

ADARUN Parameter Directory

This chapter describes all of the ADARUN parameters available for Adabas.

Parameter	Use to	Values	Default
AOSlog	Log to DDPRINT the commands issued by ADADBS OPERCOM or equivalent AOS functions that modify the active nucleus.	YES NO	NO
ARExclude	Exclude file(s) from autorestart.	1 - 5000	---
ARMname	Specify the automatic restart management (ARM) program name.	<i>arm-pgm-name</i>	---
ASSocache	Activate or deactivate controller caching for the Associator component (z/OS only).	YES NO	YES
ASYtvs	Flush buffers asynchronously based on volume serial number.	YES NO	YES
CACHe	Load ADACSH (Adabas Caching Facility control) during Adabas session initialization.	YES NO	NO
CACTivate	Control RABN activation. (Adabas Caching Facility)	YES NO	NO
CASSODsp	Cache an Associator RABN or RABN range in the data space cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CASSOExt	Cache an Associator RABN or RABN range in the extended memory cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CASSOG64	Cache an Associator RABN or RABN range backed by 2G large pages in the virtual 64-bit storage cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments.	<i>rabn rabn1 - rabnx</i>	---
CASSOHsp	Cache an Associator RABN or RABN range in the hiperspace cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CASSOL64	Cache an Associator RABN or RABN range backed by large pages in the virtual 64 storage cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments.	<i>rabn rabn1 - rabnx</i>	---
CASSOV64	Cache an Associator RABN or RABN range in the virtual 64 storage cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CASSOMaxs	Specify the unit size in which areas are allocated for Associator cache space. (Adabas Caching Facility)	80K - 2047M (extended memory, data spaces, or hiperspaces) 1M - your installation limit (virtual 64 memory)	1M (extended memory, data spaces, or hiperspaces) 1M (virtual 64 memory)
CBufno	Specify the number of read-ahead buffers for concurrent I/O processing. (Adabas Caching Facility)	0 - the NT parameter value	0
CCTimeout	Specify the cache space area inactivity time limit (in seconds) before it is released back to the system. If demand caching is in effect, use this parameter to specify the inactivity time limit (in seconds) of a RABN range, file or file range, before it is disabled. (Adabas Caching Facility)	60 - 2147483647	7200 seconds (2 hours)
CDATADsp	Cache a Data Storage RABN or RABN range in the data space cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---

Parameter	Use to	Values	Default
CDATAExt	Cache a Data Storage RABN or RABN range in the extended memory cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. Only extended memory is available for BS2000 RISC machines. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CDATAG64	Cache a Data Storage RABN or RABN range backed by 2G large pages in the virtual 64-bit storage cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CDATAHsp	Cache a Data Storage RABN or RABN range in the hiperspace cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CDATAL64	Cache a Data Storage RABN or RABN range backed by large pages in the virtual 64 storage cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CDATAV64	Cache a Data Storage RABN or RABN range in the virtual 64 storage cache. This parameter is mutually exclusive with the CFILE and CXFILE parameters. This parameter is available only in z/OS environments. (Adabas Caching Facility)	<i>rabn rabn1 - rabnx</i>	---
CDATAMaxs	Specify the unit size in which areas are allocated for Data Storage cache space. (Adabas Caching Facility)	80K - 2047M (extended memory, data spaces, or hiperspaces) 1M - your installation limit (virtual 64 memory)	1M (extended memory, data spaces, or hiperspaces) 1M (virtual 64 memory)
CDEmand	Set the lowest Adabas buffer efficiency level accepted before caching is activated. A value of "0" indicates that no demand-level caching occurs. (Adabas Caching Facility)	0 - 2147483647	0
CDIspstat	Control whether RABN range statistics display on both the operator console and DDPRINT, or only on DDPRINT. (Adabas Caching Facility)	YES (display on both) NO (display only on DDPRINT)	YES
CDXnn	Specify the user routine of collation descriptor user exits. The <i>nn</i> in the parameter name can range from 01 through 08, so up to eight collation descriptor exits may be specified (in any order).	user routine name (up to 8 characters)	---
CExclude	Identify command types that should be excluded from read-ahead caching. More than one command type may be specified by separating values with forward slashes (/). (Adabas Caching Facility)	PHYS LOGI HIST FIND	---

Parameter	Use to	Values	Default
CFile	<p>Select a file or range of files to be cached. This parameter is also used to indicate the class of caching service, the caching scope, and the cache storage type that should be implemented for the file or file range. Specify this parameter using the following syntax:</p> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> CFILE={<i>fnr</i> <i>fnr1-fnrx</i>}[<i>,class</i>][<i>,scope</i>][<i>,cachetype</i>] </div> <p>where:</p> <ul style="list-style-type: none"> • <i>fnr, fnr1, fnrx</i>: valid file numbers. When a range is specified, the lower file number must be specified first (<i>fnr1</i>). • <i>class</i>: The percentage of cache space that can be used, which effectively assigns a priority to the files, with lower priority RABNs being purged first. A value of 1 indicates the highest priority, 5 indicates the lowest priority. • <i>scope</i>: Identifies the type of caching used: Associator (A), Data Storage (D), or both (B). • <i>cachetype</i>: Identifies the cache storage type used: data space (D), extended (E), 2G large pages in virtual 64-bit storage (G), hiperspace (H), 1M large pages in virtual 64-bit storage (L), or virtual 64-bit storage (V). Virtual 64-bit storage cache is only available in z/OS environments. <p>This parameter is mutually exclusive with the CASSOxxx, CDATAxxx and CSTORAGE parameters.</p> <p>(Adabas Caching Facility)</p>	<i>fnr, fnr1, fnrx</i> : 0 - n <i>class</i> : 1 2 3 4 5 <i>scope</i> : A D B <i>cachetype</i> : D E G H L V	<i>fnr, fnr1, fnrx</i> : --- <i>class</i> : 3 <i>scope</i> : B <i>cachetype</i> : D
CLOGBmax	Specify the maximum size (bytes) of a logged buffer.	368 - 2147483647	4096
CLOGDev	Specify the device type used for multiple command log data sets. This parameter must be specified with the CLOGSIZE parameter.	Valid device types are listed in <i>Adabas Device Types and Block Sizes</i> .	value of the DEVICE parameter
CLOGLayout	Specify the format of the Adabas command log.	5 8	5
CLOGMAX	Specify the maximum size (bytes) of all of the logged buffers allowed for an Adabas command.	368 - 2147483647	16384
CLOGMRg	Indicate whether the Adabas cluster should run with an automatic CLOG merge. (Adabas Cluster Services, Adabas Parallel Services)	YES NO	NO
CLOGSize	Specify the number of blocks available for each command log in multiple command log data sets. This parameter must be specified with the CLOGDEV parameter.	16 - 16777215	---
CLUCACHEExtra	Specify 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. (Adabas Cluster Services, Adabas Parallel Services)	0 - 2147483647	1000
CLUCACHENAME	Identifies the name of the Adabas cluster cache structure/area in which the current nucleus participates. (Adabas Cluster Services)	<i>cluster-name</i>	---
CLUCACHESize	Specify the amount of storage to allocate for the global cache area that services the Adabas Parallel Services cluster in which the current nucleus participates. (Adabas Parallel Services)	128K - varies (depending on the cache type specified by the CLUCACHETYPE parameter)	---
CLUCACHETYPE	Identify the virtual storage type for the global cache area: shared dataspace (DSP), virtual 64-bit storage (V64), virtual 64-bit cache backed by page-fixed two-gigabyte (2G) large pages (G64), or virtual 64-bit cache backed by page-fixed one-megabyte (1M) large pages (L64). The G64, L64 and V64 values for this parameter are available in z/OS environments only. (Adabas Parallel Services)	DSP G64 L64 V64	DSP

Parameter	Use to	Values	Default
CLUCACHEUnchanged	Indicate whether unchanged blocks are written to global cache during Adabas Parallel Services processing. (Adabas Parallel Services)	YES NO	NO
CLUGroupname	Specify the XCF messaging group name in Adabas Cluster Services or the name of the messaging group in Adabas Parallel Services. (Adabas Cluster Services, Adabas Parallel Services)	<i>groupname</i>	none in Adabas Cluster Services; the router name in Adabas Parallel Services
CLULOCKName	Specify the lock structure name defined in the CFRM policy or the global lock area in which the current nucleus participates. (Adabas Cluster Services)	<i>lockstructurename</i>	---
CLULOCKSize	Specify the amount of storage allocated for the global lock area. Specify values for this parameter using the following syntax: <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"><code>nnn[K M G]</code></div> where <i>nnn</i> : the number of bytes (no unit specification), kilobyte (K), megabytes (M), or gigabytes (G). (Adabas Parallel Services)	131072 - 2147483147 (bytes) 128K - 2097152K 128K - 2048M 128K - 2G	---
CLUSter	Indicate whether the Adabas nucleus will participate in a cluster of nuclei working on the same database, and, if so, the type of cluster in which it will participate. Valid values are "NO" (nucleus not participating in a cluster), "LOCAL" (nucleus participating in a cluster on the same operating system image), and "SYSPLEX" (nucleus participating in a cluster on different systems in a parallel sysplex). (Adabas Cluster Services, Adabas Parallel Services)	NO LOCAL SYSPLEX	NO
CMADdr	Identify the starting address of the GETMAIN common memory pool <i>above</i> the 16M line. This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	X'1000000' - X'FE000000'	0 (X'1000000')
CMAxcsp	Specify the number of storage areas that can be allocated for ADACSH cache space in data spaces, hiperspaces, or virtual 64-bit storage. This parameter has no affect on a cache space in extended memory, where only one Associator and one Data Storage area are allowed. (Adabas Caching Facility)	1 - 16	8
CMDqmode	Indicate where the command queue memory pool should be allocated: below the 16 MB line (BELOW) or above it (ABOVE). This parameter is valid only in BS2000 environments.	ABOVE BELOW	ABOVE (BELOW for Adabas versions prior to Version 8)
CMFix	Indicate whether the GETMAIN common memory pool must have a fixed location <i>above</i> the 16 MB line. This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	YES NO	NO
CMLAddr	Identify the starting address of the GETMAIN common memory pool <i>below</i> the 16M line. This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	X'10000' - X'E00000'	0 (X'10000')
CMLFix	Indicate whether the GETMAIN common memory pool must have a fixed location <i>below</i> the 16 MB line. This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	YES NO	NO
CMLScope	Control access to the GETMAIN common memory pool <i>below</i> the 16 MB line. Valid values are "GROUP" (the pool can be accessed only by other tasks using the same user ID) and "GLOBAL" (the pool can be accessed by all user IDs in the system). This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	GROUP GLOBAL	GROUP
CMLSize	Specify whether the GETMAIN is performed in a common memory pool <i>below</i> the 16 MB line and, if so, its size in bytes. Valid values are "0" (required application memory is not obtained in a common memory pool, but in the class 6 memory of the system) or the size (in bytes) of user storage to allow in a common memory pool (<i>below</i> the 16 MB line) of the specified size, rounded up to the next megabyte. This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	0 <i>bytesize</i>	0

Parameter	Use to	Values	Default
CMSCope	Control access to the GETMAIN common memory pool <i>above</i> the 16 MB line. Valid values are "GROUP" (the pool can be accessed only by other tasks using the same user ID) and "GLOBAL" (the pool can be accessed by all user IDs in the system). This parameter is valid only in BS2000 environments.	GROUP GLOBAL	GROUP
CMSIze	Specify whether the GETMAIN is performed in a common memory pool <i>above</i> the 16 MB line and, if so, its size in bytes. Valid values are "0" (required application memory is not obtained in a common memory pool, but in the class 6 memory of the system) or the size (in bytes) of user storage to allow in a common memory pool (<i>above</i> the 16 MB line) of the specified size, rounded up to the next megabyte. This parameter is valid only in BS2000 environments running Sockets versions less than 2.2.	0 <i>bytesize</i>	0
CRetry	Specify the number of seconds between Adabas Caching Facility attempts to acquire an area of cache space if the previous GETMAIN or space allocation failed. (Adabas Caching Facility)	60 - 2147483647	900 (15 minutes)
CStorage	Identify the type of cache space to which all RABNs for the database are cached (using system defaults). This parameter overrides the use of the CASSOxxx and CDATAxxx ADARUN parameters. Valid values are "DATASPACE" (dataspace), "EXTENDED" (extended memory), "HIPERSPACE" (hiperspace), "G64" (2G large pages in virtual 64-bit storage), "L64" (1M large pages in virtual 64-bit storage), or "VIRTUAL64" (virtual 64-bit storage). "EXTENDED" is the only option available on BS2000 RISC machines; "HIPERSPACE", "G64", "L64", and "VIRTUAL64" are only valid in z/OS environments. This parameter is mutually exclusive with the CFILE and CXFILE parameters. (Adabas Caching Facility)	DATASPACE EXTENDED G64 HIPERSPACE L64 VIRTUAL64	---
CT	Specify the maximum time limit (in seconds) for interregion communication of results from Adabas to the user.	1 - 16777215	60
CWORKStorage	Identify the type of cache space used for Adabas Caching Facility work areas CWORK2FAC and CWORK3FAC. Valid values are "DATASPACE" (dataspace), "EXTENDED" (extended memory), "G64" (2G large pages in virtual 64-bit storage), "HIPERSPACE" (hiperspace), "L64" (1M large pages in virtual 64-bit storage), or "VIRTUAL64" (virtual 64-bit storage). "EXTENDED" is the only option available on BS2000 RISC machines; "HIPERSPACE", "G64", "L64", and "VIRTUAL64" are only valid in z/OS environments. (Adabas Caching Facility)	DATASPACE EXTENDED G64 HIPERSPACE L64 VIRTUAL64	---
CWORK2fac	Specify the percentage of Work part 2 to be cached. If "0" is specified, Work part 2 is not cached. (Adabas Caching Facility)	0 - 100	0
CWORK3fac	Specify the percentage of Work part 3 to be cached. If "0" is specified, Work part 3 is not cached. (Adabas Caching Facility)	0 - 100	0

Parameter	Use to	Values	Default
CXfile	<p>Select a file or range of files to be excluded from caching. Specify this parameter using the following syntax:</p> <pre>CXFILE=({fnr fnr1-fnrx}[, ,scope])</pre> <p>where:</p> <ul style="list-style-type: none"> • <i>fnr, fnr1, fnrx</i>: valid file numbers to be excluded. When a range is specified, the lower file number must be specified first (<i>fnr1</i>). • <i>scope</i>: Identifies the scope of exclusion: Associator RABNs only excluded (A), Data Storage RABNs only excluded (D), or both types of RABNs excluded (B). If a scope is specified, you must specify two commas before it. <p>This parameter is mutually exclusive with the CASSOxxx, CDATAxxx, and CSTORAGE ADARUN parameters.</p> <p>(Adabas Caching Facility)</p>	<p><i>fnr, fnr1, fnrx</i>: 0 - n</p> <p><i>scope</i>: A D B</p>	<p><i>fnr, fnr1, fnrx</i>: ---</p> <p><i>scope</i>: B</p>
DAtacache	Control whether caching control is enabled (YES) or disabled (NO) for the Data Storage data set. This parameter is valid only in z/OS environments.	YES NO	YES
DBid	Specify the physical database ID of the database. Note: If you specify the DBID parameter, you must also specify the ADARUN SVC parameter.	1 - 65535	1
DEvice	Specify the device type on which the first block of the Associator is stored. Valid device types for each platform are listed in <i>Adabas Device Types and Blocks Sizes</i> .	<i>devtype</i>	2300 (BS2000) 3390 (z/OS, z/VSE)
DIRratio	Together with the ELEMENTRATIO ADARUN parameter, define the ratio of directory entries to data elements in Adabas cluster environments. (Adabas Cluster Services, Adabas Parallel Services)	1 - 32767	1
DSF	Enable Adabas Delta Save support. (Adabas Delta Save)	YES NO	NO
DSFEx1	Specify the user routine to be given control by Adabas Delta Save. (Adabas Delta Save)	<i>exitname</i>	---
DTp	Indicate whether to enable distributed transaction processing, and, if so, whether the nucleus should be a resource or transaction manager. Valid values are NO (distributed transaction processing is not enabled), RM (the nucleus participates in distributed transaction processing as a resource manager), and TM (the nucleus participates in distributed transaction processing as a transaction manager). (Adabas Transaction Manager)	NO RM TM	NO
DUALCLD	Specify the device type to be used for dual command logging. Valid device types for each platform are listed in <i>Adabas Device Types and Blocks Sizes</i> .	<i>devtype</i>	value of the DEVICE parameter
DUALCLS	Specify the number of blocks available for each dual command log.	16 - 16777215	---
DUALPLD	Specify the device type to be used for dual protection logging. Valid device types for each platform are listed in <i>Adabas Device Types and Blocks Sizes</i> .	<i>devtype</i>	value of the DEVICE parameter
DUALPLS	Specify the number of blocks available for each dual protection log.	16 - 16777215	---
ELementratio	Together with DIRRATIO ADARUN parameter, define the ratio of directory entries to data elements in Adabas cluster environments. Adabas Cluster Services, Adabas Parallel Services	1 - 32767	1
EXcpvr	Indicate whether EXCP or EXCPVR should be used when running APF-authorized. This parameter is valid only in z/OS environments.	YES NO	YES

Parameter	Use to	Values	Default
FAstpath	Enable Adabas Fastpath support. (Adabas Fastpath)	YES NO	NO
FMxio	Set a limit on the number of I/O operations that can be started in parallel by LFIOP flush processing. The full meaning of this parameter varies, depending on the setting of the ASYTVS ADARUN parameter. If ASYTVS=YES, you can use FMXIO to set the limit on the number of I/O <ul style="list-style-type: none"> If ASYTVS=YES, you can use FMXIO to set a limit on the number of I/O operations to be started in parallel on each volume. If ASYTVS=NO, you can use FMXIO to set a limit on the number of I/O operations to be started in parallel overall. 	1 - 16 (ASYTVS=YES)	1 (ASYTVS=YES)
		1 - 100 (ASYTVS=NO)	60 (ASYTVS=NO)
FOrce	Indicate whether the nucleus or Adabas Review hub can overwrite an existing ID table entry.	YES NO	NO
GRoups	Indicate whether BS2000 interprocess communication is limited to users with the same logon ID. This parameter is available only in BS2000 environments.	YES NO	NO
HEXnn	Specify the hyperdescriptor exit(s) and associated user routine to be used by the nucleus. Specify this parameter using the following syntax: <div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> <p>HEXnn=user-routine-name</p> </div> where: <ul style="list-style-type: none"> nn: the number of the hyperdescriptor exit user-routine-name: the name (up to eight characters) of the user routine that gets control for the associated hyperdescriptor exit. 	<ul style="list-style-type: none"> nn: 01 - 31 user-routine-name: text (up to eight characters) 	<ul style="list-style-type: none"> nn: --- user-routine-name: ---
IDTName	Define the name for an alternate ID table, resulting in an alternate Adabas environment. This parameter is valid only in BS2000 environments. Valid names must be begin with the letters "ADA" and must be eight characters long.	ADAxxxxx	ADABAS5B
IDTPsup	Suppress the IDT name in Adabas output messages. This parameter is valid only in BS2000 environments.	YES NO	NO
IGNDib	Indicate whether an active nucleus entry in the data integrity block (DIB) should be ignored and deleted.	YES NO	NO
IGNDTp	Indicate whether the two-phase commit area (Work part 4) should be ignored.	YES NO	NO
INFobuffersize	Specify the size of the information buffer pool, in bytes. The information buffer pool is also known as the Adabas event log.	0 or 1024 - 2147483647	0 (buffer will not be allocated)
INTAuto	Specify the time interval (in seconds) between autorestart progress messages, which will be printed in addition to the standard messages at the end of each autorestart phase, if the phase takes longer than the interval specified. The default (0 seconds) indicates that autorestart progress messages should only be printed at the completion of each autorestart phase.	0 - 16777215	0
INTNas	Specify the time interval (in units of 1.048576 seconds) between SYNS 60 checkpoints.	1 - 16777215	3600 (approximately one hour)
LARgepage	Indicate whether the Adabas nucleus should use page-fixed one-megabyte (1M) large pages above the two-gigabyte (2G) bar (L64) or page-fixed 2G large pages above the 2G bar (G64). A value of "YES" is equivalent to specifying "L64". A value of "NO" indicates that virtual storage above the 2G bar will be backed by four-kilobyte (4K) pages. Values of "G64", "L64", and "YES" can only be specified if the ADARUN V64BIT parameter is also set to "YES". In addition, the operating system must support large pages of the required type. Note: At this time, this parameter is applicable only to z/OS systems.	G64 L64 NO YES	NO

Parameter	Use to	Values	Default
LBP	Specify the maximum number of bytes to be used for the Adabas buffer pool during a session.	80000 - 2147483647	350000
LCP	Specify the maximum number of bytes to be used for the Adabas security pool during a session.	2000 - 16777215	10000
LDEuqp	Specify the number of bytes to be used for the unique (UQ) descriptor pool during a session. Note: When you apply zap AN822111, the minimum setting size of the ADARUN LDEUQP parameter is increased from zero to 5000 bytes.	5000 - 16777215	5000
LDTp	Define the length of the distributed transaction processing area (Work part 4) or the index of Work part 4: If you have Adabas Transaction Manager Version 7.4 or earlier installed, this parameter defines the length of the distributed transaction processing area (Work part 4) for use with DTP=RM. If you have Adabas Transaction Manager Version 7.5 or later installed, this parameter defines the size (in blocks) of the Work part 4 index which is used to administer the data on Work part 4 (DDWORKR4). (Adabas Transaction Manager)	Adabas Transaction Manager 7.4 or earlier: 0 - 65535 Adabas Transaction Manager 7.5 or later: 8 - one eighth of the WORK4 size	Adabas Transaction Manager 7.4 or earlier: 0 Adabas Transaction Manager 7.5 or later: 8
LFTop	Enable asynchronous buffer flush operations and define the length of the related buffer flush pool. A value of "0" disables asynchronous buffer flush operations; only synchronous flushing occurs. A value of "1" enables asynchronous buffer flush operations, and defines the buffer pool size as 25% of the buffer pool (LBP parameter) size. Any other value (with a minimum of 80000 allowed) enables asynchronous buffer flush operations and defines the size of the pool directly (in bytes).	0 1 80000 - 18,446,744,073,709,551,615 (although this is really limited by the available amount of virtual and real memory)	0
LFP	Specify the size (in bytes) of the Adabas internal format buffer pool.	6000 - 2147483647	12000
LI	Specify the size (in bytes) allocated for the table of ISNs (TBI) used to store ISN lists (overflow ISNs or saved ISN lists).	2000 - 16777215	10000
LNKGNAME	The name of the link globals table to be employed by an Adabas 8 batch/TSO link routine. Note: If you specify the LNKGNAME parameter, you must also specify the ADARUN DBID and SVC parameters.	valid module name for an IBM z/OS or z/VSE operating system, or for a BS2000 module	LNKGBLS
LOCAl	Indicate whether an Adabas nucleus or Adabas Review hub is isolated and available for local use only. The isolated nucleus or hub will be unknown to the network.	YES NO	YES
LOGAbdx	Indicate whether Adabas buffer descriptions (ABDs) should be logged during this session. For the dependencies of the various logging parameters, please review the detailed description.	YES NO	NO
LOGCB	Indicate whether Adabas extended control blocks (ACBX) should be logged during this Adabas session. For the dependencies of the various logging parameters, please review the detailed description.	YES NO	NO
LOGCLex	Indicate whether the Adabas command log extension (CLEX) should be logged. For the dependencies of the various logging parameters, please review the detailed description.	YES NO	NO
LOGFb	Indicate whether Adabas format buffers (FB) should be logged.	YES NO	NO
LOGGing	Indicate whether command logging should be used for the Adabas session. If this is set to "NO", all other LOGxxxx parameters are ignored.	YES NO	NO
LOGIB	Indicate whether Adabas ISN buffers (IB) should be logged.	YES NO	NO
LOGIO	Indicate whether Adabas I/O activity should be logged.	YES NO	NO
LOGMb	Indicate whether Adabas multifetch buffers (MB) should be logged. For the dependencies of the various logging parameters, please review the detailed description.	YES NO	NO
LOGRb	Indicate whether Adabas record buffers (RB) should be logged.	YES NO	NO
LOGSB	Indicate whether Adabas search buffers (SB) should be logged.	YES NO	NO

Parameter	Use to	Values	Default
LOGSIze	Specify the maximum block size for the sequential command log.	100 - 32760	value of the QBLKSIZE ADARUN parameter
LOGUX	Indicate whether Adabas user exit B data should be logged. For the dependencies of the various logging parameters, please review the detailed description.	YES NO	NO
LOGVB	Indicate whether Adabas value buffers (VB) should be logged.	YES NO	NO
LOGVolio	Indicate whether the extended I/O list should be written to the command log for CLOGLAYOUT=5 and CLOGLAYOUT=8.	YES NO	NO
LOGWarn	Specify the frequency, in seconds, at which the PLOG and CLOG status is checked and resulting alert messages are produced. This parameter is optional, but a valid user exit 2 or user exit 12 must also be specified when this parameter is set to any value other than zero (0). A value of zero (0), the default, indicates that no CLOG and PLOG status checks should occur and that no alert messages should be produced. Any other valid LOGWARN value represents the interval (in seconds) at which the PLOG and CLOG status is checked and resulting alert messages are produced.	0 - 2147483647	0
LP	Specify the number of blocks to allocate to the data protection area (Work part 1). The highest value you can set for this parameter is limited by the size of the Work data set and the sizes of Work parts 2, 3, and 4.	200 - *	1000
LQ	Specify the size (in bytes) of the table of sequential commands, which contains entries required during the processing of Adabas read sequential (L2/L5, L3/L6, and L9) commands.	2000 - 16777215	10000
LRDp	Specify the size (in bytes) of the redo pool, which allows for deferred publishing. If the value of this parameter is set to 0, updated database blocks are always written to global cache at the time the update is made; no deferred publishing occurs. (Adabas Cluster Services, Adabas Parallel Services)	0 80000 - 2137483647	For Adabas Parallel Services 8.2 or later, the default is 0 (zero). For Adabas Cluster Services 8.2 or later, the default is either the value of the ADARUN LFIOP parameter or 10 megabytes, whichever is smaller. For earlier releases of Adabas Parallel Services and Adabas Cluster Services, the default is the value of the ADARUN LFIOP parameter.
LRPI	Specify the size of the Adabas or Event Replicator replication pools, in bytes. This parameter should be specified for only: <ul style="list-style-type: none"> an Adabas nucleus that also has REPLICATION=YES set. an Event Replicator database The values for the Adabas nucleus and the Event Replicator database may be different.	20000 - *	100000 (if REPLICATION=YES) 0 (if REPLICATION=NO)
LS	Specify the maximum number of bytes for internal sort processing. Note: To ensure that a single search command does not use too much of the work pool space, the maximum LS value must be less than or equal to the LWP (work pool length) divided by two minus the minimum LS value (19968) (LS = LWP / 2 - 19968).	19968 - *	49920

Parameter	Use to	Values	Default
LU	Specify the length (in bytes) of the intermediate user buffer area. The range of values varies based on the operating system on which Adabas is installed. On z/OS systems, LU cannot exceed a value greater than that produced by this calculation: $(NABvalue \times 4096)$ On z/VSE and BS2000 systems, LU cannot exceed a value greater than that produced by this calculation: $(NABvalue \times 4096) - 256$	a positive integer, minimum and maximum governed by the operating system.	65535
LWKP2	Specify the number of Work blocks to use for intermediate ISN lists Work part 2).	0 - 16777215	0
LWP	Specify the size (in bytes) of the Adabas work pool. This value must be equal to or greater than twice the sum of the specified internal sort area length (LS parameter) plus the minimum LS value (19968). In other words, the LWP value must be greater than or equal to a value calculated as follows: $(LS-value + 19968) \times 2$ The minimum value you can specify is 80000 or a value not less than the value of the NT parameter multiplied by 25000.	80000-2147483647	150000
MLwto	Specify 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. (Event Replicator for Adabas only)	YES NO	NO
MOde	Identify the Adabas operating mode. Valid operating modes include running in single-user mode (SINGLE) or running in multiuser mode (MULTI).	SINGLE MULTI	MULTI
MSGBuf	Specify the size (in kilobytes) of the message buffer area, which is used to buffer messages for use for performance and tuning and problem analysis.	0 - 128	0
MSGConsl	Specify the case of Adabas messages sent to the console.	Upper Mixed	Mixed
MSGDruck	Specify the case of Adabas messages sent to the DD/DRUCK data set.	Upper Mixed	Mixed
MSGPrint	Specify the case of Adabas messages sent to the DD/PRINT data set.	Upper Mixed	Mixed
MXCANCEL	Specify the timeout threshold, in seconds, in which processing of an intracluster cancellation request must occur. Adabas Cluster Services, Adabas Parallel Services	15 - 2147483647	The value of the MXMSG parameter.
MXCANCELWarn	Specify the timeout threshold, in seconds, for an intracluster cancellation request to occur before issuing an early warning about the failure of the cancellation request. The value of this parameter must be less than the value of the MXCANCEL parameter. A value of "0" indicates that no warnings should be issued. Adabas Cluster Services, Adabas Parallel Services	0 4 - (MXCANCEL_value - 1)	The value of the MXCANCEL parameter divided by 4.
MXMSG	Specify the timeout threshold, in seconds, in which the processing of an intracluster command must occur. Adabas Cluster Services, Adabas Parallel Services	15 - 32767	300
MXMSGWarn	Specify the timeout threshold, in seconds, for a response to an intracluster command to occur before issuing an early warning about the failure of the request. The value of this parameter must be less than the value of the MXMSG parameter. A value of "0" indicates that no warnings should be issued. Adabas Cluster Services, Adabas Parallel Services	0 4 - (MXMSG_value - 1)	The value of the MXMSG parameter divided by 4.

