GBL}
```

where

RMTALL	is not applicable
OP	is OPEN
CL	is CLOSE
LCLALL	is LOCAL ALL: all nuclei on the local image
LCL	is LOCAL: local users
GBL	is GLOBAL: all Adabas Parallel Services cluster users

The possible options settings for the SN command are described as follows:

Option	Action
RMTALL {OP CL}	not applicable
<i>image-name</i> {OP CL}	not applicable
<i>nucleus-id</i> {OP CL}	not applicable
LCLALL {OP CL} LCL	open or close all nuclei on the <i>local</i> image to local users. Information is not broadcast to other images.

Option	Action
nucleus-id {OP CL} LCL	open or close the specified <i>local</i> nucleus to local users. Information is not broadcast to other images.
LCLALL {OP CL} GBL	open or close all nuclei on the <i>local</i> image to all Adabas Parallel Services cluster users.
nucleus-id {OP CL} GBL	open or close the specified <i>local</i> nucleus to all Adabas Parallel Services cluster users.

By default, nuclei start open to users.

After the nuclei start, the SN operator commands may be used to exercise some control over the assignment of users to nuclei.

Entering z/VSE System Commands

➤ To enter operator commands in z/VSE environments, use the following steps:

- 1 Enter an MSG command for the z/VSE partition in which ADACOM is executing, using the following format:

```
MSG partition-id
```

When ready for communication, ADACOM will respond with message ADAI29. In cases where the ADAEND command causes ADACOM to end its session, no outstanding reply is presented so that an orderly shutdown can occur without the need for operator intervention.

- 2 Enter the desired operator command or commands. More than one command can be entered in a session.

You can also enter a command followed directly by a slash (/) to end any further operator communications until the next MSG command by z/VSE.

- 3 When all desired commands have been entered, close the operator communication session by entering a null command (EOB).
- 4 To enter any more commands after the session is closed, repeat this entire procedure.

Every operator command is directed to the ADACOM job and is echoed with the message:

```
PLI060 SVC=svc DBID=dbid OPERATOR COMMAND:xxx
```

- DIM - Display Image
- DN - Display Active Nuclei
- DUMP - Dump Storage Areas

- [SN - Set Nucleus Status](#)

DIM - Display Image

```
DIM [ image-name ]
```

The DIM command displays information about each active cluster nucleus on the specified (or all) operating system image. Since with Adabas Parallel Services (in contrast to Adabas Cluster Services) all cluster nuclei run on the same system image, its output for DIM is equivalent to that for the DN command. See [DN - Display Nuclei](#) for sample output.

Specifying the image name is optional.

DN - Display Active Nuclei

Use DN to display the number of commands processed and the number of currently active users for each active nucleus.

Sample Output

```
PLI060 SVC=svc DBID=dbid OPERATOR COMMAND:DN
PLI004 image-name NUCID UP LO RO -#USERS- -#CMNDS- LURA= RULA=
      jobname      00001 Y Y N 00000152 00000001
PLI004 image-name NUCID UP LO RO -#USERS- -#CMNDS- LURA= RULA=
      jobname      00002 Y Y N 00000089 00000000
```

The display uses the following indicators:

<i>image-name</i>	the name of the image
<i>jobname</i>	the job name of a cluster nucleus active on the local image
NUCID	unique cluster nucleus identifier between 1 and 65000
UP	whether (Y or N) the specified nucleus is available for normal processing
LO	whether the specified nucleus is open (Y); or closed (N) for new users
RO	not applicable to Parallel Services
#USERS	the number of users currently assigned to the specified nucleus
#CMNDS	the number of commands currently in progress in the specified nucleus
LURA=	not applicable to Parallel Services
RULA=	not applicable to Parallel Services

DUMP - Dump Storage Areas

The DUMP command displays internal information about various storage areas that Adabas Cluster Services and Adabas Parallel Services use for controlling the communication with the cluster nuclei. Use this command as directed by Software AG Support. Its output might be helpful in the diagnosis of communication-related problems with cluster databases.

```
DUMP stg-area
```

where *stg-area* can be any of the following storage areas: CLUDSP, IDTE, IDTH, IDTHPRFX, PLXCB, PLXMAP, PLXNUC, PLXUSER, or SVC.

SN - Set Nucleus Status

