

Adabas Cluster Services

Reference

Version 8.5.4

April 2020

This document applies to Adabas Cluster Services Version 8.5.4 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 © 2020 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: ALS-AALSREF-854-20210929

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 |
| CMDMGR - ADACOM Use | 8 |
| DBID - Database Identification | 9 |
| FORCE - Force Restart after Abnormal Termination | 9 |
| LOC - Specify PLXCB Location | 10 |
| NU - Number of User Table Elements | 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 |
| 5 Entire Net-Work (NETWRK) Initialization Parameters | 25 |
| Guidelines | 27 |
| CT - Command Timeout Limit | 27 |
| FORCE - Overwrite ID Table Entry | 54 |
| LU - Length of Intermediate User Buffer | 29 |
| NAB - Number of Attached Buffers | 29 |
| NC - Number of Request Queue Elements | 30 |
| SVC - Interregion Communication Number | 30 |
| TARGETID - Entire Net-Work Target Identification | 31 |
| 6 Useful ADARUN Parameters for Adabas Cluster Services | 33 |
| CLOGMRG Parameter: Cluster Command Log Merge | 34 |
| CLUCACHEEXTRA Parameter: Adjust Shared Cache Vector Size | 35 |
| CLUCACHENAME Parameter: Cluster Cache Structure/Area Name | 35 |
| CLUGROUPNAME Parameter: Cluster Group Name | 36 |
| CLULOCKNAME Parameter: Cluster Lock Structure/Area Name | 37 |
| CLUPUBLPROT Parameter: Publishing of Protection Data | 38 |
| CLUSTER Parameter: Adabas Nucleus Cluster Session Control | 39 |
| CLUWORK1CACHE Parameter: Number of WORK Blocks in Global Cache | 39 |
| DIRRATIO / ELEMENTRATIO Parameters: Cluster Cache Directory-to-Data Elements Ratio | 41 |
| LDEUQP Parameter: Length of Unique Descriptor Pool | 42 |
| LRDP Parameter: Length of the Redo Pool | 43 |
| MLWTO Parameter: Multi-Line Console Message Control | 44 |
| MXCANCEL Parameter: Timeout Threshold for Canceled Peer Nucleus | 44 |
| MXCANCELWARN Parameter: Timeout Threshold for Canceled Peer Nucleus Warning | 45 |

| | |
|--|-----|
| MXMSG Parameter: Timeout Threshold for Internucleus Command Processing | 46 |
| MXMSGWARN Parameter: Timeout Threshold for Internucleus Command Processing Warning | 47 |
| MXSTATUS Parameter: Member-Level XCF Status Monitoring Heartbeat Interval | 48 |
| MXWTOR Parameter: Self-Termination Operator Query Interval | 49 |
| NUCID Parameter: Adabas Cluster Nucleus ID | 50 |
| PROGRAM Parameter: Program to Run | 51 |
| SMFDETAIL Parameter: Selecting Detail SMF Record Sections | 52 |
| UPDATECONTROL Parameter: Control Scheduling of Update Commands During Buffer Flush | 54 |
| V64BIT Parameter: 64-Bit Virtual Storage Option | 55 |
| 7 Entire Net-Work Parameter Statements | 57 |
| Conventions | 58 |
| XCF NODE Statement | 61 |
| XCF DRIVER Statement | 71 |
| XCF LINK Statement | 74 |
| 8 Adabas Cluster Operator Commands | 79 |
| ADACOM Operator Commands | 80 |
| Adabas Cluster Nucleus Operator Commands | 87 |
| 9 Entire Net-Work Operator Commands | 111 |
| Overview of Entire Net-WorkCommands for z/OS Environments | 113 |
| Command Descriptions | 114 |
| XCF Line Driver Commands | 127 |
| 10 Entire Net-Work Diagnostic Utilities | 133 |
| NETPFIL1 Utility | 134 |
| NETPFIL2 Utility | 136 |
| Example JCL | 138 |
| Index | 139 |

Preface

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

This documentation provides information on the operator commands for ADACOM and for an Adabas cluster nucleus; it also describes the operator commands for Entire Net-Work and for its XCF line driver.

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

It also describes the parameter statements required for enabling Entire Net-Work to communicate with other nodes.

This document is organized as follows:

| | |
|--|---|
| <i>Initialization Parameters</i> | Describes the initialization parameters needed for an Adabas Cluster Services environment. |
| <i>ADACOM Initialization Parameters</i> | Describes the ADACOM initialization parameters needed for an Adabas Cluster Services environment. |
| <i>ADANUC Initialization Parameters for Cluster Nuclei</i> | Describes the ADARUN parameters needed for each nucleus in a cluster. |
| <i>Entire Net-Work (NETWRK) Initialization Parameters</i> | Describes the Entire Net-Work parameters needed for each nucleus in a cluster. |
| <i>Useful ADARUN Parameters</i> | Provides a reference for ADARUN parameters useful to Adabas Cluster Services processing. |
| <i>Entire Net-Work Parameter Statements</i> | Describes the initialization parameters needed for Entire Net-Work. |
| <i>Adabas Cluster Operator Commands</i> | Describes operator commands for ADACOM and for an Adabas Cluster Services cluster nucleus. |
| <i>Entire Net-Work Operator Commands</i> | Describes operator commands for Entire Net-Work. |
| <i>Entire Net-Work Diagnostic Utilities</i> | Describes the Entire Net-Work diagnostic utilities, provided in the Entire Net-Work load library. |

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 cluster environment. It also describes the parameters required for enabling Entire Net-Work to in cluster environments.

ADARUN parameters are used to customize the Adabas environment. For Adabas Cluster Services environments, three 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 Entire Net-Work node, the name must be NETWRK.
- To set up an ADACOM task, the name must be ADACOM (no other ADARUN parameters are recognized for the ADACOM program run).

| | |
|---|---|
| <i>ADACOM Initialization Parameters</i> | Describes the ADACOM initialization parameters required by Adabas Cluster Services. |
| <i>About ADARUN Parameters for Cluster Nuclei</i> | Describes the Adabas nucleus (ADANUC) initialization parameters required by Adabas Cluster Services. |
| <i>Entire Net-Work (NETWRK) Initialization Parameter</i> | Describes the Entire Net-Work (NETWRK) initialization parameters for Entire Net-Work in Adabas Cluster Services environments. |
| <i>Useful ADARUN Parameters for Adabas Cluster Services</i> | Provides a reference for all of the ADARUN parameters useful to Adabas Cluster Services processing. |

3 ADACOM Initialization Parameters

- CMDMGR - ADACOM Use 8
- DBID - Database Identification 9
- FORCE - Force Restart after Abnormal Termination 9
- LOC - Specify PLXCB Location 10
- NU - Number of User Table Elements 10
- SVC - Interregion Communication Number 11

An unlimited number of Adabas nucleus clusters, each with up to 32 members sharing a common database, can be defined per parallel sysplex.

The ADACOM initialization task:

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

Parameters specify the SVC/DBID combinations (sets) that the ADACOM is to manage. The DBID identifies the external physical database shared by a particular cluster of nuclei and known to the application. The SVC value may vary within a cluster and the same SVC may be used for different clusters. Other ADACOM parameters are discussed in the following sections.

Although a single ADACOM job can run all SVC/DBID sets in a sysplex environment, it is possible to run multiple ADACOM tasks simultaneously with the same, mixed, or completely different SVC/DBID sets. A subtask is attached to the ADACOM task for each SVC/DBID set occurring in the input of ADACOM. All work pertaining to an SVC/DBID set is performed by the subtask.

The COMPRINT data set prints global messages that apply to all SVC/DBID sets defined to an ADACOM task. In addition, a SYSOUT data set is dynamically allocated for command output to each SVC/DBID set. The format of this data set DD-name is *Pssdddd* where *ss* is the last two digits of the SVC and *dddd* is the DBID.

For a sample job to run ADACOM, see the section *Create a Startup Procedure for ADACOM* in *Adabas Cluster Services Installation*.

CMDMGR - ADACOM Use

| Parameter | Specify... | Possible Values | Default |
|-----------|--|--------------------------|---------|
| CMDMGR | whether ADACOM remains in use as a command manager after initialization or quiesces. | <u>Y</u> ES <u>N</u> O | YES |

The CMDMGR parameter specifies whether the ADACOM task remains in operation as a command manager for the cluster nuclei (CMDMGR=YES) or whether it runs like a batch job and terminates after processing the commands given in DDKARTE (CMDMGR=NO).

The ADACOM task used to allocate and monitor the PLXCB and subordinate data structures in common storage (ECSA) terminates after processing the commands given in DDKARTE, if CMDMGR=NO has been specified for all SVC/DBID pairs. If CMDMGR=YES (the default setting) has been specified for one or more SVC/DBID pairs, the ADACOM task stays active and accepts operator commands for displaying and controlling the multiprocessing environment.

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 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 nucleus cluster.

This number must be unique among all Adabas database IDs, NUCIDs, Natural buffer pool IDs, Entire Net-Work target IDs (suggested range 65482-65535), and internal target IDs used by Adabas Cluster Services (restricted range 65001-65481).

FORCE - Force Restart after Abnormal Termination

| Parameter | Specify... | Possible Values | Default |
|-----------|--|--------------------------|---------|
| FORCE | whether ADACOM forces a restart after an abnormal termination. | <u>Y</u> ES <u>N</u> O | NO |

If Adabas believes there are still active nuclei, ADACOMs, or users on an image after a termination, an `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 is wrong in its belief, you can use `FORCE=YES` to force a clean restart.

Note, however, that if a sysplex cluster nucleus or an ADACOM with the same SVC number 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 re-setting the environment.
2. The nucleus ADARUN parameter `FORCE` is neither seen nor used by the Adabas Cluster Services `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. | <u>C</u> SA <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. To use a PLXCB dataspace ADACOM must be started to define the dataspace before any nuclei start.
2. ADACOM must remain active as long as any nuclei or database users on the local system are active.
3. 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).
4. A mapped dataspace can be captured in SYSUDUMP.
5. 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

| Parameter | Specify... | Minimum | Maximum | Default |
|------------|--|---------|------------|---------|
| <u>N</u> U | the number of users that can be active in parallel on sysplex cluster nuclei in the local 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 sysplex cluster nuclei on the local 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 the assigned cluster nucleus otherwise deletes the user's UQE (e.g., due to timeout).

The ADACOM `NU` parameter specifies the number of concurrent users (UTEs) allowed for all cluster nuclei on the local operating system image. The first cluster nucleus or 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 sysplex cluster nuclei and ADACOM initialization tasks (if used) in the local image that have the same SVC number, 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 Cluster Services control block (PLXCB) is freed, and the ADACOM subtask for this SVC/DBID set terminates, regardless of the `CMDMGR` parameter setting. However, if the ADACOM subtask believes that there are active nuclei or users on this database, and `FORCE=YES` is not used, no action is taken and the subtask terminates abnormally with user ABEND 40 (hex). The ADACOM main task and any subtasks handling other SVC/DBID sets are not affected.

Example:

Run the Adabas cluster nuclei on the local image with 500 elements in the user table.

```
NU=500
```

SVC - Interregion Communication Number

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|---------|
| SVC | the supervisor call number to be used for the Adabas Cluster Services session. | 200-255 | 249 |

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 version 8.2 Adabas SVC (ADASVC). Valid `SVC` values are 200-255.

Example:

Execute an Adabas Cluster Services session under an OS/390 or z/OS image using SVC 202.

```
SVC=202
```


4 About ADARUN Parameters for Cluster Nuclei

- Global ADARUN Parameters 14
- ADARUN Parameter Usage in Cluster Environments 15

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

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 Cluster 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](#)
- [Useful ADARUN Parameters for Adabas Cluster 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, intracenter 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 intracenter 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 intracenter messages.

All other nuclei in the cluster receive the intracenter 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 Cluster 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 Cluster Services environments and are described in [Useful ADARUN Parameters for Adabas Cluster Services](#), elsewhere in this guide. A description of every ADARUN parameter (including those not described in the Adabas Cluster 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 Note: Please refer to Online Recovery in the <i>Operations</i> documentation for recommendations concerning the LWP formula. | | 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 | | | |

| Parameter | Usage | N/A | LF | LM | GF | GM |
|---------------|--|-----|----|----|----|----|
| MSGPRINT | Case of messages sent to DD/PRINT data set | | LF | | | |
| 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 | |

| Parameter | Usage | N/A | LF | LM | GF | GM |
|-----------------------|---|-----|----|----|----|----|
| PAMREQS | Maximum number of BS2000 PAM blocks that can be transferred in a single PAM macro request | | LF | | | |
| 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 | | | |

| Parameter | Usage | N/A | LF | LM | GF | GM |
|-----------------------|---|-----|----|----|----|----|
| RPWARNINTERVAL | Replication pool usage warning message suppression interval | | LF | | | |
| 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 | | |

| Parameter | Usage | N/A | LF | LM | GF | GM |
|-------------------------|--|-----|----|----|----|----|
| TMSGSEV ³ | Severity threshold for suppression of Adabas Transaction manager warning messages | | LF | | | |
| 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.

5

Entire Net-Work (NETWRK) Initialization Parameters

| | |
|--|----|
| ▪ Guidelines | 27 |
| ▪ CT - Command Timeout Limit | 27 |
| ▪ FORCE - Overwrite ID Table Entry | 54 |
| ▪ LU - Length of Intermediate User Buffer | 29 |
| ▪ NAB - Number of Attached Buffers | 29 |
| ▪ NC - Number of Request Queue Elements | 30 |
| ▪ SVC - Interregion Communication Number | 30 |
| ▪ TARGETID - Entire Net-Work Target Identification | 31 |

The ADARUN parameters for Entire Net-Work specify all interregion communication information for the node. They are specified with ADARUN PROGRAM=NETWRK.

Software AG recommends that you use your existing values (or the default settings) of the ADARUN parameters for each Entire Net-Work node, and then tune the values after analyzing the performance of the node. See the *Performance and Tuning* documentation for information about expected differences.

The following table summarizes the Entire Net-Work ADARUN parameters:

| Parameter | Use | Minimum | Maximum | Default |
|-----------|----------------------------------|---------|----------|---------|
| CT | Command timeout limit (sec.) | 1 | 16777215 | 60 |
| FORCE | Overwrite target ID entry | - | - | NO |
| LU | User buffer size | 4000 | 65535 | 65535 |
| NAB | Number of attached buffers | 3 | 49150 | 16 |
| NC | Number of request queue elements | 20 | 32767 | 200 |
| PROGRAM | Program to be executed | NETWRK | NETWRK | (none) |
| SVC | SVC number | 200 | 255 | 249 |
| TARGETID | Entire Net-Work target ID | 1* | 65535 | 1 |

* Software AG strongly recommends that you restrict Entire Net-Work target IDs in a sysplex cluster environment to values in the range 65482 through 65535 to avoid interference with the target IDs used for other cluster requirements.