Parameter	Use to	Values	Default
MXStatus	Specify the interval, in seconds, at which the member-level XCF status monitoring heartbeat should occur. A value of "0" indicates that no member-level status monitoring should occur. Adabas Cluster Services	0 15 - 21474836	15
MXTNa	Set the maximum non-activity time, in seconds, that can be specified in the OP command for an individual user.	1 - 65535	3600
MXTSx	Set the maximum execution time (in seconds) that can be specified for the ADARUN TLSCMD parameter and for any override execution time setting that might be specified in the Adabas control block of an OP command.	1 - 65535	3600
MXTT	Set the maximum transaction time limit, in seconds, that can be specified in the OP command for an individual user.	1 - 65535	3600
MXWtor	Specify the timeout interval, in seconds, in which an operator response to a self-termination operator query must occur before issuing a warning. A value of "0" indicates that no warning message should be issued. Adabas Cluster Services	0 15 - 64800	0
NAb	Specify the number of attached buffers to be used during the session. The maximum value you can specify for this parameter varies, depending on the amount of available virtual storage.	1 -varies	16
NC	Specify the maximum number of command queue elements (CQEs) for the session. The number of CQEs established for the session determines the maximum number of Adabas commands that may be queued or be in process at any one time during the session.	20 - 32767	200
NCLog	Specify the maximum number of command log data sets to be opened and logged in the PPT for the session. A value of "0" disables this parameter.	0 2 - 8	0
NH	Specify the maximum number of hold queue elements for the session.	20 - 16777215	500
NIsnhq	Specify the maximum number of records that can be placed in hold status (hold queues) at the same time by a single user. The maximum value that can be specified for this parameter is the smaller of 65535 or the value of the NH parameter divided by 4. The default is the larger of 20 or the value of the NH parameter divided by the value of the NU parameter with the result multiplied by two ((NH/NU) * 2).	1 - 65535 NH/4 (whichever is smaller)	20 (NH/NU) * 2 (whichever is larger)
NKsp	Indicate whether to span large sequential file records on no-keyed disk files to avoid the DMS0BBB record size error. (BS2000 environments only)	YES NO	NO
NOndes	Indicate whether to allow (YES) searches using nondescriptors in search buffers.	YES NO	YES
NPLOG	Specify the maximum number of protection log data sets to be opened and logged in the PPT. A value of "0" disables this parameter.	0 2 - 8	0
NPLOGBuffers	Specify the number of PLOG protection I/O buffers allocated in the Adabas nucleus session. PLOG protection I/O buffers are only allocated when ADARUN parameter LFIOP is set to a nonzero value.	1 - 2,147,483,687	1
NQcid	Specify the maximum number of active command IDs per user.	1 - 65535	20
NSisn	Specify the maximum number of ISNs per ISN table (TBI) element. The maximum value that can be specified varies based on the device type of the Work data set or, in the case where user-defined Work block sizes are used, the maximum value is based on a formula. For more information, refer to the more detailed NSISN documentation.	7 - *	51
NT	Specify the number of threads to be used during the Adabas session.	4 - 250	5
NU	Specify the maximum number of user queue elements for the Adabas session.	20 - 16777215	200
NUCid	Specify the ID of an Adabas nucleus in an Adabas cluster. Adabas Cluster Services, Adabas Parallel Services	0 - 65000	0

Parameter	Use to	Values	Default
NWork1buffers	Specify the number of Work part 1 protection I/O buffers allocated in the Adabas nucleus session. Work part 1 protection I/O buffers are only allocated when ADARUN parameter LFIOP is set to a nonzero value.	1 - 2,147,483,687	1
OPenrq	Indicate whether an open command (OP) must be issued (is required) as the first command of a user session.	YES NO	YES
PAmreqs	BS2000 only. Set the maximum number of PAM blocks that can be transferred in one PAM macro request. The default can be reduced to avoid the receipt of I/O errors DMS09AC -- too many simultaneous I/O requests on some machines.	1 - 1024	80
PGfix	Indicate whether pages containing I/O control blocks are released after I/O processing is completed or after the job has ended when using EXCPVR on z/OS systems. When this parameter is set to YES, pages containing the I/O control blocks are fixed for the duration of the job. When set to "NO", pages are fixed only for the duration of I/O processing.	YES NO	NO
PLOGDev	Specify the device type used for multiple protection log data sets. This parameter must be specified with the PLOGSIZE parameter.	Valid device types are listed in <i>Adabas Device Types and Block Sizes</i>	value of the DEVICE parameter
PLOGRq	Indicate whether a protection log is required for the Adabas session. For details on the meanings of the different values for this parameter, refer to the detailed PLOGRQ documentation.	YES NO SEL FORCE	YES
PLOGSIze	Specify the number of blocks available for each protection log in multiple protection log data sets. This parameter must be specified with the PLOGDEV parameter.	16 - 16777215	---
PREFEtch	Indicate whether the Adabas command-level prefetch or multifetch features should be automatically enabled. YES enables the multifetch (M) option; OLD enables the prefetch (P) option; NO disables the prefetch feature.	YES OLD NO	NO
PREFICmd	Identify the commands that should be included in prefetch or multifetch processing. To specify multiple values for this parameter, list the values separated by forward slashes (/). For example: PREFICMD=L3/L9. The maximum number of entries is 50. The same number of list items must be specified for this parameter as are specified for the PREFIFIL parameter. If you want several commands to be prefetched or multifetched for a single file, repeat the same file number in the PREFIFIL parameter. The values for this parameter are used in conjunction with the value of the PREFIFIL parameter to customize prefetch and multifetch processing during the session. You cannot specify this parameter in the same session as the PREFXCMD parameter; if you do, the values of the PREFXCMD parameter are used and the values for this parameter are ignored.	L1 - L9	---
PREFIFil	Identify the files that should be included in prefetch or multifetch processing. To specify multiple values for this parameter, list the values separated by commas. The maximum number of entries is 50. The same number of list items must be specified for this parameter as are specified for the PREFICMD parameter. If you want several commands to be prefetched or multifetched for a single file, repeat the same file number in the PREFIFIL parameter. The values for this parameter are used in conjunction with the value of the PREFICMD parameter to customize prefetch and multifetch processing during the session. You cannot specify this parameter in the same session as the PREFXFIL parameter; if you do, the values of the PREFXFIL parameter are used and the values for this parameter are ignored.	1 - 5000	---
PREFNrec	Specify the number records that can be prefetched or multifetched. The maximum number of records that can be multifetched depends on the maximum that can be held by the record and ISN buffers.	0 - *	0
PREFSBl	Specify the size, in bytes, of the internal buffer used by multifetch/prefetch processes, which determines how much record data is prefetched.	1024 - <PREFTBL value>	3000

Parameter	Use to	Values	Default
PREFSTdd	Specify the job statement label (DD name) in the nucleus startup job that defines the location to which multifetch statistics should be written.	<i>ddname</i>	---
PREFTbl	Specify the size, in bytes, of the total prefetch buffer area.	6144 - 16777215	30000
PREFXCmd	<p>Identify the commands that should be <i>excluded</i> in prefetch or multifetch processing.</p> <p>To specify multiple values for this parameter, list the values separated by forward slashes (/). For example: PREFXCMD=L3 / L9. The maximum number of entries is 50. The same number of list items must be specified for this parameter as are specified for the PREFXFIL parameter. If you want several commands to be prefetched or multifetched for a single file, repeat the same file number in the PREFXFIL parameter.</p> <p>The values for this parameter are used in conjunction with the value of the PREFXFIL parameter to customize prefetch and multifetch processing during the session. You cannot specify this parameter in the same session as the PREFICMD parameter; if you do, the values of this parameter are used and the values for the PREFICMD parameter are ignored.</p>	L1 - L9	---
PREFXFil	<p>Identify the files that should be excluded in prefetch or multifetch processing. To specify multiple values for this parameter, list the values separated by commas. The maximum number of entries is 50. The same number of list items must be specified for this parameter as are specified for the PREFXCMD parameter. If you want several commands to be prefetched or multifetched for a single file, repeat the same file number in the PREFXFIL parameter.</p> <p>The values for this parameter are used in conjunction with the value of the PREFICMD parameter to customize prefetch and multifetch processing during the session. You cannot specify this parameter in the same session as the PREFXFIL parameter; if you do, the values of the PREFXFIL parameter are used and the values for this parameter are ignored.</p>	1 - 5000	---
PROGram	Identify the program to be run.	ADACOM ADANUC ADAREV NETWRK RENTUSER USER <i>utility-name</i>	USER
QBlksize	Specify the block size, in bytes, for sequential data sets used by ADAIOR, the Adabas I/O component. A value of "0" does not indicate that 0-byte block sizes should be used, but that the block size should be determined by the media type. For more information, read the detailed QBLKSIZE documentation.	0 - 262144	0
REAdonly	Indicate whether the Adabas session is a read-only session or it should permit database updates.	YES NO	NO
REPllication	<p>Indicate whether replication should be activated for the Adabas nucleus.</p> <p>(Event Replicator for Adabas)</p>	YES NO	NO
REVFilter	<p>Indicate whether Adabas Review record filtering should be allowed during the session.</p> <p>(Adabas Review)</p>	YES NO	YES
REVIew	<p>Indicate whether Adabas Review should be run in local mode, hub mode (specifying the hub ID), or not at all.</p> <p>(Adabas Review)</p>	NO LOCAL <i>hubid</i>	NO
REVLOGBmax	<p>Specify the maximum allowable number of bytes of a logged buffer for Adabas Review.</p> <p>(Adabas Review)</p>	0 - 30000	5120
REVLOGMax	<p>Specify the maximum size of all of the logged buffers allowed for an Adabas Review command.</p> <p>(Adabas Review)</p>	2000 - 32768 (32K)	16384

Parameter	Use to	Values	Default
RPLCONNECTCount	Specify the number of connection attempts made for the Adabas or Event Replicator Server nucleus after an attempt fails.	0 - 2147483647	0
RPLCONNECTInterval	Specify the interval (in seconds) between connection attempts made for the Adabas or Event Replicator Server nucleus after an attempt fails.	0 - 2147483647	0
RPLParms	Identify the location where your replication definitions (initialization parameters) should be read from. Replication definitions can be read from the Replicator system file (FILE), from DDKARTE (PARMS), both the Replicator system file and from DDKARTE (where the DDKARTE specifications override any duplicates with the Replicator system file), or not at all (NONE). (Event Replicator for Adabas)	BOTH FILE NONE PARMS	If the Replicator system file resides on the Event Replicator database, definitions are read first from the Replicator system file and then from DDKARTE. If the Replicator system file does not reside on the Event Replicator database, the definitions are read from DDKARTE.
RPLSort	Indicate whether Event Replicator for Adabas transaction data sorting should occur. (Event Replicator for Adabas)	YES NO	YES
RPWARNINCrement	Specify the interval (in percentage of LRPL usage) at which warning messages should be sent indicating that the replication pool usage has exceeded the threshold set by the RPWARNPERCENT parameter. (Event Replicator for Adabas)	1 - 99	10
RPWARNINterval	Specify the interval, in seconds, during which replication pool usage warning messages are suppressed on the console. (Event Replicator for Adabas)	1 - 2147483647	60
RPWARNMessageLimit	Specify the number of replication pool usage warning messages that can be issued to the console before message suppression begins. (Event Replicator for Adabas)	1 - 2147483647	5
RPWARNPercent	Specify the threshold for replication pool usage (expressed as a percentage of LRPL) at which replication pool usage warning messages should be sent. A value of 0 for this parameter indicates that no warning messages are issued (Event Replicator for Adabas)	0 - 99	0
RVclient	Specify whether Adabas Review client reporting should be activated when you want to run client reports in batch environments. Note: This parameter is valid only in z/OS environments. (Adabas Review)	ACTIVE INACTIVE	INACTIVE
SMF	Specify whether or not Adabas SMF recording should be enabled.	YES NO	NO
SMF89	Specify whether or not the Adabas nucleus should register with z/OS for type 89 SMF records.	YES NO	NO

Parameter	Use to	Values	Default
SMFDETAIL	<p>Specify 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.</p> <p>When NONE or ALL are specified for this parameter, they should be specified alone; no other values can be combined with NONE or ALL.</p> <p>Valid detail section names are ALL, CMD, FILE, IODD, NONE, PARM, STG, THRD, and USER. List of section names must be separated by commas.</p> <p>In cluster environments (Adabas Cluster Services and Adabas Parallel Services), additional detail section names can be specified: CSHB, CSHF, CSHG, CSHP, LOCK, MSGB, MSGC, and MSGH.</p>	(NONE ALL <i>section-name-list</i>)	(NONE)
SMFINTERVAL	Specify whether SMF interval records should be generated and the interval at which they should be generated.	NONE GLOBAL SUBSYS <i>minutes</i>	NONE
SMFRECNO	Specify the Adabas SMF record number used for user-defined SMF records.	128 - 255	255
SMFSUBSYS	Specify the name of the IBM or user-defined SMF subsystem from PARMLIB member SMFPRMxx.	CURRENT <i>subsystem-name</i>	CURRENT
SMGt	<p>Indicate whether the error handling and message buffering facility should be enabled for the session.</p> <p>This parameter is not available on BS2000 systems.</p>	YES NO	NO
SOrtcache	Indicate whether to enable (YES) or disable (NO) controller caching for the Adabas Sort data set. This parameter is valid only in z/OS environments.	YES NO	YES
SPt	Indicate whether use of Adabas triggers and stored procedures should be enabled for the session.	YES NO	NO
SRlog	Indicate how spanned records should be logged to the protection logs.	ALL UPD PART	UPD
SUBmpsz	Specify the common memory pool size, in bytes, for subtask communication in products such as Adabas Parallel Services, Adabas Review, and Event Replicator for Adabas. This parameter is valid only in BS2000 environments.	100000 - *	1024000
SVC	Specify the Adabas SVC number or Adabas Review hub SVC number to be used for the session.	<p>z/OS: 200 - 255</p> <p>z/VSE: any free SVC value, although 45 is recommended.</p>	<p>z/OS: 249</p> <p>z/VSE: 45</p>
SWitchnr	BS2000 systems only. Specify the job switch number that is set when a nucleus or utility exits with a non-zero response code.	1 - 31	10
TAPerel	Specify how end-of-file (EOF) processing is handled for tapes and cartridges. This parameter is valid only in BS2000 environments.	KEEPUNL KEEP NO RELEASE UNLOAD	KEEPUNL
TARgetid	Specify the unique Entire Net-Work target ID for this node.	1 - 65535	1
TASkctgy	Identify the Adabas task category. This parameter is valid only in BS2000 environments.	BATCH TP	BATCH
TCPIp	Indicate whether a direct TCP/IP link to the Adabas nucleus should be activated for this session. This parameter is valid only in z/OS environments.	YES NO	NO
TCPUrl	Identify the universal resource locator (URL) for the direct TCP/IP link to the Adabas nucleus (when TCPIP is set to "YES"). This parameter is valid only in z/OS environments.	<i>url</i>	---
TEmpcache	Indicate whether to enable (YES) or disable (NO) controller caching for the Adabas Temp data set. This parameter is valid only in z/OS environments.	YES NO	YES
TFlush	Specify the time, in seconds, to allow for a synchronous buffer flush.	1 - 16777215	1
TLscmd	Specify the maximum time, in seconds, to be used to process a single Adabas S1, S2, or S4 command with complex search criteria.	1 - the setting of the ADARUN MXTSX parameter	300

Parameter	Use to	Values	Default
TMDrq	Set the number of entries allowed in the Adabas Transaction Manager's internal request queue. (Adabas Transaction Manager)	10 - 32767	10
TMEtdata	Identify the database or databases that will store Adabas Transaction Manager's ET data. (Adabas Transaction Manager)	ATM TARGETS	TARGETS
TMGtt	Set the time limit, in seconds, during which a global transaction can be open without being prepared. (Adabas Transaction Manager)	1 - 16777215	720
TMLog	Identify the logging option for the current execution of the Adabas Transaction Manager. (Adabas Transaction Manager)	YES NO NEVER	NO
TMMsgsev	Set the severity threshold for the suppression of Adabas Transaction Manager warning messages. If this parameter is set to "0", no messages are suppressed. (Adabas Transaction Manager)	0 4 8	0
TMRestart	Identify how restart processing of problematic transactions should be handled. (Adabas Transaction Manager)	NORMAL FORCE FORCEALL	NORMAL
TMSyncmgr	Indicate whether the Adabas Transaction Manager is to interact with an external transaction coordinator. (Adabas Transaction Manager)	NONE RRMS	NONE
TMTcidpref	Define the first one or two characters (the prefix) of dynamically allocated client IDs. Up to two alphanumeric characters can be specified. (Adabas Transaction Manager)	<i>prefix</i>	TM
TNAA	Specify the maximum elapsed time, in seconds, that an access-only user may be considered active without issuing an Adabas command.	1 - 16777215	900
TNAE	Specify the maximum elapsed time, in seconds, that an ET logic user may be considered active without issuing an Adabas command.	1 - 16777215	900
TNAX	Specify the maximum elapsed time, in seconds, that an exclusive control user may be considered active without issuing an Adabas command.	1 - 16777215	900
TT	Specify the maximum elapsed time, in seconds, permitted for a logical transaction issued by an ET logic user.	1 - 16777215	900
UEX nm	Activate a user exit and specify the name of the user routine for the exit. The value of nm in the parameter name must be an integer in the range from 1 through 12, inclusive that identifies the number of the user exit. The value assigned the parameter must be a one to eight-character name of the corresponding user routine that gets control. Note: User exit 2 (UEX2) and user exit 12 (UEX12) are mutually exclusive.	<i>user-routine-name</i>	---
UEXSMF	Specify the name of the SMF user exit module.	<i>module-name</i>	---
UTionly	Indicate whether the session should be restricted to Adabas utilities only.	YES NO	NO
V64bit	Indicate whether the Adabas nucleus should use virtual storage above the 2 gigabyte bar. In order for the nucleus to use 64-bit virtual storage, the operating system must also support 64-bit virtual storage. Note: This parameter is applicable only to z/OS systems.	YES NO	NO
Vlsta	Indicate whether support for Adabas Vista should be enabled. (Adabas Vista)	YES NO	NO

Parameter	Use to	Values	Default
WOrkcache	Indicate whether to enable (YES) or disable (NO) controller caching for the Adabas Work data set. This parameter is valid only in z/OS environments.	YES NO	YES

AOSLOG Parameter: Log Active Nucleus Update Commands

Parameter	Specify . . .	Possible Values	Default
<u>A</u> OSLOG	whether to log update commands from ADADBS OPERCOM or AOS that modify the active nucleus	YES NO	NO

Value	Meaning
YES	Log to DDPRINT Adabas calls that modify the active nucleus (not read/display calls) issued by ADADBS OPERCOM or the equivalent Adabas Online System (AOS) functions.
NO	The default setting. Commands are not logged.

Example

In the following example, Adabas calls that modify the nucleus are logged.

```
ADARUN PROG=ADANUC ,AOSLOG=YES
```

AREXCLUDE Parameter: Exclude Files from Autorestart Processing

Parameter	Specify . . .	Possible Values	Default
<u>A</u> REXCLUDE	the file(s) to exclude from autorestart processing.	1-5000	none

The AREXCLUDE parameter is intended for use only when autorestart processing fails (abends) due to response codes on one or more Adabas files. In this case, excluding the problem file(s) from autorestart might allow for successful autorestart processing on the other files.

Excluded files are marked as "utility use only". Software AG recommends that you restore excluded files from the latest save tape and regenerate them.

Multiple file numbers can be specified on the same parameter if separated by commas.

Example

The following example excludes files 199 and 231 from autorestart processing.

```
ADARUN PROG=ADANUC,AREXCLUDE=199,231
```

ARMNAME Parameter: Activate Automatic Restart Management (ARM) Program

Parameter	Specify . . .	Possible Values	Default
ARMNAME	and activate a particular automatic restart management (ARM) program.	<i>arm-pgm-name</i>	none

Automatic restart management (ARM) is used to automatically restart a nucleus when it abends. It is suppressed when the abend is intentional; for example, when it results from a parameter error.

ARM can be used in both cluster and noncluster nucleus environments.

ARM programs ('elements') are defined by the system programmer in the ARM 'policy'. They specify when, where, and how often a restart is to be attempted.

The ARMNAME parameter has no effect if an ARM policy has not been defined.

The name can be 1-16 characters long. In a cluster environment, the name must be unique across the cluster. Uppercase alphabetic characters, the numbers 0-9, and the special characters \$, @, and #, and the underscore (_) may be used. The first character may not be numeric. Names that start with A through I and with SYS are reserved for use by IBM.

Example

The following example activates the automatic restart management program RESTART002.

```
ADARUN PROG=ADANUC,ARMNAME=RESTART002
```

ASSOCACHE Parameter: Associator Controller Caching

This parameter applies to the z/OS operating system only.

Parameter	Specify . . .	Possible Values	Default
<u>ASSOCACHE</u>	whether to enable (YES) or disable (NO) controller caching for the Associator (ASSO) data set.	YES NO	YES

Value Meaning

YES	The default setting. ASSO blocks are preread and maintained as long as possible in the caching facilities of the disk controller.
NO	ASSO blocks are not preread.

Notes:

1. If ASSO blocks are highly scattered or the update frequency is low, controller caching can cause longer I/O response times.
2. ADARUN caching control parameters also exist for Data Storage (DATACACHE), the work area (WORKCACHE), the sort area (SORTCACHE), and the temp area (TEMPCACHE).

Example

The following example runs the nucleus without controller caching for the ASSO data set.

```
ADARUN PROG=ADANUC,ASSOCACHE=NO
```

ASYTVS Parameter: Asynchronous Buffer Flush

Parameter	Specify . . .	Possible Values	Default
<u>ASYTVS</u>	whether to flush buffers asynchronously based on volume serial number.	YES NO	YES

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.

Example

The following example disables the asynchronous buffer flush based on volume.

```
ADARUN PROG=ADANUC,ASYTVS=NO
```

CACHE Parameter: Adabas Caching Facility Control

CACHE (or CACH) specifies whether or not ADACSH is to be loaded during Adabas session initialization.

CACHE = { YES | NO }

Value	Description
YES	ADACSH is loaded and available for operation.
NO	ADACSH is not loaded. Caching is not available until the Adabas session is terminated and subsequently restarted with CACHE=YES.

Example

```
ADARUN CACHE=YES
```

ADACSH is loaded during the Adabas session initialization. ADACSH operator commands are accepted and caching is in effect.

CACTIVATE Parameter: Adabas Caching Facility RABN Activation

CACTIVATE (or CACT) specifies whether or not *all* RABN ranges are to be activated at session open.

CACTIVATE = { YES | NO }

Value	Description
YES	All RABN ranges are activated at session open. This provides for immediate use of the cache without consideration of the CDEMAND threshold level.
NO	None of the RABN ranges is activated unless: <ul style="list-style-type: none"> the CDEMAND threshold level is reached; or the CENABLE operator command is accepted by ADACSH.

Example

```
ADARUN CACTIVATE=YES
```

All RABN ranges are made available for immediate use; that is, as soon as the Adabas nucleus needs to move a block from the Adabas buffer pool into a RABN range, the RABN range is allocated.

CASSODSP / CASSOEXT / CASSOHSP / CASSOG64 / CASSOL64 / CASSOV64 Parameters: Adabas Caching Facility Associator RABN Cache Space Control

Note:

CASSOxxx parameters are mutually exclusive with the CFILE and CXFILE parameters.

The CASSOxxx parameters identify one or more Associator RABNs to be cached where xxx indicates the type of cache space.

- CASSODSP (or CASSOD) caches the Associator RABNs in data space.
- CASSOEXT (or CASSOE) caches the Associator RABNs in extended memory.
- CASSOHSP (or CASSOH) caches the Associator RABNs in hiperspace.
- CASSOG64 (or CASSOL) caches the Associator RABNs in virtual 64-bit storage backed by 2G large pages.

Note:

If "G64" is specified, we recommend you specify a cache space size (CASSOMAXS) that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.

- CASSOL64 (or CASSOL) caches the Associator RABNs in virtual 64-bit storage backed by 1M large pages.
- CASSOV64 (or CASSOV) caches the Associator RABNs in virtual 64-bit storage.

Note:

Virtual 64 storage is only available in z/OS environments.

{ CASSODSP | CASSOEXT | CASSOHSP | CASSOG64 | CASSOL64 | CASSOV64 } = { rabn | rabn - rabn }, ...

Value	Description
rabn	The Associator RABN to be cached.
rabn - rabn	The range of Associator RABNs to be cached. The range is specified as two numbers separated by a hyphen with no intervening spaces or special characters. More than one RABN range (separated by a comma) may be specified on the same line. There are no continuations. If more than one line is required, additional CASSOxxx parameters must be specified.

If CASSO_{xxx} is specified, *only* the specified RABNs are cached. Unspecified RABNs are *not* cached and are subject to normal Adabas I/O operations.

Associator RABNs 1 through 30 are not cached; they are ignored if specified by the user.

Notes:

1. If an invalid RABN range is specified in the CASSO_{xxx} parameter, subsequent RABN ranges for the CASSO_{xxx} parameter are ignored. The ignored parameters can be added after initialization by using an operator command.
2. The CSTORAGE parameter overrides any CASSO_{xxx} parameter and ADACSH will cache all Associator RABNs as defined for the database.

RABNs for the ranges specified are cached in the cache space indicated by the parameter suffix. RABN ranges defined with this parameter must therefore compete for the space available in a given cache space. When the maximum is allocated, RABNs are purged from the cache space based on a first-in-first-out (FIFO) mechanism and are reused to accommodate the latest RABNs being purged from the Adabas buffer pool.

Example

```
ADARUN CASSODSP=1330-2349,126-187,124-125,205-216
ADARUN CASSOHSP=218-252,203-204,260-275,277-335
```

ADACSH caches the specified RABN ranges. No other Associator RABNs are cached.

The first four RABN ranges defined are cached in the Associator cache space allocated in data space (suffix DSP). The second four RABN ranges defined are cached in the Associator cache space allocated in hiperspace (suffix HSP).

CASSOMAXS Parameter: Associator Cache Space Unit Size

The CASSOMAXS (or CASSOM) parameter specifies the unit size in which areas are allocated for Associator cache space.

```
CASSOMAXS = n
```

Value	Description
<i>n</i>	<p>The unit size in which areas are allocated for Associator cache space. Size may be specified in bytes, in kilobytes (K), in megabytes (M) or in gigabytes (G).</p> <p>For extended memory, data spaces, and hiperspaces:</p> <ul style="list-style-type: none"> ● the minimum unit size is 81920 bytes (or 80K); ● the maximum unit size is 2147475456 bytes (or 2097144K; or 2047M); and ● the default unit size is 1048576 bytes (or 1024K; or 1M). ● In these cases, CASSOMAXS is always rounded up to the nearest 4-kilobyte boundary. <p>Note: When using data space storage on BS2000 systems, allocation units are rounded to one-megabyte units. The Adabas Caching Facility may extend data space storage and deliver the message <code>ADAN8M dddd area-type (stg-type) size extended by nnnnnn bytes</code>. However, if the megabyte limit is reached, it will deliver the message: <code>ADAN8S area-type (stg-type) Allocate failed. Retcode=64</code>. We recommend that you set the ADARUN parameters CASSOMAXS and CDATAMAXS high enough (in the megabyte range) to prevent repeated expansion. It may be necessary to set the BS2000 user attribute ADDRESS-SPACE-LIMIT high enough to accommodate this.</p> <p>For virtual 64 memory and virtual 64 backed by large pages, virtual storage above the 2G bar is allocated in terms of memory objects, each of which comprises one or more virtual segments or pages of one megabyte each. In essence, CASSOMAXS determines the number of 1-megabyte segments to be allocated for the memory object:</p> <ul style="list-style-type: none"> ● the minimum unit size and the default unit size is 1M (1048576 bytes or 1024K). ● the maximum unit size depends on the limit set for your installation. ● For memory objects, CASSOMAXS is rounded up to the nearest megabyte boundary. For cache spaces for which storage is backed by 2G large pages ("G64"), CASSOMAXS is rounded up to the nearest multiple of two gigabytes. <p>Note: If "G64" is specified for a cache space, we recommend that you specify a CASSOMAXS value that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.</p>

The total space available is equal to:

Maximum-ASSO-cache-space-size = CASSOMAXS x CMAXCSPS

Storage areas for a cache space are allocated when required. The Adabas nucleus informs ADACSH when a RABN is about to be purged from the buffer pool. If ADACSH determines that this RABN is being cached, it attempts to allocate space in the appropriate cache space.

When no space is available in the currently allocated areas for a cache space, or no areas are allocated, a new area of storage is allocated with the size specified by the CASSOMAXS parameter. For cache storage in

- extended memory, additional storage is requested in this area;
- data space, an additional data space is allocated;
- hiperspace, an additional hiperspace is built; and
- virtual 64 memory backed by large pages, an additional memory object is allocated; and
- virtual 64 memory, an additional memory object is allocated.

This occurs until the maximum number of areas as determined by the CMAXCSPS parameter is reached. At this point, space is made for the latest RABNs to be purged from the Adabas buffer pool based on the aging algorithms.

The CASSOMAXS limit used for the allocation of new areas in a cache space is based on the last specification for this value. If no operator or Adabas Online System command is issued, this is the value specified in the startup parameters. If this value is specified by an operator command or through Adabas Online System, it takes effect the next time an additional area is allocated in an ASSO cache space.

Example

ADARUN CASSOMAXS=15000000

Cache space areas for ASSO RABNs are allocated in units of 15003648 bytes (or 14652K; or approximately 14M).

CBUFNO Parameter: Number of Read-Ahead Buffers for Concurrent I/O Processing

CBUFNO specifies the number of read-ahead buffers available for concurrent I/O processing.

CBUFNO is one of two parameters (CBUFNO and CEXCLUDE) used for read-ahead caching. Read-ahead caching performs a read-ahead of consecutive Associator and/or Data Storage blocks when processing serial Adabas commands (Read Logical, Read Physical, Histogram), and when searching using nondescriptors.

The read-ahead is performed by issuing a single EXCP to read all the consecutive Associator or Data Storage blocks that reside on a single track of the disk device. The blocks are kept in cache and are immediately available when the nucleus requests the next block in a sequence.

This feature may enhance performance by reducing the number of physical read I/Os. For example, for a 3380 Associator, a reduction factor of 18:1 may be achieved.

CBUFNO = *n*

Value	Description
<i>n</i>	<p>The number of read-ahead buffers available for concurrent I/O processing.</p> <p>The recommended setting is a value that accommodates the expected number of concurrent L2, L3, and L9 commands.</p> <p>The minimum number of read-ahead buffers is 0 (the default). If CBUFNO=0 is specified, <i>no</i> read-ahead caching is performed.</p> <p>The maximum number is the same as the number of threads specified for the nucleus parameter NT, because there cannot be more than NT number of commands active at any time.</p>

During normal cache processing, if a read I/O request qualifies for read-ahead processing but no read-ahead buffer is available, the read I/O is satisfied with a single block EXCP. In other words, no optimization is done to issue an EXCP to read the entire track.

Note:

Read-end-of-track processing is not available in Cluster and SMP environments.

Example

```
ADARUN CBUFNO=5,CEXCLUDE=FIND/HIST
```

Five buffers are available for read-ahead caching. No read-ahead processing is performed for S1 (nondesoriptor search) and L9 (histogram) commands.

CCTIMEOUT Parameter: Cache Space Area or RABN Inactivity Time

The CCTIMEOUT (or CC) parameter specifies the amount of time (in seconds) a cache space area can remain inactive before it is released back to the system.

If demand caching (CDEMAND) is in effect, CCTIMEOUT specifies the number of seconds a RABN range, file, or file range can remain inactive before it is disabled.

CCTIMEOUT = *n*

Value	Description
<i>n</i>	<p>The amount of time (in seconds) a cache space area can remain inactive before it is released back to the system.</p> <p>If demand caching (CDEMAND) is in effect, CCTIMEOUT specifies the number of seconds a RABN range, file, or file range can remain inactive before it is disabled.</p> <p>The minimum number of seconds is 60; the maximum is 2,147,483,647.</p> <p>The default number of seconds is 7200 (2 hours).</p>

If a cache space area has not been accessed for the specified period of time, ADACSH releases the storage associated with the area back to the operating system. A message is displayed on the operator console indicating that the cache space area was timed out. The cache space area can be re-allocated immediately by the caching system when required.

If a RABN range, file, or file range is not referenced for the specified period of time, and demand caching is in effect, the RABN range, file, or file range is disabled. A message is displayed on the operator console informing the user that the RABN range, file, or file range was timed out. No further caching occurs for that RABN range, file, or file range until the Adabas buffer efficiency goes below the level specified by CDEMAND. Once this occurs, the RABN range, file, or file range is enabled again and begins caching its associated RABNs.

Activity is checked at an interval of one half of the CCTIMEOUT value. For example, if CCTIMEOUT=3600 is specified, the allocated cache space areas, RABN ranges, files, or file ranges are checked every 1800 seconds (30 minutes).

A change to CCTIMEOUT using an operator command causes the check to be performed immediately, and subsequently at intervals of one half of the new CCTIMEOUT value.

Example

```
ADARUN CCTIMEOUT=3600
```

A cache space area is released (or a RABN range/file/file range is disabled if demand caching is in effect) if there is no activity associated with it for a continuous period of 1 hour (3600 seconds).

CDATADSP / CDATAEXT / CDATAHSP / CDATAG64 / CDATAL64 / CDATAV64 Parameter: Adabas Caching Facility Data Storage RABN Cache Space Control

Note:

CDATAxxx parameters are mutually exclusive with the CFILE and CXFILE parameters.

The CDATAxxx parameters identify one or more Data Storage RABNs to be cached where xxx indicates the type of cache space.

- CDATADSP (or CDATAD) caches the Data Storage RABNs in data space.
- CDATAEXT (or CDATAE) caches the Data Storage RABNs in extended memory.
- CDATAHSP (or CDATAH) caches the Data Storage RABNs in hiperspace.
- CDATAG64 (or CDATAL) caches the Data Storage RABNs in virtual 64-bit storage backed by 2G large pages.

Note:

If "G64" is specified, we recommend you specify a cache space size (CDATAMAXS) that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.

- CDATAL64 (or CDATAL) caches the Data Storage RABNs in virtual 64-bit storage backed by 1M large pages.
- CDATAV64 (or CDATAV) caches the Data Storage RABNs in virtual 64-bit storage.