```
SN {RMTALL | image-name | nucleus-id } {OP | CL }
SN {LCLALL | nucleus-id } } {OP | CL } {LCL | GBL }
```

where

RMTALL	is not applicable
OP	is OPEN
CL	is CLOSE
LCLALL	is LOCAL ALL: all nuclei on the local image
LCL	is LOCAL: local users
GBL	is GLOBAL: all Adabas Parallel Services cluster users

The possible options settings for the SN command are described as follows:

Option	Action
RMTALL {OP CL}	not applicable
<i>image-name</i> {OP CL}	not applicable
<i>nucleus-id</i> {OP CL}	not applicable
LCLALL {OP CL} LCL	open or close all nuclei on the <i>local</i> image to local users. Information is not broadcast to other images.
<i>nucleus-id</i> {OP CL} LCL	open or close the specified <i>local</i> nucleus to local users. Information is not broadcast to other images.
LCLALL {OP CL} GBL	open or close all nuclei on the <i>local</i> image to all Adabas Parallel Services cluster users.
<i>nucleus-id</i> {OP CL} GBL	open or close the specified <i>local</i> nucleus to all Adabas Parallel Services cluster users.

By default, nuclei start open to users.

After the nuclei start, the SN operator commands may be used to exercise some control over the assignment of users to nuclei.

Adabas Cluster Nucleus Operator Commands

The following console operator commands can be issued against a local Adabas cluster nucleus. In addition, the console operator commands documented in the *Adabas Operations* documentation can be issued against any local nucleus.

Command	To ...
CLUFREEUSER	Delete leftover cluster user table elements (UTEs or PLXUSERS) in common storage that are no longer associated with user queue elements (UQEs) in a nucleus
CLUPUBLPROT	Dynamically modify the setting of the ADARUN CLUPUBLPROT parameter
DMEMTB	Display information about active Adabas cluster nuclei
DNFV	Display information about current file use
DPARM	Display Adabas session parameters currently in effect
DPPT	Display information about the nucleus's own PPT block
DRES	Display the allocated pool space and the highest use level (high water mark) during current session
DTARG	Display the system image currently servicing commands routed using the Entire Net-Work DBID target.
DXCACHE	Display the primary cache-related statistics
DXFILE	Display the cache-related statistics for 1 to 5 files
DXLOCK	Display the lock-related statistics
DXMSG	Display the Adabas Parallel Services messaging performance statistics
DXSTAT	Display all cache- and lock-related statistics
DXWORK	Display the I/O statistics for WORK datasets of peer nuclei in the cluster
MXCANCEL	Dynamically change the ADARUN MXCANCEL setting
MXCANCELWARN	Dynamically change the ADARUN MXCANCELWARN setting
MXMSG	Dynamically change the ADARUN MXMSG setting
MXMSGWARN	Dynamically change the ADARUN MXMSGWARN setting
MXWTOR	Dynamically change the ADARUN MXWTOR setting
SMFDETAIL	Dynamically select the detail section types in SMF Interval and Termination records
TARGET	Route commands to a specified target.



Note: See the *Adabas Parallel Services Operations* documentation for information about issuing Adabas utility ADADBS OPERCOM commands against the local cluster nucleus, a specified cluster nucleus, or all cluster nuclei.

CLUFREEUSER Command: Delete Leftover Cluster User Table Elements

```
CLUFREEUSER=[,TNA= max-time ][,{UID | UIDPRFX} = userid ][,FORCE][,GLOBAL]
```

 **Note:** 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.

Use the CLUFREEUSER command to delete leftover cluster user table elements (UTES or PLXUSERS) in common storage that are no longer associated with user queue elements (UQEs) in a nucleus.

TNA TNA specifies a decimal number specifying the timeout value in seconds (*max-time*). PLXUSERS that are not used during the time specified may be deleted if other conditions are fulfilled. If TNA= is not specified, PLXUSERS may be deleted without regard to their recent use.

UID UID specifies a character string or hexadecimal byte string format user ID (*userid*) as follows:

<i>cccccccc</i>	where the argument is 1-8 letters, digits, or embedded '-' signs without surrounding apostrophes.
' <i>cccccccc</i> '	where the argument is 1-8 characters with surrounding apostrophes.
X' <i>xxxxxxxxxxxxxxxxxx</i> '	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 eight 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 eight bytes of a user's 28-byte communication ID match a specific user ID or user ID prefix (UIDPRFX parameter), that user's PLXUSER may be deleted if other conditions are fulfilled. If UID is not specified, PLXUSERS may be deleted regardless of their user IDs. "UID" may be abbreviated to "U".

UID and UIDPRFX are mutually exclusive; one or the other, but not both, can be specified.

UIDPRFX UIDPRFX accepts the same operands as UID. It operates in the same manner as UID, except that the operand is considered a prefix and there is no implicit padding. If the operand matches the initial bytes of the last eight bytes of the 28-byte communication ID, that PLXUSER may be deleted if other conditions are fulfilled. "UIDPRFX" may be abbreviated to "UP".

UID and UIDPRFX are mutually exclusive; one or the other, but not both, can be specified.

FORCE FORCE indicates that leftover PLXUSERS are to be deleted even if the users are due a response code 9 (ADARSP009), subcode 20. If FORCE is not specified, such PLXUSERS are not deleted. Before using the FORCE parameter, ensure that the users owning the PLXUSERS to be deleted will not expect any of their transactions to remain open.

GLOBAL GLOBAL indicates that leftover PLXUSERS throughout the Adabas cluster are to be deleted if they are no longer associated with UQEs and are eligible according to the other specified

parameters. Additionally and subject to the other rules, leftover PLXUSERS are deleted if their assigned nuclei have terminated since their last use. If GLOBAL is not specified, only PLXUSERS assigned to the local nucleus and used since the nucleus start are eligible for deletion.

CLUPUBLPROT Command: Dynamically Change CLUPUBLPROT Parameter Setting

```
CLUPUBLPROT={YES | NO}
```

Use the CLUPUBLPROT command to dynamically modify the setting of the ADARUN CLUPUBLPROT parameter. For more information, read CLUPUBLPROT Parameter: Publishing of Protection Data.

The CLUPUBLPROT command is only valid in cluster environments. It is global by definition and affects all nuclei in the cluster.

DMEMTB Command: Display Member State Table



The DMEMTB command can be issued against Adabas cluster nuclei to display the member state table, which provides information about all of the active nuclei in an Adabas cluster. The information produced by this command may be used by Software AG technical support.

The output produced by the DMEMTB command includes a third flag that indicates whether a system- or message-level XCF status monitoring exception has been encountered and whether a message was issued for the exception.

Sample Output