Following is an example of an ADARUN statement for Entire Net-Work:

```
ADARUN PROG=NETWRK,TARGETID=3333,NAB=20,NC=50,LU=65535,SVC=251
```

For this node, Entire Net-Work

- runs with a target ID of 3333;
- allocates an attached buffer pool for 20 interregion communication buffers;
- processes as many as 50 requests simultaneously;
- uses the SVC 251; and
- ensures adequate buffer size for Adabas compatibility by setting the value of the LU parameter to 65535.

Guidelines

Use the following guidelines when specifying ADARUN session parameters:

- Ensure that the correct program to be executed is specified (`PROG=NETWRK`).
- Ensure that the correct target ID is specified (`TARGETID` parameter).
- Determine which settings for the following parameters are applicable for the session:

| | |
|-------|-----------------------------|
| FORCE | (override active target ID) |
| SVC | (Adabas SVC number) |

CT - Command Timeout Limit

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|----------|---------|
| CT | the time allowed for a user program to obtain results from Entire Net-Work. | 1 | 16777215 | 60 |

The maximum number of seconds (more precisely, units of 1.048576 seconds) that can elapse from the time an Entire Net-Work request is completed until the results are retrieved by the user through the interregion communication ("16-call"/operating system-dependent).

This parameter is used to prevent a request queue element (RQE) and the attached buffer from being held indefinitely if a user with an outstanding Entire Net-Work request terminates abnormally.

Possible causes for a command timeout are

- non-dispatchable or swapped-out user address space;
- user canceled;
- low user priority in a high-activity system.

If the CT limit is exceeded, the RQE and the attached buffer are released and message NETM93 is printed. If the user attempts to access the RQE after the timeout has occurred, response code 254 is returned.

Example:

Execute an Entire Net-Work task; set the maximum permitted time to obtain the result of a command using interregion communication to approximately 30 seconds.

```
ADARUN PROG=NETWRK,CT=30
```

FORCE - Overwrite ID Table Entry

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| FORCE | whether to overwrite the target ID table entry. | YES NO | NO |

By default (FORCE=NO), error NETM98 occurs if the TARGETID parameter (see [TARGETID](#)) specifies a target ID that is already in the target ID table:

```
NETM98 targetid TARGET INITIALIZATION ERROR: DUP ID (LOCAL)
```

FORCE=NO prevents the table entry from being overwritten. FORCE=YES allows the existing target ID table entry to be overwritten by the value specified for the TARGETID= parameter.

When a target terminates, its ID table entry is normally removed from the ID table. However, if an Entire Net-Work task has terminated abnormally without the recovery mechanism receiving control, its entry can remain in the ID table. FORCE=YES permits an Entire Net-Work target to be restarted even though its ID is already present in the ID table.

In addition to Entire Net-Work, other Software AG products such as Adabas or Entire System Server (formerly Natural Process) may have entries in the ID table.

Because FORCE=YES replaces *any* entry that matches the specified TARGETID value, you must ensure that the ID table does not contain that same value for any target that is active.



Caution: If you were to mistakenly override the ID table entry of an active target, it could severely disrupt the operation of the affected service system.

For this reason, all target and database IDs must be unique, and FORCE should be issued *only* after you determine that no other target with the ID specified by TARGETID is active.

Example:

Execute an Entire Net-Work task.

If the ID table contains an entry for target ID 3333, overwrite it; that is, allow the new task to use ID table entry 3333. Disconnect any existing task with an ID of 3333 from all user access.

```
ADARUN PROG=NETWRK, FORCE=YES, TARGETID=3333
```

If FORCE=NO (the default), Entire Net-Work does not permit the use of target ID 3333 if the ID table already contains an entry for that ID. The message NETM98 is issued, indicating that a target using the ID 3333 is already active.

LU - Length of Intermediate User Buffer

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|---------|---------|
| LU | the size of the intermediate user buffer. | 4000 | 65535 | 65535 |

The size specified must be large enough to accommodate all Adabas control information (108 bytes), all user buffers (format, record, search, value, ISN) that may be required for any given Adabas command, and any user information passed from Adabas link routines to nucleus user exits.

If the PREFETCH option or a utility that needs large record/ISN buffers is to be used during the session, the setting of LU must be large enough to contain the buffers needed.

An error occurs if the LU parameter specifies a value greater than the byte count implied by the NAB (number of attached buffers) parameter. The NAB default of 16 allocates more than 64 kilobytes.

Example:

Start an Entire Net-Work session using an Adabas intermediate user buffer area of 20,000 bytes.

```
ADARUN PROG=NETWRK,LU=20000
```

NAB - Number of Attached Buffers

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|--|---------|---------|---------|
| NAB | the number of attached buffers to be used. | 3 | 49150 | 16 |

NAB is an optional parameter that defines the number of attached buffers to be used during the Entire Net-Work session. Entire Net-Work allocates an attached buffer pool whose size equals the NAB value multiplied by 4112.

An attached buffer is an internal buffer used for interregion communication. Entire Net-Work requires the use of attached buffers in all environments.



Note: A value of less than 16 prevents some Adabas utility functions from being performed remotely.

Example:

Execute an Entire Net-Work session; use 50 attached buffers.

```
ADARUN PROG=NETWRK,NAB=50
```

NC - Number of Request Queue Elements

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|--|---------|---------|---------|
| NC | the number of elements allowed in the Entire Net-Work request queue. | 20 | 32767 | 200 |

NC is an optional parameter that specifies the maximum number of user requests that can be processed simultaneously by this Entire Net-Work communicator. Any user requests originating on this node that cause the NC value to be exceeded receive response code 151.

Each user request originating from this node is assigned a request queue element (RQE), which exists for the duration of the request. The RQE is released when the user has either received the results of this request from Entire Net-Work, or has been timed out. Each RQE requires 192 bytes.

Example:

Execute an Entire Net-Work session; set the request queue to contain as many as 500 elements at one time.

```
ADARUN PROG=NETWRK,NC=500
```

SVC - Interregion Communication Number

| Parameter | Specify... | Possible Values | Default |
|-----------|-----------------------------------|-----------------|---------|
| SVC | the number of the SVC to be used. | 200-255 | 249 |

The number of the SVC to be used for the Entire Net-Work session. The SVC number is specified as an integer. The SVC is used to perform various internal functions, especially the processing of interregion communication.

SVC is an optional parameter that defines the interregion communications supervisor call (SVC) to be used for this session.

The SVC number specified must correspond to that which has been installed as the Software AG communications SVC at the user installation. The same SVC may be used for Adabas, Adabas Cluster Services, Entire Net-Work, and all other Software AG products using interregion communication.

Example:

Execute an Entire Net-Work session using SVC 205.

```
ADARUN PROG=NETWRK,SVC=205
```

TARGETID - Entire Net-Work Target Identification

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|-------------------------------------|---------|---------|---------|
| TARGETID | the ID for an Entire Net-Work node. | 1* | 65535 | 1 |

* Software AG strongly recommends that you restrict Entire Net-Work target IDs in a sysplex cluster environment to values in the range 65482 through 65535 to avoid interference with the target IDs used for other cluster requirements.

The TARGETID parameter value must uniquely identify a particular Entire Net-Work node. Each target (node) ID must be unique, and can be any valid decimal value ranging from 1 (the default) to 65535. If the default is used for one node, it cannot be used for any other node and thus, this parameter must be specified for all other nodes.

In a sysplex cluster environment, all NUCIDs, Adabas database IDs (DBIDs), internally assigned operating system image IDs, target IDs used by Entire Net-Work, Natural global buffer pool IDs, etc., must be unique throughout all Entire Net-Work nodes. In particular, the Entire Net-Work target ID *must not* coincide with any database ID used in the network. An exception to this rule allows databases to be specified as "local" targets as explained in [Special Exception \(Local\)](#), elsewhere in this section.

Example:

Execute an Entire Net-Work session using target ID 65485.

```
ADARUN PROG=NETWRK,TARGETID=65485
```

This section covers the following topics:

- [Uniqueness Required \(Global\)](#)
- [Special Exception \(Local\)](#)

Uniqueness Required (Global)

Entire Net-Work normally enforces unique target IDs throughout the network, that is, across all connected nodes. A target is not permitted to start if another target with the same ID is active anywhere in the network. The ADAM98 message (target initialization error) will result.

When a link is established between two previously unconnected partial networks, it is possible that the same target ID is used in both parts simultaneously. Because target IDs must be unique within the network, this is considered an error situation.

Entire Net-Work cannot determine which of the targets should be deactivated. Therefore the link between the two parts of the network is not established and explanatory messages (NET0108 and NET0106) are written to the operator consoles of both nodes trying to connect.

It is the responsibility of the installation to terminate one of the conflicting targets. The link can then be connected from either side by using the CONNECT link operator command.

Special Exception (Local)

The Adabas ADARUN parameter LOCAL permits the use of duplicate target IDs under special circumstances. A target ID can be used on all nodes in the network simultaneously if *all* targets with that ID are declared LOCAL=YES (see the *Adabas Operations* documentation for the appropriate ADARUN parameter). In this case, any one of the targets with that ID can be accessed *only* by users running on the same node.

Entire Net-Work does not transport requests to LOCAL targets. Response code 148 is generated if a user program calls a target declared LOCAL from a remote node.

A target ID cannot be used simultaneously for both a local target and a global target. The reason for this restriction is illustrated in the following example.

Imagine that a (global) database, DBID 35, is active on node PROD, and user programs from all nodes in the network are issuing calls to it. If a local database, also with DBID 35, starts on node TEST, it immediately draws all calls that originate on node TEST to the local database DBID 35. Users on other nodes are not affected, but all users on node TEST are switched in mid-transaction from the global database DBID 35 to the local database DBID 35, a different and probably totally unrelated database. This error situation cannot be permitted.

Entire Net-Work enforces this rule and thus prevents the confusion that would result if a local target ID coincided with a global target ID by broadcasting information about local targets to all nodes the same way it does for global targets.

6

Useful ADARUN Parameters for Adabas Cluster Services

| | |
|--|----|
| ▪ CLOGMRG Parameter: Cluster Command Log Merge | 34 |
| ▪ CLUCACHEEXTRA Parameter: Adjust Shared Cache Vector Size | 35 |
| ▪ CLUCACHENAME Parameter: Cluster Cache Structure/Area Name | 35 |
| ▪ CLUGROUPNAME Parameter: Cluster Group Name | 36 |
| ▪ CLULOCKNAME Parameter: Cluster Lock Structure/Area Name | 37 |
| ▪ CLUPUBLPROT Parameter: Publishing of Protection Data | 38 |
| ▪ CLUSTER Parameter: Adabas Nucleus Cluster Session Control | 39 |
| ▪ CLUWORK1CACHE Parameter: Number of WORK Blocks in Global Cache | 39 |
| ▪ DIRRATIO / ELEMENTRATIO Parameters: Cluster Cache Directory-to-Data Elements Ratio | 41 |
| ▪ LDEUQP Parameter: Length of Unique Descriptor Pool | 42 |
| ▪ LRDP Parameter: Length of the Redo Pool | 43 |
| ▪ MLWTO Parameter: Multi-Line Console Message Control | 44 |
| ▪ MXCANCEL Parameter: Timeout Threshold for Canceled Peer Nucleus | 44 |
| ▪ MXCANCELWARN Parameter: Timeout Threshold for Canceled Peer Nucleus Warning | 45 |
| ▪ MXMSG Parameter: Timeout Threshold for Internucleus Command Processing | 46 |
| ▪ MXMSGWARN Parameter: Timeout Threshold for Internucleus Command Processing Warning | 47 |
| ▪ MXSTATUS Parameter: Member-Level XCF Status Monitoring Heartbeat Interval | 48 |
| ▪ MXWTOR Parameter: Self-Termination Operator Query Interval | 49 |
| ▪ NUCID Parameter: Adabas Cluster Nucleus ID | 50 |
| ▪ PROGRAM Parameter: Program to Run | 51 |
| ▪ SMFDETAIL Parameter: Selecting Detail SMF Record Sections | 52 |
| ▪ UPDATECONTROL Parameter: Control Scheduling of Update Commands During Buffer Flush | 54 |
| ▪ V64BIT Parameter: 64-Bit Virtual Storage Option | 55 |

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

CLUCACHENAME Parameter: Cluster Cache Structure/Area Name

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

For Adabas Cluster Services, the CLUCACHENAME identifies the particular sysplex cache structure (defined in the CFRM policy) associated with a particular Adabas sysplex nucleus cluster.

The same name must be set in each nucleus in the cluster. CLUCACHENAME must be unique for each cluster; that is, each database.

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



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

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

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



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

Example

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

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

CLUGROUPNAME 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
```

CLULOCKNAME Parameter: Cluster Lock Structure/Area Name

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

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

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



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

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

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



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

Example

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

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

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 Cluster 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.

MXSTATUS Parameter: Member-Level XCF Status Monitoring Heartbeat Interval

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

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



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

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

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

XCF status monitoring provides a second method by which Adabas Cluster Services can warn you that a cluster nucleus might be unable to respond in a timely way to intracluster commands. The first method is, of course, via the normal intracluster communication that occurs between cluster members. If a nucleus has heartbeat exceptions (as determined by XCF status monitoring),

it most likely will be unable to process and respond to an intracluster command; if a nucleus is slow to respond to an intracluster command, it might or might not have a heartbeat monitor exception (a nucleus may appear to be active to XCF but be unable to respond to an intracluster command). If the ADARUN `MXMSGWARN` parameter for a cluster nucleus is nonzero, it produces warning messages (ADAX9B or ADAX9C) when intracluster communication with other nuclei in the cluster is too slow; likewise, when XCF status monitoring determines that a nucleus is missing its heartbeat updates, it produces warning messages (ADAX22 and ADAX04). You can use an automated mechanism set up at installation to raise an alert or take other appropriate action based on the existence of these messages, as they identify existing or potential problems in the cluster.

MXWTOR 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. |
| PARM | 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.



Notes:

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

2. The replication pool will be allocated in 64-bit virtual storage when the Event Replicator Server runs with ADARUN V64BIT=YES parameter.

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 ,LARGE PAGE=YES ,V64BIT=YES
```

7 Entire Net-Work Parameter Statements

| | |
|------------------------------|----|
| ▪ Conventions | 58 |
| ▪ XCF NODE Statement | 61 |
| ▪ XCF DRIVER Statement | 71 |
| ▪ XCF LINK Statement | 74 |

To communicate with other nodes, Entire Net-Work requires a definition of its own operating environment, access method information, and peer node characteristics. This is accomplished with the Entire Net-Work parameter statements: NODE, DRIVER, LINK.

The NODE statement specifies the global operating parameters for the Entire Net-Work node (note that all interregion communication information is specified by the ADARUN control statement described in the previous section).

Entire Net-Work control statements must always begin with the NODE statement, followed by one or more DRIVER statements (one for each driver type), and the related LINK statements. DRIVER and LINK statements may be in any order, as long as no LINK statement precedes its related DRIVER statement.