Note:

Only extended memory is available for BS2000 RISC machines. Hiperspace is an option only in z/OS environments. Virtual 64 storage is available only in z/OS environments.

```
{ CDATADSP | CDATAEXT | CDATAHSP | CDATAG64 | CDATAL64 | CDATAV64 } = { rbn | rbn - rbn }, ...
```

Value	Description
rbn	The Data Storage RABN to be cached.
rbn - rbn	The range of Data Storage RABNs to be cached. The range is specified as two numbers separated by a hyphen with no intervening spaces or special characters. More than one RABN range (separated by a comma) may be specified on the same line. There are no continuations. If more than one line is required, additional CDATAxxx parameters must be specified.

If CDATAxxx is specified, *only* the specified RABNs are cached. Unspecified RABNs are *not* cached and are subject to normal Adabas I/O operations.

Notes:

1. If an invalid RABN range is specified in the CDATAxxx parameters, subsequent RABN ranges for the CDATAxxx parameters will be ignored. The ignored parameters can be added after initialization by using an operator command.
2. The CSTORAGE parameter overrides any CDATAxxx parameter and ADACSH will cache all Data

Storage RABNs as defined for the database.

RABNs for the ranges specified are cached in the cache space indicated by the parameter suffix. RABN ranges defined with this parameter must therefore compete for the space available in a given cache space. When the maximum is allocated, RABNs are purged from the cache space based on a first-in-first-out (FIFO) mechanism and are reused to accommodate the latest RABNs being purged from the Adabas buffer pool.

Example

```
ADARUN CDATADSP=1-4900,6523-7853
ADARUN CDATAHSP=5032-6204,8123-10382
```

ADACSH caches the specified RABN ranges. No other DATA RABNs are cached.

The first two RABN ranges defined are cached in the Data Storage cache space allocated in data space (suffix DSP). The second two RABN ranges defined are cached in the Data Storage cache space allocated in hiperspace (suffix HSP).

CDATAMAXS Parameter: Data Storage Cache Space Unit Size

The CDATAMAXS (or CDATAM) parameter specifies the unit size in which areas are allocated for Data Storage cache space.

```
CDATAMAXS = n
```

Value	Description
n	<p>The unit size in which areas are allocated for Data Storage cache space. Size may be specified in bytes, in kilobytes (K), in megabytes (M) or in gigabytes (G).</p> <p>For extended memory, data spaces, and hiperspaces:</p> <ul style="list-style-type: none"> ● the minimum unit size is 81920 bytes (or 80K); ● the maximum unit size is 2147475456 bytes (or 2097144K; or 2047M); and ● the default unit size is 1048576 bytes (or 1024K; or 1M). ● In these cases, CDATAMAXS is always rounded up to the nearest 4-kilobyte boundary. <p>Note: When using data space storage on BS2000 systems, allocation units are rounded to one-megabyte units. The Adabas Caching Facility may extend data space storage and deliver the message <code>ADAN8M dddd area-type (stg-type) size extended by nnnnn bytes</code>. However, if the megabyte limit is reached, it will deliver the message: <code>ADAN8S area-type (stg-type) Allocate failed. Retcode=64</code>. We recommend that you set the ADARUN parameters CASSOMAXS and CDATAMAXS high enough (in the megabyte range) to prevent repeated expansion. It may be necessary to set the BS2000 user attribute ADDRESS-SPACE-LIMIT high enough to accommodate this.</p> <p>For virtual 64 memory and virtual 64 backed by large pages, virtual storage above the 2G bar is allocated in terms of memory objects, each of which comprises one or more virtual segments or pages of one megabyte each. In essence, CDATAMAXS determines the number of 1-megabyte segments to be allocated for the memory object:</p> <ul style="list-style-type: none"> ● the minimum unit size and the default unit size is 1M (1048576 bytes or 1024K). ● the maximum unit size depends on the limit set for your installation. ● For memory objects, CDATAMAXS is rounded up to the nearest megabyte boundary. For cache spaces for which storage is backed by 2G large pages ("G64"), CDATAMAXS is rounded up to the nearest multiple of two gigabytes. <p>Note: If "G64" is specified for a cache space, we recommend that you specify a CDATAMAXS value that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.</p>

The total space available is equal to:

$$\text{Maximum-DATA-cache-space-size} = \text{CDATAMAXS} \times \text{CMAXCSPS}$$

Storage areas for a cache space are allocated when required. The Adabas nucleus informs ADACSH when a RABN is about to be purged from its buffer pool. If ADACSH determines that this RABN is being cached, it attempts to allocate space in the appropriate cache space.

When no space is available in the currently allocated areas for a cache space, or no areas are allocated, a new area of storage is allocated with the size specified by the CASSOMAXS parameter. For cache storage in

- extended memory, additional storage is requested in this area;
- data space, an additional data space is allocated;
- hiperspace, an additional hiperspace is built; and
- virtual 64 memory backed by large pages, an additional memory object is allocated; and
- virtual 64 memory, an additional memory object is allocated.

This occurs until the maximum number of areas as determined by the CMAXCSPS parameter is reached. At this point, space is made for the latest RABNs to be purged from the Adabas buffer pool based on the aging algorithms.

The CDATAMAXS limit used for the allocation of new areas in a cache space is based on the last specification for this value. If no operator or Adabas Online System command is issued, this is the value specified in the startup parameters. If this value is specified by an operator command or through Adabas Online System, it takes effect the next time an additional area is allocated in a Data Storage cache space.

Example

```
ADARUN CDATAMAXS=15000000
```

Cache space areas for Data Storage RABNs are allocated in units of 15003648 bytes (or 14652K; or approximately 14M).

CDEMAND Parameter: Lowest Adabas Buffer Efficiency for Caching

The CDEMAND (or CDE) parameter specifies the lowest Adabas buffer efficiency level accepted before caching is activated.

Calculated by the nucleus, Adabas buffer efficiency is the number of read requests divided by the actual number of reads to storage. If the Adabas buffer efficiency level drops below the specified CDEMAND level, ADACSH activates all disabled RABN ranges. In this way, RABN ranges are dynamically activated as demand on the Adabas buffer pool increases.

The CDEMAND parameter should be used in conjunction with the CCTIMEOUT parameter to establish resource usage boundaries that permit the timely acquisition and release of system resources. In other words, CDEMAND may be used to activate RABN ranges when the demand load on Adabas is high, and CCTIMEOUT may be used to deactivate RABN ranges when Adabas is idle.

CDEMAND= { 0
n }

Value	Description
0	If CDEMAND=0 (the default), <i>no</i> demand-level caching occurs.
n	Demand-level caching is to be activated whenever the Adabas buffer efficiency drops below the value specified. An integer value between 1 and 2,147,483,647 is permitted.

If CDEMAND is not zero, the Adabas buffer efficiency is first checked approximately 30 minutes after the Adabas nucleus is initialized. Thereafter, the Adabas buffer efficiency is checked at intervals of approximately 15 minutes.

CDEMAND does not affect the activation of caching for Work parts 2 and 3 because these areas do not influence the Adabas buffer efficiency. Only Associator and Data Storage blocks reside in the Adabas buffer pool and, consequently, only the cache spaces for the Associator and Data Storage serve to increase the Adabas buffer efficiency level.

Example

ADARUN CDEMAND=7

ADACSH activates the RABN ranges if the Adabas buffer efficiency level drops below 7.

CDISPSTAT Parameter: Adabas Caching Facility RABN Range Statistics Display Control

The CDISPSTAT (or CDI) parameter specifies whether the RABN range statistics are to be displayed on both the operator console and DDPRINT, or only on DDPRINT.

The individual RABN range statistics are displayed whenever a RABN range is released due to a parallel utility operation or an inactivity timeout, or when a RABN range is dynamically added or deleted.

CDISPSTAT = { YES | NO }

Value	Description
YES	RABN range statistics are to be displayed on both the operator console and DDPRINT.
NO	RABN range statistics are to be displayed only on DDPRINT.

Example

```
ADARUN CDISPSTAT=NO
```

When the status of a RABN range is changed due to a parallel utility operation, a timeout, or an operator command to add or delete, the individual RABN range statistics are to be displayed on DDPRINT only.

CDXnn Parameter: Collation Descriptor User Exit

Parameter	Specify . . .	Possible Values	Default
CDXnn	the collation descriptor exit and the user routine to be used with it.	see text	none

Used in conjunction with the universal encoding subsystem (UES), the format of the collation descriptor user exit parameter is

```
CDXnn=exit-name
```

where

nn is the number of the collation descriptor exit, a two-digit decimal integer in the range 01-08 inclusive.

exit-name is the name of the user routine that gets control at the collation descriptor exit; the name can be up to 8 characters long.

Only one program may be specified for each collation descriptor exit. Up to 8 collation descriptor exits may be specified (in any order). See the *Adabas User, Hyperdescriptor, Collation Descriptor, and SMF Exits Manual* for more information.

Example

The following example runs the Adabas nucleus using collation descriptor exits 1 and 2. It also gives control to the user routine UCOLDE1 at collation descriptor exit 1 and gives control to the user routine UCOLDE2 at collation descriptor exit 2.

```
ADARUN PROG=ADANUC,CDX01=UCOLDE1,CDX02=UCOLDE2
```


CEXCLUDE Parameter: Adabas Caching Facility Excluded Command Types

CEXCLUDE specifies the command types to be excluded from read-ahead caching.

CEXCLUDE is one of two parameters (CBUFNO and CEXCLUDE) used for read-ahead caching. Read-ahead caching performs a read-ahead of consecutive Associator and/or Data Storage blocks when processing serial Adabas commands (Read Logical, Read Physical, Histogram), and when searching using nondescriptors.

The read-ahead is performed by issuing a single EXCP to read all the consecutive Associator or Data Storage blocks that reside on a single track of the disk device. The blocks are kept in cache and are immediately available when the nucleus requests the next block in a sequence.

This feature may enhance performance by reducing the number of physical read I/Os. For example, for a 3380 Associator, a reduction factor of 18:1 may be achieved.

```
CEXCLUDE = { PHYS | LOGI | HIST | FIND } / ...
```

Value	Description
PHYS	Exclude L2 and L5 (physical sequence) commands.
LOGI	Exclude L3 and L6 (logical sequence) commands.
HIST	Exclude L9 (histogram) command.
FIND	Exclude S1 (nondesoriptor search) command.

The exclusion list must be specified with each four-byte name separated by a slash (/). No intervening spaces or special characters, such as commas, are permitted.

The exclusion is made globally; for example, if CEXCLUDE=PHYS is specified, all L2 and L5 commands are excluded from read-ahead processing, regardless of the file to which the command was directed.

The exclusion list may be changed during the nucleus session using the CEXCLUDE and CINCLUDE operator commands.

Example

```
ADARUN CBUFNO=5,CEXCLUDE=FIND/HIST
```

Five buffers are available for read-ahead caching. No read-ahead processing is performed for S1 (nondesoriptor search) and L9 (histogram) commands.

CFILE Parameter: Adabas Caching Facility File Range

Note:

CFILE is mutually exclusive with the CASSO_{xxx}, CDATA_{xxx}, and CSTORAGE parameters. If any of these parameters is used with CFILE, they will be ignored.

The CFILE parameter defines a file or file range to be cached. It is also used to provide the settings used to control how caching is to be implemented for the file or file range.

```
CFILE = ( { fnr | fnr - fnr }, [ cls ], [ scope ], [ cache ] )
```

Value	Description
<i>fnr</i>	The file number of the file to be cached. The minimum value is 0; the maximum is the actual number of files in the database. File number 0 may be specified to cache RABNs that are not associated with any given file. Currently, this is only the DSST on the Associator data set.
<i>fnr - fnr</i>	Range of files to be cached specified by two file numbers (the lower number first) separated by a hyphen with no intervening spaces or special characters.
<i>cls</i>	Class of service. The percentage of available cache space that can be used to cache RABNs associated with the specified file(s). This term effectively assigns a priority to the file(s). RABNs belonging to the file(s) with the lowest priority are purged first. Valid values are: <ul style="list-style-type: none"> ● 1: 100% (highest priority) ● 2: 75% ● 3: 50%; the default ● 4: 25% ● 5: 10% (lowest priority)
<i>scope</i>	Caching scope. Valid values are: <ul style="list-style-type: none"> ● ASSO (or A) : Caching is to be used only for Associator for the file ● DATA (or D) : Caching is to be used only for Data Storage for the file ● BOTH (or B) : Caching is to be used for both Associator and Data Storage for the file <p>The default is BOTH.</p>

Value	Description
<i>cache</i>	<p>Specifies the type of cache storage in which RABNs for the file will be cached. Valid values are:</p> <ul style="list-style-type: none"> ● DATASPACE (or D) ● EXTENDED (or E) ● G64 (or G) ● HIPERSPACE (or H) ● L64 (or L) ● VIRTUAL64 (or V) <p>The default is DATASPACE.</p> <p>Notes:</p> <ol style="list-style-type: none"> 1. Virtual 64-bit values "G64", "L64", and "VIRTUAL64" are cache options only in z/OS environments. 2. Specify "L64" to request that cache storage be virtual storage above the 2G bar backed by page-fixed 1M large pages; specify "G64" to request that cache storage be virtual storage above the 2G bar backed by page-fixed 2G large pages. 3. If "G64" is specified, we recommend you specify a cache space size (CASSOMAXS/CDATAMAXS) that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.

Within the parentheses, the value terms are positional. This means that you can omit a term if you want to accept the default value for it, but you must enter a comma to identify the position if you are specifying a subsequent term within the parentheses.

Multiple CFILE parameters may be specified.

CFILE and CXFILE parameters together are processed in order from left to right (or top to bottom if different ADARUN statements are used). Parameters or statements that specify file ranges are processed before those that specify specific files.

Examples

```
ADARUN CACHE=YES,CFILE=(7,2)
```

File-level caching is active for file number 7. The file is permitted to use up to 75% of the maximum total available cache space (class of service 2). Both Associator and Data Storage RABNS are cached (the default) and the type of cache space used is data space (the default).

```
ADARUN CACHE=YES,CFILE=(3-6,,DATA,HIPERSPACE)
```

File-level caching is active for file numbers 3 through 6. The files are permitted to use up to 50% of the maximum total available cache space (class of service 3, the default). Only Data Storage RABNS are cached, and the type of cache space used is hiperspace.

```
ADARUN CACHE=YES,CFILE=(14-17,,,HIPERSPACE)
```

File-level caching is active for file numbers 14 through 17. The files are permitted to use up to 50% of the maximum total available cache space (class of service 3, the default). Both Associator and Data Storage RABNs are cached (the default), and the type of cache space used is hiperspace.

CLOGBMAX Parameter: Logged Buffer Size Limit

Parameter	Specify . . .	Possible Values	Default
CLOGBMAX	the maximum allowable number of bytes of a logged buffer.	368 - 2147483647	4096

If a buffer is longer than this value, the logged buffer is truncated from the point at which its size exceeds the setting of the CLOGBMAX parameter. The CLOGBMAX setting affects the ADARUN LOGGING parameter specifications for both CLOGLAYOUT=5 and CLOGLAYOUT=8. The minimum value (368) is the size of the length of the CLOGLAYOUT=8 basic record plus the length of the extended Adabas control block (ACBX).

Example

The following example runs the Adabas nucleus using a logged buffer size limit of 512. Individual logged buffers for an Adabas command cannot exceed 512 bytes.

```
ADARUN PROG=ADANUC,CLOGB=512
```

CLOGDEV Parameter: Multiple Command Log Device Type

CLOGSIZE : Multiple Command Log Size

Parameter	Specify . . .	Possible Values	Default
<u>CLOGDEV</u>	the device type to be used for multiple command logs.	see text	value of DEVICE parameter
<u>CLOGSIZE</u>	the number of blocks available for each command log.	Min 16, Max 16777215	none

The CLOGDEV and CLOGSIZE parameters specify the type of external storage device used for multiple command logs and the number of blocks available for each command log data set. These parameters must be used *together*; if you specify one parameter, you must also specify the other parameter.

You must also specify the NCLOG parameter to use multiple command logs. If you specify a user exit, it must be UEX12.

The section *Adabas Device Types and Blocks Sizes* in the Adabas Installation documentation specifies the device types that Adabas supports.

All command logging data sets must be formatted using the ADAFRM utility before they are used.

The maximum size of the data set depends on the limits of the operating system and on the requirements of the user.

Example

The following example runs the Adabas nucleus with command logging. It uses six command log data sets and a model 3380 device for command logging. The size of each command log data set is 1350 blocks, based on an allocation of 10 cylinders. (Each cylinder of a 3380 device has 15 tracks, and each track can contain 9 blocks of command log data. Thus, 10 cylinders = 10 x 15 x 9 = 1350 blocks).

```
ADARUN PROG=ADANUC,LOGGING=YES,NCLOG=6,CLOGDEV=3380,CLOGSIZE=1350
```

CLOGLAYOUT Parameter: Command Logging Format

Parameter	Specify . . .	Possible Values	Default
<u>CLOGLAYOUT</u>	the command log format.	5 8	5

This parameter specifies the format of the Adabas command log.

Value	Meaning
5	The default setting. User programs designed for Adabas 7 or earlier should specify CLOGLAYOUT=5.

Note:

CLOGLAYOUT=5 CLOG records may be caused by internal Adabas Review V4 commands. These records should be ignored by users.

For CLOGLAYOUT=5 records, the sum of all buffers must fit within the physical blocksize of the CLOG. As the buffers are moved to the output buffer to be written, the length of the buffer is added to the current used portion of the buffer; if it exceeds the physical blocksize of the CLOG, then the buffer is not included and the next buffer is checked. This continues until all buffers have been processed. A CLOG device type should be chosen that has a large enough physical blocksize to hold all the buffers when CLOGLAYOUT=5 is specified, otherwise, some buffers may not be included in the output.

You can restrict the number of bytes to be logged for each buffer by specifying a value in the ADARUN CLOGBMAX parameter. In addition, if the ADARUN CLOGMAX parameter is specified, a value equal to or less than the CLOG physical blocksize should be specified to be effective.

8	User programs designed for Adabas 8 or later should specify CLOGLAYOUT=8.
---	---

Use the command log format from Adabas 8. Both ACB and ACBX control blocks are converted to a common format. Each of the buffers are written out in much the same way as for CLOGLAYOUT=5, except that each buffer is prefixed by its corresponding Adabas buffer description (ABD). Each segmented buffer (format, record, or multifetch) is written separately and uniquely identified.

The format of CLOGLAYOUT=8 records is described in the LORECX DSECT distributed with Adabas 8.

Notes:

1. Timestamps in an Adabas 8 command log created using CLOGLAYOUT=8 are stored in machine time (GMT), whereas CLOGLAYOUT=5 timestamps are stored, as always, in local time. The LORECX record layout that describes the CLOGLAYOUT=8 command log includes a differential time field that stores the difference between machine time and local time at the time the CLOG record is written. This field allows you to calculate the local time of a command log record.
2. Because of the different timestamp formats used in CLOGLAYOUT=5 and CLOGLAYOUT=8 command logs, we do not recommend that you mix or merge command logs created using different CLOGLAYOUT settings (and, in fact, Adabas does not allow this in cluster or parallel services environments). If you do, your results will be unpredictable.

Example

The following example uses the Adabas 8 format of the command log.

```
ADARUN PROG=ADANUC,CLOGLAYOUT=8
```

CLOGMAX Parameter: Total Logged Buffer Size Limit for a Command

Parameter	Specify . . .	Possible Values	Default
CLOGMAX	the maximum size of all of the logged buffers allowed for an Adabas command.	368 - 2147483647	16384

When the sum of sizes of the logged buffers for an Adabas command reaches the value of the CLOGMAX parameter, the buffer exceeding the limit is truncated and all following buffers are omitted. The size of CLOGMAX must at least be as large as the ADARUN CLOGBMAX parameter setting in order to accommodate the largest buffer of the Adabas command. The minimum value (368) is the size of the length of the CLOGLAYOUT=8 basic record plus the length of the extended Adabas control block (ACBX).

The CLOGMAX setting affects the ADARUN LOGGING parameter specifications for both the ADARUN CLOGLAYOUT=5 and CLOGLAYOUT=8 parameter settings.

When CLOGMAX is specified with CLOGLAYOUT=5, a value equal to or less than the CLOG physical blocksize should be specified to be effective.

Example

The following example runs the Adabas nucleus using a command logged buffer size limit of 32K. The sum of all logged buffers for an Adabas command to this nucleus cannot exceed 32K.

```
ADARUN PROG=ADANUC,CLOGMAX=32768
```

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 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.
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.

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
<u>CLUCACHEEXTRA</u>	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
<u>CLUCACHENAME</u>	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
```

CLUCACHESIZE Parameter: Size of the Global Cache Area

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

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

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

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

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

The global cache area must be large enough to retain:

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

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

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

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

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

Note:

If "G64" is specified for a cache space, we recommend that you specify a CLUCACHESIZE value that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.

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

Example

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

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

CLUCACHETYPE Parameter: Global Cache Area Storage Type Option

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

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

The size of the global cache is specified in the CLUCACHESIZE parameter.

CLUCACHETYPE Value	Description
DSP	This is the default value. It indicates that a shared dataspace will be used for the global cache. In this case, the maximum cache size is two gigabytes (2G).
G64	<p>If "G64" is specified (on z/OS 2.1 or later systems running on IBM zEC12 machines), the global cache space shared by the cluster nuclei will reside in 64-bit virtual storage that is backed by page-fixed 2G large pages (provided the large page pool has been configured to a sufficient size and is available in the system). If the system cannot satisfy the allocation request with storage backed by 2G large pages, storage backed by pageable 4K pages will be allocated instead.</p> <p>Note: Allocation requests for storage to be backed by 2G large pages will be rounded up to the nearest multiple of two gigabytes. Any difference between the specified size and the allocated size will remain unused. To avoid leaving a potentially large amount of allocated storage unused when specifying CLUCACHETYPE=G64, we recommend that you supply a multiple of two gigabytes for the CLUCACHESIZE parameter.</p>
L64	If "L64" is specified (on z/OS 1.9 or later systems running on IBM z10 machines), the global cache space shared by the cluster nuclei will reside in 64-bit virtual storage that is backed by page-fixed one-megabyte (1M) large pages (provided the large page pool has been configured to a sufficient size and is available in the system). If the system cannot satisfy the allocation request with storage backed by 1M large pages, storage backed by pageable four-kilobyte (4K) pages will be allocated instead.
V64	If "V64" is specified (on z/OS systems), a storage area in 64-bit virtual storage will be used for the global cache. In this case, the cache can be made much larger than 2G, but is still subject to the MEMLIMIT value in effect for the ADACOM job for the Adabas Parallel Services database.

Notes:

1. If one ADACOM job is used to hold the shared cache and lock spaces for multiple Adabas Parallel Services clusters, L64 and G4 cache spaces of different clusters are slightly less isolated from one another than DSP or V64 cache spaces, creating a marginally higher risk that a cache-related error in one cluster might induce an error in a different cluster. This very small risk can be avoided altogether by using different ADACOM jobs for different Adabas Parallel Services clusters that use L64 or G64 cache spaces, but this is not mandatory.
2. Virtual 64-bit storage backed by page-fixed 1M large pages can only be used on IBM z10 mainframes running z/OS Version 1.9 or above and for which IBM large page support has been enabled via IBM APAR OA20902. You must also allocate the size of the large page pool (use the LFAREA parameter in the IEASYS.xx member of SYS1.PARMLIB). The LFAREA parameter

allows you to specify the amount of real storage to be used for page-fixed 1M large pages; this parameter cannot be changed dynamically and, if it is not set, page frames will remain allocated at a size of 4K.

- Virtual 64-bit storage backed by page-fixed 2G large pages can only be used on IBM zEC12 mainframes running z/OS version 2.1 or above or an earlier z/OS version for which 2G large page support has been enabled via IBM APAR. You must also allocate the size of the large page pool (use the LFAREA parameter in the IEASYSxx member of SYS1.PARMLIB). The LFAREA parameter is extended to allow you to specify the amount of real storage to be used for page-fixed 1M and 2G large pages; this parameter cannot be changed dynamically and, if it is not set, page frames will remain allocated at a size of 4K.

Example

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

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

CLUCACHEUNCHANGED Parameter: Global Cache Unchanged Block Option

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

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

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

Recommendation

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

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

Example

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

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

CLUGROUPNAME Parameter: Cluster Group Name

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

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

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

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

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

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

Example

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

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

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
```

CLULOCKSIZE Parameter: Size of the Global Lock Area

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

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

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

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

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

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

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

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
```

CMADDR and CMLADDR Parameters: GETMAIN Memory Pool Start Address

These parameters apply only to BS2000 environments running Sockets versions less than 2.2.

Parameter	Specify . . .	Minimum	Maximum	Default
	the start address of the GETMAIN memory pool			
CMADDR	above the 16MB line	0, or X'1000000'	X'FE000000'	0
CMLADDR	below the 16MB line	0, or X'10000'	X'E00000'	0

Note:

These parameters should only be used if required by an installation site, in which case the site will provide the needed value.

These parameters are involved with setting all required memory into a common memory pool for use by associated tasks on BS2000 above (CMADDR) or below (CMLADDR) the 16 megabyte line. They set the start address of this memory pool.

If the default value ("0") is used for a CMADDR or CMLADDR parameter, the start address depends on the value of the corresponding CMSIZE or CMLSIZE parameter:

- If the CMADDR or CMLADDR parameter value is "0" and the corresponding CMSIZE or CMLSIZE parameter is *not* "0", the operating system will set the start address at the next available megabyte boundary.
- If the CMADDR or CMLADDR parameter value is "0" and the corresponding CMSIZE or CMLSIZE parameter is *also* "0", required memory will not be obtained in a common memory pool but rather in class 6 memory.

Specific Product Recommendations

For Entire Net-Work 6 (mainframe), Software AG recommends setting the value of the CMADDR parameter to X'01200000'.

Examples

The following example allows for three (3) megabytes of user storage in the address range X'2000000' to X'4FFFFFF':

```
ADARUN PROG=ADANUC,CMADDR=2000000,CMSIZE=2500000
```

The following example allows for three (3) megabytes of user storage below 16 megabytes in the address range X'200000 to X'4FFFFFF':

```
ADARUN PROG=ADANUC,CMLADDR=200000,CMLSIZE=2500000
```

CMAXCSPS Parameter: Adabas Caching Facility Storage Area Quantities

The CMAXCSPS (or CM) parameter specifies the number of storage areas that can be allocated for an ADACSH cache space in data spaces, hiperspaces, or memory objects (64-bit virtual storage). It has no effect on a cache space in extended memory, where only one Associator and one Data Storage area are allowed.

In conjunction with the CASSOMAXS and CDATAMAXS parameters, this parameter is used to control both the total amount of space that can be allocated for a cache space and the granularity with which it is allocated.

Value	Description
<i>n</i>	<p>The number of storage areas that can be allocated for an ADACSH cache space in data spaces, hiperspaces, or memory objects (64-bit virtual storage). It has no effect on a cache space in extended memory, where only one ASSO and one DATA storage area are allowed.</p> <p>The minimum value permitted is 1; the maximum is 16, the default is 8.</p> <p>This means that a potential maximum of 34 storage areas can be allocated for each type of cache space used: 16 for Associator RABNs, 16 for Data Storage RABNs, and 1 each for Work parts 2 and 3.</p> <p>Note: Increasing the value of CMAXCSPS can increase the overall CPU load used by ADACSH as more storage areas must be searched when allocating storage. This value must be carefully selected in association with the CASSOMAXS and CDATAMAXS values to obtain the correct mix of granularity and CPU overhead. Experience indicates that the default CMAXCSPS setting (8) usually provides an appropriate mix.</p>

Example

```
ADARUN CMAXCSPS=6
```

Up to six Associator cache space areas of size CASSOMAXS and up to six Data Storage cache space areas of size CDATAMAXS can be allocated for a particular type of cache space.

CMDQMODE Parameter: Command Queue Mode

This parameter applies to the BS2000 operating system only.

Parameter	Specify . . .	Possible Values	Default
CMDQMODE	whether to allocate the command queue memory pool below or above the 16-MB line.	BELOW ABOVE	ABOVE (BELOW for Adabas versions prior to Version 8)

CMDQMODE specifies whether to allocate the BS2000 memory pool for the Adabas command queue below or above the 16-MB line.

Value	Meaning
BELOW	The default setting. Places the BS2000 memory pool for the Adabas command queue below the 16-MB line in one or more 64-kilobyte segments.
ABOVE	Places the BS2000 memory pool for the Adabas command queue above the 16-MB line in one or more 1-MB segments.

Example

The following example, places the Adabas command queue memory pool above the 16-MB line in 1-MB segments.

```
ADARUN PROG=ADANUC ,CMDQMODE=ABOVE
```

CMFIX or CMLFIX Parameters: GETMAIN Memory Pool Fixed Location

These parameters apply only to BS2000 environments running Sockets versions less than 2.2.

Parameter	Specify . . .	Possible Values	Default
	whether the GETMAIN memory pool must have a fixed location or not		
<u>CMFIX</u>	above the 16MB line	YES NO	NO
<u>CMLFIX</u>	below the 16MB line	YES NO	NO

Note:

Use these parameters only if required by an installation site, in which case the site will provide the needed value.

The `..FIX` parameter is ignored if the corresponding `..SIZE` parameter value is 0.

These parameters are involved with setting all required memory into a common memory pool for use by associated tasks on BS2000 above (CMFIX) or below (CMLFIX) the 16 megabyte line. This parameter specifies whether the memory pool is at a fixed location or not:

Value	The application's GETMAIN memory pool . . .
NO	need not be located at the same start address in the subtask as in the main task.
YES	must be located at the same start address in the subtask as in the main task.

Note:

A subtask depends on the application. An example of this is Adabas Review in local mode.

Specific Product Recommendations

For Entire Net-Work 6 (mainframe), Software AG recommends setting the values of both the CMCIX and CMLFIX parameters to "YES".

Examples

The following example allows for three (3) megabytes of user storage above 16 megabytes in the address range X'2000000' to X'4FFFFFF'. This address range will also be used by the subtask:

```
ADARUN PROG=ADANUC,CMADDR=2000000,CMSIZE=2500000,CMFIX=YES
```

The following example allows for three (3) megabytes of user storage below 16 megabytes in the address range X'200000' to X'4FFFFF'. This address range will also be used by the subtask:

```
ADARUN PROG=ADANUC,CMLADDR=200000,CMSIZE=2500000,CMFIX=YES
```

CMLxxxx Parameters

The discussion of these parameters is paired with related Adabas parameters as follows:

CMLADDR
CMLFIX
CMLSCOPE
CMLSIZE

CMSCOPE or CMLSCOPE Parameters: GETMAIN Memory Pool Scope

These parameters apply only to BS2000 environments running Sockets versions less than 2.2.

Parameter	Specify . . .	Possible Values	Default
	access to the GETMAIN memory pool		
<u>CMSCOPE</u>	above the 16MB line	GROUP GLOBAL	GROUP
<u>CMLSCOPE</u>	below the 16MB line	GROUP GLOBAL	GROUP

Note:

These parameters should only be used if required by an installation site, in which case the site will provide the needed value.

The ..SCOPE parameter is ignored if the corresponding ..SIZE parameter value is 0.

These parameters are involved with setting all required memory into a common memory pool for use by associated tasks on BS2000 above (CMSCOPE) or below (CMLSCOPE) the 16 megabyte line. This parameter specifies accessibility to the memory pool:

Value Meaning

GROUP The GETMAIN common memory pool can only be accessed by other tasks using the same user ID.

GLOBAL The GETMAIN common memory pool can only be accessed by all user IDs in the system.

Specific Product Recommendations

For Entire Net-Work 6 (mainframe), Software AG recommends setting the values of the CMSCOPE and CMLSCOPE parameters to "GROUP".

Examples

The following example allows for three (3) megabytes of user storage at the next available megabyte boundary above the 16 megabyte line accessible to tasks with the same user ID.

```
ADARUN  PROG=ADANUC ,CMSIZE=3M ,CMSCOPE=GROUP
```

The following example allows for three (3) megabytes of user storage at the next available megabyte boundary below the 16 megabyte line accessible to tasks with the same user ID.