```
ADAX61 00226 2009-06-24 16:53:26 Member Status Table
ADAX61 00226 2009-06-24 16:53:26 Other members:      1
ADAX61 00226 2009-06-24 16:53:26   This system:      1
ADAX61 00226 2009-06-24 16:53:26
ADAX61 00226 2009-06-24 16:53:26           NUCID:      2,261
ADAX61 00226 2009-06-24 16:53:26   Internal ID:      2
ADAX61 00226 2009-06-24 16:53:26           System: DA2F
ADAX61 00226 2009-06-24 16:53:26           Job Name: USACHPD1
ADAX61 00226 2009-06-24 16:53:26   Flags 1/2/3: 11/00/00
ADAX61 00226 2009-06-24 16:53:26
ADAX61 00226 2009-06-24 16:53:26           NUCID:      2,262
ADAX61 00226 2009-06-24 16:53:26   Internal ID:      3
ADAX61 00226 2009-06-24 16:53:26           System: DA2F
ADAX61 00226 2009-06-24 16:53:26           Job Name: USACHPD2
ADAX61 00226 2009-06-24 16:53:26   Flags 1/2/3: D7/00/00
ADAN41 00226 2009-06-24 16:53:26 Function completed
```

Flags 1, 2, and 3 in this report are hexadecimal values that provide the following information about each nucleus:

Flag	Provides
1	<p>The cluster member service summary, as a combination of one or more of the following values:</p> <ul style="list-style-type: none"> ■ X'80' – All services are active ■ X'40' – Work data set is needed for recovery ■ X'10' – Accepting and processing intracluster commands ■ X'04' – Cache services are available ■ X'02' – Lock services are available ■ X'01' – Messaging services are available

Flag	Provides
	For example, a value of 97 for Flag 1 indicates that all services are active (X'80'), the nucleus is accepting and processing intracenter commands (X'10'), and cache (X'04'), lock (X'02') and messaging (X'01') services are all available (80+10+4+2+1=97).
2	<p>The cluster member failure status, as a combination of one or more of the following values:</p> <ul style="list-style-type: none"> ■ X'80' – All services are down ■ X'40' – Canceled by this nucleus ■ X'10' – Abnormal termination of intracenter command processing occurred ■ X'04' – Abnormal exit from cache occurred ■ X'02' – Abnormal exit from lock occurred ■ X'01' – Abnormal exit from the XCF group (messaging) occurred
3	<p>The cluster XCF status monitoring summary, as a combination of one or more of the following values:</p> <ul style="list-style-type: none"> ■ X'80' – A system-level XCF status monitoring exception occurred. ■ X'40' – A member-level XCF status monitoring exception occurred. ■ X'02' – A warning message was issued for the system-level XCF status monitoring exception. ■ X'01' – A warning message was issued for the member-level XCF status monitoring exception. <p>For example, a value of 41 for Flag 3 indicates that a member-level XCF status monitoring exception occurred (X'40') and that a warning message (X'01') was issued for it (40+1=41).</p>

DNFV Command: Display Nucleus File Variables



Use the DNFV command to display information about current file use.

This command provides information about the files in use at a particular point in time. It also indicates which other nucleus has exclusive file control if, for example, a user program receives a response 148 (ADARSP148), subcode 15.

Here is some sample output:

```
ADAI29 OPER CMD: DNFV
FNR=00008  A=Y  U=Y  ID=          CA=00000  CU=00001
```

where;

FNR= <i>nnnnn</i>	is the file number
A={Y N}	(yes or no) indicates whether the file is used for access (read and/or search)
U={Y N}	(yes or no) indicates whether the file is used for update. Use for update includes use for access.
ID= <i>nucid</i>	is the ID of the nucleus that owns the file lock, if the file is locked.
CA= <i>nnnnn</i>	is the number of users on this nucleus who are currently accessing this file.
CU= <i>nnnnn</i>	is the number of users on this nucleus who are currently updating this file.

DPARM Command: Display Adabas Session Parameters



Use the `DPARM` command to display the Adabas session parameters currently in effect. Here is an example:

```
17:00:05 ADAI29 Oper cmd: DPARM
17:00:05 ADAN16 00035 2016-11-29 17:00:03 AOSLOG=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 ARNWORKBUFFERS=5
17:00:05 ADAN16 00035 2016-11-29 17:00:03 ASSOSPACEWARN=0
17:00:05 ADAN16 00035 2016-11-29 17:00:03 ASYTVS=YES
17:00:05 ADAN16 00035 2016-11-29 17:00:03 AUTOCQENV=ALL
17:00:05 ADAN16 00035 2016-11-29 17:00:03 AUTOCQTIME=0
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLOGMRG=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLUCACHESIZE=2,097,152
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLUCACHETYPE=DSP
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLUCACHEUNCHANGED=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLULOCKSIZE=5,242,880
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLUPUBLPROT=YES
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLUSTER=LOCAL
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CLUWORK1CACHE=5
17:00:05 ADAN16 00035 2016-11-29 17:00:03 CT=900
17:00:05 ADAN16 00035 2016-11-29 17:00:03 DATASPACEWARN=0
17:00:05 ADAN16 00035 2016-11-29 17:00:03 DIRRATIO=4
17:00:05 ADAN16 00035 2016-11-29 17:00:03 ELEMENTRATIO=1
17:00:05 ADAN16 00035 2016-11-29 17:00:03 FMXIO=1
17:00:05 ADAN16 00035 2016-11-29 17:00:03 INDEXCROSSCHECK=YES
17:00:05 ADAN16 00035 2016-11-29 17:00:03 INDEXUPDATE=ADVANCED
17:00:05 ADAN16 00035 2016-11-29 17:00:03 INTNAS=3,600
```

Cluster Operator Commands