A DRIVER statement must be specified for each line driver. The DRIVER statement indicates the operational parameters for the line driver's access method and default values for the related links.

A LINK statement may be specified for each link to a remote node. The LINK statement specifies the driver-specific parameters to be used for the connection. Each LINK statement must be associated with a previously specified DRIVER statement.

Conventions

This section covers the following topics:

- [Syntax Conventions](#)
- [Syntax Rules](#)

Syntax Conventions

The following table describes the conventions used in syntax diagrams of Entire Net-Work statements.

| Convention | Description | Example |
|-----------------|---|--|
| uppercase, bold | Syntax elements appearing in uppercase and bold font are keywords. When specified, these keywords must be entered exactly as shown. | <div style="border: 1px solid black; padding: 5px; display: inline-block;"> DRIVER TCPI [DRVCHAR = driver-char #] </div> <p>The syntax elements DRIVER, TCPI, and DRVCHAR are Entire Net-Work keywords.</p> |

| Convention | Description | Example |
|--------------------------------|--|---|
| lowercase, italic, normal font | Syntax elements appearing in lowercase and normal, italic font identify items that you must supply. | <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>DRIVER TCPI [DRVCHAR = driver-char #]</p> </div> <p>The syntax element <i>driver-char</i> identifies and describes the kind of value you must supply. In this instance, you must supply the special character used to designate that an operator command is directed to the TCP/IP line driver, rather than to a specific link.</p> |
| underlining | <p>Underlining is used for two purposes:</p> <ol style="list-style-type: none"> 1. To identify default values, wherever appropriate. Otherwise, the defaults are explained in the accompanying parameter descriptions. 2. To identify the short form of a keyword. | <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>DRIVER TCPI [DRVCHAR = driver-char #]</p> </div> <p>In the example above, # is the default that will be used for the DRVCHAR parameter if no other record buffer length is specified.</p> <p>Also in the example above, the short version of the DRVCHAR parameter is D.</p> |
| vertical bars () | <p>Vertical bars are used to separate mutually exclusive choices.</p> <p>Note: In more complex syntax involving the use of large brackets or braces, mutually exclusive choices are stacked instead.</p> | <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>DRIVER TCPI API = { BS2 CNS EZA HPS OES }</p> </div> <p>In the example above, you must select BS2, CNS, EZA, HPS, or OES for the API parameter. There are no defaults.</p> |
| brackets ([]) | <p>Brackets are used to identify optional elements. When multiple elements are stacked or separated by vertical bars within brackets, only one of the elements may be supplied.</p> | <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>DRIVER TCPI [DRVCHAR = driver-char #]</p> </div> <p>In this example, the DRVCHAR parameter is optional.</p> |
| braces ({ }) | <p>Braces are used to identify required elements. When multiple elements are stacked or separated by vertical bars within braces, one and only one of the elements must be supplied.</p> | <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>DRIVER TCPI API = { BS2 CNS EZA HPS OES }</p> </div> <p>In this example, one of the following values is required for the API parameter: BS2, CNS, EZA, HPS, or OES.</p> |

| Convention | Description | Example |
|-------------------------------|---|--|
| other punctuation and symbols | All other punctuation and symbols must be entered exactly as shown. | <pre data-bbox="662 281 1338 312">LINK linkname TCPI [(NETADDR = n1.n2.n3.n4) [,][-]</pre> <p data-bbox="636 390 1378 422">In this example, the periods must be specified in the IP address.</p> <p data-bbox="636 453 1378 548">In addition, options must be separated by commas and dashes should be used as needed to indicate that parameter settings continue on the next line.</p> |

Syntax Rules

The following rules apply when specifying Entire Net-Work parameter statements:

- Each Entire Net-Work parameter statement occupies positions 1 - 72 of at least one line.
- The statement type (NODE, LINK, or DRIVER) must be specified as the first nonblank item on the statement.
- The node name, driver name, translation definition function, or link name follows the statement type, separated by at least one blank (space).
- Keyword parameters may be specified following either the node name on NODE statements or the driver name on DRIVER and LINK statements. Keyword parameters are separated from their arguments by an equal (=) sign, and from other keyword parameters by at least one blank (space) or a comma (,).
- When the acceptable values for a parameter are Y and N (yes and no), any other value is treated as an N, unless there is a documented default, and processing continues without any warning.
- When the acceptable values for a parameter fall within a range (e.g., 1 - 2147483647) and a value outside the range is specified, the value is automatically reset to the maximum value within the range, unless documented otherwise for the parameter. Processing continues without any warning.
- A statement can be continued beginning in any column of the next line by specifying a dash (-) as the last nonblank character in any column of the current line, before column 73.
- Comment lines begin with an asterisk (*) in position 1 and can be inserted anywhere in the statement sequence.
- Some keywords may require a list of subparameters separated by commas; the list must be enclosed in parentheses () unless only the first subparameter is to be entered. Omitted ("defaulted") subparameters must be represented by placeholder commas if subsequent parameters are to be entered. The following are examples of correct subparameter strings:

```
KEYWORD=(value1,value2,value3)
KEYWORD=(value1,,value3)
KEYWORD=(,value3)
KEYWORD=(,value2)
KEYWORD=value1
```

- Hexadecimal keyword values can be entered by prefixing the value with an "X". For example:

```
LINK . . . ADJID=X0064, . . .
```

XCF NODE Statement



Note: This section describes the `NODE` statement and parameters for the XCF line driver. For detailed information see the sections [XCF DRIVER Statement](#) and [XCF LINK Statement](#).

The `NODE` statement must be the first Entire Net-Work control statement. It defines the node's name and operating characteristics.

Format

The following is the format of the `NODE` statement:

```
NODE nodename BUFFERS = (abufsize,lbufsize,stbufsize,pxbufsize),-
CQTIMER = seconds,-
DEFINE = {N | Y},-
DOMAIN = domain-name,-
DUMP = storage-area,-
LOG = {N | Y},-
MAXPATH = link-count,-
MSGFORM = message-form,-
NID0 = {N | Y},-
NTRACE = n,-
REPLYTIM = seconds,-
START={Y | N},-
TIMER = seconds,-
TRACE = argument-list,-
UCMSG={ N | Y},-
ULINK={ N | Y}
```

Parameters

The `NODE` statement parameters along with a description of their function and their accepted values and defaults are described in the following sections.



Note: The underlined portion of the parameter is the minimum abbreviation.

- nodename
- BUFFERS
- CQTIMER
- DEFINE
- DOMAIN
- DUMP
- LOG
- MAXPATH
- MSGFORM
- NID0
- NTRACE
- REPLYTIM
- START
- TIMER
- TRACE | TROFF | TRON
- UCMSG
- ULINK

nodename

A 1-8 character name for this node. The node name must be specified immediately after `NODE`, separated by at least one blank space. All Entire Net-Work nodes must have unique names. Choose a meaningful name. Entire Net-Work uses the node name when referring to the node for operator messages and when logging.

BUFFERS

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|-----------------------|----------|
| BUFFERS | the amount of working storage (in bytes) reserved for internal buffer management. | 0 | 2147483647 (see text) | see text |

The values required depend on the number and types of line drivers and links, and on the message volume; these values can be adjusted according to usage statistics printed at the end of each Entire Net-Work session, or from issuing the operator command "D S" (display statistics).

Specify decimal values for the operands ranging from 0 to 2147483647 bytes; or, optionally, a value followed by either the multiplier "M" (multiply by 1048576) or "K" (multiply by 1024). If followed by "M", the specified value must range from 0 to 2048; if followed by "K", the value must range

from 0 to 2097151. If a subparameter is omitted or zero is specified, no corresponding buffer pool will exist for this node. The requested storage space must be available; if the space is not available, Entire Net-Work ends with error message NET0013.

For OS/390 systems, the buffer pool manager initializes the common buffer pool with a subpool of 256K. Additional subpools can be created dynamically, and all subpools can be expanded or contracted as needed. The only size limitation for the buffer pool is the size of the region or partition. The `BUFFERS` specification on the `NODE` statement remains the same, even though the first three values (*abufsize*, *ltbufsize*, and *stbufsize*) are ignored by the new buffer pool manager. The fourth value (*pfxbufsize*) is used to set the size of the page-fixed buffer pool.

```
BUFFERS= (abufsize ,ltbufsize,stbufsize,{pfxbufsize | 0 } )
```

where

| | |
|-------------------|--|
| <i>abufsize</i> | is the asynchronous buffer pool size. This required subparameter specifies the storage space to reserve for the asynchronous buffer pool; used whenever asynchronously scheduled routines (exit routines, interrupt handlers) require storage. All requests to this pool are serialized. |
| <i>ltbufsize</i> | is the long-term buffer pool size. This required subparameter specifies the storage space to reserve for the long-term buffer pool, from which comparatively small buffers are allocated for relatively long-term use (for example, dynamically created control blocks). |
| <i>stbufsize</i> | is the short-term buffer pool size. This required subparameter specifies the storage space to reserve for the short-term buffer pool, from which comparatively large individual buffers are allocated for short time periods (for example, message and transmission buffers). |
| <i>pfxbufsize</i> | is the page-fixed buffer pool size. This optional subparameter specifies the bytes to reserve for the page-fixed buffer pool, from which all requests for page-fixed buffers are filled. The use of this pool is similar to that of the short-term buffer pool; that is, relatively large buffers are allocated for short time periods. The default '0' (zero) means that no page-fixed buffer pool is allocated. Because the segment size of the page-fixed buffer pool is identical to the page size of the operating system, more page-fixed storage than short-term storage may be required. |

CQTIMER

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|------------|---------|
| CQTIMER | the time allowed (in seconds) for a user to retrieve command results with a router 16-call. | 1 | 2147483647 | 60 |

Specify a practical decimal value, depending on the node system's environment. The maximum value is approximately 68 years: effectively, no timeout will occur. The default is approximately one minute.

The purpose of the CQTIMER= timeout is to prevent an Entire Net-Work request queue element (RQE) and the attached buffer from becoming irretrievable if the user has ended abnormally. This parameter performs the same function as the ADARUN CT parameter.



Note: This parameter can be changed during Entire Net-Work operation by the SET CQTIMER= command.

DEFINE

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| DEFINE | whether the DEFINE operator command can be used to define links during Entire Net-Work operation. | N Y | N |

If "Y" is specified, the DEFINE operator command is accepted and executed; "N" (the default) rejects the DEFINE operator command.

DOMAIN

| Parameter | Specify... | Possible Values | Default |
|-----------|---------------|--------------------|------------------------|
| DOMAIN | a domain name | 1-6 character name | blank (no domain name) |

This parameter allows you to subdivide the network into multiple domains. Using domains simplifies network management and limits administrative message traffic.

DUMP

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|---------|
| DUMP | the areas of storage to be printed after an abnormal termination of Entire Net-Work. | see text | ALL |

The information is printed to the NETPRINT file if it is open. Otherwise, it is printed to the DDPRINT file. The DUMP parameter can be used to reduce the amount of output generated during an ABEND, especially on large Entire Net-Work systems. This parameter cannot be abbreviated.

DUMP= { ALL | NONE | BLOCKS | TRACETAB | BUFFERS | LINKAREA | FORMAT }

where

| | |
|----------|--|
| ALL | All storage areas are dumped. This is the default value. |
| NONE | No storage areas are dumped. |
| BLOCKS | The major control blocks are dumped. |
| TRACETAB | The internal trace table is dumped. |
| BUFFERS | All internal buffer areas are dumped. |
| LINKAREA | All storage areas related to a driver and link are dumped. |
| FORMAT | The driver and link trace tables are formatted. |

In general, the default value of ALL should be used so that all diagnostic information is available to Software AG support.

Multiple values can be specified, separated by commas and surrounded by parentheses:

```
DUMP = (BLOCKS, TRACETAB, FORMAT)
```

If values conflict, the last value specified is used. 'NONE' is used in the following example:

```
DUMP = (BLOCKS, TRACETAB, NONE)
```



Note: This parameter can be changed during Entire Net-Work operation by the SET DUMP command.

LOG

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| LOG | whether control flow is recorded and/or selected data areas are logged to the printer data set. | see text | NO |

Normally, logging should not be used because of the extra system resources required. The LOG function is intended primarily as a diagnostic tool: use it only with the assistance of your Software AG technical support representative.

LOG is a test parameter for recording control flow and for logging selected data areas. The information is written to the NETPRNT file if it is open. Otherwise, it is written to the DDPRINT file. It can have the following values:

LOG= { ON | OFF | YES | NO | FULL | SHORT }

where

| | |
|-------------------|---|
| OFF, NO | No logging occurs during this Entire Net-Work node's session. |
| ON, YES, and FULL | Log both the node's checkpoint records and data areas. |
| SHORT | Log the checkpoint records only. |

 **Note:** This parameter can be changed during Entire Net-Work operation by the SET LOG= command.

MAXPATH

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|--|---------|---------|---------|
| MAXPATH | the maximum path length (in links) that a message can travel in the network. | 1 | 32767 | 4 |

The default (4) results in a stack large enough for four en route node IDs.

The Entire Net-Work communicator uses this optional value to build a list of two-byte entries for tracking each message. This list, called a 'node stack', is included in the message header. As the message passes through nodes on the way to its target, each node's ID is added to the stack.

If the specified MAXPATH value results in a node stack that is larger than needed, messages will be unnecessarily long. If the MAXPATH value is too small, Entire Net-Work automatically copies the message, increasing the node stack size; this causes unnecessary processor overhead.

 **Note:** This parameter can be changed during Entire Net-Work operation by the SET MAXPATH= command.

MSGFORM

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|---------|
| MSGFORM | the format of console messages and DDPRINT output. | 1 2 3 4 | 2 |

where the valid values are

| | |
|---|---|
| 1 | is compatible with the message format used by Entire Net-Work Version 5.2. For example: <code>NET0090 BUFFER USAGE STATISTICS</code> |
| 2 | provides a severity letter (I for information, W for warning, or E for error) with the message number. For example: <code>NET0090I: BUFFER USAGE STATISTICS</code> |
| 3 | provides the message number followed by the node name of the issuing Entire Net-Work node, padded with blanks to a length of 8. For example: <code>NET0090I NODE2 : BUFFER USAGE STATISTICS</code> |
| 4 | provides the message number followed by the node name, not padded, of the issuing Entire Net-Work node. For example: <code>NET0090I NODE2: BUFFER USAGE STATISTICS</code> |

**Notes:**

1. Individual line drivers may not recognize this parameter.
2. Values between 5 and 255 will be accepted, but they have no meaning and are not valid.
3. This parameter can be changed during Entire Net-Work operation by the SET MSGFORM= command.

NIDO

| Parameter | Specify... | Possible Values | Default |
|-------------|---|-----------------|---------|
| <u>NIDO</u> | whether to force a node ID of 0 (zero) for all unsolicited connections. | N Y | N |

If NIDO=Y is specified, all nodes that attempt to connect and are not explicitly defined are assigned a node ID of 0. No Adabas servers on those nodes are broadcast through the network.