```
ADARUN  PROG=ADANUC ,CMLSIZE=3M ,CMLSCOPE=GROUP
```

CMSIZE or CMLSIZE Parameters: GETMAIN Memory Pool Use and Size

These parameters apply only to BS2000 environments running Sockets versions less than 2.2.

Parameter	Specify . . .	Minimum	Maximum	Default
	whether the GETMAIN is performed in a common memory pool and if so, its size in bytes (decimal)			
<u>CMSIZE</u>	above the 16MB line	0	address-limit	0
<u>CMLSIZE</u>	below the 16MB line	0	address-limit	0

Note:

These parameters should only be used if required by an installation site, in which case the site will provide the needed value.

These parameters are involved with setting all required memory into a common memory pool for use by associated tasks on BS2000 above (CMSIZE) or below (CMLSIZE) the 16 megabyte line. These parameters determine whether a common memory pool is used and if so, its size:

- If the ..SIZE parameter value is 0, the default, the required application memory is obtained in the class 6 memory of the system.
- If the ..SIZE parameter value is *not* 0, the required application memory is obtained in a common memory pool of the specified size rounded up to the next megabyte.

Specific Product Recommendations

For Entire Net-Work 6 (mainframe) running BS2000 Sockets earlier than version 2.2, Software AG recommends setting the CMSIZE parameter to "31M" and the CMLSIZE parameter to "4096K".

Examples

The following example allows for three (3) megabytes of user storage at the next available megabyte boundary above the 16 megabyte line accessible to tasks with the same user ID.

```
ADARUN  PROG=ADANUC ,CMSIZE=3M ,CMSCOPE=GROUP
```

The following example allows for three (3) megabytes of user storage at the next available megabyte boundary below the 16 megabyte line accessible to tasks with the same user ID.

```
ADARUN  PROG=ADANUC ,CMLSIZE=3M ,CMLSCOPE=GROUP
```

CRETRY Parameter: Adabas Caching Facility Retry Time

The CRETRY (or CR) parameter specifies the amount of time (in seconds) between attempts to acquire an area of cache space if the previous GETMAIN or space allocation failed.

```
CRETRY = n
```

Value	Description
<i>n</i>	<p>The number of seconds between attempts to acquire an area of cache space if the previous GETMAIN or space allocation failed.</p> <p>The minimum number of seconds is 60; the maximum is 2147483647.</p> <p>The default number of seconds is 900 (15 minutes).</p>

If an area of cache space cannot be allocated, ADACSH waits for the specified number of seconds before attempting to allocate the area of cache space again. During this period of time, caching continues in the other areas of the cache space already allocated, but the maximum allowed cache space is not available. If, after the time period has elapsed, there is a request to cache a RABN and no space is available in the current allocation, ADACSH attempts to allocate the area of cache space again. These attempts continue until the cache space is successfully allocated.

Note:

CRETRY does not apply to an area of cache space that was released due to a timeout (see CCTIMEOUT).

Example

```
ADARUN  CRETRY=500
```

If an attempt to allocate an additional area of a cache space fails, another attempt to cache a RABN in the cache space in question is made after 500 seconds. If the currently allocated space is still full, the attempt is again made to allocate an additional area of cache space.

CSTORAGE Parameter: Adabas Caching Facility Override Cache Type

The CSTORAGE parameter can be used to indicate that all RABNs for the database are to be cached using the system defaults and the type of cache specified by this parameter.

The CSTORAGE (or CS) parameter overrides any CASSOxxx or CDATAxxx parameter that may be specified. If the CSTORAGE parameter is specified, all RABNs for the database are cached using the system defaults and the type of cache space specified by CSTORAGE.

Note:

This parameter is mutually exclusive with the CFILE and CXFILE parameters.

CSTORAGE = { DATASPACE | EXTENDED | HIPERSPACE | G64 | L64 | VIRTUAL64 }

Value	Description
DATASPACE	Data space
EXTENDED	Extended memory
HIPERSPACE	Hiperspace
G64	Virtual 64-bit storage backed by 2G large pages Note: If "G64" is specified, we recommend you specify a cache space size (CASSOMAXS/CDATAMAXS) that is a multiple of two gigabytes. The specified size will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. Specifying a multiple of two gigabytes avoids leaving a potentially large amount of storage unused.
L64	Virtual 64-bit storage backed by 1M large pages
VIRTUAL64	Virtual 64-bit storage

Extended memory is the only option available on BS2000 RISC machines. Hiperspace is an option only in z/OS environments. Virtual 64 storage is available only in z/OS environments.

Example

```
ADARUN CASSOHSP=1-1795
ADARUN CDATAHSP=81-135,1140-1298,136-170,181-230
ADARUN CSTORAGE=EXTENDED
```


ADACSH caches all Associator and Data Storage RABNs in extended memory (above the 16 megabyte line). All CASSOHSP and CDATEHSP RABN ranges are ignored.

CT Parameter: Command Timeout Limit

Parameter	Specify . . .	Minimum	Maximum	Default
CT	the maximum time (seconds) for interregion communication of results from Adabas to the user.	1	16777215	60

For Adabas Review, this is the maximum number of seconds (more precisely, units of 1.048576 seconds) that can elapse from the time an Adabas Review hub command has been completed until the results are returned to the user through the interregion communication (operating-system-dependent).

This parameter is used to prevent a command queue element and attached buffer from being held for a long period of time for a user who has terminated abnormally.

Possible causes of a command timeout are

- user region is swapped out or cannot be dispatched;
- user is canceled;
- user has low priority in high activity system.

If the CT limit is exceeded,

- the command queue element and attached buffer are released;
- a message ADAM93 is printed; and
- if the user has not terminated, response code 254 (ADARSP254) is returned to the user program.

Example

The following example permits about 30 seconds to obtain a result through interregion communication from the Adabas Review hub

```
ADARUN PROG=ADAREV,CT=30
```

CWORKSTORAGE Parameter: Adabas Caching Facility Work Area Cache Type

The CWORKSTORAGE (or CWORKS) parameter specifies the type of cache space to use for the CWORK2FAC and/or CWORK3FAC areas.

CWORKSTORAGE = { DATASPACE | EXTENDED | G64 | HIPERSPACE | L64 | VIRTUAL64 }

Value	Description
DATASPACE	Data space
EXTENDED	Extended memory
HIPERSPACE	Hiperspace
G64	Virtual 64-bit storage backed by 2G large pages Note: 'If "G64" is specified, the sizes of the cache spaces for Work parts 2 and 3 will be rounded up to the nearest multiple of two gigabytes and any difference between the specified size and the allocated size will remain unused. To avoid leaving a large amount of cache storage unused, we recommend you use "G64" only if the sizes of the Work part 2 and Work part 3 cache spaces (as per the CWORK2FAC and CWORK3FAC parameters) will not be much smaller than a multiple of two gigabytes.
L64	Virtual 64-bit storage backed by 1M large pages
VIRTUAL64	Virtual 64-bit storage

Extended memory is the only option available on BS2000 RISC machines. Hiperspace is an option only in z/OS environments. Virtual 64 storage is available only in z/OS environments.

If CWORKSTORAGE is not specified, but the CWORK2FAC and/or CWORK3FAC parameters are specified, then the default CWORKSTORAGE is as follows:

- DATASPACE for BS2000 (except RISC machines), z/OS, z/VSE
- EXTENDED for BS2000 RISC machines

Example

```
ADARUN CWORKSTORAGE=HIPERSPACE
```

For Work parts 2 and/or 3, cache space is allocated in hiperspace.

CWORK2FAC Parameter: Adabas Caching Facility Work Part 2 Percentage

The CWORK2FAC (or CWORK2) parameter specifies the percentage of Work part 2 to be cached.

Value	Description
<i>n</i>	<p>The percentage of Work part 2 to be cached.</p> <p>The minimum and the default is 0; the maximum is 100.</p> <p>If CWORK2FAC=0 is specified, Work part 2 is not cached.</p>

The number of Work part 2 RABNs to be cached is computed by multiplying the specified percentage factor by the number of RABNs allocated for Work part 2.

Each cache space allocation is based on the number of RABNs and the device block size, rounded up to the nearest 4-kilobyte boundary.

```
cache size = int ( ( ( #RABNs x blksize ) + 4095 ) x 4096 ) + 4096
```

- where

```
#RABNs = int ( WORK part 2 size x CWORK2FAC )
```

If the specified CWORK2FAC cannot accommodate the working set of Work part 2 RABNs, the RABNs beyond the computed CWORK2FAC are not cached. Normal I/O operations are performed for the non-cached Work part 2 RABNs.

For any percentage less than 100, the writes to Work are still done; however, RABNs read from Work will be obtained from the cache if the RABN is available.

Example

```
ADARUN CWORK2FAC=50
```

ADACSH caches 50 percent of Work part 2, beginning from the first block. If Work part 2 contains 1000 blocks, a cache space is allocated for 500 blocks. If the Work device type is a 3380, the cache space size required is

```
cache size = int ( ( ( #RABNs x blksize ) + 4095 ) ÷ 4096 ) x 4096
            = int ( ( ( 500 x 5492 ) + 4095 ) ÷ 4096 ) x 4096
            = 2,750,095
```

CWORK3FAC Parameter: Adabas Caching Facility Work Part 3 Percentage

The CWORK3FAC (or CWORK3) parameter specifies the percentage of Work part 3 to be cached.

Value	Description
<i>n</i>	The percentage of Work part 3 to be cached. The minimum and the default is 0; the maximum is 100. If CWORK3FAC=0 is specified, Work part 3 is not cached.

The number of Work part 3 RABNs to be cached is computed by multiplying the specified percentage factor by the number of RABNs allocated for Work part 3.

Each cache space allocation is based on the number of RABNs and the device block size, rounded up to the nearest 4-kilobyte boundary.

```
cache size = int ( ( ( #RABNs x blksize ) + 4095 ) ÷ 4096 ) x 4096
```

- where

```
#RABNs = int ( WORK part 3 size x CWORK3FAC )
```

If the specified CWORK3FAC cannot accommodate the working set of Work part 3 RABNs, the RABNs beyond the computed CWORK3FAC are not cached. Normal I/O operations are performed for the non-cached Work part 3 RABNs.

For any percentage less than 100, the writes to Work are still done; however, RABNs read from Work will be obtained from the cache if the RABN is available.

Example

```
ADARUN CWORK3FAC=50
```

ADACSH caches 50 percent of Work part 3, beginning from the first block of Work part 3. If Work part 3 contains 1000 blocks, a cache space is allocated for 500 blocks. If the Work device type is a 3380, the cache space size required is

```
cache size = int ( ( ( #RABNs x blksize ) + 4095 ) ÷ 4096 ) x 4096
            = int ( ( ( 500 5492 ) + 4095 ) ÷ 4096 ) x 4096
            = 2,750,095
```

CXFILE Parameter: Adabas Caching Facility Excluded Files

The CXFILE parameter is used to exclude files from caching.

Note:

CXFILE is mutually exclusive with CASSO_{xxx}, CDATA_{xxx}, and CSTORAGE parameters. If any of these parameters is used with CXFILE, they will be ignored.

```
CXFILE = ( { fnr | fnr - fnr } [ ,, scope ] )
```

Value	Description
<i>fnr</i>	The file to be excluded specified by file number. The minimum value is 0; the maximum is the actual number of files in the database. File number 0 may be specified to exclude RABNs that are not associated with any given file. Currently, this is only the DSST on the Associator data set.
<i>fnr</i> - <i>fnr</i>	A range of files to be excluded specified by two file numbers (the lower number first) separated by a hyphen with no intervening spaces or special characters.
<i>scope</i>	<p>Specifies the scope of exclusion:</p> <ul style="list-style-type: none"> ● ASSO (or A): Only Associator RABNs are to be excluded ● DATA (or D): Only Data Storage RABNs are to be excluded ● BOTH (or B): Both Associator and Data Storage RABNs are to be excluded <p>If the scope term is specified, it must be preceded by two commas to maintain its position within the parentheses. See the CFILE parameter for more information about the positional terms used within the parentheses. The <i>cls</i> and <i>cache</i> terms are not used for the CXFILE parameter but the position of the <i>scope</i> term is nevertheless maintained.</p>

Multiple CXFILE parameters may be specified.

CFILE and CXFILE parameters together are processed in order from left to right (or top to bottom if different ADARUN statements are used).

Examples

```
ADARUN CACHE=YES , CFILE=( 13-19 ) , CXFILE=( 17 )
```

Files 13 through 19 are cached using the standard defaults for class of service (3), scope (BOTH), and cache (DATASPACE) with the exception of file 17, which is not cached.

```
ADARUN CACHE=YES ,CFILE=( 13-19 ) ,CXFILE=( 17 ) ,CXFILE=( 16 , ,DATA)
```

Files 13 through 19 are cached using the standard defaults; however, file 17 is not cached and only the associated Associator RABNs are cached for file 16 since the Data Storage RABNs are excluded.

DATACACHE Parameter: Data Storage Controller Caching

This parameter applies to the z/OS operating system only.

Parameter	Specify . . .	Possible Values	Default
DATACACHE	whether to enable (YES) or disable (NO) controller caching for Data Storage (DATA) data set.	YES NO	YES

Value Meaning

YES	The default setting. Data Storage blocks are pre-read and maintained as long as possible in the caching facilities of the disk controller.
NO	Data Storage blocks are not pre-read.

Notes:

1. If Data Storage blocks are highly scattered or the update frequency is low, controller caching can cause longer I/O response times.
2. ADARUN caching control parameters also exist for the Associator (ASSOCACHE), the work area (WORKCACHE), the sort area (SORTCACHE), and the temp area (TEMPCACHE).

Example

The following example runs the nucleus without controller caching for the Data Storage data set.

```
ADARUN PROG=ADANUC ,DATACACHE=NO
```

DBID Parameter: Database Identification

Parameter	Specify . . .	Minimum	Maximum	Default
DBID	the physical ID of the database.	1	65535	1

Note:

If you specify the DBID parameter, you must also specify the ADARUN SVC parameter.

Example

The following example runs the ADAINV utility against physical database 2 using SVC 241.

```
ADARUN PROG=ADAINV,DBID=2,SVC=241
```

DEVICE Parameter: Device Type

Parameter	Specify . . .	Possible Values	Default
<u>DEVICE</u>	the device type on which the first block of the Associator is stored.	see text	2300 (BS2000) 3390 (other)

The device type of the external storage device on which the first block of the Associator is stored.

This device type is also used as default for most device type parameters.

The Adabas Installation documentation specifies the device types supported by Adabas.

Some device type designations have been reserved for special use such as VSAM data set support on z/OS systems. The following is a partial list of reserved device types and their uses:

Device	Use
5555	VSAM DD/xxxxR5
6666	VSAM DD/xxxxR4
7777	VSAM DD/xxxxR3
8888	VSAM DD/xxxxR2
9999	VSAM DD/xxxxR1 - (default for self-defined VSAM block sizes)

For more information about VSAM device usage, see *Device and File Considerations* in the Adabas Installation documentation.

Example

The following example runs the ADALOD utility and stores the first block of the Associator on a 3390 device.

```
ADARUN PROG=ADALOD,DEVICE=3390
```

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
<u>DIRRATIO</u>	the directory portion of the directory-to-data element ratio	1	32767	4
<u>ELEMENTRATIO</u>	the element portion of the directory-to-data element ratio	1	32767	1

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

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

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

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

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

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

You can use the Adabas Online System cache structure size calculator described in *Estimating Sizes for the Cache Structure in a Cluster Environment* in the *Adabas Parallel Services Operations Guide* or 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.

DSF Parameter: Adabas Delta Save Facility Control

Parameter	Specify . . .	Possible Values	Default
DSF	whether to enable support for the (optional) Delta Save Facility.	YES NO	NO

The DSF parameter controls the Adabas Delta Save Facility for an Adabas nucleus or utility. If Adabas Delta Save is not installed on your system, specifying this parameter may cause an error.

Value Meaning

YES Loads certain Delta Save Facility program modules and runs the nucleus or utility (as specified in the ADARUN PROG parameter) in DSF mode. Delta Save functions are enabled: the DLOG area can be installed or removed, Delta Save operations can be performed, full and delta save tapes can be merged. DSF=YES is required for Delta Save functions of the nucleus and the ADASAV, ADARES, and ADAULD utilities; it is optional for other functions and utilities. It is required for ADASAV SAVE database (if a DLOG area exists), SAVE DELTA, MERGE, and RESTORE DELTA; for ADARES PLCOPY (only if online saves are performed), for COPY (only for rebuilding the DSIM data set); and for ADAULD UNLOAD SAVETAPE, if the inputs are full and delta save tapes produced by DSF.

NO (default) Runs the nucleus or utility in non-DSF mode. If DSF=NO is specified or the parameter is omitted, Delta Save functions cannot be performed. Specifying DSF=NO after specifying DSF=YES switches the nucleus from DSF to non-DSF mode. If a DLOG area is defined when the nucleus is started and DSF=NO is specified, the nucleus removes the DLOG area. If a DLOG area is defined when the nucleus is started, the DSF parameter may not be omitted. You can manually remove the DSF logging area from the nucleus using the Remove DSF command through Adabas Online System.

Example

The following example runs the Adabas nucleus in DSF mode; performing Adabas Delta Save functions is possible for this nucleus.

```
ADARUN PROG=ADANUC,DSF=YES
```

DSFEX1 Parameter: Delta Save Facility Exit

Parameter	Specify . . .	Possible Values	Default
<u>DSFEX1</u>	the user routine to be given control by the Delta Save Facility.	see text	none

Used in conjunction with the Adabas Delta Save Facility, the format of the user exit parameter is

```
DSFEX1=exit-name
```

where

`exit-name` is the name of the user routine that gets control at the DSF user exit; the name can be up to 8 characters long.

Only one program may be specified for the DSF user exit. Only one DSF user exit may be specified. See the Adabas Delta Save Facility documentation for more information.

Example

This example indicates that the Adabas nucleus should use the Delta Save Facility with the user exit routine DSAVEPGM.

```
ADARUN PROG=ADANUC,DSF=YES,DSFEX1=DSAVEPGM
```

DTP Parameter: Distributed Transaction Processing Control

Parameter	Specify . . .	Possible Values	Default
<u>DTP</u>	whether to enable distributed transaction processing and, if so, whether the nucleus is to be a resource or transaction manager.	NO RM TM	NO

The DTP parameter controls the distributed transaction processing status of the Adabas nucleus. If Adabas Transaction Manager is not installed on a separate nucleus in your system, specifying this parameter may cause an error.

Value	Meaning
NO	(default) The nucleus does not support distributed transaction processing.
RM	<p>The nucleus is to participate in distributed transaction processing as a resource manager (RM) with other resource managers (databases) coordinated by a transaction manager. This setting is only possible when</p> <ul style="list-style-type: none">- MODE=MULTI- READONLY=NO <p>If you have Adabas Transaction Manager Version 7.4 or earlier installed and if DTP=RM, the LDTP parameter specifies the size of the Work part 4 area that the nucleus uses to keep information about distributed transactions.</p> <p>If you have Adabas Transaction Manager Version 7.5 or later installed and if DTP=RM, the LDTP parameter specifies the size (in blocks) of the WORK4 index which is used to administer the data on DDWORKR4.</p>
TM	<p>The nucleus supports distributed transaction processing as a transaction manager (TM). Only one such nucleus is allowed per operating system instance. This setting is only possible when</p> <ul style="list-style-type: none">- MODE=MULTI- the nucleus is not part of a cluster (Parallel Services or Cluster Services)- LOCAL=NO- READONLY=NO

Example

The following example indicates that the nucleus will participate as a resource manager in a distributed transaction processing system coordinated by a transaction manager.

```
ADARUN PROG=ADANUC,DTP=RM
```

DUALCLD Parameter: Dual Command Log Device Type

DUALCLS Parameter: Dual Command Log Size

Parameter	Specify . . .	Possible Values	Default
DUALCLD	the device type to be used for dual command logging.	see text	value of DEVICE parameter
DUALCLS	the number of blocks available for each dual command log.	Min 16, Max 16777215	none

The DUALCLD and DUALCLS parameters specify the type of external storage device used for dual command logging and the number of blocks available for each dual command log data set. These parameters must be used *together*; if you specify one parameter, you must also specify the other parameter. If you specify a user exit, it must be UEX2.

The Adabas Installation documentation specifies the device types that Adabas supports.

Both dual command logging data sets must be formatted using the ADAFRM utility before they are used.

The maximum size of the data set depends on the limits of the operating system and on the requirements of the user.

Example

The following example runs the Adabas nucleus with command logging and uses a model 3380 device for dual command logging. The size of each command log data set is 1350 blocks, based on an allocation of 10 cylinders. (Each cylinder of a 3380 device has 15 tracks, and each track can contain 9 blocks of command log data. Thus, 10 cylinders = 10 x 15 x 9 = 1350 blocks).

```
ADARUN PROG=ADANUC,LOGGING=YES,DUALCLD=3380,DUALCLS=1350
```

DUALPLD Parameter: Dual Protection Log Device Type

DUALPLS : Dual Protection Log Size

Parameter	Specify . . .	Possible Values	Default
DUALPLD	the device type to be used for dual protection logging.	see text	value of DEVICE parameter
DUALPLS	the number of blocks available for each dual protection log.	Min 16, Max 16777215	none

The DUALPLD and DUALPLS parameters specify the type of external storage device used for dual data protection logging and the number of blocks available for each log data set. These parameters must be used *together*; if you specify one parameter, you must also specify the other parameter. If you specify a user exit, it must be UEX2.

When these parameters are specified, dual protection logging is activated for the session. See *Adabas Restart and Recovery* for information about dual protection logging.

The Adabas Installation documentation specifies the device types that Adabas supports.

Both dual protection log data sets must be formatted using the ADAFRM utility before they are used. The appropriate JCL for the data sets is required.

Example

The following example runs the Adabas nucleus with dual data protection logging and uses a model 3380 device; the size of each dual log data set is 1200 blocks, based on an allocation of 10 cylinders. (Each cylinder of a 3380 device has 15 tracks, and each track contains 8 blocks of protection log data. Thus, 10 cylinders = 10 * 15 * 8 = 1200 blocks).

```
ADARUN PROG=ADANUC ,DUALPLD=3380 ,DUALPLS=1200
```

EXCPVR Parameter: Control EXCP or EXCPVR Use

Parameter	Specify . . .	Possible Values	Default
EXCPVR	whether EXCP or EXCPVR should be used when running APF-authorized, in z/OS environments.	YES NO	YES

The EXCPVR parameter controls whether EXCP or EXCPVR should be used when from an APF-authorized load library in z/OS environments. If, EXCPVR=YES is specified, EXCPVR is used; if EXCPVR=NO is specified, EXCP is used.

Note:

If EXCPVR=YES is specified when running from a non-APF-authorized load library, this parameter is ignored.

Example

The following example runs Adabas APF-authorized using EXCP.

```
ADARUN PROG=ADANUC ,EXCPVR=NO
```

FASTPATH Parameter: Adabas Fastpath Control

Parameter	Specify . . .	Possible Values	Default
<u>F</u> ASTPATH	whether to enable Adabas Fastpath support.	YES NO	NO

The FASTPATH parameter controls the optional Adabas Fastpath product. If Adabas Fastpath is installed on your system, FASTPATH=YES enables the full functionality of the product.

If Adabas Fastpath is not installed on your system, FASTPATH=YES invokes the Adabas Fastpath command analysis sampler AFPLOOK that is provided with Adabas.

The demo version of Adabas Online System that is supplied with Adabas provides access to the Adabas Fastpath online system SYSAFP so that the AFPLOOK sample can be viewed online exactly as it would appear if Adabas Fastpath were installed. For complete information about AFPLOOK, read *AFPLOOK*.

Example

The following example runs Adabas Fastpath with Adabas.

```
ADARUN PROG=ADANUC,FASTPATH=YES
```

FMXIO Parameter: Parallel LFIOP I/O Operations

Parameter	Specify . . .	ASYTVS=	Minimum	Maximum	Default
<u>F</u> MXIO	the maximum number of parallel I/O operations by LFIOP flush processing.	YES	1	16	1
		NO	1	100	60

The LFIOP parameter enables asynchronous buffer flush operation and sets the I/O pool size. The FMXIO parameter sets the limit on the number of I/O operations that can be started in parallel by LFIOP flush processing.

Parameter Dependencies

The precise meaning of the FMXIO parameter depends on the setting of the ASYTVS parameter):

- When ASYTVS=YES (perform buffer flushes by volume), FMXIO specifies the number of I/Os to be started in parallel *on each volume*. The minimum and default number is 1; the maximum number is 16. If the specified number exceeds that, it is automatically reduced to 16.
- When ASYTVS=NO (perform buffer flushes in ascending RABN sequence, without regard to the distribution of the blocks over volumes), FMXIO specifies the number of I/Os to be started in parallel overall. The minimum number is 1; the maximum value is 100; and the default is 60.

Dynamic Modification

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

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.

Example

The following example specifies that a maximum of 30 I/O operations for asynchronous buffer flushes can be started in parallel:

```
ADARUN  PROG=ADANUC ,ASYTVS=NO ,FMXIO=30
```

FORCE Parameter: Allow Nucleus Database ID or Review Hub Table Entry Overwrite

Parameter	Specify . . .	Possible Values	Default
<u>FORCE</u>	whether the nucleus or Adabas Review hub can overwrite an existing ID table entry.	YES NO	NO

If running Adabas Review, this indicates whether the Adabas Review hub can overwrite an existing ID table entry. When a Review hub starts up, ADARUN scans the ID table to ensure that no entry exists for the Review hub. You can use the FORCE parameter to indicate whether the Review hub can overwrite an existing ID table entry.

The ID table entry is derived from the database ID and the job name. For Adabas Review, the ID table entry is derived from the Review hub ID (REVIEW=). The ID table entry is deleted when the nucleus terminates normally.

The FORCE parameter allows the nucleus or Adabas Review hub to overwrite the existing ID table entry and access the database.



Warning:

Do not use the FORCE parameter unless absolutely necessary, or the integrity of the database could be lost. Ensure that no nucleus or Review hub is active for the ID table entry being overwritten.

Value	Meaning
YES	The nucleus or Adabas Review hub that is starting can overwrite an existing ID table entry. FORCE=YES is required when restarting a session that terminated abnormally with an ADAM98 message. In this case, the ID table still contains an active entry for the nucleus or Review hub. Overwriting the existing entry by specifying FORCE=YES prevents further communication to the overwritten nucleus or hub and causes loss of cross-memory environment resources, which cannot be restored until the next IPL.
NO	(default) If the ID table contains an entry for the nucleus or Adabas Review hub that is starting, the nucleus is denied access to the database or the Review hub is not permitted to start.

Note:

In an Adabas Cluster Services or Adabas Parallel Services environment, the FORCE parameter applies to the NUCID, rather than the database ID, because a cluster nucleus builds an ID table entry for the NUCID.

A data integrity block (DIB) entry will only be removed once the ID Table initialization has been successful. Therefore, you must set IGNDIB and FORCE both to "YES" if either of the following occur:

- You receive a PARM ERROR 26 after parameter settings IGNDIB=NO and FORCE=YES were applied;
- You receive a PARM ERROR 23 after parameter settings IGNDIB=YES and FORCE=NO were applied.

Examples

The following example specifies that if the ID table contains an active entry for DBID 7, overwrite the entry.

```
ADARUN PROG=ADANUC, FORCE=YES, DBID=7
```

The following example specifies that if the ID table contains an active entry for the Adabas Review hub, overwrite the entry.

```
ADARUN PROG=ADAREV, FORCE=YES, REVIEW=202
```

GROUPS Parameter: User Group Interprocess Communication

This parameter applies to the BS2000 operating system only.

Parameter	Specify . . .	Possible Values	Default
<u>G</u> ROUPS	whether BS2000 interprocess communication is limited to users with the same logon ID.	YES NO	NO

This parameter limits the accessibility of ADARER and its ID table (IDT) to a group of users with the same logon ID.

Value	Meaning
YES	Allows the user(s) access to the IDT that starts under the same logon ID as the Adabas nucleus. When initialized, the IDT is set to GROUPS=YES; all users (nuclei, utilities, Entire Net-work, and all user tasks) must therefore also specify GROUPS=YES in either the ADARUN statement or using the ADALNK parameters.
NO	The default setting. All users on one machine have access to this IDT, even if they are logged on under a different logon ID.

Note:

More than one IDT with the same name can be used if one is set to GROUPS=NO (the default) and the rest specify GROUPS=YES.

Example

The following example allows BS2000 users access to the IDT that starts under the same logon ID as this Adabas nucleus.

```
ADARUN PROG=ADANUC,DBID=11,GROUPS=YES
```

HEXnn Parameter: Hyperdescriptor User Exit

Parameter	Specify . . .	Possible Values	Default
HEXnn	the hyperdescriptor exit and the user routine to be used with it.	see text	none

Used in conjunction with the hyperdescriptor facility, the format of the parameter is

```
HEXnn=exit-name
```

where

nn is the number of the hyperdescriptor exit, a two-digit decimal integer in the range 01-31 inclusive.

exit-name is the name of the user routine that gets control at the hyperdescriptor exit; the name can be up to 8 characters long.

Only one program may be specified for each hyperdescriptor exit. Up to 31 hyperdescriptor exits may be specified (in any order). See the *Adabas User, Hyperdescriptor, Collation Descriptor, and SMF Exits Manual* for more information.

Example

The following example runs the Adabas nucleus using hyperdescriptor exits 1 and 2. It also gives control to the user routine UHYPER1 at hyperdescriptor exit 1 and gives control to the user routine UHYPER2 at hyperdescriptor exit 2.

```
ADARUN PROG=ADANUC,HEX01=UHYPER1,HEX02=UHYPER2
```

IDTNAME Parameter: Define ID Table Name

This parameter applies to the BS2000 operating system only.

Parameter	Specify . . .	Possible Values	Default
IDTNAME	an alternate ID table.	ADAccccc	ADABAS5B

Defines a name for an (alternate) Adabas environment. The name must be eight characters long and must begin with the characters "ADA".

Note:

All users of the new ID table must include the IDTNAME in either the ADARUN statement or ADALNK parameters.

Specific Product Recommendations

For Entire Net-Work 6 (mainframe), Software AG recommends setting the value of the IDTNAME parameter to "ADABAS5B".

Example

The following example defines the name ADAOURDB for the second Adabas environment.

```
ADARUN PROG=ADANUC, IDTN=ADAOURDB
```

IDTPSUP Parameter: Suppress IDTNAME in Adabas Output Messages

This parameter applies to the BS2000 operating system only.

Parameter	Specify . . .	Possible Values	Default
IDTPSUP	whether the IDTNAME should be suppressed in the Adabas nucleus output messages.	YES NO	NO

This parameter allows you to suppress the IDTNAME in Adabas messages such as:

```
ADAM83 dbid NEW IDT CREATED, NAME=idtname, GROUPS=NO
```

Example

The following specification produces output messages with X's replacing the IDT name.

```
ADARUN IDTPSUP=YES
```

For example, the ADAM83 message would print as:

```
ADAM83 dbid NEW IDT CREATED, NAME=XXXXXXXX, GROUPS=NO
```

IGNDIB Parameter: Ignore DIB Entry for Adabas Nucleus

Parameter	Specify . . .	Possible Values	Default
IGNDIB	whether to ignore and delete an active nucleus entry in the data integrity block.	YES NO	NO

The DIB is the data *integrity* block in the Associator. The DIB contains an entry that indicates whether an Adabas session is active for the database. This feature prevents an update nucleus session from being started while another one is running on the same database; it is especially important when several CPUs share the same database DASD devices. The DIB entry is automatically released when the Adabas session is terminated normally.

This parameter indicates whether this update nucleus session should ignore and delete the entry in the data integrity block (DIB) in the Associator.

Value	Meaning
YES	Adabas ignores and deletes the existing DIB entry and initiates a new session.
NO	The default setting. If the DIB entry indicates that an Adabas session is already active for the database, Adabas rejects an attempt to start an update session against the database. This is true even when the session that the DIB indicates is active has terminated abnormally.

Note:

Software AG recommends that you use IGNDIB=YES only once and immediately remove it from the JCL. If no DIB entry exists when IGNDIB=YES is specified, the nucleus terminates abnormally.

A data integrity block (DIB) entry will only be removed once the ID Table initialization has been successful. Therefore, you must set IGNDIB and FORCE both to "YES" if either of the following occur:

- You receive a PARM ERROR 26 after parameter settings IGNDIB=NO and FORCE=YES were applied;
- You receive a PARM ERROR 23 after parameter settings IGNDIB=YES and FORCE=NO were applied.

Example

The following example runs an Adabas update session and ignores and deletes the DIB entry for an active Adabas update nucleus.