```
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LARGEPAGE=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LBP=104,858,624
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LCP=10,000
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LDEUQP=7,000
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LFIOP=90,000
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LFP=100,000
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LI=4,000
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGCB=YES
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGFB=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGGING=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGIB=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGIO=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGMB=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGRB=NO
17:00:05 ADAN16 00035 2016-11-29 17:00:03 LOGSB=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LOGUX=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LOGVB=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LOGVOLIO=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LOGWARN=0
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LP=1,200
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LQ=30,000
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LRDP=10,485,760
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LS=19,968
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LU=655,350
17:00:06 ADAN16 00035 2016-11-29 17:00:03 LWP=5,242,880
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXCANCEL=6,000
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXCANCELWARN=1,500
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXMSG=6,000
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXMSGWARN=1,500
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXSTATUS=6,000
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXTNA=3,600
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXTSX=3,600
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXTT=3,600
17:00:06 ADAN16 00035 2016-11-29 17:00:03 MXWTOR=0
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NC=500
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NCLOG=0
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NH=8,000
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NISNHQ=1,999
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NONDES=YES
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NPLOG=0
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NPLOGBUFFERS=5
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NQCID=250
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NSISN=100
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NT=20
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NU=500
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NUCID=10,001
17:00:06 ADAN16 00035 2016-11-29 17:00:03 NWORK1BUFFERS=50
17:00:06 ADAN16 00035 2016-11-29 17:00:03 READONLY=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 REFSTPRT=YES
17:00:06 ADAN16 00035 2016-11-29 17:00:03 REVIEW=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 RIAFTERUPDATE=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 SECUID=ALLOW
```

```

17:00:06 ADAN16 00035 2016-11-29 17:00:03 SMF=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 TLSCMD=20
17:00:06 ADAN16 00035 2016-11-29 17:00:03 TNAA=1,800
17:00:06 ADAN16 00035 2016-11-29 17:00:03 TNAE=1,800
17:00:06 ADAN16 00035 2016-11-29 17:00:03 TNAX=1,800
17:00:06 ADAN16 00035 2016-11-29 17:00:03 TT=1,800
17:00:06 ADAN16 00035 2016-11-29 17:00:03 UPDATECONTROL=NODELAY
17:00:06 ADAN16 00035 2016-11-29 17:00:03 UTIONLY=NO
17:00:06 ADAN16 00035 2016-11-29 17:00:03 V64BIT=NO
17:00:06 ADAN41 00035 2016-11-29 17:00:03 Function completed

```

Additional Adabas add-on products and nucleus components may display more parameters than are shown in this sample.

DPPT Command: Display Parallel Participant Table (PPT)



Use the DPPT command to display the parallel participant table (PPT) block for a nucleus.

This command produces internal information for use by Software AG technical support.

Sample Output

```

ADAI29 Oper cmd: DPPT
ADAN24 00199 2009-01-27 22:31:35 Display PPT RABNs 00000083 to 000000A2
ADAN24 00199 2009-01-27 22:31:35
ADAN24 00199 2009-01-27 22:31:35           PPT RABN: 00000083
ADAN24 00199 2009-01-27 22:31:35 Number of entries: 05
ADAN24 00199 2009-01-27 22:31:35 Nucleus indicator: E2
ADAN24 00199 2009-01-27 22:31:35           NUCID: 0000
ADAN24 00199 2009-01-27 22:31:35           Session number: 0005
ADAN24 00199 2009-01-27 22:31:35           Last PLOG block: 00000000
ADAN24 00199 2009-01-27 22:31:35           PLOG block ind: 00
ADAN24 00199 2009-01-27 22:31:35 PPT Entry length: 0023
ADAN24 00199 2009-01-27 22:31:35           Entry ID: E6
ADAN24 00199 2009-01-27 22:31:35 Dataset=/SAGUID/PLX2/V7/WORKR1/
ADAN24 00199 2009-01-27 22:31:35 PPT Entry length: 0023
ADAN24 00199 2009-01-27 22:31:35           Entry ID: 61
ADAN24 00199 2009-01-27 22:31:35 Dataset=/ SAGUID /PLX2/V7/PLOGR1/
ADAN24 00199 2009-01-27 22:31:35 PPT Entry length: 0023
ADAN24 00199 2009-01-27 22:31:35           Entry ID: 62
ADAN24 00199 2009-01-27 22:31:35 Dataset=/ SAGUID /PLX2/V7/PLOGR2/
ADAN24 00199 2009-01-27 22:31:35 PPT Entry length: 0023
ADAN24 00199 2009-01-27 22:31:35           Entry ID: 41
ADAN24 00199 2009-01-27 22:31:35 Dataset=/ SAGUID /PLX2/V7/CLOGR1/
ADAN24 00199 2009-01-27 22:31:35 PPT Entry length: 0023

```