NTRACE

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|------------|---------|
| NTRACE | the number of 32-byte entries to be allocated for the internal trace table. | 0 | 2147483647 | 0 |

Specify a practical decimal value, depending on the node's system configuration. The maximum may be limited by the system. If the default (0) is used, no trace table is allocated for this node and no tracing takes place, whether TRACE= or TRON= is specified or not.

REPLYTIM

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|------------|---------|
| REPLYTIM | the time allowed (in seconds) for a user request to complete. | 1 | 2147483647 | 60 |

A request is considered complete when the originating node receives a reply.

Specify a practical decimal value, depending on the node system's operation. The maximum value is approximately 68 years: effectively, no timeout will occur. The default is approximately one minute.

In the event that a message is "stranded" (that is, a reply cannot be returned to the originating node), REPLYTIM= specifies a time after which a response code 224 is returned to the user.



Note: This parameter can be changed during Entire Net-Work operation by the SET REPLYTIM= command.

START

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| START | whether Entire Net-Work starts normal operations automatically. | N Y | Y |

The values for START are interpreted as follows:

| | |
|---|--|
| Y | Entire Net-Work automatically starts all line drivers and initiates connections for all links that have the parameter ACQUIRE=Y specified. |
| N | Entire Net-Work initializes line drivers but does not start them; nor does it connect any links. Line drivers can be started individually by using the START operator command. |

TIMER

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|--|---------|------------|---------|
| TIMER | the interval (in seconds) between handling of time-dependent requests. | 1 | 2147483647 | 30 |

Each time the specified interval expires, Entire Net-Work scans its tables for any time-dependent action that needs to be taken. The `TIMER=` value determines the precision of *all* time-dependent Entire Net-Work services.

Specify a practical decimal value depending on the node operation. The maximum value is approximately 68 years: effectively, no timing supervision will occur, even if other timing parameters, such as `REPLYTIM`, `CQTIMER`, or `ADARUN CT` are set.

There is an interaction between `TIMER=` and other timing parameters. If the `TIMER` interval is greater than the individual `CQTIMER=` and `REPLYTIM=` intervals, the specified action may not be started until the `TIMER=` interval has expired.

TRACE | TROFF | TRON

| Parameter | Specify... | Possible Values | Default |
|----------------------|---|------------------|---------|
| TRACE TROFF TRON | trace control parameters for performing program traces. | trace (see text) | none |

Tracing should not be active during normal operation. Tracing is intended as a diagnostic tool; it is recommended that you use tracing only with the assistance of your Software AG technical support representative.

`TRACE=` and `TRON=` are synonyms to either start or resume tracing of the specified events. `TROFF=` stops tracing. If this parameter is not specified, no tracing will occur.

Note that the `NTRACE` parameter must specify a nonzero value in order to permit tracing.

`{TRACE | TROFF | TRON} = { trace | (trace, ...)}`

One or more trace values can be specified. If more than one value is specified, the values must be enclosed in parentheses and separated by commas:

```
(trace,trace,... )
```

The value 'trace' can be one or more of the following:

| | |
|------|----------------------------------|
| MAIN | mainline trace |
| BPM | buffer pool manager trace |
| TQM | transmission queue manager trace |
| RQM | receive queue manager trace |
| XCF | XCF line driver trace |

An XCF line driver trace can be requested for installed line drivers on the local node only.

 **Note:** Values set by this parameter can be changed during Entire Net-Work operation by the SET TRACE=, TROFF=, or TRON=. . . commands.

UCMSG

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| UCMSG | whether messages are issued in uppercase (Y) or mixed case (N). | N Y | N |

 **Notes:**

1. Individual line drivers may not recognize this parameter.
2. This parameter can be changed during Entire Net-Work operation by the SET UCMSG command.

ULINK

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| ULINK | whether multiple links are allowed between two Entire Net-Work nodes. | N Y | N |

The values for ULINK are interpreted as follows:

| | |
|---|---|
| N | (the default) multiple links between two Entire Net-Work nodes are allowed. |
| Y | Entire Net-Work ensures that each connection to an adjacent node is unique; incoming connection requests from adjacent nodes that are already known as active are rejected. |

In networks with many PCs, two PCs may be assigned the same node name and ID by mistake. If both PCs are simultaneously connected to Entire Net-Work, they are perceived as one Entire Net-Work node that is connected by two different links. As a result, one of the PCs may receive a reply to a call that originated on the other PC.

To avoid this type of situation, specify ULINK=YES. When the second PC tries to connect, it is rejected. The integrity of the network is maintained and the duplicate node name and ID can be identified.

 **Note:** This parameter can be changed during Entire Net-Work operation by the SET ULINK command.

XCF DRIVER Statement

The Entire Net-Work DRIVER control statement defines the line driver type (that is, XCF) to be loaded. Only one DRIVER statement may be specified for a given line driver type.

DRIVER statement parameters are read from a sequential file during system startup, and can be modified after startup using the ALTER operator command. Some parameters can be modified when the line driver is open or closed; others can be modified only when the line driver is closed. See the operator commands ALTER and CLOSE in *XCF Line Driver Commands*, elsewhere in this guide. The open/closed requirement for each parameter is included in its description.

The XCF DRIVER statement and its parameters are used to activate and define the characteristics of the local sysplex node. The access method name 'XCF' or 'XCFD' instructs Entire Net-Work to load the line driver module NETXCF, which must exist in the libraries defined for the related job step.



Note: The software accepts both XCF and XCFD as the access method name. However, the name used in the DRIVER statement must also be used for the LINK statements.

This section covers the following topics:

- [Format](#)
- [Parameters](#)

Format

The XCF DRIVER statement has the following format:

```

DRIVER XCFD ACCEPTUI = { Y | N },-
GROUP = group-name,-
LARGMSG = { message-size | 8192 },-
PSTATS = { Y | N },-
RSTATS = { Y | N },-
SMALLMSG = { message-size | 1024 },-
STATINT = { statistics-interval | 3600 },-
TRACESIZ = { internal-trace-buffer-size | 8192 }

```

where 'XCFD' is the four-character access method name for the XCF driver type:

| Driver | Module Name | Access Method |
|--------|-------------|---|
| XCFD | NETXCF | parallel sysplex cross-system coupling facility |

Parameters



Note: The underlined portion of the parameter is the minimum abbreviation.

- ACCEPTUI
- GROUP
- LARGMSG
- PSTATS
- RSTATS
- SMALLMSG
- STATINT
- TRACESIZ

ACCEPTUI

| Parameter | Specify... | Possible Values | Default |
|-----------------|--|-----------------|---------|
| <u>ACCEPTUI</u> | whether this node will connect to other nodes in the XCF group that are not defined by a LINK statement. | N Y | Y |

The ACCEPTUI parameter can be modified when the driver is open or closed.

GROUP

| Parameter | Specify... | Possible Values | Default |
|--------------|---------------------------------|-----------------|----------|
| <u>GROUP</u> | the Entire Net-Work group name. | group-name | WCPXCFV5 |

All Entire Net-Work nodes in the SYSPLEX must use the same unique XCF group name. The group name must be the same on all Entire Net-Work nodes; it must also be different from the name of any other product or system using XCF. Results are unpredictable if Entire Net-Work nodes are assigned to the same group as other types of systems such as DB2 or CICS. The GROUP parameter specifies the Entire Net-Work group name, which may be defined in the COUPLE_{xx} member in SYS1.PARMLIB (see step *Modify COUPLE_{xx} and the CFRM Policy* in the *Installation* documentation). The GROUP parameter can be modified only when the driver is closed.

LARGMSG

| Parameter | Specify... | Default |
|-----------|--|-----------|
| LARGMSG | the minimum size (in bytes) of a large message for statistics reporting. | 8192 (8K) |

Any data messages larger than this size are counted as large messages. Any data messages smaller than this size are counted as medium or small messages, depending on the value specified by the SMALLMSG parameter described below.

For example, using the default values for SMALLMSG and LARGMSG:

- Data messages between 1 and 1024 bytes long are reported as small messages.
- Data messages between 1025 and 8192 bytes long are reported as medium messages.
- Data messages larger than 8192 bytes long are reported as large messages.

PSTATS

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|---------|
| PSTATS | whether statistics are printed automatically at regular intervals. | N Y | N |

This parameter is only used to set the default for all links.

RSTATS

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|---------|
| RSTATS | whether statistics are reset automatically at regular intervals. | N Y | Y |

This parameter is only used to set the default for all links.

SMALLMSG

| Parameter | Specify... | Default |
|-----------|--|---------|
| SMALLMSG | the maximum size (in bytes) of a small message for statistics reporting. | 1024 |

Any data messages smaller than or equal this size are counted as small messages. Any data messages larger than this size are counted as medium or large messages, depending on the value specified by the LARGMSG parameter described above.

STATINT

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|--|---------|------------|---------|
| STATINT | the amount of time, in seconds, before statistics are automatically reported or reset. | 0 | 2147483647 | 3600 |

This parameter is only used to set the default for all links.

TRACESIZ

| Parameter | Specify... | Default |
|-----------|---|-----------|
| TRACESIZ | the size (in bytes) of the internal trace buffer. | 8192 (8K) |

The TRACESIZ parameter can be modified when the driver is open or closed.

XCF LINK Statement

The LINK statement is currently not required for the XCF DRIVER statement. During the initialization process, the XCF line driver requests a list of all other members in its XCF group and then, if the DRIVER statement parameter ACCEPTUI=Y is set, logically connects to each member.

Although the LINK statement is optional, it can be added for documentation or security purposes.



Note: If two or more Entire Net-Work nodes are active in the sysplex within the same group, the Entire Net-Work operator command `D L` (display links) can be used to display the list of dynamically added links and parameter-defined links.

Each link to another node can be defined with a LINK statement. Each link uses the XCF communications access method as defined on an XCF DRIVER statement. LINK statements specify the XCF driver by name, and follow the XCF DRIVER statement in the Entire Net-Work statement order.



Note: The software accepts both XCF and XCFD as driver names. However, the name used in the DRIVER statement must also be used for the LINK statements.

The LINK statement parameters are read from a sequential file during system startup, and can be modified after startup using the ALTER operator command. Some parameters can be modified when the link is open or closed; others can be modified only when the link is closed. See the operator commands ALTER and CLOSE in *XCF Line Driver Commands*, elsewhere in this guide. The open/closed requirement for each parameter is included in its description.

- [Format](#)

- Parameters

Format

The XCF LINK statement has the following format:

```
LINK linkname XCFD PSTATS = { Y | N },-
                   RSTATS = { Y | N },-
                   SAF = { Y | L | N },-
                   STATINT = statistics-interval,-
                   WEIGHT = { n | 256 }
```

where 'XCFD' is the four-character name of the XCF line driver. The XCF DRIVER statement must precede the LINK statement.

Parameters



Note: The underlined portion of the parameter is the minimum abbreviation.

- linkname
- PSTATS
- RSTATS
- SAF
- STATINT
- WEIGHT

linkname

This required parameter specifies the name by which this link is to be known; it is a one- to eight-character name identifying the link and must be unique on the node. It is positional and must be specified immediately after the LINK keyword and immediately before the driver name (XCFD).



Note: If more than 8 characters are entered for linkname, only the first 8 characters are used. The connection is issued correctly and no error message is generated.

PSTATS

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|----------|
| PSTATS | whether statistics are printed automatically at regular intervals. | N Y | see text |

The default value is the value of the PSTATS parameter on the DRIVER statement (see *XCF DRIVER Statement* parameter [PSTATS](#)).

RSTATS

| Parameter | Specify... | Possible Values | Default |
|-----------|--|-----------------|----------|
| RSTATS | whether statistics are reset automatically at regular intervals. | N Y | see text |

The default value is the value of the PSTATS parameter on the DRIVER statement (see *XCF DRIVER Statement* parameter [RSTATS](#)).

SAF

| Parameter | Specify... | Possible Values | Default |
|-----------|---|-----------------|---------|
| SAF | whether Entire Net-Work will call the SAF Interface for all incoming requests on this link. | N L Y | N |

Failure to load the Interface is considered a security violation and Entire Net-Work will shut down. If SAF=L, the calls are traced and the output directed to DDPRINT. An error code is transmitted to the user if access to SAF is denied. The SAF parameter can be modified when the link is open or closed.

STATINT

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|--|---------|------------|----------|
| STATINT | the elapsed time (in seconds) before statistics are automatically reported or reset. | 0 | 2147483647 | see text |

The default value is the value of the STATINT parameter on the DRIVER statement (*XCF DRIVER Statement* parameter [STATINT](#)).

WEIGHT

| Parameter | Specify... | Minimum | Maximum | Default |
|-----------|---|---------|---------|---------|
| WEIGHT | the weight of this link with respect to other links going to the same node. | 1 | 999999 | 256 |

If a given target can be reached by more than one path (chain of connected links), the path with the lowest weight is used. Slow or expensive links should be given a higher value than fast or inexpensive links. The WEIGHT parameter can be modified only when the link is closed.

8

Adabas Cluster Operator Commands

- ADACOM Operator Commands 80
- Adabas Cluster Nucleus Operator Commands 87

This part of the documentation provides information on the operator commands for ADACOM and for an Adabas 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 covers the following topics:

- [ADACOM Environment](#)
- [ADACOM Command Syntax and Processing](#)
- [Controlling Subtasks](#)
- [Controlling PLXCBs](#)
- [DIM - Display Images](#)
- [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.

Cluster Services (ALS) nuclei do not require ADACOM SVC/DBID subtasks to be active for normal operation. However, operator commands other than those that start the subtasks or terminate ADACOM cannot be issued unless the subtasks are active.

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. If there is no PLXCB, the first Cluster Services nucleus will allocate one using the nucleus ADARUN NU value and the defaults for the other parameters.

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. It is also done automatically if, when the first Cluster Services nucleus starts, there are no ADACOM SVC/DBID subtasks running and NU is different from the existing 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 Images



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

This command is serviced by an SVC/DBID subtask.

Use DIM to display the number of commands processed and the number of currently active users for each active nucleus on every image known to the local image.

If the image name is specified, DIM displays the information for the specified image only.

Sample Output

```

PLI060 SVC=svc DBID=dbid OPERATOR COMMAND:DIM
PLI004 image-name NUCID UP LO RO -#USERS- -#CMNDS- LURA= RULA=
      jobname      00001 Y  Y  N  00000152 00000001
PLI007 image-name NUCID UP LO RO -#USERS- -#CMNDS- LURA= RULA=
      jobname      00002 Y  ** NN 00000089 00000000
    
```

Nuclei on the local image start open; nuclei on remote images start closed to the local image. The display uses the following indicators:

| | |
|------------|--|
| image-name | the name of the image: in PLI004 displays, this is the local image; in PLI007 displays, this is a remote image |
| jobname | the job name of a cluster nucleus active on that image |
| NUCID | the 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 on the local image and open for new users (Y); on the local image and closed for new users (N); or not on the local image (**) |
| RO | whether the specified local nucleus (PLI0004 message) is closed for new users from remote images (N) or open for new users from all images (G); or whether the specified remote nucleus (PLI007 message) is closed for new users from remote (as seen from that nucleus) images (NN), open for new users from this image (where DIM was issued) but not necessarily other images (LN), or open for new users from all images (NG and LG) |
| #USERS | the number of users currently assigned to the specified nucleus |
| #CMNDS | the number of commands currently in progress in the specified nucleus |
| LURA= | the number of local users currently assigned to a remote image |
| RULA= | the number of remote users currently assigned to the local image |

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 on the local image only.