```
ADARUN PROG=ADANUC,IGNDIB=YES
```

IGNDTP Parameter: Ignore Two-Phase Commit Area (Work Part 4)

Parameter	Specify . . .	Possible Values	Default
IGNDTP	whether to ignore the two-phase commit area (Work part 4).	YES NO	NO

**Warning:**

This parameter is for emergency use only; contact your Software AG technical support representative before using it. Whenever the information on Work part 4 is ignored, the integrity of the incomplete global transactions that are related to that information cannot be guaranteed.

IGNDTP=YES can be specified with DTP=NO or DTP=RM.

When DTP=NO, the nucleus will not start if there is data in Work part 4 related to the recovery of previously incomplete global transactions. In this situation, specifying IGNDTP=YES forces the nucleus to ignore the data in Work part 4 and start anyway.

When DTP=RM, the nucleus also ignores any data on Work part 4 and reinitializes the area. This can be useful during a recovery when the state and correctness of the incomplete global transactions that are stored on Work part 4 has already been established.

Example

The following example runs an Adabas update session without support for distributed transactions; it ignores data in the Work part 4 area left over from the previous session with support for distributed transaction processing.

```
ADARUN PROGRAM=ADANUC ,DTP=NO ,LDTP=8 ,IGNDTP=YES
```

INFOBUFFERSIZE Parameter: Information Buffer Pool (Adabas Event Log) Size

Parameter	Specify . . .	Possible Values	Default
<u>INFOBUFFERSIZE</u>	the size of the information buffer pool, in bytes. The information buffer pool is also known as the Adabas event log.	0 or 1024 - 2147483647	0 (no buffer pool is allocated)

The information buffer pool is a wraparound buffer in memory that represents the size of the new Adabas event log. At this time, the Adabas event log is used to log each response code 145 (ADARSP145) event. The INFOBUFFERSIZE parameter identifies the size of the Adabas event log. Each entry in the event log is currently 128 bytes, although this may change in later Adabas releases. When the Adabas event log fills up, the oldest entries in the log are overwritten.

Setting INFOBUFFERSIZE to zero (0) indicates that no buffer pool (Adabas event log) is allocated. The minimum size of the pool is 1024 bytes.

Note:

Each response code 145 (ADARSP145) event is logged to the command log, if CLOGLAYOUT=8 is specified, regardless of the setting of the INFOBUFFERSIZE parameter.

The information in the Adabas event log can be printed using the ADADBS DEVENTLOG utility function. For more information, refer to the Adabas Utilities Manual.

Example

The following example specifies an information buffer pool (Adabas event log) that is 50,000 bytes in size.

```
ADARUN PROGRAM=ADANUC,INFOBUFFERSIZE=50000
```

INTAUTO Parameter: Time Interval Between Autorestart Progress Messages

Parameter	Specify . . .	Minimum	Maximum	Default
<u>INTAUTO</u>	the time interval (in seconds) between autorestart progress messages.	1	16777215	0

When performing session autorestart to recover from a preceding failure, Adabas prints autorestart progress messages (in message ADAN56) to indicate how much repair and recovery work has already been done and how much remains to be done. By default, these messages occur when each phase of the autorestart (backward repair, forward repair, autobackout) completes.

The INTAUTO parameter allows you to set the time interval (in seconds) at which autorestart progress messages will be printed, in addition to the standard messages printed at the end of each phase (assuming a phase takes longer than the specified time interval). The default (0 seconds) indicates that autorestart progress messages should only be printed at the completion of each autorestart phase.

Example

The following ADARUN INTAUTO parameter indicates that Adabas should print progress messages every ten seconds when performing session autorestart processing.

```
ADARUN PROGRAM=ADANUC,INTAUTO=10
```

INTNAS Parameter: Time Interval Between SYNS 60 Checkpoints

Parameter	Specify . . .	Minimum	Maximum	Default
<u>INTNAS</u>	the time interval between SYNS 60 checkpoints.	1	16777215	3600

The SYNS 60 statistic checkpoint is generated by the Adabas nucleus and written to the Adabas Checkpoint system file. It may be displayed using the CPLIST function of the ADAREP utility or Adabas Online System. For more information about checkpoints see the ADAREP utility discussion in the Adabas Utilities Manual.

INTNAS allows you to set the time interval (in units of 1.048576 seconds) between SYNS 60 checkpoints. The default (3600) is approximately 1 hour meaning that nucleus statistics will be recorded once each hour.

Example

The following example sets Adabas so that nucleus statistics are recorded every 30 minutes.

```
ADARUN PROGRAM=ADANUC,INTNAS=1800
```

LARGEPAGE Parameter: Large Page Support Option

Parameter	Specify . . .	Possible Values	Default
LARGEPAGE	whether or not the Adabas nucleus should use page-fixed one-megabyte (1M) large pages of real storage above the two-gigabyte (2G) bar (LARGEPAGE=L64) or page-fixed 2G large pages above the 2G bar (LARGEPAGE=G64). A value of "YES" is equivalent to specifying a value of <i>L64</i> . Large pages can only be used if your operating system supports them and 64-bit virtual storage.	G64 L64 NO YES	NO

This parameter indicates whether or not the Adabas nucleus should use large pages. The default is NO.

Note:

At this time, this parameter is applicable only to z/OS systems.

Value	Meaning
G64	<p>The Adabas nucleus will attempt to acquire virtual storage above the 2G bar backed by page-fixed 2G large pages. The virtual storage will be backed by large pages if the operating system supports them and if sufficient large pages are available in the system. If insufficient large pages are available, the virtual storage above the 2G bar will be backed by pageable four-kilobyte (4K) pages.</p> <p>You can specify <code>LARGEPAGE=G64</code> only if the <code>ADARUN V64BIT</code> parameter is set to "YES".</p> <p>The <code>LARGEPAGE=G64</code> setting does not apply to the flush I/O pool, Work and PLOG I/O buffers, or cast-out directory buffer (in a cluster environment).</p> <p>Note: Allocation requests for storage to be backed by 2G large pages are rounded up to the nearest multiple of two gigabytes. Any difference between the specified size and the allocated size will remain unused. For this reason, we recommend that you specify a multiple of two gigabytes for any storage area to which the <code>LARGEPAGE=G64</code> setting applies, to avoid leaving a potentially large amount of allocated storage unused.</p>
L64 or YES	<p>The Adabas nucleus will attempt to acquire virtual storage above the 2G bar backed by page-fixed 1M large pages. The virtual storage will be backed by large pages if the operating system supports them and if sufficient large pages are available in the system. If insufficient large pages are available, the virtual storage above the 2G bar will be backed by pageable 4K pages. A value of "YES" is equivalent to specifying a value of "L64".</p> <p>You can specify <code>LARGEPAGE=L64</code> or <code>LARGEPAGE=YES</code> only if the <code>ADARUN V64BIT</code> parameter is set to "YES".</p>
NO	<p>Virtual storage above the 2G bar will be backed by 4K pages. This is the default value.</p>

The amount of storage in your operating system available for page-fixed 1M or 2G large pages is set using the `LFAREA` parameter in `PARMLIB` member `IEASYSxx`. For complete information on the `LFAREA` parameter, refer to the *IBM z/OS Extended Addressability Guide*.

Example

The following example requests that the Adabas nucleus should make use of virtual storage above the 2G bar and should use 1M large pages.

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


LBP Parameter: Length of Buffer Pool

Parameter	Specify . . .	Minimum	Maximum	Default
<u>LBP</u>	the maximum size of the Adabas buffer pool.	80 000	2147483647	350000

The maximum number of bytes to be used for the Adabas buffer pool during a session.

Adabas allocates the number of bytes specified to the Adabas buffer pool if sufficient memory is available. A minimum of 80000 bytes must be available to successfully initiate a session. For performance reasons, however, Software AG recommends that you specify at least 50000 bytes per thread. Ensure that sufficient memory is available to accommodate the LBP size specified. If asynchronous buffer flush management is used (that is, the LFIOP parameter specifies a nonzero value), the LBP value must be greater than that specified by LFIOP.

In environments running in 31-bit addressing mode, the buffer pool space is allocated above the 16-megabyte line.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

The buffer efficiency statistic can be used to determine the optimal size of the parameter.

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

The following example allocates a maximum of 800,000 bytes to the Adabas buffer pool for this session.

```
ADARUN PROG=ADANUC,LBP=800000
```

LCP Parameter: Length of Security Pool

Parameter	Specify . . .	Minimum	Maximum	Default
<u>LCP</u>	the maximum size of the Adabas security pool.	2000	16777215	10000

The maximum number of bytes to be used for the Adabas security pool during a session. The pool is used to keep security information in memory in order to optimize performance.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

Example

The following example allocates a maximum of 50,000 bytes to the Adabas security pool for this session.

```
ADARUN PROG=ADANUC,LCP=50000
```

LDEUQP Parameter: Length of Unique Descriptor Pool

Parameter	Specify . . .	Minimum	Maximum	Default
LDEUQP	the size of the unique descriptor pool.	5000	16777215	5000

Note:

Zap AN822111 has increased the minimum setting size of the ADARUN LDEUQP parameter from zero to 5000 bytes.

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
```

LDTP Parameter: Length of Distributed Transaction Processing Area (Work Part 4)

Parameter	Specify . . .	Minimum	Maximum	Default
<u>LDTP</u>	the number of blocks to allocate to Work part 4 if you have Adabas Transaction Manager Version 7.4 or earlier installed.	0	65535	0
<u>LDTP</u>	the size (in blocks) of the WORK4 index which is used to administer the data on DDWORKR4 if you have Adabas Transaction Manager Version 7.5 or later installed.	8	one eighth of the WORK4 size	8

If you have Adabas Transaction Manager Version 7.4 or earlier installed, this parameter specifies the length of the distributed transaction processing area, Work part 4:

- If DTP=RM and LDTP is either unspecified or specified as '0', the length of Work part 4 is equivalent to that specified for Work part 1 in the ADARUN LP parameter.
- If specified, the ADARUN LDTP must be at least 200 blocks. A safe size of LP/4 is a good value to start with. If you do not specify an LDTP value, it is set to the value of the LP parameter.
- When Work part 4 contains data about incomplete transactions, it may be increased in size but not decreased. Work part 4 must be empty before it can be made smaller.

If you have Adabas Transaction Manager Version 7.5 or later installed, this parameter specifies the size (in blocks) of the WORK4 index which is used to administer the data on DDWORKR4:

- If DTP=RM and LDTP is either unspecified or specified as "0", a default of 8 blocks is used.
- If specified, the ADARUN LDTP must be at least 8 blocks. The maximum size that can be specified is the size of the DDWORKR4 data set divided by 8.

DDWORKR4 is a container data set used for the same purpose as Work part 4 of DDWORKR1 was used previously, with the difference that DDWORKR4 can be used in parallel by all members in a cluster. The DDWORKR4 data set should be allocated and formatted in the normal way, using a block size greater than or equal to DDWORKR1. It should be at least as large as the cluster's LP parameter of the database or cluster.

Example

In the following example, the LP parameter specifies 2000 blocks for the data protection area (Work part 1). This example also uses 2500 blocks for the distributed transaction processing area on the Adabas Work data set (part 4).

```
ADARUN PROG=ADANUC,DTP=RM,LP=2000,LDTP=2500
```

LFIOP Parameter: Asynchronous Buffer Pool Control

Parameter	Specify . . .	Minimum	Maximum	Default
<u>LFIOP</u>	whether to enable asynchronous buffer flush operation and set the I/O pool size.	0 1 80000	18,446,744,073,709,551,615 (although this is really limited by the available amount of virtual and real memory)	0

The asynchronous buffer flush I/O pool enables asynchronous Adabas buffer flush operation and defines the size of the related buffer. Specifying a non-zero value for LFIOP enables asynchronous buffer flush operation; the value specified determines the buffer size, in bytes.

If LFIOP is not specified or zero is specified (the default), synchronous buffer flushing is performed. Note that flushing is always synchronous in single-user mode (MODE=SINGLE).

Value	Meaning
0	No asynchronous buffer flush (the default); flushing is synchronous.
1	Asynchronous buffer flush operations occur. The flush pool is sized at 25% of the buffer pool (LBP=) size.
n (minimum of 80000)	Asynchronous buffer flush operations occur. The flush pool size is defined by the LFIOP setting (in bytes).

When LFIOP=1 or LFIOP=n and the specified or calculated flush pool size is less than 80000 bytes, Adabas uses a size of 80000 bytes.

In environments running in 31-bit addressing mode, the I/O pool is allocated above the 16-megabyte line.

In environments supporting 64-bit virtual storage, the flush I/O pool is allocated above the 2 gigabyte bar when the V64BIT ADARUN parameter is set to "YES".

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

The following example enables asynchronous buffer flush operation for this session and allocates an I/O pool of 200,000 bytes (25% of the LBP size).

```
ADARUN PROG=ADANUC,LFIOP=1,LBP=800000
```

LFP Parameter: Length of Internal Format Buffer Pool

Parameter	Specify . . .	Minimum	Maximum	Default
LFP	the size of the Adabas internal format buffer pool.	6000	2147483647	12000

This pool stores internal format buffers used with Adabas read and update commands as described in the *Adabas Command Reference Guide*.

Each internal format buffer pool entry requires approximately $60 + (64 * n)$ bytes, where n equals the number of fields referenced in the command.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

In environments running in 31-bit addressing mode, the LFP pool is allocated above the 16-megabyte line.

Specific Product Recommendations

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

The following example allocates 8,000 bytes for the Adabas internal format buffer pool for this session.

```
ADARUN PROG=ADANUC,LFP=8000
```

LI Parameter: Length of Table of ISNs (TBI)

Parameter	Specify . . .	Minimum	Maximum	Default
LI	the size of the TBI.	2000	16777215	10000

LI specifies the number of bytes allocated for the table of ISNs (TBI), which is used to store ISN lists (overflow ISNs or saved ISN lists). Each ISN list is identified by a command ID (CID) and stored in a TBI element.

You must also specify a value for the NSISN parameter, which sets the maximum number of ISNs per TBI element. Each TBI element has a length of $64 + (NSISN * 4)$ bytes.

The number of TBI elements required determines the length of the TBI. For complex searches, the Adabas nucleus builds temporary TBI elements in addition to the permanent TBI element; for example, for holding the results of descriptor searches, for reading the result records in Data Storage and collecting matches, or for holding the results of the partial criteria associated with the R operator. Although the temporary TBI elements are deleted at the end of a search, the TBI must be large enough to hold all of the temporary and permanent TBI elements in parallel during search processing.

The Adabas session statistics or Adabas Online System can be used to set this parameter for the next session.

When running with Adabas expanded files, a larger-than-normal LI value is recommended.

In environments running in 31-bit addressing mode, the LI pool is allocated above the 16-megabyte line.

Example

The following example allocates 30,000 bytes for the table of ISN list IDs for this session.

```
ADARUN PROG=ADANUC,LI=30000
```

LNKNAME Parameter: Link Globals Table Name

Parameter	Specify . . .	Possible Values	Default
LNKNAME	the module name of the link globals table to be employed by an Adabas 8 batch/TSO link routine. The link globals table is used to obtain default information and for invoking any linked exits.	valid module name for an IBM z/OS or z/VSE operating system, or for a BS2000 module	LNKGBLS

The parameter name LNKNAME may be abbreviated to LGN.

Notes:

1. If you specify the LNKNAME parameter, you must also specify the ADARUN DBID and SVC parameters.
2. To use this and other keywords for a batch or TSO job step, the application must be linked with the ADAUSER load module and call Adabas to invoke the ADAUSER load module.
3. If the job step specifies PROG=USER or PROG=RENTUSER with no LNKNAME keyword, either the ADALNK or ADALNKR module will be loaded by ADARUN. This is intended to allow the Adabas 7 batch/TSO link routines to be loaded and executed for the job step. Take care to ensure the proper load library concatenation is provided on the job step so the Adabas 7 ADALNK or ADALNKR module will be loaded. This should allow gradual migration of job steps from Adabas 7 to Adabas 8 as required.

Example:

In the following example, the link globals table in module MYGBLS is used.

```
ADARUN PROG=USER, LNKGNAME=MYGBLS
```

LOCAL Parameter: Local Nucleus or Adabas Review Hub

Parameter	Specify . . .	Possible Values	Default
<u>L</u> OCAL	whether an Adabas nucleus or Adabas Review hub is isolated and available for local use only. At this time, the only valid value for an Adabas Review hub is "YES".	YES NO	YES

Defines an isolated Adabas nucleus or Adabas Review hub that is only available locally. The nucleus or hub is unknown to Entire Net-Work. A nucleus or Adabas Review hub specifying LOCAL=YES (the default) can have the same database ID or hub ID as another database nucleus or Review hub on another network node.

Value	Meaning
YES	Isolates this nucleus or Adabas Review hub (that is, makes it unaddressable) from other Entire Net-Work nodes. This is the only value valid at this time for Adabas Review hubs; hubs are not currently supported in an Entire Net-Work environment.
NO	Allows the nucleus or Adabas Review hub to receive calls from other Entire Net-Work nodes. This value is not currently supported for Adabas Review hubs.

Example

In the following example, the Adabas nucleus is isolated and cannot be addressed by other Entire Net-Work nodes.

```
ADARUN PROG=ADANUC, LOCAL=YES
```

LOGGING / LOGxxx Parameters: Command Logging Control

Parameter	Specify . . .	Possible Values	Default
<u>L</u> OGGING	whether to log Adabas commands.	YES NO	NO

The LOGGING parameter specifies whether to use command logging for the Adabas session. The LOGxxxx parameters specify the type of information to be logged. Valid values are YES and NO for all of these parameters.

Note:

User exit 4 is still called even if LOGGING=NO. The only way to disable user exit 4 is to remove the ADARUN UEX4 parameter from the Adabas run.

If you specify LOGGING=YES, you must also specify YES for one or more of the following ADARUN LOGxxxx parameters. By default, each type of information is *not* logged.

Note:

Some parameters are only valid when CLOGLAYOUT=5; other parameters are only valid when CLOGLAYOUT=8.

Parameter	Specify whether to log the ...	Valid for CLOGLAYOUTs...	Possible Values	Default
LOGABDX	Adabas buffer descriptions (ABDs)	8	YES NO	NO
LOGCB	extended Adabas control block ¹	8 ²	YES NO	NO
LOGCLEX	command log extension (CLEX) ²	5, 8	YES NO	NO
LOGFB	format buffer	5, 8	YES NO	NO
LOGIB	ISN buffer	5, 8	YES NO	NO
LOGIO	I/O activity	5, 8	YES NO	NO
LOGMB	multifetch buffer	8	YES NO	NO
LOGRB	record buffer	5, 8	YES NO	NO
LOGSB	search buffer	5, 8	YES NO	NO
LOGUX	user exit B data	5, 8	YES NO	NO
LOGVB	value buffer	5, 8	YES NO	NO
LOGVOLIO	extended I/O list ³	5, 8	YES NO	NO

¹ You can specify the LOGCB parameter with CLOGLAYOUT=5, however its value will be ignored because when CLOGLAYOUT=5 is specified, the Adabas control block is always written to the command log file.

² LOGCLEX=YES requires REVIEW=NO although the Adabas Review modules must be available. Command log records with CLEX can be viewed using the batch program REVIEWB.

³ If the LOGIO parameter is set to "NO", no I/O buffer element is logged at all. If LOGIO is set to "YES" but LOGVOLIO is set to "NO", only the standard I/O buffer element is logged. The extended I/O list is only logged if both the LOGIO and LOGVOLIO parameter are set to "YES".

You can use Adabas operator commands or Adabas Online System to modify logging parameters during session execution.

Example

The following example runs the Adabas nucleus with command logging and logs the Adabas control block.

```
ADARUN PROG=ADANUC,LOGGING=YES,LOGCB=YES
```

LOGSIZE Parameter: Maximum Command Log Size

Parameter	Specify . . .	Minimum	Maximum	Default
LOGSIZE	the maximum block size for the sequential command log DD/LOG.	100	32760	value of QBLKSIZE

This parameter controls the maximum block size for the sequential command log DD/LOG. The default value is identical to the value of the QBLKSIZE parameter. If dual or multiple command logging is used, this parameter is ignored.

Example

In the following example, the block size of the command log is set as 2,000 bytes.

```
ADARUN PROG=ADANUC,LOGSIZE=2000
```

LOGWARN Parameter: CLOG and PLOG Alert Message Interval

Parameter	Specify . . .	Minimum	Maximum	Default
LOGWARN	the frequency, in seconds, at which the PLOG and CLOG status is checked and resulting alert messages are produced.	0	2147483647	0

This parameter is optional, but a valid user exit 2 or user exit 12 must also be specified when this parameter is set to any value other than zero (0). A value of zero (0), the default, indicates that no CLOG and PLOG status checks should occur and that no alert messages should be produced. Any other valid LOGWARN value represents the interval (in seconds) at which the PLOG and CLOG status is checked and resulting alert messages are produced.

After checking the status of the CLOG and PLOG data sets, alert messages are written when Adabas detects that a PLOG or CLOG data set has not been copied within a specified time interval. Thereafter, an alert message is written at regular intervals as long as there is a PLOG or CLOG data set waiting to be copied. These alert messages are issued until the PLOG or CLOG data set is copied or until no logs

remain to be copied. Your user exits 2 or 12 can be set up so that the messages appear providing a comprehensive trace of activity so you can resolve and problems with the PLOG or CLOG.

This feature helps you prevent the potential problem of Adabas overwriting an older CLOG or PLOG data set without waiting for direction from the user. Ordinarily, when a PLOG or CLOG data set fills up, the Adabas nucleus examines the return code from user exit 2 or user exit 12 to determine whether it should wait before overwriting an older PLOG or CLOG data set with new information. If there is no user exit 2 or 12, Adabas does not wait, but simply overwrites the older log. If this LOGWARN parameter is set to a value other than zero, the statuses of the CLOG and PLOG data sets are analyzed at the frequency specified by this parameter, warning about their status and providing you with the information you need to prevent any overwriting of older logs.

Example

In the following example, the CLOG and PLOG statuses are checked every hour.

```
ADARUN PROG=ADANUC,LOGWARN=3600
```

LP Parameter: Length of Data Protection Area (Work Part 1)

Parameter	Specify . . .	Minimum	Maximum	Default
LP	the number of blocks to allocate to Work part 1.	200	see text	1000

The LP parameter specifies the number of blocks to allocate to the Work data protection area, which is part 1 of the Work data set. This area must be large enough to accommodate the data protection information for the current transaction for all ET logic users. Details can be found in *Adabas Restart and Recovery*.

The LP parameter is no longer limited to 65535 blocks; however, it is still limited by the size of the Work data set and the sizes of Work parts 2, 3, and 4.

Note:

For any pending autorestart processing, Adabas retains and uses the value of the LP parameter from the previous session. If the LP parameter for the new session is specified with a different value than in the previous session, the nucleus will process any autorestart using the old value of LP, and establish the new LP size at the end of the autorestart.

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

The following example uses 2,000 blocks to store data protection information on the Adabas Work data set.

```
ADARUN PROG=ADANUC,LP=2000
```

LQ Parameter: Length of Sequential Command Table

Parameter	Specify . . .	Minimum	Maximum	Default
LQ	the size of the table of sequential commands.	2000	16777215	10 000

The table of sequential commands contains internal entries required during the processing of the Adabas read sequential (L2/L5, L3/L6, and L9) commands.

Each sequential pass of an Adabas file requires one entry in this table. Each table entry requires approximately 60-100 bytes, depending on the command type. Additionally, a variable portion (descriptor value) is required in the table entry for the L3/L6/L9 commands. An entry is deleted when an end-of-file condition is detected during sequential processing or when an RC or CL command is issued.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session. If using expanded files, the LQ value may need to be set to a higher-than-normal value.

The LQ value also determines the maximum number of command IDs that each user can specify. See the NQCID parameter description, later in this chapter, for more information.

In environments running in 31-bit addressing mode, the LQ pool is allocated above the 16-megabyte line.

Example

The following example allocates 30,000 bytes for the table of sequential commands for this session.

```
ADARUN PROG=ADANUC,LQ=30000
```

LRDP Parameter: Length of the Redo Pool

Parameter	Specify . . .	Minimum	Maximum	Default
<u>LRDP</u>	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.

LRPL Parameter: Adabas or Event Replicator Replication Pool Size

Parameter	Specify . . .	Possible Values	Default
LRPL	The size of the Adabas or Event Replicator replication pools.	minimum 20,000	100,000 (if REPLICATION=YES) or 0 (if REPLICATION=NO)

Notes:

1. The LRPL parameter should be specified for an Adabas nucleus running with ADARUN parameter REPLICATION=YES.
2. The LRPL parameter should be specified for an Event Replicator Server.
3. The LRPL parameter values in the Adabas nucleus and the Event Replicator Server may be different.
4. Software AG recommends setting LRPL to a relatively large value (e.g. LRPL=40M or greater). To determine what this setting should be for your site, you could run Adabas in a test environment with replication turned on, simulate the activity in the production environment on this test system, and use the Adabas DRES operator command to view the current and high-water mark of replication pool usage.

Example:

```
ADARUN PROG=ADANUC,REPLICATION=YES,LRPL=50000K
```

LS Parameter: Length of Internal Sort Area

Parameter	Specify . . .	Minimum	Maximum	Default
LS	the maximum number of bytes to be used for internal sort processing.	19968	see note	49920

Note:

To ensure that a single search command does not use too much of the work pool space, the maximum LS value must be less than or equal to the LWP (work pool length) divided by two minus the minimum LS value (19968) ($LS = LWP / 2 - 19968$). See the *LWP parameter*.

The LS parameter specifies the maximum number of bytes within the Work pool to be used for sort work space during the processing of a single command of the following type:

- S2 and S9 commands;
- complex search commands, which require the storage of bit strings.

The following formula may be used to estimate the number of Adabas ISNs that may be sorted in the work space provided (S2 or S9 commands):

If the TOPISN in the file is less than 65536:

Using	Formula to determine the number of Adabas ISNs that may be sorted in the work space provided
using 1 descriptor	LS/12
using 2 descriptors	LS/16
using 3 descriptors	LS/20

Notes:

1. TOPISN is the highest ISN currently in use for the file, less (MINISN - 1). TOPISN is not controlled by a parameter.
2. 4-byte ISNs are used for the internal sort area, regardless of the ISNSIZE specified for the file.

If the TOPISN in the file is greater than 65536:

Using	Formula to determine the number of Adabas ISNs that may be sorted in the work space provided
using 1 descriptor	LS/14
using 2 descriptors	LS/20
using 3 descriptors	LS/26

For complex searches, it is desirable for performance reasons to set LS value to the next integer multiple of 256 bytes greater than TOPISN / 8. For example, if TOPISN = 480000, then TOPISN / 8 = 60000, so LS should be set to the next higher multiple of 256, namely 60160.

Example

In the following example, the maximum amount of internal sort space during this session is set at 60,160 bytes.

```
ADARUN PROG=ADANUC,LS=60160
```

LU Parameter: Length of Intermediate User Buffer Area

Parameter	Specify . . .	Minimum	Maximum	Default
LU	the size of the intermediate user buffer area.	none	none (see note 1)	65535 (see note 2)

The LU parameter is optional. The size specified must be large enough to accommodate all Adabas control information (204 bytes), all user buffers (format, record, search, value, ISN) that may be required for any given Adabas command, plus any user information passed from Adabas link routines to nucleus user exits. If the ADALNK has user exits, the addresses of the ACB and each buffer passed will be stored in this space. Also, if the ADALNK has declared LUSIZE and has user exits, this buffer size plus 64 bytes

will also be stored in this space.

If the multifetch/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.

Notes:

1. An error occurs if the LU parameter specifies a value greater than the byte count implied by the NAB (number of attached buffers) parameter. On z/OS and VM systems, LU cannot exceed a value greater than that produced by the following calculation: $(NABvalue \times 4096)$; on z/VSE and BS2000 systems, LU cannot exceed a value greater than that produced by the following calculation: $(NABvalue \times 4096) - 256$. For more information about the NAB parameter, read *NAB: Number of Attached Buffers*.
2. Due to the length of the record buffer of the utilities that need the nucleus, e.g., ADAULD, the default value is set to 65,535. If the value of LU is less than 65,535 for an Adabas session, a response code will occur when such a utility is running.

The LU parameter syntax is:

```
LU={ n | 65535 }
```

Specific Product Recommendations

- For Event Replicator Server databases running with Adabas 8, the LU parameter must be greater than or equal to 167,000. In addition, if data is sent through Entire Net-Work from one or more Adabas nuclei to an Event Replicator Server, the Entire Net-Work LU parameter must be greater than or equal to the LU parameter setting for the Event Replicator Server itself.
- If data is sent through Entire Net-Work from one or more Adabas nuclei to an Event Replicator Server, the Entire Net-Work LU parameter setting must be greater than or equal to 164,000.

Example

The following example runs the Adabas nucleus with an Adabas intermediate user buffer area of 20,000 bytes.

```
ADARUN PROG=ADANUC,LU=20000
```

LWKP2 Parameter: Length of ISN List Processing Area (Work Part 2)

Parameter	Specify . . .	Minimum	Maximum	Default
LWKP2	the number of Work blocks to use for intermediate ISN lists.	0	16777215	0

This parameter specifies the size (in blocks) of the Adabas Work part 2, which holds intermediate ISN lists during Sx commands.

If LWKP2=0 (the default), Adabas calculates the Work part 2 size automatically using the following formula:

$$\text{WORK2} = 22 + 2 * (4 * \text{RECORDS} / (\text{BLKSIZE} - 16))$$

where

WORK2 is the Work part 2 space requirement in blocks.

RECORDS is the number of records in the file with the most records. This number equals $\text{TOPISN} - \text{MINISN} + 1$, where **TOPISN** is the highest ISN currently used in the file. **MINISN** is the lowest ISN used in the file. The **MINISN** value is specified with the ADACMP/ADALOD parameter **MINISN**; 1 is the default. You can use the ADAREP utility to display the **TOPISN** and **MINISN** values for the files in a database.

BLKSIZE is the block size of the device where the Work data set is stored (see the Adabas Installation documentation).

If you choose not to use LWKP2=0, the minimum value is LWKP2=20 blocks.

Note:

An Adabas internal table requires one byte of storage for each Work part 2 block.

Example:

For example, for the following Adabas session, the size of the Adabas Work part 2 is 50,000 blocks.

```
ADARUN PROG=ADANUC,LWKP2=50000
```

LWP Parameter: Length of Adabas Work Pool

Parameter	Specify . . .	Minimum	Maximum	Default
LWP	the size of the Adabas work pool.	80 000 but not less than NT * 25000	2147483647	150 000

Note:

The LWP (work pool length) parameter value must be equal to or greater than the LS (internal sort area length) value plus its minimum value (19968) multiplied by 2; that is, $(\text{LS} + 19968) * 2$.

The Adabas work pool area is used to store the following:

- descriptor value tables (DVT);
- intermediate user buffer areas (see parameter LU);
- sort work space areas (see parameter LS);
- work space during session and command execution.

In environments running in 31-bit addressing mode, the work pool space is allocated above the 16-megabyte line.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session. If expanded files are being used, the LWP value may need to be increased above the normal setting.

Specific Product Recommendations

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

The following example runs the Adabas nucleus with a work pool of 900,000 bytes.

```
ADARUN PROG=ADANUC,LWP=900000
```

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 or z/VM environments has no effect on the console messages.

Example:

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

MODE Parameter: Mode of Operation

Parameter	Specify . . .	Possible Values	Default
<u>MODE</u>	the Adabas operating mode.	SINGLE MULTI	MULTI

Value	Meaning
SINGLE	Run in single-user mode: Load the Adabas nucleus and the user program or Adabas utility into the same region.
MULTI	Run in multiuser mode: Load only the ADALNK routine into the same region as the program/utility. Load the Adabas nucleus into a different region. ADALNK will manage all interregion communication between the program or utility and the nucleus.