```
ADAN24 00199 2009-01-27 22:31:35      Entry ID: 42
ADAN24 00199 2009-01-27 22:31:35 Dataset=/ SAGUID /PLX2/V7/CLOGR2/
ADAN41 00199 2009-01-27 22:31:35 Function completed ↵
```

DRES Command: Display Allocated Pool Space and High Water Marks



Use the DRES command to display the allocated pool space and the highest use level (high water mark) during current session. Here is an example:

```
17:32:07 ADAI29 Oper cmd: DRES ↵
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Resource      Size      Current  ↵
Highwater
17:32:07 ADAN28 00035 2016-11-29 17:32:05 AB  -Pool      1228800   N/A      ↵
1536
17:32:07 ADAN28 00035 2016-11-29 17:32:05 CQ  -Pool      96000    0        ↵
192
17:32:07 ADAN28 00035 2016-11-29 17:32:05 DUQ -Pool      7000     0        ↵
0
17:32:07 ADAN28 00035 2016-11-29 17:32:05 FI  -Pool      100000   0        ↵
160
17:32:07 ADAN28 00035 2016-11-29 17:32:05 HQ  -Pool      416000   0        ↵
5200
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Redo-Pool  10485760  0        ↵
176
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Sec  -Pool      10000    0        ↵
0
17:32:07 ADAN28 00035 2016-11-29 17:32:05 TBI  -Pool      4000     0        ↵
0
17:32:07 ADAN28 00035 2016-11-29 17:32:05 TBS  -Pool      30000    0        ↵
0
17:32:07 ADAN28 00035 2016-11-29 17:32:05 UQ   -Pool      183456   1092     ↵
1456
17:32:07 ADAN28 00035 2016-11-29 17:32:05 UQF  -Pool      48384    72       ↵
144
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Wrk1-IO B    50       0        ↵
5
17:32:07 ADAN28 00035 2016-11-29 17:32:05 WORK-Pool    5242880  128     ↵
34584
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Work Part1   1200     0        ↵
7
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Work Pt 1b   4        4        ↵
4
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Work Part2   20       0        ↵
0
```

```
17:32:07 ADAN28 00035 2016-11-29 17:32:05 Work Part3      270      0      ↵
0
17:32:07 ADAN41 00035 2016-11-29 17:32:05 Function completed
```



Note: The usage of the protection data extension area (resource name `Work Pt 1b`) reflects the number of WORK blocks that the nucleus is keeping/has kept in the global cache. The protection data extension area provides space for these WORK blocks for the case that the nucleus occasionally does not participate in a buffer flush performed by a peer nucleus in the cluster.

DTARG Command: Display System Image



Use the `DTARG` command to display the system image currently servicing commands routed using the Entire Net-Work DBID target. This command can be issued against Adabas cluster nuclei.

The Entire Net-Work DBID target is used to route commands from users on remote systems with no active PLXCB. Such systems may be part of the z/OS sysplex or other Entire Net-Work nodes outside the sysplex.

DXCACHE Command: Display Cache Statistics



Issue the `DXCACHE` command against cluster nuclei to display the primary cache-related statistics. Statistics are displayed in messages on the console. The kinds of statistics displayed are described in *Global Cache Statistics (Cluster Nucleus Only)*, in the *Adabas Parallel Services Operations Guide*.

The full set of statistics shown in this output is displayed only for users who have the selectable unit Adabas Online System (AOS) installed.

If you have installed only the demo version of AOS delivered with Adabas, only the statistics in the sections *Totals*, *Data Storage*, and *Normal Index* are displayed using this command.

DXFILE Command: Display Cache Statistics for Files

```
DXFILE= fnr [, fnr1 ] ...
```

Issue the `DXFILE` command against cluster nuclei to display cache-related statistics for one to five files.

Statistics are displayed in messages on the console. The kinds of statistics displayed are described in *Global Cache Statistics (Cluster Nucleus Only)*, in the *Adabas Parallel Services Operations Guide*.

DXLOCK Command: Display Lock Statistics

A rectangular box with a thin blue border containing the word "DXLOCK" in a bold, blue, sans-serif font.

Issue the DXLOCK command against cluster nuclei to display lock-related statistics. Statistics are displayed in messages on the console. The kinds of statistics displayed are described in *Global Lock Statistics (Cluster Nucleus Only)*, in the *Adabas Parallel Services Operations Guide*.

DXMSG Command: Display Messaging Performance Statistics

A rectangular box with a thin blue border containing the word "DXMSG" in a bold, blue, sans-serif font.

The DXMSG command can be used to display the messaging performance statistics available for Adabas Cluster Services or Adabas Parallel Services. Statistics are displayed in messages on the console. The kinds of statistics displayed are described in *Messaging Statistics*, in the *Adabas Cluster Services Operations Guide* or *Adabas Parallel Services Operations Guide*.

The statistics produced by the DXMSG command can be used to determine the impact of messaging on system performance and to determine how to set the ADARUN MXMSG parameter.

DXSTAT Command: Display Cache and Lock Statistics



Use the `DXSTAT` command against cluster nuclei to display all cache- and lock-related statistics as well as I/O statistics for `WORK` datasets in the cluster. Statistics are displayed in messages on the console. The kinds of statistics displayed are described in *Global Cache Statistics (Cluster Nucleus Only)*, *Global Lock Statistics (Cluster Nucleus Only)* and *WORK Data Set I/O Statistics (Cluster Nucleus Only)*, in the *Adabas Parallel Services Operations Guide*.

The full set of external cache statistics shown in this output is displayed only for users who have the selectable unit Adabas Online System (AOS) installed.

If you have installed only the demo version of AOS delivered with Adabas, only the external cache statistics in the sections *Totals*, *Data Storage*, and *Normal Index* are displayed using this command. All file cache statistics for files and all external lock statistics are displayed.

DXWORK Command: Display I/O Statistics for WORK Datasets**DXWORK**

Use the `DXWORK` command to display the read and write I/O statistics for the accesses to the `WORK` datasets of peer nuclei by a particular nucleus in the cluster. The `DXWORK` command may or may not display I/O statistics for a nucleus' own `WORK` dataset. Here is an example:

