

Adabas Cluster Services

Parameters and Parameter Statements

Version 7.4.2

September 2009

This document applies to Adabas Cluster Services Version 7.4.2 and to all subsequent releases.

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

Copyright © Software AG 2009. All rights reserved.

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

Table of Contents



1 Parameters and Parameter Statements	1
2 Initialization Parameters	3
ADARUN Initialization Parameters	29
ADACOM Initialization Parameters	4
ADANUC Initialization Parameters for Cluster Nuclei	7
NETWRK Initialization Parameters for Entire Net-Work	13
3 Entire Net-Work Parameter Statements	21
Statement Format	22
XCF NODE Statement	23
XCF DRIVER Statement	33
XCF LINK Statement	36
Index	41

1 Parameters and Parameter Statements

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.

The documentation is organized in the following parts:

 Initialization Parameters
 Entire Net-Work Parameter Statements

2 Initialization Parameters

- ADARUN Initialization Parameters 29
- ADACOM Initialization Parameters 4
- ADANUC Initialization Parameters for Cluster Nuclei 7
- NETWRK Initialization Parameters for Entire Net-Work 13

This section describes the initialization parameters needed for an Adabas Cluster Services environment.

This chapter covers the following topics:

ADARUN Initialization Parameters

ADARUN parameters are used to customize the Adabas environment. See the *Adabas Operations* documentation for information about ADARUN parameters and the format of ADARUN control statements.

For Adabas Cluster Services environments, three types of initialization parameters must be set based on the ADARUN parameter PROGRAM:

PROGRAM - Execute an ADARUN Program

Parameter	Specify...	Possible Values	Default
PROGRAM	the ADARUN program	NETWRK ADANUC ADACOM	none

The PROGRAM parameter is used to specify the name of the program being run, and must always be specified. There is no default.

- 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).

ADACOM Initialization Parameters

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 dataset prints global messages that apply to all SVC/DBID sets defined to an ADACOM task. In addition, a SYSOUT dataset is dynamically allocated for command output to each SVC/DBID set. The format of this dataset 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 Start-up Procedure for ADACOM* in the *Installation* documentation.

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 DD/KARTE (CMDMGR=NO).

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.	YES NO	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.