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  00000000 00000000

```

Nuclei on the local image start open. The display uses the following indicators:

| | |
|------------|--|
| image-name | the name of the image: in PLI004 displays, this is the local image name |
| jobname | the job name of the cluster nucleus active on the local image |
| NUCID | the 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 local nucleus is open (Y) or closed for new users (N) |
| RO | whether the specified nucleus is closed for new users from remote images (N) or open for new users from all images (G) |
| #USERS | the number of users currently assigned to the specified nucleus |
| #CMNDS | the number of commands currently in progress in the specified nucleus |
| LURA= | the number of local users currently assigned to a remote image |
| RULA= | the number of remote users currently assigned to the local image |

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, IDTHPRFX, 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 REMOTE ALL: all nuclei on all remote images |
| 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 sysplex cluster users |

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

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

By default, nuclei start open to local users and closed to remote users.

After the nuclei start, the SN operator commands may be used to exercise some control over the assignment of users to nuclei. For example, in a node with no active local nuclei, it is necessary to open at least one remote nucleus to users from the local node in order to run user programs in that node.

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 *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 Cluster 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 |
| MXSTATUS | Dynamically change the ADARUN MXSTATUS 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 Cluster 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 | The cluster member failure status, as a combination of one or more of the following values: <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 | The cluster XCF status monitoring summary, as a combination of one or more of the following values: <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=nnnnn | 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=nucid | is the ID of the nucleus that owns the file lock, if the file is locked. |
| CA=nnnnn | is the number of users on this nucleus who are currently accessing this file. |
| CU=nnnnn | 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:

```
15:54:43 ADAI29 Oper cmd: DPARM
15:54:43 ADAN16 00035 2016-11-29 15:54:42 AOSLOG=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 ARNWORKBUFFERS=5
15:54:43 ADAN16 00035 2016-11-29 15:54:42 ASSOSPACEWARN=0
15:54:43 ADAN16 00035 2016-11-29 15:54:42 ASYTVS=YES
15:54:43 ADAN16 00035 2016-11-29 15:54:42 AUTOCQENV=ALL
15:54:43 ADAN16 00035 2016-11-29 15:54:42 AUTOCQTIME=0
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLOGMRG=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLUCACHENAME=SAG1_CACHE
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLUGROUPNAME=SAG1CLU
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLULOCKNAME=SAG1_LOCK
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLUPUBLPROT=YES
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLUSTER=SYSPLEX
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CLUWORK1CACHE=5
15:54:43 ADAN16 00035 2016-11-29 15:54:42 CT=900
15:54:43 ADAN16 00035 2016-11-29 15:54:42 DATASPACEWARN=0
15:54:43 ADAN16 00035 2016-11-29 15:54:42 DIRRATIO=4
15:54:43 ADAN16 00035 2016-11-29 15:54:42 ELEMENTRATIO=1
15:54:43 ADAN16 00035 2016-11-29 15:54:42 FMXIO=1
15:54:43 ADAN16 00035 2016-11-29 15:54:42 INDEXCROSSCHECK=YES
15:54:43 ADAN16 00035 2016-11-29 15:54:42 INDEXUPDATE=ADVANCED
15:54:43 ADAN16 00035 2016-11-29 15:54:42 INTNAS=3,600
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LARGEPAGE=NO
```

```
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LBP=104,858,624
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LCP=10,000
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LDEUQP=7,000
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LFIOP=90,000
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LFP=100,000
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LI=4,000
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGCB=YES
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGFB=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGGING=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGIB=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGIO=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGMB=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGRB=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGSB=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGUX=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGVB=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGVOLIO=NO
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LOGWARN=0
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LP=1,200
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LQ=30,000
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LRDP=10,485,760
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LS=19,968
15:54:43 ADAN16 00035 2016-11-29 15:54:42 LU=655,350
15:54:44 ADAN16 00035 2016-11-29 15:54:42 LWP=5,242,880
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXCANCEL=6,000
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXCANCELWARN=1,500
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXMSG=6,000
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXMSGWARN=1,500
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXSTATUS=6,000
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXTNA=3,600
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXTSX=3,600
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXTT=3,600
15:54:44 ADAN16 00035 2016-11-29 15:54:42 MXWTOR=0
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NC=500
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NCLOG=0
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NH=8,000
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NISNHQ=1,999
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NONDES=YES
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NPLOG=0
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NPLOGBUFFERS=5
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NQCID=250
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NSISN=100
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NT=20
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NU=500
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NUCID=10,001
15:54:44 ADAN16 00035 2016-11-29 15:54:42 NWORK1BUFFERS=50
15:54:44 ADAN16 00035 2016-11-29 15:54:42 READONLY=NO
15:54:44 ADAN16 00035 2016-11-29 15:54:42 REFSTPRT=YES
15:54:44 ADAN16 00035 2016-11-29 15:54:42 REVIEW=NO
15:54:44 ADAN16 00035 2016-11-29 15:54:42 RIAFTERUPDATE=NO
15:54:44 ADAN16 00035 2016-11-29 15:54:42 SECUID=ALLOW
15:54:44 ADAN16 00035 2016-11-29 15:54:42 SMF=NO
```

```

15:54:44 ADAN16 00035 2016-11-29 15:54:42 TLSCMD=20
15:54:44 ADAN16 00035 2016-11-29 15:54:42 TNAA=1,800
15:54:44 ADAN16 00035 2016-11-29 15:54:42 TNAE=1,800
15:54:44 ADAN16 00035 2016-11-29 15:54:42 TNAX=1,800
15:54:44 ADAN16 00035 2016-11-29 15:54:42 TT=1,800
15:54:44 ADAN16 00035 2016-11-29 15:54:42 UPDATECONTROL=NODELAY
15:54:44 ADAN16 00035 2016-11-29 15:54:42 UTIONLY=NO
15:54:44 ADAN16 00035 2016-11-29 15:54:42 V64BIT=NO
15:54:44 ADAN41 00035 2016-11-29 15:54:42 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-I0 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

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

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 *External Cache Statistics (Cluster Nucleus Only)*, in the *Adabas Cluster 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 *External Cache Statistics (Cluster Nucleus Only)*, in the *Adabas Cluster Services Operations Guide*.

DXLOCK Command: Display Lock Statistics



DXLOCK

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 *External Lock Statistics (Cluster Nucleus Only)*, in the *Adabas Cluster Services Operations Guide*.

DXMSG Command: Display Messaging Performance Statistics



DXMSG

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 *External Cache Statistics (Cluster Nucleus Only)*, *External Lock Statistics (Cluster Nucleus Only)* and *WORK Data Set I/O Statistics (Cluster Nucleus Only)*, in the *Adabas Cluster 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



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*.

MXSTATUS Command: Dynamically Change MXSTATUS Parameter Setting



The MXSTATUS command can be used to dynamically change (while the Adabas nucleus is running) the setting of the MXSTATUS ADARUN parameter for a cluster nucleus. This command can be used to specify the number of seconds between member-level XCF status monitoring heartbeats. Valid values range from 15 through 21474836.



Note: The MXSTATUS command is only used for Adabas Cluster Services and not for Adabas Parallel Services. Adabas Parallel Services does not use XCF and ignores the setting.

The MXSTATUS command can only be used to change the value of the MXSTATUS parameter if the ADARUN MXSTATUS parameter is not zero (0). This command cannot be used to turn status monitoring on (by specifying a nonzero value if the ADARUN MXSTATUS setting is zero) or to turn it off (by specifying a value of zero when the ADARUN MXSTATUS setting is nonzero). However, setting MXSTATUS to a very high value is effectively the same as turning status monitoring off (monitoring will occur very infrequently at least).

For more information about the MXSTATUS setting for a nucleus, read about the ADARUN MXSTATUS 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.

9 Entire Net-Work Operator Commands

- Overview of Entire Net-WorkCommands for z/OS Environments 113
- Command Descriptions 114
- XCF Line Driver Commands 127

This section contains operator commands for Entire Net-Work, and for its XCF line driver.

Although Entire Net-Work operates automatically, there are operator commands available during execution to display or modify the status of the network and control the local Entire Net-Work node.

In OS/390 environments, the following operator commands are used:

```
MODIFY  
F
```

```
STOP  
P
```

The STOP (or P) command serves as an alternative to the NETEND command and terminates Entire Net-Work. Its synonyms are described as part of the following example:

Entering the following long form MODIFY commands results in the following status displays (assuming that NETWK is the name of the started task running Entire Net-Work):

```
MODIFY NETWK,D STATS  
NET0090I: BUFFER USAGE STATISTICS:  
NET0091I: ASYNCH. BUFFERS:      000016 (= 24.2 %) OF 000064 K USED  
NET0091I: LONG TERM BUFFERS:   000000 (=  0.4 %) OF 000064 K USED  
NET0091I: SHORT TERM BUFFERS:  000000 (=  6.1 %) OF 001025 K USED  
NET0091I: ATTACHED BUFFERS:    000000 (= 11.9 %) OF 000080 K USED  
NET0091I: REQUEST QUEUE:       000000 (=  6.0 %) OF 000050 K USED  
NET0087I: 0000010847 REQUESTS FROM LOCAL RQ
```

The following two commands are equivalent ways to terminate the Entire Net-Work session:

```
F NETWK,NETEND  
P NETWK
```

The Entire Net-Work commands described in this section are similar to Adabas operator commands. The underlined portion of the command is the minimum abbreviation.

Entire Net-Work line drivers have the ability to process operator commands that are directed to a specific link or directly to the driver. The operator commands that are specific to the XCF line driver are described in the section [XCF Line Driver Commands](#).

Overview of Entire Net-Work Commands for z/OS Environments

The following table summarizes the Entire Net-Work operator commands for z/OS:


| Command | Argument | Action |
|---------------|-----------|--|
| ADAEND | | Terminate Entire Net-Work session. |
| CLOSE | driver | Disconnect all links of a driver, then close the driver. |
| CLOSE NETPRNT | | Close the NETPRNT file and route all trace and snap output to DDPRINT. |
| CONNECT | link | Reconnect a link after a disconnect or handshake error. |
| DEFINE | link | Dynamically define a new link. |
| DISABLE | link | Disable a link (link cannot accept connects). |
| DISCONNECT | link | Disconnect a link. |
| DISPLAY | parameter | Display link, nodes, targets, paths, or statistics. |
| DUMP | | Snap data areas, then terminate the Entire Net-Work session. |
| ENABLE | link | Enable a link (the link can accept connects). |
| END | | Terminate Entire Net-Work session. |
| FORCE | node | Broadcast a "node down" message. |
| HALT | | Terminate Entire Net-Work session. |
| HELP | | List available operator commands. |
| NETEND | | Terminate Entire Net-Work session. |
| OPEN | driver | Reopen a driver after a close or access method failure. |
| OPEN NETPRNT | | Open the NETPRNT file and route all trace and snap output to the NETPRNT file. |
| PROBE | node | Send a probe message to a node. |
| RESUME | link | Resume sending messages using this link. |
| SET | parameter | Change the values of Entire Net-Work parameters. Note: The minimum abbreviation for SET is the null string (zero characters long). |
| SNAP | | Snap data areas to DDPRINT. |
| START | driver | Restart driver (then connect all links of the driver). |
| STOP | taskid | Terminate Entire Net-Work session. |
| SUSPEND | link | Stop sending messages on this link. |
| TERMINATE | | Terminate Entire Net-Work session. |

Command Descriptions

This section covers the following topics:

- [Commands to Terminate an Entire Net-Work Session](#)
- [CLOSE](#) - Terminate Line Driver Activity
- [CLOSE NETPRNT](#) - Close the NETPRNT File
- [CONNECT](#) - Connect a Link
- [DEFINE LINK](#) - Dynamically Define a New Link
- [DISABLE](#) - Disallow Link Connection
- [DISCONNECT](#) - Disconnect a Link
- [DISPLAY](#) - Display Information About a Network Component
- [DUMP](#) - Terminate Entire Net-Work Session with Snap Dump
- [ENABLE](#) - Enable a Previously Disabled Link
- [FORCE](#) - Broadcast a "Node Down" Message
- [HELP](#) - List the Available Entire Net-Work Operator Commands
- [OPEN](#) - Reopen a Stopped or Closed Line Driver
- [OPEN NETPRNT](#) - Open the NETPRNT File
- [PROBE](#) - Verify Node Availability
- [RESUME](#) - Allow Link to Resume Sending Queued Messages
- [SET](#) - Dynamically Change NODE Statement Parameters
- [SNAP](#) - Snap Dump Selected Data Areas
- [START](#) - Restart a Stopped or Closed Line Driver
- [SUSPEND](#) - Stop Link from Sending Queued Messages

Commands to Terminate an Entire Net-Work Session



Any one of the above commands can be used to terminate an Entire Net-Work session normally. The `STOP` operator command (for example, `STOP taskid` or `P taskid`) can be used in the OS/390 environment.

Once the termination command has been accepted by Entire Net-Work, no more requests are selected from the request queue. Message NET0999 is displayed on the operator console confirming that normal termination procedures have been started.



Note: The `DUMP` command also ends Entire Net-Work operation after performing a snap dump of pertinent data areas.

CLOSE - Terminate Line Driver Activity

```
CLOSE drivername
```

Terminate all activities of the line driver by disconnecting and closing all links related to the driver, then closing the driver itself.

The driver name must be the same as was specified on the `XCF DRIVER` statement (that is, either `XCF` or `XCFD`). See section [XCF DRIVER Statement](#), elsewhere in this guide, for more information.

The effect of this command can be reversed by issuing the `OPEN` or `START` command for the driver, and `CONNECT` commands for the links (as appropriate).

CLOSE NETPRNT - Close the NETPRNT File

```
CLOSE NETPRNT
```

Close the `NETPRNT` file and route all trace and snap output to the `DDPRINT` file. When the `NETPRNT` file is closed, the data set can be copied for sending to Software AG support without shutting down Entire Net-Work. The file must be allocated `SHR`. This command cannot be abbreviated.

CONNECT - Connect a Link

```
CONNECT linkname
```

Attempt to connect link 'linkname'. The link name specified must match that used on the `LINK` statement. If the link was disconnected after a "handshaking" conflict, the `CONNECT` command can be used to retry the procedure. If the link is disabled, the `CONNECT` command can be used to enable it.