```
17:04:58 ADAI29 Oper cmd: DXWORK
17:04:58 ADAX61 00035 2016-12-01 17:04:57 WORK dataset I/O statistics:
17:04:58 ADAX61 00035 2016-12-01 17:04:57
17:04:58 ADAX61 00035 2016-12-01 17:04:57 Internal NUCID :          2
17:04:58 ADAX61 00035 2016-12-01 17:04:57 External NUCID :        10002
17:04:58 ADAX61 00035 2016-12-01 17:04:57 Reads          :          3
17:04:58 ADAX61 00035 2016-12-01 17:04:57 Writes         :          0
17:04:58 ADAX61 00035 2016-12-01 17:04:57
17:04:58 ADAX61 00035 2016-12-01 17:04:57 Internal NUCID :          3
17:04:58 ADAX61 00035 2016-12-01 17:04:57 External NUCID :        10001
17:04:58 ADAX61 00035 2016-12-01 17:04:57 Reads          :          1
17:04:58 ADAX61 00035 2016-12-01 17:04:57 Writes         :          0
17:04:58 ADAN41 00035 2016-12-01 17:04:57 Function completed
```

MXCANCEL Command: Dynamically Change MXCANCEL Parameter Setting



MXCANCEL

The `MXCANCEL` command can be used against cluster nuclei to dynamically change (while the Adabas nucleus is running) the setting of the `MXCANCEL ADARUN` parameter for a nucleus. This command can be used to specify the maximum number of seconds for which one cluster nucleus waits for the termination of another nucleus in the cluster that it has canceled for not responding to an internal, intracluster command. Valid values range from 15 through 2147483647.

If the new `MXCANCEL` value is less than or equal to the current `MXCANCELWARN` setting, the `MXCANCELWARN` setting is automatically set to zero (0). (You can then use the `MXCANCELWARN` command to dynamically set the `MXCANCELWARN` setting appropriately.)

For more information about the `MXCANCEL` setting for a nucleus, read about the `ADARUN MXCANCEL` parameter in the *Adabas Operations Manual*.

MXCANCELWARN Command: Dynamically Change MXCANCELWARN Parameter Setting

MXCANCELWARN

The MXCANCELWARN command can be used to dynamically change (while the Adabas nucleus is running) the setting of the MXCANCELWARN ADARUN parameter for a nucleus. This command can be used to specify the number of seconds after an intracluster cancellation request is made when the cluster nucleus requesting the cancellation should issue warnings about the inability of the canceled peer nucleus to terminate. Valid values can be zero (0) or an integer ranging from 4 through the value of MXCANCEL minus 1 (MXCANCEL-1).

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

For more information about the MXCANCELWARN setting for a nucleus, read about the ADARUN MXCANCELWARN parameter in the *Adabas Operations Manual*.

MXMSG Command: Dynamically Change MXMSG Parameter Setting



The MXMSG command can be used to dynamically change (while the Adabas nucleus is running) the setting of the `MXMSG ADARUN` parameter for a cluster nucleus. This command can be used to specify the maximum number of seconds in which one cluster nucleus waits for the response of another nucleus in the cluster to an internal, intracluster command. Valid values range from 15 through 32767.

If the new MXMSG value is less than or equal to the current MXMSGWARN setting, the MXMSGWARN setting is automatically set to zero (0). (You can then use the MXMSGWARN command to dynamically set the MXMSGWARN setting appropriately.)

For more information about the MXMSG setting for a nucleus, read about the `ADARUN MXMSG` parameter in the *Adabas Operations Manual*.

MXMSGWARN Command: Dynamically Change MXMSGWARN Parameter Setting

MXMSGWARN

The MXMSGWARN command can be used to dynamically change (while the Adabas nucleus is running) the setting of the MXMSGWARN ADARUN parameter for a cluster nucleus. This command can be used to specify the number of seconds after an intracenter request is made when the cluster nucleus issuing the intracenter request should issue warnings about the failure of the target cluster nucleus to respond. Valid values can be zero (0) or an integer ranging from 4 through the value of MXMSG minus 1 (MXMSG-1).

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