NU - Number of User Table Elements

Parameter	Specify...	Minimum	Maximum	Default
NU	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 7.4 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
```

ADANUC Initialization Parameters for Cluster Nuclei

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. See the *Performance and Tuning* documentation for information about expected differences.

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

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.

Changing Parameter Values at Nucleus Startup

After the first nucleus in an Adabas cluster starts, XCF messages are used to communicate the ADARUN parameter settings of the first nucleus to all subsequent cluster nuclei. Each following nucleus receives this information during initialization and determines whether its global non-changeable 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 XCF messages and a corresponding message is printed.

Changing Parameter Values during a Session

On a running system, a cluster nucleus may want to modify one or more of the "changeable" global parameters. This nucleus acquires a "parameter change lock", makes the changes in its local parameter area, and communicates the changes to the other cluster nuclei using XCF messages.

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

Parameter Types

A cluster nucleus requires

- *global* parameters. Adabas Cluster Services enforces the same value for all nuclei in a cluster. Some of these parameters are modifiable (GM) during a session using an operator command or AOS (NISNHQ, NONDES, and AOSLOG are only modifiable using AOS); others are fixed (GF) and cannot be modified.
- *local* parameters, which can be different for each nucleus. Some of these parameters are modifiable (LM) using an operator command or AOS; others are fixed (LF) and cannot be modified.

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

Parameter Directory



Note: All Adabas Caching Facility ADARUN parameters are supported as local, modifiable parameters.

Parameter	Usage	No*	LF	LM	GF	GM
AQSLOG	Log to DDPRINT commands issued by AOS or ADADBS OPERCOM that modify the active nucleus			Y		
AREXCLUDE	Exclude file(s) from autorestart		Y			
ARMNAME	Name used to activate ARM		Y			
ASSOCACHE	Controller caching control for the Associator component		Y			
ASYTVS	Asynchronous buffer flush based on vol-ser			Y		
CDXnn	Collation descriptor user exit(s)				Y	
CLOGDEV	Multiple command log device		Y			
CLOGLAYOUT	Define command log format		Y			
CLOGMRG	Automatic command log merge control in a cluster environment					Y
CLOGSIZE	Multiple command log size (blocks)		Y			
CLUCACHENAME	Cluster cache structure name (Adabas Cluster Services only)				Y	
CLUCACHESIZE	Cluster cache area size (Parallel Services only)				Y	
CLUGROUPNAME	Cluster group name (Adabas Cluster Services only)				Y	
CLULOCKNAME	Cluster lock structure name (Adabas Cluster Services only)				Y	
CLULOCKSIZE	Cluster lock area size (Parallel Services only)				Y	
CLUSTER	Adabas cluster session control				Y	
CT	Command time limit (seconds)					Y
DATACACHE	Controller caching control for the Data Storage component		Y			
DIRRATIO/ELEMENTRATIO	Ratio of directory entries to data elements in a cluster cache structure/area				Y	
DBID	Database ID (physical)				Y	
DEVICE	Device type of the first ASSO extent				Y	
DSF	Delta Save Facility control				Y	
DSFEX1	Delta Save Facility user exit				Y	
DTP	Distributed transaction processing control	Y			Y	
DUALCLD	Dual command log device		Y			
DUALCLS	Dual command log size (blocks)		Y			

Initialization Parameters

Parameter	Usage	No*	LF	LM	GF	GM
DUALPLD	Dual protection log device		Y			
DUALPLS	Dual protection log size (blocks)		Y			
FASTPATH	Adabas Fastpath control				Y	
FMXIO	Limit parallel I/O operations by LFIOP flush processing			Y		
FORCE	Overwrite IDTE		Y			
HEXnn	Hyperdescriptor exit(s)				Y	
IGNDIB	Ignore DIB entry		Y			
IGNDTP	Ignore distributed transaction processing area (Work part 4)	Y				
INTNAS	Interval between nucleus statistic checkpoints (SYNS 60)				Y	
LBP	Length of buffer pool		Y			
LCP	Length of security pool		Y			
LDEUQP	Length of unique (UQ) descriptor pool		Y			
LDTP	Length of distributed transaction processing area (Work part 4)	Y				
LFIOP	Length of asynchronous flush pool		Y			
LFP	Length of internal format buffer pool		Y			
LI	Length of ISN list table (TBI)		Y			
LOCAL	Nucleus (cluster) unknown to the network				Y	
LOGCB	Log control block			Y		
LOGCLEX	Log command log extension (CLEX)			Y		
LOGFB	Log format buffer			Y		
LOGGING	Logging of Adabas commands			Y		
LOGIB	Log ISN buffer			Y		
LOGIO	Log I/O activity			Y		
LOGRB	Log record buffer			Y		
LOGSB	Log search buffer			Y		
LOGSIZE	Maximum command log size		Y			
LOGUX	Log user exit B data			Y		
LOGVB	Log value buffer			Y		
LP	Length of data protection area (Work part 1)		Y			
LQ	Length of sequential command table		Y			
LRDP	Length of the deferred caching redo pool in cluster environments.		Y			
LS	Length of sort area		Y			
LU	Length of intermediate user buffer				Y	

Parameter	Usage	No*	LF	LM	GF	GM
LWKP2	Length of ISN list processing area (Work part 2)		Y			
LWP	Length of Adabas work pool		Y			
MODE	Mode of operation	Y				
MSGBUF	Size of the message buffer		Y			
MXMSG	Maximum message reply time between cluster nuclei in Cluster Services or Parallel Services.		Y			
MXTNA	Maximum inactivity time limit override for a user				Y	
MXTSX	Maximum Sx execution time limit override for a user				Y	
MXTT	Maximum transaction time limit override for a user				Y	
NAB	Number of attached buffers		Y			
NC	Number of command queue elements		Y			
NCLOG	Number of command logs		Y			
NH	Number of hold queue elements		Y			
NISNHQ	Number of ISNs in hold queue for user					Y
NPLOG	Number of protection logs		Y			