DEFINE LINK - Dynamically Define a New Link

```
DEFINE {LINK statement | LINK linkname LIKE linkname}
```

Defines a link during Entire Net-Work operation. The link statement must adhere to the format described in [XCF LINK Statement](#), elsewhere in this guide. The following example applies to the Entire Net-Work XCF line driver:

```
DEFINE LINK SYS01 XCF MEMBER=NODE01,PSTATS=Y,-  
RSTATS=N,STATINT=3600,TRACESIZ=4096
```

The LIKE linkname clause can be used instead of the LINK statement to define a link by copying the parameters specified for a previously defined link. For example:

```
DEFINE LINK SYS02 LIKE SYS01
```



Note: DEFINE LINK is permitted only when DEFINE=Y is specified on the NODE statement.

DISABLE - Disallow Link Connection

```
DISABLE linkname
```

Instructs the specified link not to accept any connections from other Entire Net-Work nodes. If the link is connected, it is disconnected and then disabled.

DISCONNECT - Disconnect a Link

```
DISCONNECT linkname
```

Disconnect the link 'linkname' connected to this node. The link name specified must be the same as that used on the LINK statement.

DISPLAY - Display Information About a Network Component

```
DISPLAY { ALINKS | LINKS | NODES | PATHS | STATS | TARGETS | ZAPS }
        [ name | string* ]
```

Displays current information about the specified network component. Only one component type (link, node, path, or target) can be specified in a single `DISPLAY` command. The information is displayed in the form of Entire Net-Work messages. Refer to the Entire Net-Work messages and codes documentation, provided in *Adabas Cluster Services Messages and Codes* for information about specific messages.

The optional second parameter serves to qualify the display request, thereby limiting the information displayed. At the same time, additional information is displayed for qualified `DISPLAY LINK` or `DISPLAY NODES` requests.

The possible qualifier values and their meanings depend on the type of request. A link name, node name, or (numeric) target ID may be specified. Alternatively, a string ending in a "wild card" character (*) may be used to indicate all links or nodes whose names start with the specified string. The asterisk (*) alone may be used to produce a display of all links or nodes, but additional information is shown only for qualified display requests.

DISPLAY ALINKS

The following is an example of `DISPLAY ALINKS` output. `DISPLAY ALINKS` lists currently active links only:

```
F NETWK,D AL
NET0120I: VTAM LINK LNKE   TO NODE ENODE   STAT=ACTIVE
NET0120I: VTAM LINK LNKA   TO NODE ANODE   STAT=ACTIVE
```

DISPLAY LINKS

The following is an example of `DISPLAY LINKS` output:

```
F NETWK,D L
NET0120I: VTAM LINK LNKALS TO NODE ALSNODE STAT=DISC
NET0120I: VTAM LINK LNKE   TO NODE ENODE   STAT=ACTIVE
NET0120I: VTAM LINK LNKA   TO NODE ANODE   STAT=ACTIVE
NET0120I: VTAM LINK LNKVM  TO NODE UNKNOWN STAT=OPEN
```

The following is an example of `DISPLAY LINKS` output for all links whose names begin with "TO":

```
F NET1,D L TO*
NET0120I: VTAM LINK TOSIX      TO NODE SIX      STAT=ACTIVE
NET0112I:          2 MSGS;          2 TR.BLKS
NET0120I: VTAM LINK TOTWO     TO NODE TWO     STAT=ACTIVE
NET0112I:          3 MSGS;          3 TR.BLKS
NET0120I: VTAM LINK TONINE    TO NODE UNKNOWN STAT=CONSTA
NET0112I:          0 MSGS;          0 TR.BLKS
```

DISPLAY NODES

The following is an example of DISPLAY NODES output:

```
F NETWK,D N
NET0122I: NODE FNODE      (50752) LOCAL
NET0122I: NODE ALSNODE   (54080) DIST 000040 VIA LINK LNKE
NET0122I: NODE ANODE     (49472) DIST 000020 VIA LINK LNKA
NET0122I: NODE ENODE     (50496) DIST 000020 VIA LINK LNKE
```

A qualifier is used in the following example:

```
F NETWK,D N A*
NET0122I: NODE ALSNODE   (54080) DIST 000040 VIA LINK LNKE
NET0123I: TARGETS: 00025 00171 00194 00175 00173 00018 00009
NET0123I: TARGETS: 00177
NET0122I: NODE ANODE     (49472) DIST 000020 VIA LINK LNKA
NET0123I: TARGETS: 00125 00192
```

DISPLAY PATHS

The following is an example of DISPLAY PATHS output:

```
F NETWK,D P
NET0122I: NODE ALSNODE   (54080) DIST 000080 (001) VIA LINK LNKA
NET0122I: NODE ALSNODE   (54080) DIST 000040 (002) VIA LINK LNKE
NET0122I: NODE ANODE     (49472) DIST 000020 (001) VIA LINK LNKA
NET0122I: NODE ANODE     (49472) DIST 000040 (002) VIA LINK LNKE
NET0122I: NODE ENODE     (50496) DIST 000040 (002) VIA LINK LNKA
NET0122I: NODE ENODE     (50496) DIST 000020 (001) VIA LINK LNKE
```

DISPLAY STATS

The `DISPLAY STATS` command produces the same type of information found in the statistics displayed at the end of an Entire Net-Work session. A qualifier parameter, if given, would have no effect. The buffer usage statistics displayed depend on the operating system being used.

The following is an example of `DISPLAY STATS` output for an OS/390 system. It includes a `NETB001I` and a `NETB009I` for each active buffer pool, a set of `NETB008I`, `NETB010I`, and `NETB012I` for each subpool within the buffer pools, and a `NETB013I` for each operator command issued.

```
F NETWK,D STATS
NETB000I:
-----
NETB001I:      Statistics For Buffer Pool COMN Loc = ANY
NETB000I:
-----
NETB008I: Req =(      13,      0,      10,      0)
NETB010I: ELM =(      512,      512,      512,      512), Sz = 512 B
NETB011I: Str =(      256,      256,      254,      252 ) K
NETB012I: Exp =(      0,      1,      0,      0)
NETB000I:
-----
NETB008I: Req =(      1,      0,      0,      0)
NETB010I: Elm =(      10,      10,      10,      2), Sz = 1 K
NETB011I: Str =(      15,      15,      13,      13) K
NETB012I: Exp =(      0,      1,      0,      0)
NETB000I:
-----
NETB008I: Req =(      1,      0,      0,      0)
NETB010I: Elm =(      1,      1,      1,      1), Sz = 14K
NETB011I: Str =(      14,      14,      0,      0) K
NETB012I: Exp =(      0,      1,      0,      0)
NETB000I:
-----
NETB009I: High Allc= 285 Curr Allc = 285 Curr Avail = 267 K
NETB000I:
-----
NETB001I:      Statistics For Buffer Pool PGFX Loc = ANY
NETB000I:
-----
NETB008I: Req =(      0,      0,      0,      0)
NETB010I: Elm =(      64,      64,      64,      0), Sz = 4K
NETB011I: Str =(      256,      256,      256,      256) K
NETB012I: Exp =(      0,      0,      0,      0)
NETB000I:
-----
NETB009I: High Allc= 256 Curr Allc = 256 Curr Avail = 256 K
NETB000I:
-----
NETB013I: Combined Buffer Pools Size                      541 K
```

```
NETB000I:
-----
```

DISPLAY TARGETS

The following is an example of `DISPLAY TARGETS` output:

```
F NETWK,D T
NET0124I: TARGET 09777 (C-N) ACTIVE ON NODE ALSNODE
NET0124I: TARGET 00009 (I-T) ACTIVE ON NODE ALSNODE
NET0124I: TARGET 02048 (L-N) ACTIVE ON NODE ANODE
NET0124I: TARGET 00237 (I-N) ACTIVE ON NODE ANODE
NET0124I: TARGET 00238 (I-N) ACTIVE ON NODE ANODE
NET0124I: TARGET 09888 (C-N) ACTIVE ON NODE ANODE
NET0124I: TARGET 00234 (I-N) ACTIVE ON NODE ANODE
NET0124I: TARGET 55769 (C-N) ACTIVE ON NODE ANODE
```

DISPLAY ZAPS

The `DISPLAY ZAPS` command lists, for each Entire Net-Work module, its name, assembly date, system maintenance level, and ZAP level. If ZAPs were applied after initial shipment, their numbers are listed as "Additional ZAPs". The following is an excerpt from a `DISPLAY ZAPS` example:

```
F NETWK,D Z
NET0037I: NETCLF (1998/11/10 SM=0001) ZAP LEVEL 0000
```

DUMP - Terminate Entire Net-Work Session with Snap Dump



Issue a snap dump, then end the Entire Net-Work session. `DUMP` is equivalent to the `SNAP` command followed by an `ADAEND` (or synonymous) command.

ENABLE - Enable a Previously Disabled Link



Revokes a previously entered `DISABLE` command. The specified link is instructed to accept incoming connect requests. Enabling a disconnected link does not connect the link.

FORCE - Broadcast a "Node Down" Message

```
FORCE {nodename | nodeid }
```

Broadcasts a control message through the network to notify all Entire Net-Work nodes that the specified node is no longer available. This command is provided for diagnosis and exception handling, and should be used only on the advice of your Software AG technical support representative.

HELP - List the Available Entire Net-Work Operator Commands

```
HELP
```

Lists the available Entire Net-Work operator commands with a short explanation of their function.

OPEN - Reopen a Stopped or Closed Line Driver

```
OPEN drivename
```

Reopen an installed and defined line driver that was stopped due to an access method or other network or system failure, or by the `CLOSE` operator command.

The driver name must be the same as was specified on the `XCF DRIVER` statement (that is, either `XCF` or `XCFD`). Read *XCF DRIVER Statement*, elsewhere in this guide, for more information.

Note that this command is currently a synonym for the `START` command. See the explanation in *XCF DRIVER Statement*, elsewhere in this guide, for more information.

OPEN NETPRNT - Open the NETPRNT File**OPEN NETPRNT**

Open the NETPRNT file and route all trace and snap output to NETPRNT. This command is necessary only after a `CLOSE NETPRNT` command has been used. It opens the NETPRNT file when Entire Net-Work is initialized. If the file is allocated SHR or OLD it will be erased when opened. This command cannot be abbreviated.

PROBE - Verify Node Availability**PROBE { *nodename* | *nodeid* } [*nnnn*]**

The `PROBE` command verifies that the specified node is available and can be reached. Entire Net-Work issues internal probe commands for the same purpose during normal operation. `PROBE` routes an internal message to the specified node and back. If the node cannot be reached, this information is sent to all active nodes, updating the node status.

The optional second parameter specifies that *nnnn* bytes of random user data (64512 bytes maximum) are to be appended to the actual probe message. The exact length of the message sent can be calculated as follows:

$$70 + (\textit{nodestack-size}) + \textit{nnnn}$$

where *nodestack-size* is twice the number specified by the `NODE` statement parameter `MAXPATH`, rounded up to the next multiple of 4. For example, if `MAXPATH=4` (the default value), the following command results in a message of length 1078:

```
PROBE nodename 1000
```

The result of the operation is displayed on the operator console, as shown in the following example:

```
F NET1,PROBE TWO
NET0136I: PROBE MESSAGE SENT
NET0135I: PROBE FOR NODE TWO      (0001.711 SEC)
NET0120I: NODE TWO      (62194) DIST 000030 VIA LINK TOFIVE
NET0140I: VERSION 5.7.1 (1999/11/10)
```

RESUME - Allow Link to Resume Sending Queued Messages

```
RESUME linkname
```

Revokes a `SUSPEND` command for the specified link. The link's status changes to "active" and the link resumes sending queued messages.

SET - Dynamically Change NODE Statement Parameters

```
SET parameter = value, ...
```

The `SET` command can be used to change Entire Net-Work parameter settings dynamically without interrupting network operations. Multiple parameters can be specified with one `SET` command.

The `SET` command keyword itself may be omitted. For example:

```
F NODEA,SET CQTIMER=180, TRACE=OFF
```

is equivalent to

```
F NODEA,CQTIMER=180, TRACE=OFF
```

The parameters allowed for the `SET` command are a subset of those defined on the Entire Net-Work `NODE` statement. They are as indicated below:

SET CQTIMER - Set Time Allowed to Retrieve Command Results

```
SET CQTIMER = seconds
```

The approximate waiting time, in seconds, allowed for a user or application to retrieve command results with a router 16-call before timeout occurs. For more information, see the [CQTIMER](#) parameter description, elsewhere in this guide.

SET DUMP - Set Storage Areas for a Dump

SET DUMP = { ALL | NONE | BLOCKS | TRACETAB | BUFFERS | LINKAREA | FORMAT }

The storage areas to be included in a dump when Entire Net-Work terminates abnormally. The information is printed to the NETPRNT file if it is open; otherwise, it is printed to the DDPRINT file. SET DUMP can be used to reduce the amount of output generated during an ABEND, especially on large Entire Net-Work systems. This command cannot be abbreviated.

In general, the default value of ALL should be used so that all diagnostic information is available to Software AG support.

Multiple values can be specified, separated by commas and surrounded by parentheses. For example:

```
SET DUMP = (BLOCKS, TRACETAB, FORMAT)
```

If conflicting values are specified, the last value specified is used. In the following, for example, the value used is NONE:

```
SET DUMP = (BLOCKS, TRACETAB, NONE)
```

| Value | Description |
|----------|--|
| ALL | All storage areas are dumped. This is the default value. |
| NONE | No storage areas are dumped. |
| BLOCKS | The major control blocks are dumped. |
| TRACETAB | The internal trace table is dumped. |
| BUFFERS | All internal buffer areas are dumped. |
| LINKAREA | All storage areas related to a driver and link are dumped. |
| FORMAT | The driver and link trace tables are formatted. |

SET LOG - Set Extent of Logging Activity

SET LOG= { ON | OFF | YES | NO | FULL | SHORT }

Regulates control flow and logging of selected data areas to the printer data set. For more information, refer to the [LOG](#) parameter description, elsewhere in this guide.

SET MAXPATH - Set Maximum Number of Links in Message Path

```
SET MAXPATH = linkcount
```

The maximum path link, specified in number of links, that a message from users on this node is expected to travel. For more information, refer to the [MAXPATH](#) parameter description, elsewhere in this guide.

SET MSGFORM - Set Format of Messages for Console Display and Print

```
SET MSGFORM = message-format
```

The message format of console messages and DDPRINT output. For more information, see the [MSGFORM](#) parameter, elsewhere in this guide.

SET REPLYTIM - Set Time Allowed for Reply to User Request

```
SET REPLYTIM = seconds
```

The time, in seconds, that this node is to wait for a reply to a user request before timing out. For more information, refer to the [REPLYTIM](#) parameter description, elsewhere in this guide.