For more information about the MXMSGWARN setting for a nucleus, read about the ADARUN MXMSGWARN parameter in the *Adabas Operations Manual*.

MXWTOR Command: Dynamically Change MXWTOR Parameter Setting

MXWTOR

The optional MXWTOR command can be used to dynamically change (while the Adabas nucleus is running) the setting of the `MXWTOR ADARUN` parameter for a cluster nucleus. This command can be used to specify the maximum number of seconds that a cluster nucleus should wait for an operator response to a self-terminating operator query. Valid values range from 15 through 64800 or zero (0).

If MXWTOR is set to zero, the operator query does not occur (no ADAX9J messages are written) and no operator intervention is expected.

For more information about the MXWTOR setting for a nucleus, read about the `ADARUN MXWTOR` parameter in the *Adabas Operations Manual*.

SMFDETAIL Command: Select the Detail Section Types in SMF Interval and Termination Records

SMFDETAIL = { ALL | NONE | *section-name-list* }

Adabas SMF records can contain a variable set of detail sections in Interval and Termination records (subtypes 2 and 3). This command allows you to override the setting of the `SMFDETAIL ADARUN` parameter defined for this Adabas session. Using this command you can select the detail section types in Interval and Termination records that should be included in the Adabas SMF records. Unlike the `SMFDETAIL ADARUN` parameter, the value or values you specify for the `SMFDETAIL` command do *not* need to be enclosed in parentheses. If you specify `NONE` or `ALL`, they should be specified alone. However, you can specify one or more of the other detail section names (`CMD`, `CSHB`, `CSHF`, `CSHG`, `CSHP`, `FILE`, `IODD`, `LOCK`, `MSGB`, `MSGC`, `MSGH`, `PARM`, `STG`, `THRD`, `ZIIP` or `USER`) in one `SMFDETAIL` parameter, separating each value with a comma.

The following table describes the meaning of the possible detail section names that can be used in the `SMFDETAIL` command:

Detail Section Name	Description
ALL	Generate all possible detail sections in the Adabas SMF records. If this value is specified, no others can be specified on the SMFDETAIL parameter.
CMD	Generate Adabas command activity detail sections in the Adabas SMF records.
CSHB ¹	Generate global cache activity by block detail sections in the Adabas SMF records.
CSHF ¹	Generate global cache activity by Adabas file number detail sections in the Adabas SMF records.
CSHG ¹	Generate global cache activity detail sections in the Adabas SMF records.
CSHP ²	Generate Adabas Parallel Services cache activity detail sections in the Adabas SMF records.
FILE	Generate Adabas file activity detail sections in the Adabas SMF records.
IODD	Generate I/O activity by DD statement detail sections in the Adabas SMF records.
LOCK ¹	Generate global lock activity detail sections in the Adabas SMF records.
MSGB ¹	Generate internucleus messaging control block activity detail sections in the Adabas SMF records.
MSGC ¹	Generate internucleus messaging count detail sections in the Adabas SMF records.
MSGH ¹	Generate internucleus messaging service time histogram detail sections in the Adabas SMF records.
NONE	Generate no detail sections in the Adabas SMF records. If this value is specified, no others can be specified on the SMFDETAIL parameter.
PARAM	Generate ADARUN parameter value detail sections in the Adabas SMF records.
STG	Generate Adabas storage pool detail sections in the Adabas SMF records.
THRD	Generate thread activity detail sections in the Adabas SMF records.
ZIIP	Generate zIIP statistics detail sections in the Adabas SMF records.
USER	Generate user-defined details sections in the Adabas SMF records. If USER is specified, a value for the UEXSMF parameter should also be specified to identify the user exit to be invoked to generate the user-defined detail section.

¹ Available only in cluster environments (Adabas Cluster Services or Adabas Parallel Services must be installed).

² Available only in cluster environments with Adabas Parallel Services installed.

TARGET Command: Route Commands to Specified Target