NONDES	Non-descriptor searches					Y
NQCID	Number of active command IDs per user					Y
NSISN	Number of ISNs per ISN table element		Y			
NT	Number of threads		Y			
NU	Number of user queue elements		Y			
NUCID	Custer nucleus ID		Y			
OPENRQ	Open command required				Y	
PLOGDEV	Multiple protection log device		Y			
PLOGRQ	Protection log required				Y	
PLOGSIZE	Multiple protection log size (blocks)		Y			
PREFETCH**	Prefetch/multifetch feature control (see note below)	Y				
PREFICMD**	Include command from prefetch/multifetch (see note below)	Y				
PREFIFIL**	Include file from prefetch/multifetch (see note below)	Y				
PREFNREC**	Multifetch record count (see note below)	Y				
PREFSBL**	Prefetch single buffer length (see note below)	Y				
PREFTBL**	Prefetch total buffer length (see note below)	Y				
PREFXCMD**	Exclude command from prefetch/multifetch (see note below)	Y				
PREFXFIL**	Exclude file from prefetch/multifetch (see note below)	Y				
QBLKSIZE	Sequential dataset block size (optimized by ADAIOR)		Y			
READONLY	Read-only session control	Y				

Parameter	Usage	No*	LF	LM	GF	GM
REVIEW	Adabas Review control			Y		
SMGT	Error handling (PIN) facility control		Y			
SORTCACHE	Controller caching control for the Adabas sort area component		Y			
SPT	Adabas triggers and stored procedures control				Y	
SVC	SVC number		Y			
TCPIP	TCP/IP access control			Y		
TCPURL	TCP/IP universal resource locator (URL)			Y		
TEMPCACHE	Controller caching for the Adabas temp area component		Y			
IFLUSH	Synchronous buffer flush time	Y				
ILSCMD	Time limit for S1, S2, and S4 complex searches (seconds)				Y	
TNAA	Non-activity time limit (access-only users)					Y
TNAE	Non-activity time limit (ET logic users)					Y
TNAX	Non-activity time limit (exclusive update users)					Y
TT	Transaction time limit					Y
UEXn	User exits: 1, 3, 4, 5, 8		Y			
UEXn	User exits: 2, 12				Y	
UEXn	User exits: 6, 9 (for utilities)	Y				
UTIONLY	Utilities-only session				Y	
VISTA	Adabas Vista control				Y	
WORKCACHE	Controller caching for the Adabas work area component		Y			



Note: * Adabas Cluster Services 7.4 does not support DTP=TM or DTP=RM. It does support DTP=ET, which is a parameter setting that must be the same for all cluster nuclei and cannot be changed online (global, fixed). ** The PREFxxx parameters are used with application programs (PROGRAM=USER) making Adabas calls. They have no effect when specified for an Adabas nucleus.

Specifying ADARUN Parameters for Cluster Nuclei

When specifying ADARUN session parameters for Adabas cluster nuclei:

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

The parameters CLOGMRG, CLUCACHENAME, CLUGROUPNAME, CLULOCKNAME, CLUSTER, DIRRATIO/ELEMENTRATIO, LRDP, MXMSG, and NUCID are used to set up the Adabas Cluster Services sysplex cluster nucleus and its environment.

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

If command logs are used, each nucleus must have its own. If command logs are to be merged, they must be dual or multiple logs and all nuclei must use them. All cluster nuclei must have the same CLOGMRG setting.

The remaining Adabas cluster nucleus parameters are the same as those of a standard Adabas nucleus. See the *Adabas Operations* documentation for more information.

NETWRK Initialization Parameters for Entire Net-Work

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 on page 25.

Example:

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

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

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.

3 Entire Net-Work Parameter Statements

- Statement Format 22
- XCF NODE Statement 23
- XCF DRIVER Statement 33
- XCF LINK Statement 36

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.

This chapter covers the following topics:

Statement Format

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

1. Each Entire Net-Work parameter statement occupies positions 1-72 of at least one line.
2. The statement type (NODE, LINK, or DRIVER) must be specified as the first nonblank item on the statement.
3. The node name, driver name, or link name follows the statement type, separated by at least one blank (space). Specific statement naming conventions are described in the following sections.
4. 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 comma. Refer to the statement descriptions in the following sections for specific parameter information.
5. In the syntax representation of arguments, those shown in lowercase are variables and should be replaced with an appropriate value as described in the text. Uppercase arguments are constants and must be specified as shown. Default values that are in effect if the parameter is not specified are underlined.
6. When the acceptable values for a parameter are Y and N (yes and no), any other value is treated as an N and processing continues without any warning.

7. When the acceptable values for a parameter fall within a range (for example, 1 - 2147483647) and a value outside the range is specified, the value is automatically reset to the maximum value within the range. Processing continues without any warning.
8. When the acceptable values for a parameter fall within a range (e.g., 1 - 2147483647) and a value outside the range is specified, it is automatically reset to the maximum value within the range. Processing continues without any warning.
9. In the discussion of keyword parameters, the underlined characters indicate the minimum abbreviation to identify a keyword.
10. 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.
11. Comment lines begin with an asterisk (*) in position 1 and can be inserted anywhere in the statement sequence.
12. Some keyword parameters (see for example *BUFFERS*) 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)
```