Note the following:

1. The MODE parameter does not affect parallel utility execution.
2. The nucleus cannot run in single-user mode (MODE=SINGLE cannot be specified) if it is to run
 - with Adabas Parallel Services or Adabas Cluster Services (NUCID 1)
 - with support for distributed transactions (DTP=RM)
 - as an Adabas Transaction Manager (DTP=TM)

In addition, a valid Adabas SVC should be specified when running in single-user (MODE=SINGLE) mode.

3. When running several nuclei against the same database, only one update nucleus can be used.

Example

The following example runs the ADAINV utility in single-user mode.

```
ADARUN PROG=ADAINV,MODE=SINGLE,SVC=234
```

MSGBUF Parameter: Size of Message Buffer

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MSGBUF</u>	the size of the message buffer area.	0	128	0

MSGBUF specifies the size (in kilobytes) of the area used to buffer messages for use for performance and tuning and problem analysis. Adabas messages buffered in this area can be reviewed even when online access to the console or DDPRINT messages is unavailable. While the buffer is intended for messages written to the system log or DDPRINT, it can be used for any purpose by any caller.

MSGBUF presupposes that ADARUN SMGT=YES.

MSGBUF=0 (the default) means that message buffering is not active for the session.

To activate message buffering, set MSGBUF to a value greater than 0. Software AG recommends setting MSGBUF to 10 (kilobytes).

Message buffering can be temporarily disabled and subsequently reactivated using the SMGT operator command.

See the *Adabas DBA Tasks Manual* for more information about message buffering.

Example

The following example activates message buffering when the Adabas nucleus is initialized with a message buffer size of 10 KB.

```
ADARUN PROGRAM=ADANUC
ADARUN SMGT=YES
ADARUN MSGBUF=10
```

MSGCONSL Parameter: Case of Messages Sent to the Console

MSGDRUCK Parameter: Case of Messages Sent to DD/DRUCK

MSGPRINT Parameter: Case of Messages Sent to DD/PRINT

Parameter	Specify the case of Adabas . . .	Possible Values	Default
<u>MSGCONSL</u>	messages sent to the console	<u>U</u> PPER <u>M</u> IXED	<u>M</u> IXED
<u>MSGDRUCK</u>	messages sent to the DD/DRUCK data set		
<u>MSGPRINT</u>	messages sent to the DD/PRINT data set		

Adabas messages are displayed in mixed upper and lowercase. You can use one or more of these ADARUN parameters to display different message types in all uppercase.

Notes:

1. The first message line in DDPRINT will not be influenced by the value of parameter MSGPRINT.
2. Lowercase fields will not display correctly (they will be converted to uppercase) if you use the ADARUN parameter settings MSGCONSL=U, MSGDRUCK=U, or MSGPRINT=U.

Example

The following example will display messages in the DDDRUCK and DDPRINT data sets in all uppercase:

```
ADARUN MSGDRUCK=U,MSGPRINT=U
```

MXCANCEL Parameter: Timeout Threshold for Canceled Peer Nucleus

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXCANCEL</u>	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
<code>MXCANCELWARN</code>	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	<code>MXCANCEL</code> - 1	<code>MXCANCEL</code> / 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
MXMSGWARN	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	MXMSG - 1	MXMSG / 4

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, 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
<u>MXSTATUS</u>	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.

MXTNA Parameter: Maximum Non-activity Time

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXTNA</u>	the maximum non-activity time limit that can be specified in the OP command for an individual user.	1	65535	3600

The TNAX parameter specifies the maximum elapsed time (in units of 1.048576 seconds) that users may be considered active without issuing an Adabas command. Limits may be set for three types of user:

- the TNAA parameter for access-only users
- the TNAE parameter for ET logic users

- the TNAX parameter for exclusive control users

You can override these parameter settings for an individual user by specifying a different value in the Adabas control block of an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

For these parameters, the MXTNA parameter sets a limit on the amount of time that can be specified in the OP command for an individual user.

If the nonactivity time limit is exceeded, action is taken as described in *Adabas Timeout Feature*.

The value of the MXTNA parameter can only be set when the nucleus is started. It cannot be changed using operator commands or the ADADBS OPERCOM utility function.

Example

For example, during the following Adabas session, the nonactivity time limit for exclusive control users is about 5 minutes (300 seconds). This limit cannot be extended for an individual user beyond 8 minutes (480 seconds).

```
ADARUN  PROG=ADANUC ,TNAX=300 ,MXTNA=480
```

MXTSX Parameter: Maximum Sx Command Execution Time

Parameter	Specify . . .	Minimum	Maximum	Default
MXTSX	the maximum execution time (in seconds) that can be specified for the ADARUN TLSCMD parameter and for any override execution time setting that might be specified in the Adabas control block of an OP command.	1	65535	3600

The TLSCMD parameter specifies the maximum time (in units of 1.048576 seconds) to be used to process a single Adabas S1, S2, or S4 command with a complex search criterion.

If the estimated time to process a command exceeds the TLSCMD time limit, the command is not processed and the user who issued the command receives a nonzero response code.

If the estimated time to process a command does not exceed the time limit, the command is processed to completion, whether the actual time required exceeds the time limit specified by the TLSCMD or not.

You can override the TLSCMD parameter setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

The MXTSX parameter sets the maximum time that can be specified for the TLSCMD parameter, including any override execution time setting specified in the Adabas control block of an OP command.

The value of the MXTSX parameter can only be set when the nucleus is started. It cannot be changed using operator commands or the ADADBS OPERCOM utility function.

Example

For example, during the following Adabas session, the Sx command execution time limit is about 1 hour 20 minutes (4800 seconds). This limit cannot be extended for an individual user beyond 2 hours (7200 seconds).

```
ADARUN  PROG=ADANUC,TLSCMD=4800,MXTSX=7200
```

MXTT Parameter: Maximum Transaction Time

Parameter	Specify . . .	Minimum	Maximum	Default
MXTT	the maximum transaction time limit that can be specified in the OP command for an individual user.	1	65535	3600

The TT parameter specifies the maximum elapsed time (in units of 1.048576 seconds) permitted for a logical transaction issued by an ET logic user. See ADARUN parameter TT for more information.

You can override this parameter setting for an individual user by specifying a different value in the Adabas control block of an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

The MXTT parameter sets a limit on the amount of time that can be specified in the OP command for an individual user.

The time measurement for a logical transaction starts when the first command that results in a record being placed in hold status is issued. It ends when an ET, BT, or CL command is issued.

If the transaction time limit is exceeded, action is taken as described in *Adabas Timeout Feature*.

The value specified for MXTT directly influences the required size of the Adabas parameter LP (length of the data protection area).

Software AG recommends that you set the MXTT parameter lower than the MXTNA parameter.

The value of the MXTT parameter can only be set when the nucleus is started. It cannot be changed using operator commands or the ADADBS OPERCOM utility function.

Example

For example, during the following Adabas session, the transaction time limit for ET logic users is about 3 minutes (180 seconds) and cannot be extended for an individual user beyond 5 minutes (300 seconds).

```
ADARUN  PROG=ADANUC,TT=180,MXTT=300
```

MXWTOR Parameter: Self-Termination Operator Query Interval

Parameter	Specify . . .	Minimum	Maximum	Default
<u>MXWTOR</u>	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.

NAB Parameter: Number of Attached Buffers

Parameter	Specify . . .	Minimum	Maximum	Default
<u>NAB</u>	the number of attached buffers to be used.	1	varies, depending on the amount of available virtual storage	16

The NAB parameter defines the number of attached buffers to be used during the session. An attached buffer is an internal buffer used for interregion communication. It is required in all environments. Adabas allocates an attached buffer pool with a size equal to the value of NAB multiplied by 4096 bytes.

You may specify as many attached buffers as fit into the available virtual storage.

In environments running in 31-bit addressing mode, the attached buffer pool space is allocated above the 16-MB line.

The NAB parameter syntax is:

```
NAB={ n | 16 }
```

Specific Product Recommendations

- For Event Replicator Server databases, set parameter NAB to a value greater than or equal to: 41 * 10 *
the-number-of-Adabas-nuclei-sending-data-to-the-Event-Replicator-Server.

For example, if one Adabas nucleus will be sending data to the Event Replicator Server, set the NAB parameter greater than or equal to 410 (for example NAB=420).

- If data is sent through Entire Net-Work from one or more Adabas nuclei to an Event Replicator Server, the Entire Net-Work NAB parameter must also be set to a value greater than or equal to: 41 * 10 *
the-number-of-Adabas-nuclei-sending-data-to-the-Event-Replicator-Server.
- 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.
- Users of the Adabas Review hub should read *Storage Requirements* in the *Adabas Review Concepts Manual* for more information about the space requirements of the Command Queue for Adabas Review.

Example

The following example runs the Adabas Review hub nucleus with 50 attached buffers.

```
ADARUN PROG=ADAREV,NAB=50
```

NC Parameter: Number of Command Queue Elements

Parameter	Specify . . .	Minimum	Maximum	Default
NC	the maximum number of command queue elements.	20	32767	200

The number of command queue elements (CQEs) established for the Adabas or Review hub session determines the maximum number of Adabas commands that may be queued or be in process at any one time during the session.

Each call from the Adabas nucleus is assigned a CQE. The CQE is released when the user receives the results of the command, the Adabas Review hub has processed the command, or the user has been timed out..

192 bytes are required for each CQE.

Software AG recommends that you set NC high enough to allow one command per active user for possible synchronization during execution of the online SAVE database function of the ADASAV utility.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

For more information about the space requirements of the Command Queue for Adabas Review, refer to *Storage Requirements* in the *Adabas Review Concepts Manual*.

Specific Product Recommendations

- For Event Replicator Server databases, set parameter NC to a value greater than or equal to: $10 * \text{the-number-of-Adabas-nuclei-sending-data-to-the-Event-Replicator-Server}$. For example, if one Adabas nucleus will be sending data to the Event Replicator Server, set the NC parameter greater than or equal to 10 (for example NC=20).
- If data is sent through Entire Net-Work from one or more Adabas nuclei to an Event Replicator Server, the Entire Net-Work NC parameter must also be set to a value greater than or equal to: $10 * \text{the-number-of-Adabas-nuclei-sending-data-to-the-Event-Replicator-Server}$.
- 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:

Run the Adabas nucleus with a maximum of 500 elements in the command queue.

```
ADARUN PROG=ADANUC,NC=500
```

The following example runs the Adabas Review hub nucleus with a maximum of 500 elements in the command queue.

```
ADARUN PROG=ADAREV,NC=500
```

NCLOG Parameter: Number of Command Log Data Sets

Parameter	Specify . . .	Minimum	Maximum	Default
NCLOG	the maximum number of command log data sets to be opened and logged in the PPT.	0, 2	8	0

This parameter specifies the maximum number of command log data sets to be opened and logged in the parallel participant table (PPT) for the current nucleus. Valid values are 0 (the default) and 2-8. Other values produce a parameter error. A zero value disables the parameter.

If more than the number of log data sets specified in the NCLOG parameter are allocated in the startup JCL, only the number specified in the NCLOG parameter are opened and logged; any additional allocated data sets are activated using the dynamic allocation process.

NCLOG must be specified with a nonzero value in conjunction with the CLOGDEV, CLOGSIZE, and UEX12 parameters.

NCLOG is not valid if DUALCLD, DUALCLS, DUALPLD, DUALPLS, or UEX2 are specified.

UEX2 and UEX12 are mutually exclusive.

If a nonzero NCLOG, CLOGDEV, and CLOGSIZE are specified, but LOGGING=NO, Adabas attempts to open and log the number of data sets specified by NCLOG; however, no logging occurs until the LOGGING parameter has been set to YES using either an operator command or Adabas Online System.

Example

The following example runs the Adabas nucleus with command logging, uses six command log data sets, and uses a model 3380 device for command logging. The size of each command log data set is 1350 blocks, based on an allocation of 10 cylinders. (Each cylinder of a 3380 device has 15 tracks, and each track can contain 9 blocks of command log data. Thus, 10 cylinders = 10 15 9 = 1350 blocks).

```
ADARUN  PROG=ADANUC ,LOGGING=YES ,NCLOG=6 ,CLOGDEV=3380 ,CLOGSIZE=1350
```

NH Parameter: Number of Hold Queue Elements

Parameter	Specify . . .	Minimum	Maximum	Default
NH	the maximum number of hold queue elements.	20	16777215	500

One hold queue element is required for each record (ISN) placed in hold status. A hold queue element is released when the record is released from hold status.

The hold queue is normally used to hold ISNs for the duration of a transaction. Therefore NH should be large enough to contain at least the average number of ISNs held per transaction multiplied by the maximum number of transactions running in parallel.

One hold queue element takes 52 bytes of storage. In environments running in 31-bit addressing mode, the hold queue is allocated above the 16-megabyte line.

Each shared hold request with a different command ID (CID), as well as a (shared or exclusive) hold request without a CID, is counted against the NH limit.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

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

The following example runs the Adabas nucleus with a maximum of 2000 elements in the hold queue.

```
ADARUN PROG=ADANUC,NH=2000
```

NISNHQ Parameter: Number of ISNs in Hold Queue per User

Parameter	Specify . . .	Minimum	Maximum	Default
NISNHQ	the maximum number of records that can be placed in hold status at the same time by one user.	1	(NH / 4) or 65535, whichever is smaller	(NH / NU) * 2 or 20, whichever is larger

NISNHQ determines the maximum number of held ISNs the hold queue (see the NH parameter description) can contain from one user. The most ISNs allowed on hold for a single user is one quarter (1/4) of the hold queue size, but not more than 65,535 ISNs. The default is the greater of either 20 or the number of ISNs the hold queue can contain divided by the number of users, multiplied by two.

If a user attempts to place more records in hold status than permitted, they receive a nonzero response code even though there may still be space in the hold queue.

Each shared hold request with a different command ID (CID), as well as a (shared or exclusive) hold request without a CID, is counted against the NISNHQ limit.

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.

Overriding the Parameter Setting

You can override this parameter setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

Examples

In the following example, the maximum number of records that may be in hold status for a single user is set at 100.

```
ADARUN PROG=ADANUC,NISNHQ=100,NH=450...
```

In the following example, NISNHQ is not specified. The maximum number of records that may be in hold status for a single user is 40 (the NISNHQ default of (NH / NU) * 2, which is greater than the alternative default of 20).

```
ADARUN PROG=ADANUC,NU=50,NH=2000...
```

NKSP Parameter: Span Large Sequential File Records on No-Keyed Disks

Parameter	Specify . . .	Possible Values	Default
<u>NKSP</u>	whether to span large sequential file records on no-keyed disk files to avoid the DMS0BBB record size error.	YES NO	NO

This parameter pertains to BS2000 environments only.

The DMS0BBB I/O error occurs when using utilities with large data block sizes such as BS2000 2008 devices with no-keyed disks. This is because the system requires 16 bytes for the record key with these disks. The maximum sequential record length is 32752 and not 32768. However, once written in this format, I/O interfaces other than the Adabas I/O interface cannot read the records correctly. To resolve this problem, run Adabas with the NKSP parameter set to "YES".

Example

In the following example, Adabas is set to write or read large sequential file records on no-keyed disks:

```
ADARUN PROGRAM=ADAORD,MODE=MULTI,NKSP=YES
```


NONDES Parameter: Nondesoriptor Searches

Parameter	Specify . . .	Possible Values	Default
NONDES	whether to allow (YES) or disallow (NO) searches using nondescriptors in search buffers.	YES NO	YES

By default, you can perform searches on Adabas fields that have not been defined as descriptors. However, nondesoriptor searches carry significant overhead.

NONDES=NO prevents searches on nondescriptors. An attempt to search on a nondesoriptor field produces response code 68 (ADARSP068).

Note:

The behavior of nondesoriptor searches in Adabas databases differs between mainframe and open systems in regards to null suppression in the fields. In open systems, nondesoriptor searches do not return records with null values in a field if the field is null-suppressed (NU); on mainframe systems, the null-suppression (NU) of fields is ignored during nondesoriptor searches. At this time, to resolve this problem, we recommend that you remove the null suppression option (NU) for open systems fields, if the fields must be used for a nondesoriptor search.

Example

The following example disallows searches using nondescriptors.

```
ADARUN  PROG=ADANUC, NONDES=NO
```

NPLOG Parameter: Number of Protection Log Data Sets

Parameter	Specify . . .	Minimum	Maximum	Default
NPLOG	the maximum number of protection log data sets to be opened and logged in the PPT.	0, 2	8	0

This parameter specifies the maximum number of protection log data sets to be opened and logged in the parallel participant table (PPT) for the current nucleus. Valid values are 0 (the default) and 2-8. Other values produce a parameter error. A zero value disables the parameter.

If more than the number of log data sets specified in the NPLOG parameter are allocated in the startup JCL, only the number specified in the NPLOG parameter are opened and logged; any additional allocated data sets are activated using the dynamic allocation process.

NPLOG must be specified with a nonzero value in conjunction with the PLOGDEV, PLOGSIZE, and UEX12 parameters.

NPLOG is not valid if DUALCLD, DUALCLS, DUALPLD, DUALPLS, or UEX2 are specified.

UEX2 and UEX12 are mutually exclusive.

Example

The following example runs the Adabas nucleus with six data protection log data sets and uses a model 3380 device. The size of each log data set is 1200 blocks, based on an allocation of 10 cylinders. (Each cylinder of a 3380 device has 15 tracks, and each track contains 8 blocks of protection log data. Thus, 10 cylinders = 10 * 15 * 8 = 1200 blocks).

```
ADARUN PROG=ADANUC, NPLOG=6, PLOGDEV=3380, PLOGSIZE=1200
```

NPLOGBUFFERS Parameter: PLOG Protection I/O Buffer Control

Parameter	Specify . . .	Minimum	Maximum	Default
NPLOGBUFFERS	the number of PLOG protection I/O buffers allocated in the Adabas nucleus session	1	2,147,483,687	1

PLOG protection I/O buffers are only allocated when ADARUN parameter LFIOP is set to a nonzero value. The realistic maximum setting of the NPLOGBUFFERS parameter is limited by the available amount of virtual and real memory. The number of protection I/O buffers allocated is limited to the number of PLOG blocks available on the dual/OCTO PLOG data sets (see the related ADARUN parameters DUALPLS and PLOGSIZE).

To determine the setting for NPLOGBUFFERS, look at the related high-water mark shown in the output of the DRES operator command or in the nucleus shutdown statistics. If the related high-water mark hits 100%, consider increasing the size of NPLOGBUFFERS

The use of the NPLOGBUFFERS parameter may considerably improve performance in update intensive applications that were formerly constrained by writes to the PLOG. When NPLOGBUFFERS is greater than 1, multiple PLOG blocks may be written in one write I/O to the PLOG.

In environments running in 31-bit addressing mode, the PLOG protection I/O buffers are allocated above the 16-megabyte line.

In environments supporting 64-bit virtual storage, the PLOG protection I/O buffers are allocated above the 2 gigabyte bar when the V64BIT ADARUN parameter is set to "YES".

Example

The following example sets the number of PLOG protection I/O buffers to 5.

```
ADARUN PROG=ADANUC, NPLOGBUFFERS=5
```

NQCID Parameter: Number of Active Command IDs per User

Parameter	Specify . . .	Minimum	Maximum	Default
NQCID	the maximum number of active command IDs per user.	1	65535	20

The table of sequential commands contains the control blocks for sequential read operations. Each operation is identified by a command ID (CID). NQCID defines the total number of CIDs allowed per user, and therefore the maximum number of table elements the user has available at any point in time.

Note:

The NQCID value specified does not ensure that all users have the specified count of CIDs available at all times. The actual number of CIDs available to a single user depends on the totals of table elements available and sequential operations active at a given time.

Overriding the Parameter Setting

You can override this parameter setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

Example

The following example sets the maximum number of parallel sequential processes per user to 25.

```
ADARUN PROG=ADANUC ,NQCID=25 ,NU=200 ,LQ=3000
```

Since one element in the sequential command table occupies approximately 100 bytes (varies depending on the command type; see the description of the LQ parameter), NQCID=25 limits each user to approximately 2500 bytes of the table.

NSISN Parameter: Number of ISNs per ISN Table Element

Parameter	Specify . . .	Minimum	Maximum	Default
NSISN	the maximum number of ISNs per element in the table of ISNs (TBI).	7	see text	51

The parameter NSISN specifies the maximum number of ISNs kept in a TBI element in memory (list of resulting ISN lists). The more ISNs that are kept in memory, the fewer the Work I/Os needed to read additional resulting ISN lists from the Work data set. This parameter influences the performance of the L1/4 command with the N (GET NEXT) option.

The highest value that can be specified for NSISN depends on the Work device. For Adabas-specified Work block sizes, the maximum values are shown in the device tables provided later in this section. However, the tables do not apply to *user-defined* Work block sizes; in this case, the maximum is

determined as follows, where WORK is the block size of the Work data set:

- $(\text{WORK} - 6) / 4$
- The TBI element has a length of $64 + (\text{NSISN} * 4)$ bytes in the TBI pool (see the LI parameter).

Although supported, the BS2000 device type 2007 is not recommended for use with Adabas. Support for the 2007 will be removed in a later Adabas release.

IBM and Compatible Devices

Device Type	Max. Value	Device Type	Max. Value
3310	1022	8345	5728
3330	1061	8350	2358
3340	877	8380	2267
3350	1155	8381	2867
3359	1155	8385	5865
3370	1278	8390	2674
3375	1022	8391	3418
3379	1022	8392	4611
3380	1371	8393	6995
3389	1371	9332	1278
3390	1429	9335	1278
		9345	2785

BS2000 Devices

Device Type	Max. Value	Device Type	Max. Value
2000	1022	2010	4096
2001	1022	2200	4096
2002	2046	2201	3072
2003	1534	2202	4096
2004	2558		
2005	2046		
2006	2558		
2007	7678		
2008	8188		
2009	8188		

Overriding the Parameter Setting

You can override this parameter setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

Example

In the following example, the maximum number of ISNs per TBI element during the Adabas session is set to 400.

```
ADARUN PROG=ADANUC,NSISN=400
```

NT Parameter: Number of Threads

Parameter	Specify . . .	Minimum	Maximum	Default
NT	the number of user threads to be used during the Adabas session.	4	250	5

A thread is used to execute a single command for any user (access or update). Each thread requires 8192 bytes.

In environments running in 31-bit addressing mode, the threads are allocated above the 16-megabyte line.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

For Event Replicator Server databases, this parameter must be greater than or equal to 15.

Specific Product Recommendations

- For Event Replicator Server databases, the NT parameter must be greater than or equal to 15.
- 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

The following example runs the Adabas nucleus with six (6) user threads.

```
ADARUN PROG=ADANUC,NT=6
```

NU Parameter: Number of User Queue Elements

Parameter	Specify . . .	Minimum	Maximum	Default
NU	the maximum number of user queue elements.	20	16777215	200

Adabas assigns a user queue element (UQE) to each active user when the user issues an OP command or (if the user did not issue an OP command) the first Adabas command.

Adabas releases the UQE when the user issues a CL command or when there is a timeout for an ET logic user who did not specify a restart ID in the control block (Additions 1 field) for the OP command.

Currently, $272 * (NU + 4)$ bytes are allocated to the user queue. The length of each UQE is 288 bytes.

In environments running in 31-bit addressing mode, the user queue is allocated above the 16-megabyte line.

Because each user's file list can be a different size, the number of users who can be active at the same time may be less than the NU value. Therefore, you may need to specify an NU value that is greater than the maximum number of users desired.

The Adabas session statistics or Adabas Online System can be used to tune this parameter for the next session.

Specific Product Recommendations

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

The following example runs the Adabas nucleus with a maximum of 500 elements in the user queue.

```
ADARUN PROG=ADANUC,NU=500
```

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 65,001-65,479). 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, ...
```

NWORK1BUFFERS Parameter: Work Part 1 Protection I/O Buffer Control

Parameter	Specify . . .	Minimum	Maximum	Default
NWORK1BUFFERS	the number of Work part 1 protection I/O buffers allocated in the Adabas nucleus session.	1	2,147,483,687	1

Work part 1 protection I/O buffers are only allocated when ADARUN parameter LFIOP is set to a nonzero value. The realistic maximum setting of the NWORK1BUFFERS parameter is limited by the available amount of virtual and real memory. The number of protection I/O buffers allocated is limited to the number of Work part 1 blocks available (see the related ADARUN LP parameter).

To determine the setting for NWORK1BUFFERS, look at the related high-water mark shown in the output of the DRES operator command or in the nucleus shutdown statistics. If the related high-water mark hits 100%, consider increasing the size of NWORK1BUFFERS

The use of the NWORK1BUFFERS parameter may considerably improve performance in update intensive applications that were formerly constrained by writes to Work part 1. When NWORK1BUFFERS is greater than 1, multiple Work part 1 blocks may be written in one write I/O to the Work data set.

In environments running in 31-bit addressing mode, the Work part 1 protection I/O buffers are allocated above the 16-megabyte line.

In environments supporting 64-bit virtual storage, the Work part 1 protection I/O buffers are allocated above the 2 gigabyte bar when the V64BIT ADARUN parameter is set to "YES".

Example

The following example sets the number of Work part 1 protection I/O buffers to 5.

```
ADARUN PROG=ADANUC,NWORK1BUFFERS=5
```

OPENRQ Parameter: Open Command Required

Parameter	Specify . . .	Possible Values	Default
OPENRQ	whether an open command is required.	YES NO	YES

This parameter may be used to indicate whether an open (OP) command must be issued as the first command of a user session.

Value	Meaning
YES	The default setting. Any user attempting to issue commands to Adabas without previously issuing an OP command receives a response code.
NO	Users are not required to issue an OP as the first command of the session.

Example

The following example specifies that the first session command issued by users does not have to be an OP command.

```
ADARUN PROG=ADANUC,OPENRQ=NO
```

PAMREQS Parameter: Maximum PAM Blocks for One PAM Macro

This parameter is valid on BS2000 systems only.

Parameter	Specify . . .	Possible Values	Default
<u>PAMREQS</u>	the maximum number of PAM blocks that can be transferred in one PAM macro request. This parameter is valid on BS2000 systems only.	1 - 1024	80

Using this parameter, you can set a reduced default to avoid the receipt of I/O errors DMS09AC (too many simultaneous I/O requests on some machines).

Example

The following example specifies that the maximum number of PAM blocks that can be transferred in a PAM macro request is 50.

```
ADARUN PROG=ADANUC ,PAMREQS=50
```

PGFIX Parameter: EXCPVR Page Fixing Control

Parameter	Specify . . .	Possible Values	Default
<u>PGFIX</u>	whether pages containing I/O control blocks are released after I/O processing is completed or after the job has ended when using EXCPVR on z/OS systems.	YES NO	NO

This parameter is only valid for z/OS users using EXCPVR (for example, when running the Adabas nucleus or utility jobs APF-authorized). It allows you to control how page-fixing occurs when EXCPVR is in use. When PGFIX=NO is specified, pages containing the I/O control blocks are fixed only for the duration of the I/O processing. When PGFIX=YES is specified, pages containing the I/O control blocks are fixed for the duration of the job.

PLOGDEV Parameter: Multiple Protection Log Device Type

PLOGSIZE Parameter: Multiple Protection Log Size

Parameter	Specify . . .	Possible Values	Default
<u>PLOGDEV</u>	the device type to be used for multiple protection logs.	see text	value of DEVICE parameter
<u>PLOGSIZE</u>	the number of blocks available for each protection log data set.	Min 16, Max 16777215	none

The PLOGDEV and PLOGSIZE parameters specify the type of external storage device used for multiple data protection log data sets and the number of blocks available for each log data set. These parameters must be used *together*; if you specify one parameter, you must also specify the other parameter.

You must also specify the NPLOG parameter to use multiple protection logs. If you specify a user exit, it must be UEX12. See *Adabas Restart and Recovery*, for information about multiple data set protection logging.

The Adabas Installation documentation specifies the device types that Adabas supports.

All protection log data sets must be formatted using the ADAFRM utility before they are used. The appropriate JCL for the data sets is required.

Example

The following example runs the Adabas nucleus with six (6) data protection log data sets and uses a model 3380 device. The size of each log data set is 1200 blocks, based on an allocation of 10 cylinders. (Each cylinder of a 3380 device has 15 tracks, and each track contains 8 blocks of protection log data. Thus, 10 cylinders = $10 * 15 * 8 = 1200$ blocks).

```
ADARUN  PROG=ADANUC ,NPLOG=6 ,PLOGDEV=3380 ,PLOGSIZE=1200
```

PLOGRQ Parameter: Protection Log Required

Parameter	Specify . . .	Possible Values	Default
PLOGRQ	whether a protection log is required for this session.	YES NO SEL FORCE	YES

The valid values are as follows:

Value	Meaning
YES	The default setting. Any attempt to start an Adabas nucleus without a protection log causes the Adabas initialization to terminate with an error message.
NO	In this Adabas session: <ul style="list-style-type: none"> - multiple PLOG data sets are used if a PLOGSIZE parameter is provided; - dual PLOG data sets are used if a DUALPLS parameter is provided; - a SIBA data set is used if JCL for a SIBA data set is provided; - otherwise, no protection log data set is used.
SEL	Selective protection logging is switched on only when running an ADASAV SAVE (database) or SAVE FILE= ... utility operation with an active nucleus. Protection logging is switched off when the ADASAV SAVE operation completes successfully.
FORCE	Any attempt to start an Adabas nucleus without UEX2 or UEX12, or without a protection log causes the Adabas initialization to terminate with an error message. Additionally, an error occurs if the initialization process determines that different PLOG data sets have been specified for the session and the old PLOG data sets have not yet been copied.

Example

The following example indicates that a data protection log (SIBA or PLOG) is not mandatory during the Adabas session but will be used if provided.

```
ADARUN PROG=ADANUC ,PLOGRQ=NO
```

PREFETCH Parameter: Activate Prefetch Feature

Parameter	Specify . . .	Possible Values	Default
PREFETCH	whether to enable a prefetch option automatically.	YES OLD NO	NO

This parameter determines whether the Adabas prefetch feature is enabled automatically; if it is, execution is transparent to the user program.

There are two Adabas prefetch options: the multifetch option is supported across Adabas platforms; the prefetch option is available only on mainframe Adabas. For information about these two options, see the *Adabas Command Reference Guide*.

Records can be prefetched and stored in a prefetch buffer for the following commands:

- L1/L4 with I or N option (read by ISN, find with GETNEXT)
- L2/L5 (read physical)
- L3/L6 (read logical by descriptor)
- L9 (histogram)

Sequential commands of the types listed above can retrieve the next record from the prefetch buffer, which saves cross-region/cross-network calls to the nucleus.

The PREFETCH parameter has the following values:

Value	Meaning
YES	Adabas enables the command-level multifetch (M) option.
OLD	Adabas enables the command-level prefetch (P) option.
NO	Adabas does not enable the prefetch feature.

Note:

Specify PREFETCH=YES or PREFETCH=OLD only when MODE=MULTI (multiuser mode). Do not specify these settings when running Adabas utilities. Dynamic reversal of read direction for L3/L6/L9 is not supported with these settings.

Examples

The following example runs a user program using multifetch.

```
ADARUN PROG=USER, PREFETCH=YES
```

The following example runs a user program using prefetch.