```
TARGET = { * | nucid | sysname }
```

Use this command against an Adabas Cluster Services nucleus to specify the system on which the Entire Net-Work DBID target for the cluster database is to be established. Specify a system name or the NUCID of a cluster nucleus on that system. The DBID target can be established on any system where a nucleus of the cluster is active. It is used to route commands from users on remote systems with no active PLXCB. Such systems may be part of the z/OS sysplex or other Entire Net-Work nodes outside the sysplex.

The command instructs Entire Net-Work to route requests to an Entire Net-Work node on an implicitly or explicitly specified system image. All future commands will be routed to the new destination. If there are existing users who have been routed using the DBID target and the target is relocated to another Entire Net-Work node, they will be reassigned to any cluster nuclei on the new node. These users lose their existing sessions with their previously assigned nuclei. As a consequence, they may receive response code 9 (ADARSP009) on the first command after the target is moved. Furthermore, they may receive response code 21 (ADARSP021) if they attempt to reuse previously defined command IDs (logical read sequences or search results).

Several possible operands can be used to specify the desired DBID target assignment, as described in the following table.

Operand	Description
*	Specify an asterisk (*) to assign the DBID target to the Entire Net-Work node associated with the nucleus receiving the command.
<i>nucid</i>	Specify the nucleus ID (as defined in the ADARUN NUCID parameter) of an active cluster nucleus to assign the DBID target to the Entire Net-Work node associated with that nucleus.
<i>sysname</i>	Specify a system name (LPAR name) to assign the DBID target to the Entire Net-Work node associated with an active cluster nucleus on that system.



Note: The DBID target identifies an Entire Net-Work node, not a specific nucleus on a system. The DBID target is jointly serviced by all nuclei on that Entire Net-Work node.

Index

Symbols

64-bit virtual storage option, 51

A

Adabas Caching Facility
 adjusting shared cache size for, 29
Adabas Cluster Services
 ADARUN parameter to enable, 35-36
Adabas Parallel Services
 ADARUN parameter to enable, 35-36
Adabas SMF records
 selecting detail SMF record sections, 48
Adabas sysplex cluster
 nucleus ID
 parameter to specify, 46
ADACOM
 initialization parameters, 7
ADACOM operator commands
 DIM
 display active images, 58, 61
 DN
 display active local nuclei, 58, 61
 DUMP, 59, 62
 SN, 59, 63
ADACOM parameters
 DBID, 9
 FORCE, 9
 LOC, 10
 NU
 maximum number of users, 10
 SVC
 supervisory call number for interregion communication,
 11
ADARUN
 directory of parameters, 15
ADARUN parameters, 13
 determining correct settings, 14
 MLWTO, 41

C

cache
 shared cache, 33
cache structure
 ratio
 directory entries to elements, 38

 statistics display, 75, 78
cache vector size, 29
CLOGMRG
 ADARUN parameter, 28
CLUCACHEEXTRA
 ADARUN parameter, 29
CLUCACHESIZE
 ADARUN parameter, 29
CLUCACHETYPE
 ADARUN parameter, 31
CLUCACHEUNCHANGED
 ADARUN parameter, 33
CLUGROUPNAME
 ADARUN parameter, 34
CLULOCKSIZE
 ADARUN parameter, 34
CLUPUBLPROT
 ADARUN parameter, 35
CLUSTER
 ADARUN parameter, 36
cluster
 cache vector size, 29
 group name, 34
 redo updates
 parameter to set size of pool, 40
CLUWORK1CACHE
 ADARUN parameter, 36
Command log
 parameter to enable merge, 28
command routing, 86
communication
 cluster nuclei
 timeout threshold for canceled peer nucleus, 41
 timeout threshold for command, 43
console messages
 multi-line WTO messages, 41
Control scheduling of update commands during buffer flush,
50

D

DBID parameter, 9
descriptor
 unique
 parameter to set size of pool, 39
detail sections
 SMF records, 84
DIM command
 display active images, 58, 61
DIRRATIO / ELEMENTRATIO

- ADARUN parameter, 38
- displaying Adabas session parameters, 69
- DMEMTB command
 - display member state table, 67
- DN command
 - display active local nuclei, 58, 61
- DNFV command
 - displaying nucleus file variables, 68
- DPARAM operator command, 69
- DPPT command
 - displaying a nucleus PPT block, 71
- DTARG command
 - display system image, 74
- DUMP command
 - dump storage areas, 59, 62
- DXCACHE command
 - display lock statistics, 75
- DXFILE command
 - display file statistics, 76
- DXLOCK command
 - display lock statistics, 77
- DXSTAT command
 - display cache and lock statistics, 78
- dynamically change MXCANCEL setting, 80
- dynamically change MXCANCELWARN setting, 81
- dynamically change MXMSGWARN setting, 83
- dynamically change MXWTOR setting, 84

F

- file
 - statistics display, 76
- FORCE parameter, 9

G

- global cache area
 - construct type, 31
 - ratio
 - directory entries to elements, 38
 - size, 29
- Global lock area
 - statistics display, 77
- global lock area
 - size, 34

I

- Initialization parameters
 - ADACOM, 7

L

- LBP
 - ADARUN parameter, 30
- LDEUQP
 - ADARUN parameter, 39
- LOC parameter, 10
- local buffer pool
 - size parameter, 30
- lock structure
 - statistics display, 78
- LRDP
 - ADARUN parameter, 40

M

- member state table
 - command to display, 67
- MLWTO parameter, 41
- multi-line console messages, 41
- multiprocessing
 - nucleus ID for
 - parameter to specify, 46
- MXCANCEL
 - ADARUN parameter, 41
- MXCANCEL operator command, 80
- MXCANCELWARN
 - ADARUN parameter, 42
- MXCANCELWARN operator command, 81
- MXMSG
 - ADARUN parameter, 43
- MXMSGWARN
 - ADARUN parameter, 44
- MXMSGWARN operator command, 83
- MXWTOR
 - ADARUN parameter, 45
- MXWTOR operator command, 84

N

- NU parameter, 10
- NUCID
 - ADARUN parameter, 46
- nucleus
 - ID for multiprocessing
 - parameter to specify, 46
- nucleus file variables
 - displaying, 68

O

- Operator commands, 54
 - ADACOM, 54
 - cluster nucleus, 64
 - DXLOCK
 - display lock statistics, 77
- operator commands
 - DMEMTB
 - display member state table, 67
 - DNFV, 68
 - DPARAM, 69
 - DPPT, 71
 - DTARG, 74
 - DXCACHE
 - display lock statistics, 75
 - DXFILE
 - display file statistics, 76
 - DXSTAT
 - display cache and lock statistics, 78
 - MXCANCEL, 80
 - MXCANCELWARN, 81
 - MXMSGWARN, 83
 - MXWTOR, 84
 - TARGET, 86

P

- PPT

displaying a nucleus block, 71
PROGRAM
ADARUN parameter, 47

R

redo pool
parameter to set size of, 40

S

session
displaying current parameters, 69
shared cache
unchanged block option, 33
SMF Interval records, 84
SMF records
detail sections, 84
selecting detail sections, 48
SMF Termination records, 84
SMFDETAIL
ADARUN parameter, 48
SN command
set nucleus status, 59, 63
SVC parameter, 11
system image display, 74

T

TARGET command
route commands to specified target, 86

U

unique descriptor pool
parameter to set size of, 39
UPDATECONTROL
ADARUN parameter, 50
Users
maximum number of, 10

V

V64BIT parameter, 51
virtual storage
64-bit option, 51

X

XCF group
name, 34