13. Hexadecimal keyword values can be entered by prefixing the value with an "X".

Example:

```
LINK . . . TRACESIZ=X1000, . . .
```

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,ltbufsize,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

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 (`pxbufsize`) is used to set the size of the page-fixed buffer pool.

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

where

<code>abufsize</code>	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.
<code>ltbufsize</code>	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).
<code>stbufsize</code>	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).
<code>pxbufsize</code>	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 NETPRNT 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 dataset.	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 enroute 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: NET0090 BUFFER USAGE STATISTICS
2	provides a severity letter (I for information, W for warning, or E for error) with the message number. For example: NET0090I: BUFFER USAGE STATISTICS
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: NET0090I NODE2 : BUFFER USAGE STATISTICS
4	provides the message number followed by the node name, not padded, of the issuing Entire Net-Work node. For example: NET0090I NODE2: BUFFER USAGE STATISTICS



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
NIDO	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 the section *XCF Line Driver Commands* in the *Operator Commands* documentation. 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.

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

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
GROUP	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 COUPLExx member in SYS1.PARMLIB (see step *Modify COUPLExx 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 the section *XCF Line Driver Commands* in the *Operator Commands* documentation. The open/closed requirement for each parameter is included in its description.


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

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.t	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.

Index

A

- Accepting unknown requests, 34
 - XCF DRIVER statement, 34
- ACCEPTUI parameter, 34
- ACCEPTUI= parameter
 - XCF DRIVER statement, 34
- ADARUN parameters
 - CT
 - command timeout limit, 26
 - determining correct settings, 7
 - example statement, 14
 - summary for Entire Net-Work, 13
- ALTER operator command, 33, 37

B

- Buffers
 - length of intermediate user buffer, 16
 - number of attached buffers, 16
 - parameter for defining, 25
 - size of asynchronous pool, 25
 - size of long-term pool, 25
 - size of page-fixed pool, 25
 - size of short-term pool, 25
- BUFFERS parameter, 25

C

- CLOSE operator command, 33, 37
- Control statements
 - Entire Net-Work, 22
 - order of, 22
- CQTIMER parameter, 26, 31
- CT parameter, 31