```
ADARUN PROG=USER, PREFETCH=OLD
```

PREFICMD Parameter: Include Commands in Prefetch/Multifetch Processing

PREFIFIL Parameter: Include Files in Prefetch/Multifetch Processing

Parameter	Specify . . .	Possible Values	Default
PREFICMD	the commands to include in prefetch processing.	L1-L9	none
PREFIFIL	the files to include in prefetch processing.	1-5000	none

Together, PREFIFIL and PREFICMD specify one or more files to be included in prefetch/multifetch processing. PREFIFIL specifies one or more files by file number; PREFICMD specifies a command for each file number.

You can either include list of files using PREFIFIL/PREFICMD or exclude lists of files using the parameters PREFXFIL/PREFXCMD), depending on which list is shorter. You cannot both include and exclude files at the same time.

If prefetching/multifetching has been specified and

- files/commands have been neither included nor excluded, the prefetching/multifetching applies to all files by default.
- files/commands have been included using the ADARUN parameters PREFIFIL/PREFICMD, the prefetching/multifetching applies only to the files on the inclusion list.
- both PREFIFIL/PREFICMD and PREFXFIL/PREFXCMD are specified, Adabas recognizes the exclusion parameters PREFXFIL/PREFXCMD and ignores the inclusion parameters PREFIFIL/PREFICMD.

Multiple Inclusions

Specify multiple file numbers in multiple PREFIFIL statements or in a list, separated by commas. Specify multiple commands in multiple PREFICMD statements or in a list, separated by slashes (/). The maximum number of entries for each is 50.

If you use lists, they must have the same number of entries. There is a one-to-one correspondence, by position, between a file number in the PREFIFIL list and a command in the PREFICMD list.

Eligible Commands and Command Pairs

The following commands can be included in prefetch/multifetch processing for a file:

- L1/L4
- L2/L5
- L3/L6
- L9

If either command in one of the pairs above is specified in a PREFICMD statement, *both* commands are included. For example, if PREFICMD=L1 is specified, both L1 and L4 are included.

Examples

The following example includes the L3 command for file 8 in prefetch processing.

```
ADARUN PREFIFIL=8 ,PREFICMD=L3
```

The following example includes the L3 command for file 8 and the L3 and L9 commands for file 12 in prefetch/multifetch processing. As this example shows, you can specify the same file number several times, with a different command each time. Likewise, you can specify the same command code several times, with a different file number each time.

```
ADARUN PREFIFIL=8 ,12 ,12 ,PREFICMD=L3/L3/L9
```

or

```
ADARUN PREFIFIL=8 ,PREFICMD=L3
ADARUN PREFIFIL=12 ,PREFICMD=L3
ADARUN PREFIFIL=12 ,PREFICMD=L9
```

The following example executes a user program with multifetch and allocates 20,000 bytes for the total prefetch buffer area. It includes the L3 and L6 commands for files 8 and 12 in multifetch processing and it includes *all* commands for file 7. No PREFSBL value is specified, so the buffer length for a single prefetch sequence defaults to 3,000 bytes; up to 6 prefetch sequences can be handled.

```
ADARUN PROG=USER ,PREFETCH=YES ,PREFTBL=20000
ADARUN PREFIFIL=8 ,12 ,PREFICMD=L3/L3
ADARUN PREFIFIL=7 ,7 ,7 ,7 ,PREFICMD=L1/L2/L3/L9
```

PREFNREC Parameter: Multifetch Record Count

Parameter	Specify . . .	Minimum	Maximum	Default
PREFNREC	the maximum number of records that can be prefetched.	0	see text	0

This parameter specifies the maximum number of records that the multifetch feature (specified by the PREFETCH=YES parameter) obtains.

The maximum that can be specified depends on the maximum that can be held by the record and ISN buffers, each of which is assumed to be less than 32 kilobytes.

If no value or zero (the default) is specified, the sizes of the ISN and record buffers determine the number of records that are prefetched. Specifying zero (0) *does not* disable multifetch operation.

Example

The following example enables the multifetch feature and allows a maximum of 32 records to be prefetched.

```
ADARUN PROG=USER, PREFETCH=YES, PREFNREC=32
```

PREFSBL Parameter: Prefetch Buffer Length per Call

Parameter	Specify . . .	Minimum	Maximum	Default
PREFSBL	the prefetch/multifetch buffer length (in bytes) for one prefetch.	1024	<PREFTBL value>	3000

The PREFSBL parameter specifies the size of an internal buffer that the multifetch/prefetch feature uses to prefetch data for a single sequential read command. By setting of this buffer size you control the maximum amount record data that is prefetched. See also the description of the PREFTBL parameter, which specifies the total buffer space used by prefetch/multifetch.

When prefetch or multifetch is active, Adabas makes a buffer available to the user program for storing prefetched records on the client side. If a required record is already in a prefetch buffer, the record is taken directly from the buffer.

Note:

Do not specify this operand unless PREFETCH=YES or PREFETCH=OLD also is specified.

The following formula can be used to estimate the approximate maximum number of records (n) that will be prefetched:

$$600 + FBL + ((RBL * 3) * n) = PREFSBL$$

In this calculation, *FBL* is the length of the format buffer and *RBL* is the length of the record buffer.

When an application uses the ACB direct call interface, the 32K buffer size limit will limit the maximum number of prefetched records to:

$$MAX(n) = 32768 / RBL$$

For long sequential processes, this parameter should be as large as possible unless the sequential processes are frequently interrupted with value start (L3/L6) or RC commands; in that case, a large value often causes records to be prefetched needlessly.

Example

The following example executes a user program using multifetch and uses a 15,000-byte prefetch buffer.

```
ADARUN PROG=USER , PREFETCH=YES , PREFSBL=15000
```

PREFSTDD Parameter: Multifetch Statistics DD Name

Parameter	Specify . . .	Valid Values	Default
PREFSTDD	the job control statement label (DD name) in the JCL that defines the location to which multifetch statistics should be written when a CL command is issued.	DD name or other job statement label	none

The PREFSTDD parameter specifies the job control statement label (DD name) in the nucleus startup job that defines the location to which multifetch statistics should be written when multifetch statistics are printed. The length of the job statement label (DD name) is limited to seven (7) characters.

When PREFSTDD is specified, the ADARUN PREFETCH parameter must be set to "YES".

Note:

Any data set defined by the job statement label referenced by PREFSTDD must have an LRECL of at least 80 bytes and must be a fixed block data set. If it does not, the PREFSTDD setting is ignored and no statistics are printed.

Multifetch statistics are written when a CL command is issued. If no CL command is issued, no multifetch statistics are produced. If multiple CL commands are issued, multifetch statistics are produced up until the first CL command completes and then no more statistics are generated.

Example

The following example executes a user program using multifetch, uses a 15,000-byte prefetch buffer, and writes multifetch statistics to the output location identified by the STATDD job control statement:

```
ADARUN PROG=USER , PREFETCH=YES , PREFSBL=15000 , PREFSTDD=STATDD
```

The STATDD job control statement that accompanies these ADARUN parameters might simply indicate that multifetch statistics should be written to SYSOUT:

```
//STATDD DD SYSOUT=*
```

Alternatively, the STATDD job control statement that accompanies these ADARUN parameters might specify a specific data set:

```
//STATDD DD DISP=(NEW,CATLG),DSN=ADABAS.MLF.STATS,
// UNIT=SYSDA,SPACE=(TRK,(1,1),RLSE),
// DCB=(RECFM=FB,LRECL=121,BLKSIZE=1210)
```


On BS2000, the STATDD must be a file. This would be defined in the job as follows:

```
/DELETE-FILE ADABAS.MLF.STATS
/SET-JOB-STEP
.
.
/SET-FILE-LINK STATDD,ADABAS.MLF.STATS,REC-FORM=F(REC-SIZE=80)
```

The following is an example of the multifetch statistics output produced:

```
ADAML1 +-----+
ADAML1 +
ADAML1 + A D A M L F   S T A T I S T I C S
ADAML1 +
ADAML1 + Percentage of PREFTBL used           :           50.00
ADAML1 + Number of calls to ADAMLF           :           1918
ADAML1 + Number of cmds to ADABAS           :             56
ADAML1 + Number of cmds returned from ADAMLF :           1875
ADAML1 + Number of over reads                :             2
ADAML1 +
ADAML1 +-----+
```

The following table describes these statistics:

Statistic	Meaning
Percentage of PREFTBL used	The percentage of the ADARUN parameter PREFTBL used for this execution.
Number of calls to ADAMLF	The number of times that ADAMLF was entered from the client program.
Number of cmds to ADABAS	The total number of commands that were passed to Adabas for this execution.
Number of cmds returned from ADAMLF	The number of commands that were returned from the ADAMLF buffers.
Number of over reads	The number of records that were prefetched that were not passed back to the caller.

PREFTBL Parameter: Total Size of Prefetch Buffer

Parameter	Specify . . .	Minimum	Maximum	Default
PREFTBL	the size of the total prefetch buffer area.	6144	16777215	30000

PREFTBL specifies the size of the total prefetch buffer area for this operation. This buffer length should be large enough to hold the PREFSBL-specified value, multiplied by the number of command-ID-controlled sequential prefetches for this operation. This allows a user program to perform multiple sequential operations requiring prefetches, each with a prefetch buffer of the size specified in the PREFSBL parameter.

If this user program has a single prefetch sequence only, the PREFTBL value must be equal to or greater than the PREFSBL value. The PREFTBL-specified space is always allocated to the program.

Note:

Do not specify this operand unless PREFETCH=YES or PREFETCH=OLD is also specified.

Example

The following example executes a user program using multifetch. No PREFSBL value is specified, so this example uses the default buffer length for a single prefetch (3,000 bytes). It also allocates 30,000 bytes for the total prefetch buffer area; this size provides buffer space for up to 10 prefetch sequences.

```
ADARUN PROG=USER , PREFETCH=YES , PREFTBL=30000
```

PREFXCMD (PREFXFIL) Parameters: Exclude Commands (Files) from Prefetch/Multifetch Processing

Parameter	Specify . . .	Possible Values	Default
PREFXCMD	the commands to exclude from prefetch processing.	L1-L9	none
PREFXFIL	the files to exclude from prefetch processing.	1-5000	none

Together, PREFXFIL and PREFXCMD specify one or more files to be excluded from prefetch/multifetch processing. PREFXFIL specifies one or more files by file number; PREFXCMD specifies a command for each file number.

You can either include list of files using PREFIFIL/PREFICMD) or exclude lists of files using the parameters PREFXFIL/PREFXCMD, depending on which list is shorter. You cannot both include and exclude files at the same time.

If prefetching/multifetching has been specified and

- files/commands have been neither included nor excluded, the prefetching/multifetching applies to all files by default.
- files/commands have been excluded using PREFXFIL/PREFXCMD, the prefetching/ multifetching applies only to the files on the exclusion list.
- both PREFIFIL/PREFICMD and PREFXFIL/PREFXCMD are specified, Adabas recognizes the exclusion parameters PREFXFIL/PREFXCMD and ignores the inclusion parameters PREFIFIL/PREFICMD.

Multiple Exclusions

Specify multiple file numbers in multiple PREFXFIL statements or in a list, separated by commas. Specify multiple commands in multiple PREFXCMD statements or in a list, separated by slashes (/). The maximum number of entries for each is 50.

If you use lists, they must have the same number of entries. There is a one-to-one correspondence, by position, between a file number in the PREFXFIL list and a command in the PREFXCMD list.

Eligible Commands and Command Pairs

The following commands can be excluded from prefetch/multifetch processing for a file:

- L1/L4
- L2/L5
- L3/L6
- L9

If either command in one of the pairs above is specified in a PREFXCMD statement, *both* commands are excluded. For example, if PREFXCMD=L1 is specified, both L1 and L4 are excluded.

Examples

The following example excludes the L3 command for file 8 from prefetch processing.

```
ADARUN PREFXFIL=8 , PREFXCMD=L3
```

The following example excludes the L3 command for file 8 and the L3 and L9 commands for file 12 from prefetch/multifetch processing. As this example shows, you can specify the same file number several times, with a different command each time. Likewise, you can specify the same command code several times, with a different file number each time.

```
ADARUN PREFXFIL=8 , 12 , 12 , PREFXCMD=L3 / L3 / L9
```

or

```
ADARUN PREFXFIL=8 , PREFXCMD=L3
ADARUN PREFXFIL=12 , PREFXCMD=L3
ADARUN PREFXFIL=12 , PREFXCMD=L9
```

The following example executes a user program with multifetch and allocates 20,000 bytes for the total prefetch buffer area. It excludes the L3 and L6 commands for files 8 and 12 from multifetch processing and it excludes *all* commands for file 7. No PREFSBL value is specified, so the buffer length for a single prefetch sequence defaults to 3,000 bytes; up to 6 prefetch sequences can be handled.

```
ADARUN PROG=USER,PREFETCH=YES,PREFETBL=20000
ADARUN PREFXFIL=8,12,PREFXCMD=L3/L3
ADARUN PREFXFIL=7,7,7,7,PREFXCMD=L1/L2/L3/L9
```

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

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.

```
ADARUN PROGRAM=NETWRK
```

QBLKSIZE Parameter: Block Size for Sequential Data Sets

Parameter	Specify . . .	Minimum	Maximum	Default
QBLKSIZE	the block size for sequential data sets used by ADAIOR.	0	262144	0 (see text)

QBLKSIZE specifies the block size, in bytes, for the *sequential data sets* used by ADAIOR, the Adabas I/O component. Use this parameter to override the default block sizes. For more information about sequential files used by Adabas utilities and their operating system dependencies, read *Adabas Sequential Files*.



Warning:

If you choose to write tape files with large block sizes, for example for database backups, these files will not be transportable to systems where support for large blocks is not available. This might include a site being used as a backup facility for disaster recovery.

If QBLKSIZE=0 (the default), the following default block sizes are used:

Data Set Block Size

Tape the optimum selected for the device by the operating system

CKD-Disk track-size / k, where k = smallest value such that track-size / k <= 32760

FBA-Disk 32760

Under all operating systems except z/VSE, this block size may be provided in the JCL or data set label.

Under BS2000, disk block sizes are always rounded up to the next integer multiple of 2048 bytes.

Examples

For example, a block size of 65,536 bytes for the Adabas sequential data sets is used by the ADAULD utility:

```
ADARUN PROG=ADAULD,QBLKSIZE=65536
```

In the following example, a block size of 32,760 bytes for the Adabas sequential data sets is used by the ADAORD utility:

```
ADARUN PROG=ADAORD,QBLKSIZE=32760
```

READONLY Parameter: Read-Only Session Control

Parameter	Specify . . .	Possible Values	Default
<u>READONLY</u>	whether to permit database updates during this session.	YES NO	NO

READONLY determines whether the Adabas session is a read-only or a read/write session.

Value	Meaning
YES	All database update commands issued by this nucleus are rejected.
NO	The default setting. The session allows both reading and updating of the database.

Guidelines

- When executing the ADAREP (database report) utility, you can specify READONLY=YES to issue an access OPEN instead of an update OPEN.
- When using READONLY=YES, ensure that the Work data set does not contain a pending autorestart; otherwise, the nucleus cannot start.
- Each nucleus being run must have its own Work data set.
- When a nucleus is started with READONLY=YES, it is *not* possible to switch the nucleus to READONLY=NO mode later in the session. When a nucleus is started with READONLY=NO, it *can* be switched back and forth between the READONLY=YES and READONLY=NO modes.
- A database save performed while the nucleus is in read-only mode is equivalent to an offline save (which is taken while the nucleus is down). You can avoid shutting down the nucleus for an offline save by switching the nucleus into read-only mode before the save and back to update mode after the save.

Example

The following example starts an Adabas session as read-only.

```
ADARUN PROG=ADANUC,READONLY=YES
```

REPLICATION Parameter: Replication Control

Parameter	Specify . . .	Possible Values	Default
REPLICATION	Whether to activate replication for the Adabas nucleus.	YES NO	NO

Valid values are described in the following table.

Value	Meaning
YES	Activate replication for the Adabas nucleus.
NO	The default setting. Do not activate replication for the Adabas nucleus.

When REPLICATION=NO is specified and there are replicated files on the database, replication will be turned off for all replicated files. A message is issued during nucleus initialization indicating which files were defined with replication.

The REPLICATION parameter should not be specified for the Event Replicator Server. REPLICATION=NO must be explicitly specified if:

- You want to start the Adabas nucleus without replication active, *and*
- REPLICATION=YES was specified in the previous Adabas nucleus session, *and*
- Replication has been defined for one or more files on the database.

REPLICATION=NO must be explicitly specified for utility execution if:

- A utility requiring an active nucleus is executed with MODE=SINGLE, *and*
- REPLICATION=YES was specified in the previous Adabas nucleus session, *and*
- Replication has been defined for one or more files on the database.

Example

The following example activates replication for an Adabas session.

```
ADARUN PROG=ADANUC,REPLICATION=YES
```

REVFILTER Parameter: Review Record Filtering Control

Parameter	Specify . . .	Possible Values	Default
REVFILTER	whether to allow Adabas Review record filtering during the session.	YES NO	YES

REVFILTER determines whether command log record filtering may be activated. Filtering can decrease the number of command log records passed to Review for report processing.

Value	Meaning
YES	The default setting. Database command log records may be filtered from report processing, depending upon Review report rules.
NO	All command log records will be passed to Review for report processing.

Examples

In the following example, Adabas Review's record filtering may be activated during the Adabas nucleus session.

```
ADARUN PROG=ADANUC,REVFILTER=YES
```

In the following example, Adabas Review's record filtering will not be in effect for the Adabas nucleus session.

```
ADARUN PROG=ADANUC,REVFILTER=NO
```

REVIEW Parameter: Adabas Review Control

Parameter	Specify . . .	Possible Values	Default
REVIEW	whether to run Adabas Review in local or hub mode specifying the hub ID, or not at all.	NO LOCAL dbid	NO

Note:

The parameter name REVIEWHUBID is a synonym for REVIEW, provided to ensure downward compatibility with past Adabas releases. We recommend that you use the parameter name REVIEW

instead, wherever possible.

REVIEW controls the use of the Adabas Review product:

Value	Meaning
NO	The default setting. Adabas Review is not started.
<u>LOCAL</u>	<p>Adabas Review is started in local mode running as an extension to ADALOG.</p> <p>In local mode, Adabas Review job control statements should be added to the Adabas nucleus startup JCL.</p> <p>Note: If an Adabas Review load library is not included in the startup JCL, the REVIEW parameter is automatically changed from LOCAL to NO.</p>
dbid	<p>Adabas Review is started in hub mode. The physical database ID that you specify for the hub identifies</p> <ul style="list-style-type: none"> ● the hub (server) itself (with PROGRAM=ADAREV) that is being started; or ● from an Adabas nucleus (client), the hub that is the target for Adabas Review processing for that nucleus (with PROGRAM=ADANUC). <p>In hub mode, Adabas Review job control statements should be added to the Adabas Review hub startup JCL.</p>

Dynamic Modification

The setting of the ADARUN PROG=ADANUC,REVIEW=dbid parameter can be changed dynamically using the REVIEWHUBID command from the operator console, the ADADBS OPERCOM REVIEWHUBID function, or the Modify Parameter function of Adabas Online System.

Examples

The following example starts hub 202 for the Adabas Review hub (server) installation.

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

The following example starts the Adabas nucleus that will log to Adabas Review hub 202 for the Adabas Review (client) installation.

```
ADARUN PROGRAM=ADANUC,REVIEW=202
```

REVLOGBMAX Parameter: Logged Buffer Size Limit for Review

Parameter	Use	Values	Default
REVLOGBMAX	Specify the maximum allowable number of bytes of a logged buffer for Review.	Any integer ranging from 0 to 30000	5120

If a buffer is longer than this value, the logged buffer is truncated from the point at which its size exceeds the setting of the REVLOGBMAX parameter. The REVLOGBMAX setting affects the ADARUN LOGGING parameter specifications only for both CLOGLAYOUT=8. The minimum value (368) is the size of the length of the CLOGLAYOUT=8 basic record plus the length of the extended Adabas control block (ACBX).

Example

The following example runs the Adabas nucleus using a logged buffer size limit of 512. Individual logged buffers, such as the format buffer for an Adabas command, will be truncated if they exceed 512 bytes.

```
ADARUN PROG=ADAREV,LOGBMAX=512
```

REVLOGMAX Parameter: Total Logged Buffer Size Limit for a Review Command

Parameter	Use	Values	Default
REVLOGMAX	Specify the maximum size of all of the logged buffers allowed for an Adabas Review command.	Any integer ranging from 2000 to 32768 (32K).	16384

When the sum of sizes of the logged buffers for an Adabas Review command reaches the value of the REVLOGMAX parameter, the buffer exceeding the limit is truncated and all following buffers are omitted. The size of REVLOGMAX must at least be as large as the REVLOGBMAX + 2000 in order to accommodate the largest buffer of the Adabas command. The minimum value (2000) is the size of the length of the CLOGLAYOUT=8 basic record plus the length of the extended Adabas control block (ACBX) and the CLEX information.

The CLOGMAX setting affects the ADARUN LOGGING parameter specifications only for CLOGLAYOUT=8.

Example

The following example runs the Adabas nucleus using a command logged buffer size limit of 32000 bytes. The sum of all logged buffers for an Adabas Review command to this nucleus cannot exceed 10000 bytes.

```
ADARUN PROG=ADAREV,REVLOGMAX=10000
```

RPLCONNECTCOUNT Parameter: Replication Connection Attempt Count

Parameter	Specify . . .	Possible Values	Default
RPLCONNECTCOUNT	the number of connection attempts made for the Adabas or Event Replicator Server nucleus after an attempt fails (response 148, ADARSP148, is issued).	0 - 2,147,483,647	0

A value of zero indicates that no connection attempts should occur; a value of zero makes the most sense in situations where the Adabas database and the Event Replicator Server execute together on the same logical partition (LPAR). If the Adabas database and the Event Replicator Server execute on different LPARs, however, setting a real value using this command helps avoid errors that might arise if network problems occur because the network is not started or a network connection between the Adabas database and the Event Replicator Server is lost.

RPLCONNECTINTERVAL Parameter: Replication Connection Attempt Interval

Parameter	Specify . . .	Possible Values	Default
RPLCONNECTINTERVAL	the interval (in seconds) between connection attempts made for the Adabas or Event Replicator Server nucleus after an attempt fails (response 148, ADARSP148, is issued).	0 - 2,147,483,647	0

A value of zero indicates that no connection attempts should occur; a value of zero makes the most sense in situations where the Adabas database and the Event Replicator Server execute together on the same logical partition (LPAR). If the Adabas database and the Event Replicator Server execute on different LPARs, however, setting a real value using this command helps avoid errors that might arise if network problems occur because the network is not started or a network connection between the Adabas database and the Event Replicator Server is lost.

RPLPARMS Parameter: Replication Parameter Location

Parameter	Specify . . .	Possible Values	Default
RPLPARMS	Where your replication definitions (initialization parameters) should be read from.	BOTH FILE NONE PARMS	See below.

This parameter is only used when starting the Event Replicator database.

If the RPLPARMS parameter is not specified, the following logic is used to determine where to read the replication definitions:

- If the Replicator system file is loaded on the Event Replicator database, the replication definitions are read from that file.
- If the Replicator system file is not loaded on the Event Replicator database, the replication definitions are read from DDKARTE.

Valid values are described in the following table.

Value	Meaning
BOTH	The replication definitions are read first from the Replicator system file and secondly from DDKARTE. Separate entries are built for each and errors are generated if there are duplicate entries.
FILE	The replication definitions are read from the Replicator system file, which is maintained using the Adabas Event Replicator Subsystem in Adabas Online System (AOS).
NONE	No replication definitions are read. The Event Replicator Server will establish itself with a null set of replication definitions and will not communicate or handshake with related Adabas nuclei. This mode allows maintenance to occur to the Replicator system file, without the Event Replicator Server communicating with related Adabas nuclei.
PARMS	The replication definitions are read from DDKARTE.

RPLSORT Parameter: Event Replicator for Adabas Transaction Data Sorting Control

Parameter	Specify . . .	Possible Values	Default
RPLSORT	Whether Event Replicator for Adabas transaction data sorting should occur or not.	YES NO	YES

This parameter is only used for the Adabas database from which transactions will be replicated. It should not be used for the Event Replicator Server database.

By default (RPLSORT=YES), Event Replicator for Adabas sorts transaction data by a key consisting of the file number, ISN, and relative update number. If you do not want this sort to occur, specify RPLSORT=NO.

It is important to note that the RPLSORT setting will affect how the Event Replicator consolidates transaction modifications and the order in which they are replicated. When RPLSORT=YES, all modifications for a given ISN in a file are consolidated before replication occurs. For complete information on regular modification consolidation, read *Detailed Adabas Nucleus Processing in the Event Replicator for Adabas Concepts*. When RPLSORT=NO is specified, the modifications are still consolidated, but their referential integrity is maintained. In other words, the chronological order of the updates is maintained for each ISN in a file.

Example

For example, assume the following modifications occur, in the following order:

```
Insert File 77 ISN 80 Seq 1
Update File 77 ISN 80 Seq 2
Insert File 77 ISN 1010 Seq 3
Update File 77 ISN 80 Seq 4
Update File 77 ISN 1010 Seq 5
```

If RPLSORT=YES, these modifications are consolidated and replicated as follows:

```
Insert File 77 ISN 80 Seq 1 + Seq 2 + Seq 4
Insert File 77 ISN 1010 Seq 3 + Seq 5
```

If RPLSORT=NO, these modifications are consolidated and replicated as follows:

```
Insert File 77 ISN 80 Seq 1 + Seq 2
Insert File 77 ISN 1010 Seq 3
Update File 77 ISN 80 Seq 4
Update File 77 ISN 1010 Seq 5
```

The end result is the same, regardless of whether RPLSORT is set to YES or NO, but multiple modifications to the same ISN in a file may be necessary when RPLSORT is set to NO.

RPWARNINCREMENT Parameter: Replication Pool Warning Message Interval

Parameter	Specify . . .	Possible Values	Default
RPWARNINCREMENT	The interval (in percent of LRPL usage) at which warning messages should be sent that indicate that the replication pool usage has exceeded the threshold set by the RPWARNPERCENT parameter.	1-99	10

This ADARUN parameter applies to both the Adabas and Event Replicator Server replication pools.

RPWARNINTERVAL Parameter: Replication Pool Usage Interval

Parameter	Specify . . .	Possible Values	Default
RPWARNINTERVAL	The interval, in seconds, during which replication pool usage warning messages are suppressed on the console.	1-2147483647	60

This ADARUN parameter applies to both the Adabas and Event Replicator Server replication pools.

When the number of replication pool usage warning messages sent to the console exceeds the limit specified by the RPWARNMESSAGELIMIT parameter, all such messages are suppressed for the interval of time specified by this parameter. When the interval specified by the RPWARNINTERVAL parameter has passed, warning messages resume.

This parameter, along with the RPWARNMESSAGELIMIT parameter, is provided to avoid flooding the console with replication pool usage warning messages.

RPWARNMESSAGELIMIT Parameter: Replication Pool Usage Warning Message Count

Parameter	Specify . . .	Possible Values	Default
RPWARNMESSAGELIMIT	The number of replication pool usage warning messages that can be issued to the console before message suppression begins.	1-2147483647	5

This ADARUN parameter applies to both the Adabas and Event Replicator Server replication pools.

When the threshold specified by this parameter is reached, all replication pool usage warning messages are suppressed until the message interval specified by the RPWARNINTERVAL parameter has passed.

This parameter, along with the RPWARNINTERVAL parameter, is provided to avoid flooding the console with replication pool usage warning messages.

RPWARNPERCENT Parameter: Replication Pool Usage Threshold

Parameter	Specify . . .	Possible Values	Default
RPWARNPERCENT	The threshold for replication pool usage (expressed as a percentage of LRPL) at which warning messages should be sent. If "0" is specified or taken as the default there are no warning messages	0-99	0

This ADARUN parameter applies to both the Adabas and Event Replicator Server replication pools.

The usage level should be expressed as a percentage. Warning messages will be printed when usage of the replication pool exceeds this percentage (threshold).

RVCLIENT Parameter: Adabas Review Client Reporting Activation

Parameter	Specify . . .	Possible Values	Default
RVclient	whether Adabas Review client reporting should be activated when you want to run client reports in batch environments.	ACTIVE INACTIVE	INACTIVE

This ADARUN parameter allows you to activate Adabas Review client reporting when you want to run client reports in batch environments. Specify "ACTIVE" to activate it; specify "INACTIVE" (or specify no setting, since INACTIVE is the default) if you do not want client reporting activated.

Note:

This ADARUN parameter is valid only in z/OS environments.

Example

In the following example, client reporting is activated.

```
ADARUN PROGRAM=USER,RVCLIENT=ACTIVE
```

SMF Parameter: Enable Adabas SMF Recording

Parameter	Specify . . .	Valid Values	Default
SMF	whether or not Adabas SMF recording should be enabled.	YES NO	NO

This parameter allows you to enable and disable Adabas SMF recording. SMF is available only in z/OS environments.

- NO is the default and disables all SMF processing for the duration of the nucleus.
- YES enables all SMF processing for the duration of the nucleus. Adabas SMF records will be generated and written when the SMF parameter is set to YES.

This parameter *cannot* be altered by operator commands while the nucleus is running. If you specify SMF=YES, note that APF authorization is required to write to system SMF files.

Example

In the following example, the Adabas SMF records are generated and written.

```
ADARUN PROG=ADANUC,SMF=YES
```

SMF89 Parameter: Enable SMF Type 89 Records

Parameter	Specify . . .	Valid Values	Default
SMF89	whether or not the Adabas nucleus should register with z/OS for type 89 SMF records.	YES NO	NO

This allows you to have the Adabas nucleus register with z/OS for type 89 SMF records. SMF89 is available only in z/OS environments.

- NO is the default and prevents registration.
- YES causes the nucleus to register with z/OS for type 89 SMF records..

This parameter *cannot* be altered by operator commands while the nucleus is running.

Example

In the following example, the nucleus will register with z/OS for type 89 SMF records.

```
ADARUN PROG=ADANUC,SMF89=YES
```


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, IODD, LOCK, MSGB, MSGC, MSGH, PARM, STG, THRD, 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.

Detail Section Name	Description
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.
PARM	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.
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)
```

SMFINTERVAL Parameter: Enabling Adabas SMF Interval Records

Parameter	Specify . . .	Valid Values	Default
SMFINTERVAL	whether SMF interval records should be generated and the interval at which they should be generated.	NONE GLOBAL SUBSYS minutes	NONE

Adabas SMF records can be generated at specific intervals. This parameter enables and disables interval recording and specifies the interval or the source from which the interval can be derived.

The interval can be one of several values from z/OS specifications or an explicit interval in minutes. Interval records contain a product detail section and whatever detail sections are selected in the SMFDETAIL parameter, except for the ADARUN parameter detail section which are not included in SMF interval records.

The following table describes the meaning of the specifications that can be used in the SMFINTERVAL parameter:

Valid Value	Description
GLOBAL	Generate all interval SMF records at the rate established by the INTVL and SYNCVAL SMF parameters in PARMLIB member SMFPRMxx.
NONE	Generate no interval SMF records. This is the default.
SUBSYS	Generate interval SMF records at the rate specified in PARMLIB member SMFPRMxx for the subsystem identified in the ADARUN SMFSUBSYS parameter.
<i>minutes</i>	Generate interval SMF records at the specified interval, in minutes. Valid minute values can range from 1 through 9999..

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

Example

In the following example, interval SMF records are generated at the rate specified in PARMLIB member SMFPRMxx for the subsystem identified in the ADARUN SMFSUBSYS parameter.

```
ADARUN PROG=ADANUC,SMFINTERVAL=SUBSYS
```

In the following example, interval SMF records are generated every 6 hours (3600 minutes)

```
ADARUN PROG=ADANUC,SMFINTERVAL=3600
```

SMFRECNO Parameter: Selecting User-Defined Adabas SMF Record Number

Parameter	Specify . . .	Valid Values	Default
SMFRECNO	the Adabas SMF record number used for user-defined SMF records.	128 - 255	255

IBM designates the range of SMF numbers from 128 through 255 for user-defined records.

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

Example

In the following example, user-defined SMF records will be assigned number 200.

```
ADARUN PROG=ADANUC,SMFRECNO=200
```

SMFSUBSYS Parameter: Specifying SMF Subsystem Name

Parameter	Specify . . .	Valid Values	Default
SMFSUBSYS	the name of the IBM or user-defined SMF subsystem from PARMLIB member SMFPRMxx.	CURRENT <i>subsystem-name</i>	CURRENT

This parameter associates the Adabas SMF records with an IBM or user-defined subsystem. The z/OS systems staff can provide different sets of SMF system parameters for IBM-defined subsystems. Up to eight user-specified subsystems can be defined, each with their own set of system parameters. This parameter allows you to optionally identify one of these subsystem names to be used with Adabas SMF recording, if the one you want is different from the current one.