SET TRACE, SET TROFF, SET TRON - Set Extent of Tracing Activity

```
SET {TRACE | TROFF | TRON} = {trace | (trace,...)}
```

Sets the trace control parameters for program traces. For more information, see the [TRACE](#) statement parameters, elsewhere in this guide.

SET UCMSG - Set Message Case

```
SET UCMSG= { N | Y }
```

Controls whether messages are issued in uppercase or mixed case. For more information, see the [UCMSG](#) parameter, elsewhere in this guide.

SET ULINK - Control Links to an Adjacent Node

```
SET ULINK= { N | Y }
```

Allows or disallows multiple links to an adjacent Entire Net-Work node. For more information, see the [ULINK](#) parameter, elsewhere in this guide.

SNAP - Snap Dump Selected Data Areas

```
SNAP { BPH | CQ | CURRMSG | MAIN | MYBLK | TRACE | UBQ }
```

Issue a snap dump of selected data areas to the DDPRINT file and continue processing.



Note: Under certain circumstances, a snap dump is performed internally at either normal or abnormal session end.

The optional parameters are used to snap one or more specific data areas:

| Parameter | Area |
|-----------|---|
| BPH | Buffer pool headers |
| CQ | Command queue |
| CURRMSG | Message that Entire Net-Work mainline is currently working on |
| MAIN | Header of mainline module |
| MYBLK | Central control block |
| TRACE | Internal trace table. |
| UBQ | User buffer queue (currently active Adabas commands) |

START - Restart a Stopped or Closed Line Driver

```
START drivername
```

Restart the installed line driver that was stopped due to an access method or other network or system failure, or by the `CLOSE` operator command.

The driver name must be the same as was specified on the XCF DRIVER statement (that is, either XCF or XCFD). Read [XCF DRIVER Statement](#), elsewhere in this guide, for more information.

The `START` command is a synonym for the `OPEN` command.

SUSPEND - Stop Link from Sending Queued Messages

```
SUSPEND linkname
```

Instructs the specified link to stop sending messages. However, Entire Net-Work can still queue messages on this link. The `SUSPEND` command is valid only if the link is active.

XCF Line Driver Commands

The Entire Net-Work XCF line driver has the ability to process operator commands that are directed to a specific link, to all links, or directly to the driver.

Format

The XCF line driver operator commands have the following format:

```
XCFD target cmd1 cmd2 ... cmdx
```

where

- 'XCFD' tells Entire Net-Work that the command is destined for the XCF line driver. Read [XCF DRIVER Statement](#), elsewhere in this guide, for more information.

- 'target' identifies the target of the command to XCF and has the following format:

```
target = { * | # | linkname }
```

XCF line driver operator commands depend on whether the target is the driver of one or more links. Specify

- "*" if the target is all links.
 - "#" if the target is the driver itself.
 - the link name if the target is a specific link.
- 'cmd1 cmd2 cmdx' are the operator commands to be issued.

Multiple commands can be specified in a single command statement. When the ALTER command is specified, it must be the last command in the statement, because everything following the ALTER command is treated as a DRIVER or LINK statement parameter.

Command Overview

The Entire Net-Work XCF line driver supports the commands listed in the following table. All commands are available for both link and driver targets. The underlined portion of the command is the minimum abbreviation.

| Command | Action |
|--------------|---|
| <u>ALTER</u> | Dynamically changes the driver or link configuration. |
| <u>HELP</u> | Lists available operator commands. |
| <u>RESET</u> | Resets all statistics for the driver or link. |
| <u>SHOW</u> | Displays parameter settings for the driver or link. |
| <u>SNAP</u> | For driver targets, snaps driver-specific data areas. For link targets, snaps all control blocks specific to the link. |
| <u>STATS</u> | For driver targets, displays/prints driver statistics immediately. For link targets, causes the immediate printing of statistics. |
| <u>TRACE</u> | For driver targets, formats and prints the driver-specific trace table. For link targets, prints the link-specific trace table. |

Command Descriptions

ALTER - Dynamically Change Configuration

ALTER *configuration-parameters*

Dynamically changes the driver or link configuration. The ALTER command is followed by the driver or link configuration parameters to be altered, as appropriate. The driver configuration parameters are the same as those specified in the DRIVER statement (read [XCF DRIVER Statement](#), elsewhere in this guide); the link configuration parameters are the same as those specified on the LINK statement (read [XCF LINK Statement](#), elsewhere in this guide).

For example:

| | |
|----------------|----------------------------------|
| driver target: | XCFD # ALTER ACCEPTUI=Y,MAXRU=8A |
| link target: | XCFD linkname ALTER PSTAT=Y |

HELP - List Available Operator Commands

HELP


Lists all the commands available to display or alter the XCF driver or an XCF link.

RESET - Reset All Statistics

RESET

Resets all statistics for the XCF driver or link target. Statistics are printed only if the STATS command precedes the RESET command.

SHOW - Display Parameter Settings



Causes the immediate printing of the current driver or link parameter settings.

SNAP - Snap Dump Selected Data Areas



Causes all driver- or link-specific control blocks and the driver- or link-specific trace table to be snapped (printed in hexadecimal). Entire Net-Work-specific control blocks are not snapped.

- For driver targets, link-specific control blocks are not snapped.
- For link targets, driver-specific control blocks are not snapped.

STATS - Display and Print Statistics



Causes the immediate printing of statistics and restarts the statistics interval. To print and reset statistics, specify `RESET` immediately after the `STATS` command.

For example:

| | |
|----------------|---------------------------|
| driver target: | XCFD # RESET |
| link target: | XCFD linkname STATS RESET |

TRACE - Format and Print Trace Table

Causes the driver- or link-specific trace table to be formatted and printed. The trace table is also formatted and printed automatically when the `SNAP` command is processed.

10 Entire Net-Work Diagnostic Utilities

| | |
|--------------------------|-----|
| ▪ NETPFIL1 Utility | 134 |
| ▪ NETPFIL2 Utility | 136 |
| ▪ Example JCL | 138 |

This section describes the Entire Net-Work diagnostic utilities, which are provided in the Entire Net-Work load library.

The Adabas modules required by the utilities are provided in the `WALvrs` library distributed with Entire Net-Work. See the section *Product Code Description* in *Entire Net-Work Release Notes* and the *Software AG Product Delivery Report*.

NETPFIL1 Utility

The NETPFIL1 utility is used to select the information to be printed from the Entire Net-Work NETPRNT file. The Entire Net-Work NETPRNT file contains tracing, logging, and dump output. The NETPRNT file should be used as input to the NETPFIL1 utility for the NETFILE file.

NETPFIL1 can be used to select

- only those records that fall within a certain date and time frame,
- only logging and dump records that match the title or any desired portion of the title,
- only trace records that match the title or any desired portion of the title, or
- any combination of the above.

The output from NETPFIL1 can be used as input to the [NETPFIL2 Utility](#) to perform additional filtering.

NETPFIL1 Parameters

This section describes the parameters for the NETPFIL1 utility.

DATE

```
DATE=yyyymmdd
```

Only records created on the specified date are written to DDPRINT. The default value is all dates.

STARTTIME

```
STARTTIME=hhmmssth
```

Records created before the specified time are not written to DDPRINT and are not checked for other selection criteria. The default is no restriction on the start time.

ENDTIME

```
ENDTIME=hhmmssst
```

Records created after the specified time are not written to DDPRINT and are not checked for other selection criteria. The default is no restriction on the end time.

LOG

```
LOG=log-title
```

Log and dump records with a title that matches the specified log title are written to DDPRINT.

The log title must be specified exactly as it appears in the NETPRINT file; the value specified can be limited to the number of characters necessary to select the desired records. All characters following the last nonblank character are considered wild cards. Spaces are valid in the log title, but the check is performed only up to the last nonblank character.

LOG= selects all log and dump records. The default value is that no log or dump records are written.

TRACE

```
TRACE=trace-title
```

Trace records that have a title matching the specified trace title are written to DDPRINT.

The trace title must be specified exactly as it appears in the NETPRINT file; the value specified can be limited to the number of characters necessary to select the desired records. All characters following the last nonblank character are considered wild cards. Spaces are valid in the trace title, but the check is performed only up to the last nonblank character.

TRACE= selects all trace records. The default value is that no trace records are written.

Anything else is written to the output file as a comment along with the parameters specified to the beginning of the DDPRINT file.

An asterisk (*) in column 1 causes this parameter record to be ignored.

Examples

The following selects all records for a five minute period with a comment:

```
ALL INFORMATION FROM 2:23 to 2:29
STARTTIME=14230000
ENDTIME=14290000
LOG=
TRACE=
```

The following selects log records for IDDDATA and IDDBLK using implied wild cards:

```
LOG=IDD
```

The following selects log records for IDDDATA only:

```
LOG=IDDDATA
```

The following selects dump records for M Y B L K, D R I V B L K, and L I N K B L K:

```
LOG=M Y B L K
LOG=D R I V B L K
LOG=L I N K B L K
```

NETPFIL2 Utility

The NETPFIL2 utility is used to search for a control block or storage area snapped by either the logging function or a dump. When the desired dump entries are found, they are printed.

Optionally, NETPFIL2 can also print all non log entries, including the trace and time stamp records. NETPFIL2 is therefore suitable as a second step in a job where NETPFIL1 is run first.

To summarize, NETPFIL2 can be used to

- find and print a control block or data area that contains a specified value at a specified offset,
- optionally print all non logging and dump records, and
- find multiple control blocks and multiple values.

The Entire Net-Work NETPRNT file contains tracing, logging, and dump output. The NETPRNT file should be used as input to the NETPFIL2 utility for the NETFILE file.

Parameters

This section describes the parameters for the NETPFIL2 utility.

NONLOG

```
NONLOG={Y | N }
```

The NONLOG parameter specifies whether all non log and dump records should be written to DDPRINT.

| | |
|------------|--|
| NONLOG=Yes | Writes all non log and dump records. |
| NONLOG=No | Ignores all non log and dump records. This is the default value. |

FIND

```
FIND=log-title (OFFSET=xxx {VALUE=characters | VALUE=X`hex values'})
```

Log and dump records that have a log title matching the specified log title are evaluated to determine whether the hexadecimal offset matches either the character or hexadecimal value specified. If it matches, the whole dumped area is written to DDPRINT.

The log title must be specified exactly as it appears in the NETPRNT file; the value specified can be limited to the number of characters necessary to select the desired records. All characters following the last non blank character are considered wild cards. Spaces are valid in the log title, but the evaluation is performed only up to the last non blank character. To select all log and dump records, use the LOG= parameter.

OFFSET=xxx must be specified as a valid hexadecimal number. This number specifies the first byte of the location in the dumped area to compare. The comparison is performed in such a way that the value must start at this location. Each additional character of the value is checked at the next logical position in the dump, even if the next logical position is on the next record.

VALUE=characters specifies the right hand portion of the value, which is in character format. The character form of the value is compared to the interpreted part of the dumped records.

VALUE=X`hex values' specifies the left hand portion of the value, following the address and offset. The hexadecimal form of the value is compared to the hexadecimal part of the dumped records.

The following example finds the driver block (D R I V B L K) for the TCP/IP (NETTCPI) driver.

```
FIND=D R I V B L K,OFFSET=10,VALUE=NETTCPI
FIND=D R I V B L K,OFFSET=10,VALUE=X`D5C5E3E3C3D7C9'
```

Example JCL

Running in z/OS Environments

The following example JCL can be used to run NETPFIL1 or NETPFIL2 in a z/OS environment.

```
//YOURJOB JOB (0),'NET-WORK',
//          CLASS=A,MSGCLASS=X,MSGLEVEL=(1,1)
//FILTER1 EXEC PGM=NETPFIL1                                <--- see Note 1 //STEPLIB ←
//          DD DSN=NETWRK.vrs.LOAD,DISP=SHR                <--- see Note 2
//          DD DSN=WAL.vrs.LOAD,DISP=SHR                  <--- see Note 3
//NETFILE DD DSN=NETWRK.vrs.NETPRNT,DISP=SHR              <--- see Note 4
//DDCARD DD *                                             <--- see Note 5
//          THIS COMMENT WILL BE PRINTED AT THE BEGINNING OF THE OUTPUT
//          DATE=20000316
//          STARTTIME=12570000
//          ENDTIME=12580000
//          LOG=T R A C E
//          * LOG=IDDDATA
//          TRACE=
//          * TRACE=SENDOUT
//DDPRINT DD SYSOUT=*,DCB=(LRECL=121,BLKSIZE=1210,RECFM=FBA)
/*
```



Notes:

1. Specify the utility program you want to run.
2. This is the same load library you use to run Entire Net-Work.
3. This is the same Adabas or WAL load library you use to run Entire Net-Work.
4. This is the file that was created by Entire Net-Work with DDNAME NETPRNT.
5. This is the parameter file and may be any sequential file with LRECL=80.