D

- D L command, 37
- D S command, 25
- DEFINE LINK operator command, 26
- DEFINE parameter, 26
- DOMAIN parameter, 26
- DRIVER statement
 - defined, 22
 - overview, 33
- DRIVER statement parameters
 - ACCEPTUI

- whether to connect to nodes not defined in LINK statement, 34
- GROUP
 - Entire Net-Work group name, 35
- DUMP parameter, 27

E

- Entire Net-Work
 - ADARUN parameters for, 13
 - request queue, 17

G

- GROUP parameter, 35
- GROUP= parameter
 - XCF DRIVER statement, 35

L

- LINK statement
 - defined, 22
 - defining for XCF, 36
 - overview, 37
- LINK statement parameters
 - linkname, 38
 - PSTATS
 - whether statistics are printed automatically, 38
 - STATINT
 - time allowed before statistics are reported or reset, 39
 - WEIGHT
 - relative weight of this link, 39
- LOG parameter, 28
- LOG= parameter (NODE statement), 28
- Logging checkpoint and data areas
 - LOG= parameter, 28

M

- MAXPATH parameter, 28
- MAXPATH= parameter (NODE statement), 28
- MSGFORM parameter, 29

N

- NID0 parameter, 30
- Node
 - name, specifying the, 24
- NODE statement

- defined, 22
- defining for XCF, 23
- NODE statement parameters
 - BUFFERS
 - space reserved for internal buffer management, 25
 - CQTIMER
 - time allowed to retrieve command results, 26
 - DEFINE
 - accept or reject DEFINE operator command, 26
 - DOMAIN, 26
 - DUMP, 27
 - LOG
 - enable/disable logging of node statistics, 28
 - MAXPATH
 - max. path length for message in network, 28
 - MSGFORM
 - define format of console messages and DDPRINT output, 29
 - NIDO, 30
 - nodename, 24
 - NTRACE
 - specify size of internal trace table, 30
 - REPLYTIM
 - time allowed for user request to complete, 30
 - START
 - enable/disable automatic start, 31
 - TIMER
 - interval between time-dependent requests, 31
 - TRACE, TROFF, TRON
 - enable/disable specific program traces, 31
 - UCMSG
 - set message for uppercase or mixed-case issue, 32
 - ULINK, 32
- NTRACE parameter, 30, 31
- NTRACE= parameter (NODE statement), 30

O

- Operator commands
 - ALTER, 33, 37
 - CLOSE, 33, 37
 - D L
 - display link, 37
 - D S
 - display statistics, 25
 - DEFINE LINK, 26
 - START
 - start individual line driver, 31

P

- Path length
 - defining the maximum, 28
- Path length, defining the maximum
 - MAXPATH= parameter (NODE statement), 28
- PSTATS parameter, 38

R

- REPLYTIM parameter, 30, 31

S

- SET CQTIMER= command, 26

- SET LOG= command, 28
- SET MAXPATH= command, 29
- SET REPLYTIM= command, 30
- SET TRACE=, TROFF=, TRON= command
 - determine extent of tracing activity, 32
- SET UCMSG command, 32
- SET ULINK command, 33
- START operator command, 31
- START parameter, 31
- STATINT parameter, 39
- Syntax
 - DRIVER statement, 33
 - LINK statement, 37

T

- Target IDs
 - allowing duplicates, 19
 - overwriting, 15
- Time
 - waiting time for results (CQTIMER parameter), 24
- TIMER parameter, 31
- Timing parameters
 - interaction between, 31
- Trace control parameters, 31
- TRACE parameter, 31
- Trace table size, 30
- Trace table size, specifying the (NTRACE= parameter), 30
- TRACESIZ= parameter
 - XCF
 - DRIVER statement, 36
- Tracing line driver activity, 31
 - XCF driver, 36
- TROFF parameter, 31
- TRON parameter, 31

U

- UCMSG parameter, 32
- ULINK parameter, 32

W

- WEIGHT parameter, 39

X

- XCF
 - LINK statement, 36
 - NODE statement, 23