The possible values for this parameter are:

- **CURRENT:** This is the default value and associates the Adabas SMF record with the IBM-defined subsystem under which the Adabas nucleus is active. The subsystem may be TSO for a nucleus running in a TSO session, STC for a starting task or the name of the job entry subsystem under which a batch job is running, JES2 or JES3.
- An explicit subsystem name can be specified. It may be either an IBM-defined name or a user-defined name from PARMLIB member SMFPRMxx. The subsystem name is one to four characters long. The first character must be alphabetic or national (#, @, or \$) and the remaining characters can be either alphanumeric or national characters. Contact your z/OS administrator for more information about the subsystem names available in your z/OS environment.

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

SMGT Parameter: Error Handling and Message Buffering Facility

Parameter	Specify . . .	Possible Values	Default
<u>SMGT</u>	whether to enable the error handling and message buffering facility.	YES NO	NO

Note:

This parameter is not available on BS2000 systems.

The meanings of the possible values are described below:

Value	Meaning
YES	ADARUN loads the module ADAMXI during initialization. If you want to use the message buffering functions, you must also specify the ADARUN MSGBUF parameter with a value greater than zero.
NO	(default) Error handling and message buffering is not enabled for the session.

This parameter can only be used when starting the nucleus. It cannot be used in utilities or user programs.

For more information about the error handling and message buffering facility, see the *Adabas DBA Tasks Manual*.

For information about temporarily turning the facility off and back on, see the SMGT operator command.

Example

The following example activates the error handling and message buffering facility with a buffer size of 36 kilobytes when the Adabas nucleus is initialized.

```
ADARUN PROGRAM=ADANUC
ADARUN SMGT=YES
ADARUN MSGBUF=36
```

SORTCACHE Parameter: Sort Area Controller Caching

This parameter is valid for the z/OS operating environment only.

Parameter	Specify . . .	Possible Values	Default
<u>SORTCACHE</u>	whether to enable (YES) or disable (NO) controller caching for the Adabas Sort data set.	YES NO	YES

SORTCACHE controls caching for the sort area component of Adabas.

Value	Meaning
YES	The default setting. Sort blocks are preread and maintained as long as possible in the caching facilities of the disk controller.
NO	Sort blocks are not preread or stored.

Guidelines

- If sort blocks are highly scattered or the update frequency is low, controller caching can cause longer I/O response times.
- ADARUN caching control parameters also exist for the Associator (ASSOCACHE), Data Storage (DATACACHE), the work area (WORKCACHE), and the temp area (TEMPCACHE).

Example

The following example runs the nucleus without controller caching for the sort data set.

```
ADARUN PROG=ADANUC , SORTCACHE=NO
```

SPT Parameter: Activate Adabas Triggers and Stored Procedures

Parameter	Specify . . .	Possible Values	Default
<u>SPT</u>	whether to enable use of triggers and stored procedures.	YES NO	NO

When the Adabas nucleus starts and SPT=YES has been specified, it passes control to the Adabas trigger driver (ADATSP) to allow it to initialize.

This parameter can only be used when starting the nucleus and only in multiuser mode. It cannot be used for utilities or user programs and it cannot be used in single-user mode.

For more information, see the Adabas Triggers and Stored Procedures Manual.

Example

The following example sets Adabas to use the triggers and stored procedures functions.

```
ADARUN PROGRAM=ADANUC ,MODE=MULTI ,SPT=YES
```

SRLOG Parameter: Controlling Spanned Record Logging

Parameter	Specify . . .	Possible Values	Default
<u>SR</u> LOG	how spanned records should be logged to the protection logs.	ALL UPD PART	UPD

When SRLOG=ALL is specified, the entire spanned record is written to the PLOG. When SRLOG=UPD is specified, only the modified spanned record segments are written to the PLOG. When SRLOG=PART is specified, all spanned records up to and including the modified segments are written to the PLOG; the remainder of the spanned record segments are not written to the PLOG.

Spanned records are logical records that are segmented into two to five data storage records and chained together. Spanned records may be segmented at the field or byte level.

Note:

Standalone secondary spanned records encountered in the PLOG will be rejected from further processing. A warning message is issued. Likewise, decompression of a partial field at the end of a spanned record will be skipped. All fields up to the partial field will be decompressed and the remaining fields on the spanned record are not available.

Examples

The following example requests complete spanned records should be logged to the PLOGs.

```
ADARUN PROGRAM=ADANUC ,SRLOG=ALL
```

The following example requests that the complete spanned record, up to and including the modified segments, should be logged to the PLOGs.

```
ADARUN PROGRAM=ADANUC ,SRLOG=UPD
```

SUBMPSZ Parameter: GETMAIN Memory Pool for Subtasks

Parameter	Specify . . .	Possible Values	Default
SUBMPSZ	the common memory pool size, in bytes, for subtask communication in products such as Adabas Review, Adabas Parallel Services, and Event Replicator for Adabas.	100000 - address-limit	1,024,000

This parameter is required for BS2000 nuclei that run subtasks. These may be Adabas Review, Adabas triggers and stored procedures, or Event Replicator for Adabas running EntireX Broker. Recommended values with any of these subtasks running are shown in the table below:

Subtask	Recommended SUBMPSZ Value
Adabas Review	14 M
Adabas triggers and stored procedures	20 M
Event Replicator for Adabas running webMethods EntireX Broker	200M (or larger)

Notes:

1. Setting this parameter for Adabas Review replaces an optional zap for increasing the subtask common memory.
2. This parameter must be set to the recommended value for the Adabas Review hub and the Adabas Review nuclei, irrespective of the value of the REVIEW parameter.
3. For Adabas triggers and stored procedures in BS2000 environments, make sure that the nucleus is started with the ADARUN parameter SUBMPSZ set to "20M" (or larger). Otherwise, the Natural subtask will deliver a "ADAI2S - 04000004 no mother task common memory" error.

Example

The following example allows for four (4) megabytes of common memory pool storage for use in the communication between the Adabas nucleus and the subtasks.

```
ADARUN PROG=ADANUC, SUBMPSZ=4096000
```

SVC Parameter: SVC Number

This parameter applies to the operating environments z/OS and z/VSE only.

Parameter	Specify . . .	Possible Values	Default
<u>SVC</u>	the Adabas SVC number or Adabas Review hub SVC number to be used for the session.	see text	45 (z/VSE) 249 (z/OS)

The SVC number is specified as an integer. It must correspond to the number used for the Adabas SVC at your installation.

The Adabas SVC or Adabas Review hub SVC are used to perform various Adabas internal functions under z/OS and z/VSE.

Valid SVC values are as follows:

z/OS 200-255

z/VSE 45 is recommended; any free SVC value can be used. See the Adabas Installation documentation for information about finding free values for z/VSE.

Example

The following example runs an Adabas session under z/OS using SVC 202 for the Adabas SVC.

```
ADARUN PROG=ADANUC,SVC=202
```

The following example runs an Adabas Review hub session under z/VSE using SVC 45 for the Adabas Review hub SVC.

```
ADARUN PROG=ADAREV,SVC=45
```

SWITCHNR Parameter: Setting the Job Switch Number

This parameter is valid on BS2000 systems only.

Parameter	Specify . . .	Possible Values	Default
<u>SWITCHNR</u>	the job switch number that is set when a nucleus or utility exits with a non-zero response code.	1 - 31	10

Example

The following example sets the job switch number to 15.

```
ADARUN PROG=ADANUC, SWITCHNR=15
```

TAPEREL Parameter: Tape Handling Control (BS2000 Only)

Parameter	Specify . . .	Possible Values	Default
TAPEREL	end-of-file (EOF) handling for tapes and cartridges.	KEEPUNL KEEP NO RELEASE UNLOAD	KEEPUNL

Value Meaning

KEEPUNL (default) Retain the I/O unit assignment but unload the tape or cartridge.

KEEP Retain the I/O unit assignment without unloading.

NO Do not perform the RELEASE macro function.

RELEASE Do not keep the I/O unit assignment; do not unload the tape or cartridge.

UNLOAD Do not keep the I/O assignment; unload the tape/cartridge.

Example

In the following example, the I/O assignment remains and the tape/cartridge remains loaded after the ADAULD utility is executed.

```
ADARUN PROG=ADAULD, TAP=KEEP
```

TARGETID Parameter: Entire Net-Work Target ID

Parameter	Specify . . .	Minimum	Maximum	Default
TARGETID	the unique Entire Net-Work target ID for this node.	1	65535	1

The TARGETID parameter is an optional ADARUN parameter that specifies the unique Entire Net-Work target ID of a node. It is synonymous with the Adabas ADARUN DBID parameter.

All target IDs used by Entire Net-Work, Adabas (database IDs), Natural global buffer pools, 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, with the exception of isolated databases that are

defined for local availability only and are therefore unknown to Entire Net-Work. For more information about global target IDs, read *Target ID Handling in the Network*, in the *Entire Net-Work Administration Guide*.

Example

The following example specifies an Entire Net-Work target ID of 12 for a node.

```
ADARUN TARGETID=12
```

TASKCTGY Parameter: Adabas Batch/TP Task Category Control

This parameter applies in BS2000 operating environments only.

Parameter	Specify . . .	Possible Values	Default
TASKCTGY	the Adabas task category.	BATCH TP	BATCH

BS2000 TP (interactive processing) environments favor TP-processing tasks over batch tasks. By default, Adabas has batch status. You can use the TASKCTGY parameter to assign TP-processing priority to the Adabas nucleus.

Value	Meaning
BATCH	(default) The nucleus retains batch status.
TP	Issues a TINF macro to force the Adabas nucleus to TP status.

Example

The following example forces the nucleus to TP-processing status.

```
ADARUN PROG=ADANUC ,TASKCTGY=TP
```

TCPIP Parameter: TCP/IP Access Control (z/OS Only)

Parameter	Specify . . .	Possible Values	Default
TCPIP	whether to activate a direct TCP/IP link to the Adabas nucleus.	YES NO	NO

For UES-enabled databases, Adabas version 7 supports a direct TCP/IP link to the z/OS Adabas nucleus from web-based applications or from PC-based applications such as Natural for Windows.

See the Adabas Installation documentation for information about installing direct TCP/IP database access.

The TCPIP parameter is used to activate a direct TCP/IP link to the Adabas nucleus.

If TCPIP=YES, you can specify the universal resource locator (URL) required to activate the direct TCP/IP link in the ADARUN parameter TCPURL.

Example

The following example runs Adabas with a direct TCP/IP link to the nucleus.

```
ADARUN  PROG=ADANUC,TCPIP=YES,TCPURL=HPS://STACKNAME:12345
```

TCPURL Parameter: TCP/IP Universal Resource Locator

Parameter	Specify . . .	Possible Values	Default
TCPURL	the universal resource locator (URL) for the TCP/IP link.	(see text)	none

If TCPIP=YES, you can specify the information required to activate the direct TCP/IP link to the Adabas nucleus. The parameter value is a 20-byte address that conforms to the RFC specification for universal resource locators (URLs):

```
TCPURL=api-name://stackid:port-number
```

where

api-name is a 1-3 character value identifying the application programming interface (API) to use. Both APIs for the IBM TCP/IP stack (HPS, OES) and the API for the Computer Associates' Unicenter TCPaccess Communications Server (ILK) are currently supported.

stackid is a 1-8 character value identifying the stack to use. For the HPS API, this is the name of the TCP/IP started task. For the OES API, no value is needed. For the ILK API, this is the subsystem identifier.

port-number is a 1-5 character number in decimal notation.

Optionally, you can specify this and additional URLs using the operator command TCPIP.

Example

The following three examples run Adabas with a direct TCP/IP link to the nucleus.

```
ADARUN  PROG=ADANUC ,TCPIP=YES ,TCPURL=OES : // : 12216
ADARUN  PROG=ADANUC ,TCPIP=YES ,TCPURL=HPS : // TCPLPAR1 : 12213
ADARUN  PROG=ADANUC ,TCPIP=YES ,TCPURL=ILK : //ACCA : 12215
```

TEMPCACHE Parameter: Temp Area Controller Caching

This parameter is valid for z/OS operating environments only.

Parameter	Specify . . .	Possible Values	Default
TEMPCACHE	whether to enable (YES) or disable (NO) controller caching for the Adabas Temp data set.	YES NO	YES

Value	Meaning
YES	The default setting. Temp blocks are pre-read and maintained as long as possible in the caching facilities of the disk controller.
NO	Temp blocks are not pre-read or stored.

Guidelines

- If temp blocks are highly scattered or the update frequency is low, controller caching can cause longer I/O response times.
- ADARUN caching control parameters also exist for the Associator (ASSOCACHE), Data Storage (DATACACHE), the work area (WORKCACHE), and the sort area (SORTCACHE).

Example

The following example runs the nucleus without controller caching for the temp data set.

```
ADARUN  PROG=ADANUC ,TEMPCACHE=NO
```

TFLUSH Parameter: Synchronous Buffer Flush Duration

Parameter	Specify . . .	Minimum	Maximum	Default
TFLUSH	the time allowed for a synchronous buffer flush.	1	16777215	1

The TFLUSH value specifies the time allowed for a synchronous buffer flush, in seconds.

If the LFIOP parameter specifies a non-zero value, the TFLUSH value has no effect because asynchronous buffer flush control is active.

If LFIOP=0 (the default) and if TFLUSH is not specified or has the value 1, one second is allowed for each synchronous buffer flush.

Example

For example, during the following Adabas session, 6 seconds are allowed for each synchronous buffer flush.

```
ADARUN PROG=ADANUC,TFLUSH=6,LFIOP=0, ...
```

TLSCMD Parameter: Time Limit for S1, S2, and S4 Complex Searches

Parameter	Specify . . .	Minimum	Maximum	Default
TLSCMD	the time limit (in seconds) for executing a single S1, S2, or S4 command with complex search criteria.	1	the setting of the ADARUN MXTSX parameter	300

The TLSCMD parameter specifies the maximum time (in units of 1.048576 seconds) to be used to process a single Adabas S1, S2, or S4 command with complex search criteria.

If the estimated time to process a command exceeds the TLSCMD time limit, the command is not processed and the user who issued the command receives a nonzero response code.

If the estimated time to process a command does not exceed the time limit, the command is processed to completion, whether the actual time required exceeds the time limit specified by the TLSCMD or not.

If the TLSCMD setting exceeds the maximum time permitted by the ADARUN MXTSX parameter, the maximum MXTSX time is used.

Overriding the Parameter Setting

You can override this parameter setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.

Example

For example, during the following Adabas session, the time limit for an S1, S2, or S4 command with a complex search criterion is about 50 minutes.

```
ADARUN PROG=ADANUC,TLSCMD=3000
```

TMDRQ Parameter: Size of ATM Request Queue

Parameter Type	Use	Possible Values	Default
Runtime	Sets the number of entries allowed in the transaction manager's internal request queue.	10 - 32767	10

The ATM transaction manager uses an internal request queue for scheduling some of the tasks it needs to perform. The TMDRQ parameter determines the total number of entries in the queue that is used to service certain incoming, asynchronous requests. This queue is currently used only when the ATM RRMS interface is active.

RRMS Considerations

When the RRMS interface is in use, the ATM dynamic request queue (DRQ) is used for communication between ATM RRMS exits and the ATM manager. The value specified for the TMDRQ parameter should therefore be increased if RRMS is to be used. Consider the number of global transactions that may reach a syncpoint at the same time, and increase TMDRQ by this number. You can specify a higher value, for safety, without incurring any performance degradation.

You can use Online Services to monitor usage of the DRQ. If you specify too small a value and the DRQ becomes full, backouts resulting in response code 9 (ADARSP009) are likely to occur.

TMETDATA Parameter: ET Data Control

Parameter Type	Use	Possible Values	Default
Runtime	Location of ET data	ATM TARGETS	TARGETS

The TMETDATA parameter determines the database or databases that will store ET data. The same TMETDATA parameter value must be specified for all Adabas Transaction Manager (ATM) instances on the network.

Possible values:

Value	Explanation
ATM	<p>ET data is always stored in and read from the ATM transaction manager's recovery file, without regard to the database ID indicated by the Adabas command.</p> <ul style="list-style-type: none"> ● treats ET data as belonging to a global transaction rather than to a database; ● there is just one copy of a client's ET data in a system; ● eliminates confusion resulting from the existence of different ET data in different databases under the same ETID; ● overrides the Natural ETDB parameter; ● should not be used if the EmergencySerialETs runtime control will be set to YES or FORCE for any client environment in which ET data is stored or read.
TARGETS	<p>ET data is stored during the commit process in all changed databases that are running with DTP=RM.</p> <p>ET data that is stored by an ET or CL command is always stored in the database that is the target of the command, whether or not that database runs with DTP=RM. If the database runs with DTP=NO, the ET data is written to the database only after any open global transaction has been committed.</p> <p>If the client storing ET data is at global transaction status, ATM is not required to carry out any commit processing; the data is simply stored in the database indicated by the Adabas command. ET data is always read from the database specified in the Adabas command.</p> <ul style="list-style-type: none"> ● treats ET data as belonging to a database or a local database transaction rather than to a global transaction; ● a client's ET data might exist in several databases in a system; ● a client can have several different ET data values in different databases at the same time; ● means that 3GL application programs using ET data must have knowledge of the database that holds ET data for a given transaction; ● honors the Natural ETDB parameter.

Notes:

1. The recommended setting for this parameter is the default, TARGETS. This setting eliminates the dangers associated with running in serial mode when the transaction manager is unavailable.
2. See also the section entitled *ET Data Storage* in your Adabas Transaction Manager operations documentation.

TMGTT Parameter: Global Transaction Time Limit

Parameter Type	Use	Possible Values	Default
Runtime	Sets time limit during which a global transaction can be open without being prepared.	1 - 16777215	720

The TMGTT parameter approximately specifies the maximum elapsed time, in units of 1.048576 seconds, that a global transaction can remain open without being prepared. If the limit is exceeded, ATM automatically backs out the global transaction from all affected databases. The next time the client issues a transactional Adabas command, a response code 9 (ADARSP009) is returned.

It is recommended that TMGTT be smaller than the lowest TT parameter of any database running with DTP=RM.

The TMGTT parameter does not supersede the TT parameter settings of other databases. If a database times out a client who has an unprepared global transaction open, the global transaction is subsequently backed out. A response code 9 (ADARSP009) is returned if the client then issues a transactional command.

Note:

Transaction timeout can happen when a client session stops issuing commands for a period of time. In these circumstances the transaction manager has no opportunity to give a response code to the client, to indicate that the timeout has taken place. In such cases, the pending response code will be preserved and returned to the client at the first possible opportunity. For further details, please refer to *Pending Response Codes* in the Adabas Transaction Manager introduction.

TMLOG Parameter: ATM Manager Logging Control

Parameter Type	Use	Possible Values	Default
Runtime	Controls the logging option for the current execution of the ATM manager.	YES NO NEVER	NO

The ATM manager can log internal events for diagnostic purposes using a dual logging system. It can

- maintain a log buffer area in memory that contains details of the latest logged events;
- maintain transaction-related event details in storage areas related to specific global transactions;
- write every log record to dual sequential ATMLOG data sets.

Possible values:

Value	Description
YES	The ATM manager writes every record to an ATMLOG data set. The ATM manager keeps a record of the latest internal events in its global log buffer area. It also maintains the contents of log buffer areas related to open global transactions. Logging to the ATMLOG data sets can be deactivated at any time using the operator command <code>ATM NOLOG</code> .
NO	The ATM manager will not write log records to the ATMLOG data sets. Logging to the ATMLOG data sets can be activated at any time using the operator command <code>ATM LOG</code> . The ATM manager keeps a record of the latest internal events in its global log buffer area; it also maintains the contents of log buffer areas related to open global transactions.
NEVER	ATM logging is disabled for the duration of this execution of the ATM manager.

On z/OS and VSE systems, if the ATM manager writes log records to the ATMLOG (TMLOG=YES or ATM LOG issued) and the data set currently in use becomes full, the manager closes that data set and switches to the other data set. Automatic log switching currently does not take place on BS2000 systems. A system message might appear in the console log, indicating that the first log data set has become full. You may also switch between the two ATMLOG data sets by using the operator command `ATM FEOFLOG`. Whenever the ATMLOG data sets are switched, a message is displayed on the console indicating the data set that is currently in use.

TMMSGSEV Parameter: Severity Threshold for Warning Messages

Parameter Type	Use	Possible Values	Default
Runtime	Sets severity threshold for the suppression of warning messages.	0 4 8	0

Every message that the ATM manager sends to the console has a severity level. Message severity levels are explained in the Adabas Transaction Manager *Messages and Codes*.

Using the TMMSGSEV parameter, you can prevent messages with low severity levels from being written to the console by the ATM manager. This parameter has no effect on messages that are issued by the ATM client proxy, or during the early part of the transaction manager's initialization.

Possible values:

Value	Description
0	No messages are suppressed.
4	Messages of severity less than 4 are not sent to the console.
8	Messages of severity less than 8 are not sent to the console.

TMRESTART Parameter: ATM Restart Control

Parameter Type	Use	Possible Values	Default
Runtime	Controls restart handling of problematic transactions.	NORMAL FORCE FORCEALL	NORMAL

When an ATM manager restarts, it rebuilds its global transaction list as nearly as possible to its latest state when it was last executing. It then attempts to complete or back out any incomplete global transactions. The TMRESTART parameter determines the manager's restart action with regard to global transactions and transaction branches that remain incomplete at the end of this process.

Possible values:

Value	Description
NORMAL	Any incomplete transactions remain in an incomplete state until such time as they can be completed.
FORCE	ATM transfers to the suspect transaction journal (STJ) the details of every incomplete global transaction that has its root local to this ATM manager. The details of these transactions are deleted from ATM's recovery information, the originating clients are closed, and any related internal resources are freed. ATM can no longer guarantee integrity for such global transactions.
FORCEALL	Details of all unresolved global transactions and transaction branches are transferred to the STJ. The details are deleted from ATM's recovery information, the originating clients are closed, and any related internal resources are freed. ATM can no longer guarantee integrity for the affected global transactions.

TMSYNCMGR Parameter: Support for External Transaction Coordinators

Parameter Type	Use	Possible Values	Default
Runtime	Indicates whether or not the ATM manager is to interact with an external transaction coordinator.	NONE RRMS	NONE

The TMSYNCMGR parameter determines whether or not the ATM manager registers and interacts with an external transaction coordinator.

Possible values:

Value	Description
NONE	The ATM manager will not interact proactively with an external transaction coordinator. This setting does not preclude the use of the CICS Syncpoint Manager, which interfaces directly with the ATM's client proxy component and not with the ATM manager.
RRMS	The ATM manager will register with the IBM Recoverable Resource Management Services so that it can participate in transactions that involve other RRMS-enabled resource managers. This parameter value is valid only for z/OS systems in which RRS (a component of RRMS) is active. The current version of ATM offers participation in RRMS-coordinated two-phase commit for single-user, single-TCB batch applications and for applications running under Com-plete or IMS TM.

TMTCIDPREF Parameter: Dynamic Client ID Prefix

Parameter Type	Use	Possible Values	Default
Runtime	Defines the first one or two characters of dynamically allocated Client IDs.	up to 2 alphanumeric characters	TM

To aid interpretation of Adabas Online Services displays and diagnostic logs, the ATM manager dynamically allocates an 8-byte client identifier (TCID) for each client session that engages in transactional activity. These client identifiers are simply labels, and have no operational significance.

The format of a dynamically allocated TCID is

ppnnxxxx

—where

pp is the value of the TMTCIDPREF parameter, left-justified and with trailing blanks if required.

nn is the 2-byte binary Node ID of the Adabas System Coordinator daemon within which the ATM transaction manager is running as a service.

xxxx is a 4-byte binary suffix. Suffixes are allocated in sequence and are reused. The first TCID to be allocated has a suffix of zeros.

TNAA Parameter: Nonactivity Time Limit for Access-Only Users

Parameter	Specify . . .	Minimum	Maximum	Default
TNAA	the nonactivity time limit for access-only users.	1	16777215	900

The TNAA parameter specifies the maximum elapsed time (in units of 1.048576 seconds) that an access-only user may be considered active without issuing an Adabas command.

If an access-only user exceeds the nonactivity time limit, action is taken as described in *Adabas Timeout Feature*.

If the NUCID parameter specifies a value greater than zero (that is, a multiprocessing nucleus), you may have to specify a larger TNAA value than for a normal nucleus. The chances of data contention are greater, and therefore more time should be allowed.

Overriding the Parameter Setting

You can override this parameter setting for an individual user or for all users:

- Override the setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.
- Override the setting for *all* users with an Adabas TNAA operator command or with the OPERCOM function of the ADADBS utility (described in the *Adabas Utilities Manual*).

Example

For example, during the following Adabas session, the nonactivity time limit for access-only users is about 30 minutes.

```
ADARUN  PROG=ADANUC ,TNAA=1800
```

TNAE Parameter: Nonactivity Time Limit for ET Logic Users

Parameter	Specify . . .	Minimum	Maximum	Default
TNAE	the nonactivity time limit for ET logic users.	1	16777215	900

The TNAE parameter specifies the maximum elapsed time (in units of 1.048576 seconds) an ET Logic user may be considered active without issuing an Adabas command.

If an ET logic user exceeds the nonactivity time limit, action is taken as described in *Adabas Timeout Feature*.

If the NUCID parameter specifies a value greater than zero (that is, a multiprocessing nucleus), you may have to specify a larger TNAE value than for a normal nucleus. The chances of data contention are greater, and therefore more time should be allowed.

Overriding the Parameter Setting

You can override this parameter setting for an individual user or for all users:

- Override the setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.
- Override the setting for all users with an Adabas TNAE operator command or with the OPERCOM function of the ADADBS utility (described in the *Adabas Utilities Manual*).

Example

For example, during the following Adabas session, the nonactivity time limit for ET logic users is about 30 minutes.

```
ADARUN PROG=ADANUC ,TNAE=1800
```

TNAX Parameter: Nonactivity Time Limit for Exclusive Control Users

Parameter	Specify ...	Minimum	Maximum	Default
TNAX	the nonactivity time limit for exclusive control users.	1	16777215	900

The TNAX parameter specifies the maximum elapsed time (in units of 1.048576 seconds) an exclusive control user (EXU) may be considered active without issuing an Adabas command.

If an exclusive control user exceeds the nonactivity time limit, action is taken as described in *Adabas Timeout Feature*.

If the NUCID parameter specifies a value greater than zero (that is, a multiprocessing nucleus), you may have to specify a larger TNAX value than for a normal nucleus. The chances of data contention are greater, and therefore more time should be allowed.

Overriding the Parameter Setting

You can override this parameter setting for an individual user or for all users:

- Override the setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.
- Override the setting for all users with an Adabas TNAX operator command or with the OPERCOM function of the ADADBS utility (described in the *Adabas Utilities Manual*).

Example

For example, during the following Adabas session, the nonactivity time limit for exclusive control users is about 30 minutes.

```
ADARUN  PROG=ADANUC, TNAX=1800
```

TT Parameter: Transaction Time Limit

Parameter	Specify . . .	Minimum	Maximum	Default
TT	the transaction time limit for ET logic users.	1	16777215	900

The TT parameter specifies the maximum elapsed time (in units of 1.048576 seconds) permitted for a logical transaction issued by an ET logic user.

The time measurement for a logical transaction starts when the first command that results in a record being placed in hold status is issued. It ends when an ET, BT, or CL command is issued.

If the transaction time limit is exceeded, action is taken as described in *Adabas Timeout Feature*.

The value specified for TT directly influences the required size of the Adabas parameter LP (length of the data protection area).

Software AG recommends that you set the TT parameter lower than the TNAE parameter.

Overriding the Parameter Setting

You can override this parameter setting for an individual user or for all users:

- Override the setting for an individual user by specifying a different value in the Adabas control block for an OP command. See the discussion of the OP command in the *Adabas Command Reference Guide*.
- Override the setting for all users with an Adabas TT operator command or with the OPERCOM function of the ADADBS utility (described in the *Adabas Utilities Manual*).

Example

For example, during the following Adabas session, the transaction time limit for ET logic users is about 5 minutes (300 seconds).

```
ADARUN  PROG=ADANUC, TT=300
```

UEXn Parameter: User Exit

Parameter	Specify . . .	Possible Values	Default
UEXn	the user exit and the user routine to be given control.	see text	none

This parameter is used in conjunction with the user exit facility. It specifies one or more optional user exits and their names in the following format:

```
UEXn=exit-name
```

where

n is a digit in the range 1-12 inclusive

exit-name is the name of a user routine that gets control at the user exit; the name can be up to 8 characters long.

Note:

User exit 2 and user exit 12 are mutually exclusive; if you specify one, you may not specify the other.

The specified user exit routine(s) must be loadable at execution time.

The Adabas user exits discussion describes the purpose and function of each user exit as well as the calling sequence.

Examples

In the following example, control is given to the user routine SECURE at user exit 11 during the Adabas session.

```
ADARUN PROG=ADANUC,UEX11=SECURE
```

During the following Adabas session, control is given to the user routine SUBR2 at user exit 2 and to the user routine SUBR4 at user exit 4.

```
ADARUN PROG=ADANUC,UEX2=SUBR2,UEX4=SUBR4
```

The following example runs the ADACMP utility and gives control to the user routine SUBR6 at user exit 6.

```
ADARUN PROG=ADACMP,UEX6=SUBR6
```


UEXSMF Parameter: Specifying SMF User Exit Module Name

Parameter	Specify . . .	Valid Values	Default
UEXSMF	the name of the SMF user exit module.	<i>module-name</i>	---

If user details sections are to be generated in the Adabas SMF record (if the ADARUN parameter SMFDETAIL is set to USER or if USER is added as a detail section via the SMFDETAIL operator command), this parameter can be used to identify the name of the SMF user exit that should be used to generate the user detail sections. This exit will be invoked for each Adabas SMF record.

This parameter *cannot* be altered by operator commands while the nucleus is running.

Example

In the following example, the SMF user exit name is SMFEXIT

```
ADARUN PROG=ADANUC,UEXSMF=SMFEXIT
```

UTIONLY Parameter: Utilities-Only Session Control

Parameter	Specify . . .	Possible Values	Default
UTIONLY	whether to restrict the session to Adabas utilities.	YES NO	NO

Use this parameter to restrict an Adabas session to the execution of Adabas utilities.

Value	Meaning
YES	Any attempt to execute a program other than an Adabas utility is rejected.
NO	(default) The session is available for both utilities and user programs.

If you specify UTIONLY=YES to reserve the session for a utility program, you can use the operator command UTIONLY=NO to make the nucleus available for user programs again when the utility is completed.

Example

In the following example, only utilities can be run during this Adabas session.

```
ADARUN PROG=ADANUC,UTIONLY=YES
```

V64BIT Parameter: 64-Bit Virtual Storage Option

Note:

This parameter is applicable only to z/OS 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.

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.

Example

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

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

VISTA Parameter: Adabas Vista Control

Parameter	Specify . . .	Possible Values	Default
<u>V</u> ISTA	whether to enable support for Adabas Vista.	YES NO	NO

The VISTA parameter enables the nucleus to support Adabas Vista processing. If Adabas Vista is not installed on your system, specifying this parameter invokes the AVILOOK analysis tool delivered with Adabas. For complete information about AVILOOK, read *AVILOOK*.

Example

The following example runs Adabas with Adabas Vista.

```
ADARUN PROG=ADANUC,VISTA=YES
```

WORKCACHE Parameter: Work Area Controller Caching

This parameter is valid for z/OS operating environments only.

Parameter	Specify . . .	Possible Values	Default
<u>W</u> ORKCACHE	whether to enable (YES) or disable (NO) controller caching for the Adabas Work data set.	YES NO	YES

Value	Meaning
YES	The default setting. Work blocks are preread and maintained as long as possible in the caching facilities of the disk controller.
NO	Work blocks are not preread or stored.

Guidelines

- If Work blocks are highly scattered or the update frequency is low, controller caching can cause longer I/O response times.

- ADARUN caching control parameters also exist for the Associator (ASSOCACHE), Data Storage (DATACACHE), the sort area (SORTCACHE), and the temp area (TEMPCACHE).

Example

The following example runs the nucleus without controller caching for the Work data set.

```
ADARUN PROG=ADANUC,WORKCACHE=NO
```