Index

Symbols

64-bit virtual storage option, 55

A

accepting unknown requests, 72
 XCF DRIVER statement, 72
ACCEPTUI parameter, 72
 XCF DRIVER statement, 72
Adabas Caching Facility
 adjusting shared cache size for, 35
Adabas Cluster Services
 ADARUN parameter to enable, 38-39
Adabas Parallel Services
 ADARUN parameter to enable, 38-39
Adabas SMF records
 selecting detail SMF record sections, 52
Adabas sysplex cluster
 nucleus ID
 parameter to specify, 50
ADACOM
 initialization parameters, 7
ADACOM operator commands
 DIM
 display active images, 83
 DN
 display active local nuclei, 84
 DUMP, 85
 SN
 set nucleus status, 86
ADACOM parameters
 CMDMGR, 8
 DBID, 9
 FORCE, 9
 LOC, 10
 NU
 maximum number of users, 10
 SVC
 supervisory call number for interregion communication,
 11
ADAEND operator command
 terminate an Entire Net-Work session, 114, 120
ADARUN
 directory of parameters, 15
ADARUN parameters
 CT
 command timeout limit, 27, 64
 determining correct settings, 14

 example statement, 26
 FORCE
 overwrite ID table entry, 28
 LU
 length of intermediate user buffer, 29
 MLWTO, 44
 NAB
 number of attached buffers, 29
 NC
 number of request queue elements, 30
 setting values, 25
 summary for Entire Net-Work, 26
 SVC
 supervisory call number for interregion communication,
 30
 TARGETID
 Entire Net-Work node number, 31
ALTER operator command, 71, 74
 change line driver configuration, 128
 XCF
 link, 129
 XCF line driver, 129

B

bold, 58
braces ({}), 59
brackets ([]), 59
buffers
 length of intermediate user buffer, 29
 number of attached buffers, 29
 parameter for defining, 62
 size of asynchronous pool, 63
 size of long-term pool, 63
 size of page-fixed pool, 63
 size of short-term pool, 63
BUFFERS parameter, 62

C

cache structure
 name, 35
 ratio
 directory entries to elements, 41
 statistics display, 98, 101
cache vector size, 35
choices in syntax, 59
CLOGMRG
 ADARUN parameter, 34
CLOSE operator command, 71, 74

terminate line driver activity, 115, 121, 127
 CLUCACHEEXTRA
 ADARUN parameter, 35
 CLUCACHENAME
 ADARUN parameter, 35
 CLUGROUPNAME
 ADARUN parameter, 36
 CLULOCKNAME
 ADARUN parameter, 37
 CLUPUBLPROT
 ADARUN parameter, 38
 CLUSTER
 ADARUN parameter, 39
 cluster
 cache vector size, 35
 group name, 36
 redo updates
 parameter to set size of pool, 43
 CLUWORK1CACHE
 ADARUN parameter, 39
 CMDMGR parameter, 8
 Command log
 parameter to enable merge, 34
 command routing, 110
 communication
 cluster nuclei
 timeout threshold for canceled peer nucleus, 44
 timeout threshold for command, 46
 CONNECT operator command
 attempt to connect a named link, 115
 console messages
 multi-line WTO messages, 44
 Control scheduling of update commands during buffer flush,
 54
 control statements
 Entire Net-Work, 58
 order of, 58
 CQTIMER parameter, 64, 69
 CT parameter, 27, 69

D

D L command, 74
 D S command, 62
 DBID parameter, 9
 default parameter values, 59
 DEFINE LINK operator command, 64
 define an Entire Net-Work link, 116
 DEFINE parameter, 64
 descriptor
 unique
 parameter to set size of pool, 42
 detail sections
 SMF records, 108
 DIM command
 display active images, 83
 DIRRATIO / ELEMENTRATIO
 ADARUN parameter, 41
 DISABLE operator command
 disconnect/disable link, 116
 DISCONNECT operator command
 disconnect a link, 116
 DISPLAY ALINKS operator command
 display active link information, 117

DISPLAY LINKS operator command
 display link information, 117
 DISPLAY NODES operator command
 display node information, 117-118
 DISPLAY operator command, 117
 DISPLAY PATHS operator command
 display path information, 118
 DISPLAY STATS operator command
 display usage statistics, 119
 DISPLAY TARGETS operator command
 display target information, 120
 DISPLAY ZAPS operator command
 display ZAP level, 120
 displaying Adabas session parameters, 92
 DMBMTB command
 display member state table, 90
 DN command
 display active local nuclei, 84
 DNFV command
 displaying nucleus file variables, 91
 DOMAIN parameter, 64
 DPARM operator command, 92
 DPPT command
 displaying a nucleus PPT block, 94
 DRIVER statement
 defined, 58
 defining for XCF, 71
 overview, 71
 parameters, 72
 DRIVER statement parameters
 ACCEPTUI
 whether to connect to nodes not defined in LINK state-
 ment, 72
 GROUP
 Entire Net-Work group name, 72
 DTARG command
 display system image, 97
 DUMP command
 dump storage areas, 85
 DUMP operator command
 terminate an Entire Net-Work session with dump, 115
 terminate an Entire Net-Work session with snap dump, 120
 DUMP parameter, 64
 DXCACHE command
 display lock statistics, 98
 DXFILE command
 display file statistics, 99
 DXLOCK command
 display lock statistics, 100
 DXSTAT command
 display cache and lock statistics, 101
 dynamically change MXCANCEL setting, 103
 dynamically change MXCANCELWARN setting, 104
 dynamically change MXMSGWARN setting, 106
 dynamically change MXWTOR setting, 108

E

ENABLE operator command
 enable a previously disabled link, 120
 END operator command
 terminate an Entire Net-Work session, 114
 Entire Net-Work
 ADARUN parameters for, 26

request queue, 30
 Entire Net-Work parameters
 using SET operator command to change settings, 123

F

file
 statistics display, 99
 FORCE parameter, 28
 format
 Entire Net-Work operator commands, 112
 Entire Net-Work statements, 58
 NODE statement, 61
 XCF DRIVER statement, 71
 XCF line driver operator commands, 127
 XCF LINK statement, 75

G

global cache area
 name, 35
 ratio
 directory entries to elements, 41
 Global lock area
 statistics display, 100
 global lock area
 name, 37
 GROUP parameter
 XCF DRIVER statement, 72

H

HALT operator command
 terminate an Entire Net-Work session, 114
 HELP operator command
 XCF line driver, 129

I

initialization parameters
 ADACOM, 7
 italic, 59

L

LDEUQP
 ADARUN parameter, 42
 LINK statement
 defined, 58
 defining for XCF, 74
 overview, 74
 parameters, 75
 LINK statement parameters
 linkname, 75
 PSTATS
 whether statistics are printed automatically, 76
 STATINT
 time allowed before statistics are reported or reset, 76
 WEIGHT
 relative weight of this link, 77
 lock structure
 name, 37
 statistics display, 101

LOG parameter, 65
 LOG parameter (NODE statement), 65
 logging checkpoint and data areas
 LOG parameter, 65
 LOG= parameter, 65
 LOG= subcommand, 124
 lowercase, 59
 LRDP
 ADARUN parameter, 43
 LU parameter, 29

M

MAXPATH parameter, 66
 MAXPATH parameter (NODE statement), 66
 member state table
 command to display, 90
 minimum keywords, 59
 MLWTO parameter, 44
 MSGFORM parameter, 66
 multi-line console messages, 44
 multiprocessing
 nucleus ID for
 parameter to specify, 50
 MXCANCEL
 ADARUN parameter, 44
 MXCANCEL operator command, 103
 MXCANCELWARN
 ADARUN parameter, 45
 MXCANCELWARN operator command, 104
 MXMSG
 ADARUN parameter, 46
 MXMSGWARN
 ADARUN parameter, 47
 MXMSGWARN operator command, 106
 MXSTATUS
 ADARUN parameter, 48
 MXWTOR
 ADARUN parameter, 49
 MXWTOR operator command, 108

N

NAB parameter, 29
 NC parameter, 30
 NETEND command
 terminate Entire Net-Work, 112, 114
 NID0 parameter, 67
 node
 name, specifying the, 62
 operation, verifying, 122
 NODE statement
 defined, 58
 defining for XCF, 61
 format, 61
 parameters, 62
 NODE statement parameters
 BUFFERS
 space reserved for internal buffer management, 62
 CQTIMER
 time allowed to retrieve command results, 64
 DEFINE
 accept or reject DEFINE operator command, 64
 DOMAIN, 64

- DUMP, 64
 - LOG
 - enable/disable logging of node statistics, 65
 - MAXPATH
 - max. path length for message in network, 66
 - MSGFORM
 - define format of console messages and DDPRINT output, 66
 - NID0, 67
 - nodename, 62
 - NTRACE
 - specify size of internal trace table, 68
 - REPLYTIM
 - time allowed for user request to complete, 68
 - START
 - enable/disable automatic start, 68
 - TIMER
 - interval between time-dependent requests, 69
 - TRACE, TROFF, TRON
 - enable/disable specific program traces, 69
 - UCMSG
 - set message for uppercase or mixed-case issue, 70
 - ULINK, 70
 - normal font, 59
 - NTRACE parameter, 68-69
 - NTRACE parameter (NODE statement), 68
 - NU parameter, 10
 - NUCID
 - ADARUN parameter, 50
 - nucleus
 - ID for multiprocessing
 - parameter to specify, 50
 - nucleus file variables
 - displaying, 91
- O**
- OPEN operator command
 - reopen a stopped/closed line driver, 115, 121, 127
 - Operator commands
 - DXLOCK
 - display lock statistics, 100
 - operator commands, v, 112
 - ADACOM, 80
 - ADAEND
 - terminate Entire Net-Work session, 114, 120
 - ALTER, 71, 74
 - change line driver configuration, 128
 - CLOSE, 71, 74
 - terminate all line driver activity, 115, 121, 127
 - cluster nucleus, 87
 - CONNECT
 - attempt to connect a named link, 115
 - D L
 - display link, 74
 - D S
 - display statistics, 62
 - DEFINE LINK, 64
 - define an Entire Net-Work link, 116
 - DISABLE
 - disconnect/disable link, 116
 - DISCONNECT
 - disconnect a link, 116
 - DISPLAY ALINKS
 - display active link information, 117
 - DISPLAY LINKS
 - display link information, 117
 - DISPLAY NODES
 - display node information, 117-118
 - DISPLAY PATHS
 - display path information, 118
 - DISPLAY STATS
 - display usage statistics, 119
 - DISPLAY TARGETS
 - display target information, 120
 - DISPLAY ZAPS
 - display ZAP level, 120
 - DMEMTB
 - display member state table, 90
 - DNFV, 91
 - DPARM, 92
 - DPPT, 94
 - DTARG, 97
 - DUMP
 - terminate Entire Net-Work session with dump, 115
 - terminate Entire Net-Work session with snap dump, 120
 - DXCACHE
 - display lock statistics, 98
 - DXFILE
 - display file statistics, 99
 - DXSTAT
 - display cache and lock statistics, 101
 - ENABLE
 - enable a previously disabled link, 120
 - END
 - terminate Entire Net-Work session, 114
 - Entire Net-Work, 112
 - summary list for z/OS, 113
 - HALT
 - terminate Entire Net-Work session, 114
 - MXCANCEL, 103
 - MXCANCELWARN, 104
 - MXMSGWARN, 106
 - MXWTOR, 108
 - NETEND
 - terminate Entire Net-Work session, 114
 - OPEN
 - reopen a stopped/closed line driver, 115, 121, 127
 - PROBE
 - verify node availability, 122
 - RESET
 - resets all statistics, 129-130
 - RESUME
 - allows link to resume sending queued messages, 123
 - SET CQTIMER
 - time allowed to retrieve command results, 123
 - SET DUMP
 - storage areas to include, 124
 - SET LOG
 - determine extent of logging activity, 124
 - SET MAXPATH
 - maximum number of links in message path, 125
 - SET MSGFORM
 - message format for display and print, 125
 - SET REPLYTIM
 - time allowed for reply to user request, 125
 - SET TRACE, SET TROFF, SET TRON
 - determine extent of tracing activity, 125

- SHOW
 - display parameter settings, 130
 - SNAP
 - snap dump selected data areas, 120, 126, 130-131
 - START
 - restart a stopped/closed line driver, 115, 121, 127
 - start individual line driver, 68
 - STATS
 - displays/prints statistics, 129-130
 - STOP
 - terminate Entire Net-Work session, 114
 - SUSPEND
 - stops link from sending queued messages, 123, 127
 - TARGET, 110
 - TERMINATE
 - terminate Entire Net-Work session, 114
 - to terminate an Entire Net-Work session, 114
 - TRACE
 - formats/prints trace table, 131
 - XCF line driver, 127
 - optional syntax elements, 59
- ## P
- parameter
 - syntax conventions, 58
 - syntax rules, 60
 - path length
 - defining the maximum, 66
 - SET MAXPATH command, 125
 - path length, defining the maximum
 - MAXPATH= parameter (NODE statement), 66
 - PPT
 - displaying a nucleus block, 94
 - PROBE operator command
 - verify node availability, 122
 - PROGRAM
 - ADARUN parameter, 51
 - PSTATS parameter, 76
 - punctuation and symbols in syntax, 60
- ## R
- redo pool
 - parameter to set size of, 43
 - REPLYTIM parameter, 68-69
 - required syntax elements, 59
 - RESET operator command
 - resets all statistics, 129-130
 - RESUME operator command
 - allows link to resume sending queued messages, 123
- ## S
- session
 - displaying current parameters, 92
 - SET CQTIMER command
 - time allowed to retrieve command results, 123
 - SET CQTIMER= command, 64
 - SET DUMP command
 - storage areas to include, 124
 - SET LOG command, 66
 - determine extent of logging activity, 124
 - SET MAXPATH command, 66
 - maximum number of links in message path, 125
 - SET MSGFORM command
 - format for message display and print, 125
 - SET operator command
 - change Entire Net-Work parameter settings, 123
 - change NODE statement parameter settings, 123
 - SET REPLYTIM command, 68
 - time allowed for reply to user request, 125
 - SET TRACE, TROFF, TRON commands
 - determine extent of tracing activity, 70, 125
 - SET UCMMSG command, 70
 - SET ULINK command, 70
 - SHOW operator command
 - display parameter settings, 130
 - SMF Interval records, 108
 - SMF records
 - detail sections, 108
 - selecting detail sections, 52
 - SMF Termination records, 108
 - SMFDETAIL
 - ADARUN parameter, 52
 - SN command
 - set nucleus status, 86
 - SNAP operator command
 - snap dump selected data areas, 120, 126, 130-131
 - START operator command, 68
 - restart a stopped/closed line driver, 115, 121, 127
 - START parameter, 68
 - statement
 - syntax conventions, 58
 - syntax rules, 60
 - STATINT parameter, 76
 - STATS operator command
 - displays/prints statistics, 129-130
 - STOP operator command
 - terminate an Entire Net-Work session, 114
 - SUSPEND operator command
 - stops link from sending queued messages, 123, 127
 - SVC parameter, 11, 30
 - syntax
 - conventions, 58
 - DRIVER statement, 71
 - LINK statement, 74
 - rules, 60
 - syntax conventions
 - bold, 58
 - braces ({}), 59
 - brackets ([]), 59
 - defaults, 59
 - italic, 59
 - lowercase, 59
 - minimum keywords, 59
 - mutually exclusive choices, 59
 - normal font, 59
 - optional elements, 59
 - punctuation and symbols, 60
 - required elements, 59
 - underlining, 59
 - uppercase, 58
 - vertical bars (|), 59
 - system image display, 97

T

TARGET command
 route commands to specified target, 110

target IDs
 allowing duplicates, 32
 overwriting, 28

TARGETID parameter, 31

TERMINATE operator command
 terminate an Entire Net-Work session, 114

time
 waiting time for results (CQTIMER parameter), 62

TIMER parameter, 69

timing parameters
 interaction between, 69

trace control parameters, 69

TRACE operator command
 formats/prints trace table, 131

TRACE parameter, 69

trace table size, 68

trace table size, specifying the (NTRACE= parameter), 68

TRACESIZ= parameter
 XCF
 DRIVER statement, 74

tracing line driver activity, 69
 XCF driver, 74

TROFF parameter, 69

TRON parameter, 69

U

UCMSG parameter, 70

ULINK parameter, 70

underlining, 59

unique descriptor pool
 parameter to set size of, 42

UPDATECONTROL
 ADARUN parameter, 54

uppercase, 58

users
 maximum number of, 10

utility control statement
 parameter values
 default, 59
 syntax conventions, 58
 syntax rules, 60

V

V64BIT parameter, 55

verify remote node operation, 122

vertical bars (|), 59

virtual storage
 64-bit option, 55

W

WEIGHT parameter, 77

X

XCF
 DRIVER statement, 71
 LINK statement, 74

 NODE statement, 61

XCF DRIVER statement
 format, 71
 parameters, 72

XCF group
 name, 36

XCF line driver
 operator commands
 descriptions, 129
 overview, 128

XCF LINK statement
 format, 75
 parameters, 75