

Adabas Session Execution

An Adabas session involves the execution of the Adabas nucleus which controls access/update to a single database. This chapter describes the job control statements needed when executing an Adabas session under each supported operating system. For examples of the Adabas utility jobs, see the Adabas Utilities documentation.

This chapter covers the following topics:

- Program Synchronization Using Operating System Services
 - BS2000 Session
 - z/OS Session
 - z/VSE Session
 - Single-User Mode
 - User Profile
-

Program Synchronization Using Operating System Services

Adabas version uses operating system services to synchronize the start and end of nucleus and utility executions. Only one program can modify the data integrity block (DIB) at a time.

The operating system services used are as follows:

Operating System	Service
z/OS	systems-wide ENQ/DEQ macros (SCOPE=SYSTEMS) with major name (QNAME) 'ADABAS'
BS2000	system-wide ENQAR/DEQAR macros (SCOPE=GLOBAL)
z/VSE	system-wide LOCK/UNLOCK macros

This feature reliably and efficiently guarantees proper synchronization of DIB updates within a single operating-system image.

If your database resides on disks that are shared among multiple images of the operating system and you run nucleus or utility jobs against the same database on more than one of the system images, you need to ensure that

- the system images are installed in such a way that synchronization is effective on all systems where nucleus and utility jobs execute; or
- nucleus and utility jobs do not execute concurrently on different system images.

Consult your system programmer for the needed information.



Warning:

If different nucleus or utility jobs updating the same file are allowed to start or terminate on different system images at the same time without proper synchronization, a DIB update may be lost. If this happens, a lock in the DIB may be violated, thereby opening the file to the possibility of destruction due to concurrent unsynchronized updates by utilities.

BS2000 Session

The following table contains all data sets which are used when executing an Adabas session under BS2000.

Data Set	Link Name	Logical Unit	Storage Medium	Additional Information
ADARUN parameters	DDCARD	SYSDTA	disk	note 1
ADARUN / Adabas messages	DDPRINT	SYSOUT	disk	note 2
Associator	DDASSOR n or DDASSO nn		disk	note 3
Data Storage	DDDATAR n or DDDATA nn		disk	note 3
Work	DDWORKR1 DDWORKR4		disk	note 4
Recovery Aid log	DDRLOGR1		disk	note 5
Protection log multiple log 1 multiple log 2	DDSIBA DDPLOGR1 DDPLOGR2		tape disk disk disk	note 6 note 7 note 7
Command log multiple log 1 multiple log 2	DDLOG DDCLOGR1 DDCLOGR2		tape disk disk disk	note 8 note 9 note 9
ECS encoding objects	DDECSOJ		tape/disk	note 10
Time zone objects	TZINFO		disk	note 11

- Example of Adabas Session Job Control (BS2000)

- JCL Required for UES Support (BS2000)
- BS2000 Job Notes

Example of Adabas Session Job Control (BS2000)

This job includes multiple protection logging, multiple command logging, Recovery Aid logging, and universal encoding support (UES):

In SDF Format:

```

/.ADANUC LOGON
/MODIFY-TEST-OPTIONS DUMP=YES
/REMARK *
/REMARK * ADABAS NUCLEUS
/REMARK *
/ASS-SYSLST L.NUC
/ASS-SYSDTA *SYSCMD
/SET-FILE-LINK DDLIB,ADAvrs.MOD
/SET-FILE-LINK BLSLIB01,$.SYSLNK.LMS
/SET-FILE-LINK DDASSOR1,ADAYyyyy.ASSO,SHARE-UPD=YES
/SET-FILE-LINK DDDATAR1,ADAYyyyy.DATA,SHARE-UPD=YES
/SET-FILE-LINK DDWORKR1,ADAYyyyy.WORK
/SET-FILE-LINK DDPLOGR1,ADAYyyyy.PLOGR1,SHARE-UPD=YES
/SET-FILE-LINK DDPLOGR2,ADAYyyyy.PLOGR2,SHARE-UPD=YES
/SET-FILE-LINK DDCLOGR1,ADAYyyyy.CLOGR1,SHARE-UPD=YES
/SET-FILE-LINK DDCLOGR2,ADAYyyyy.CLOGR2,SHARE-UPD=YES
/SET-FILE-LINK DDRLOGR1,ADAYyyyy.RLOGR1,SHARE-UPD=YES
/SET-FILE-LINK TZINFO,ADAvrs.TZ00
/START-PROGRAM *M(ADA.MOD,ADARUN),PR-MO=ANY,RUN-MODE=A(ALT-LIB=Y)
ADARUN PROG=ADANUC,DB=yyyyyy
ADARUN LBP=600000
ADARUN LWP=320000
ADARUN LS=80000
ADARUN LP=400
ADARUN NAB=24
ADARUN NC=50
ADARUN NH=2000
ADARUN NU=100
ADARUN TNAE=180,TNAA=180,TNAX=600,TT=90
ADARUN NPLOG=2,PLOGSIZE=1800,PLOGDEV=dddd
ADARUN NCLOG=2,CLOGSIZE=1800,CLOGDEV=dddd
/REMARK
/LOGOFF SYS-OUTPUT=DEL

```

In ISP Format:

```

/.ADANUC LOGON
/OPTION MSG=FB,DUMP=YES
/REMARK *
/REMARK * ADABAS NUCLEUS
/REMARK *
/SYSFILE SYSLST=L.NUC

/FILE ADAvrs.MOD,LINK=DDLIB
/FILE EXAMPLE.ADAYyyyy.ASSOR1,LINK=DDASSOR1,SHARUPD=YES
/FILE EXAMPLE.ADAYyyyy.DATAR1,LINK=DDDATAR1,SHARUPD=YES
/FILE EXAMPLE.ADAYyyyy.WORKR1,LINK=DDWORKR1
/FILE EXAMPLE.ADAYyyyy.PLOGR1,LINK=DDPLOGR1,SHARUPD=YES
/FILE EXAMPLE.ADAYyyyy.PLOGR2,LINK=DDPLOGR2,SHARUPD=YES

```

```

/FILE EXAMPLE.ADAyyyyy.CLOGR1, LINK=DDCLOGR1, SHARUPD=YES
/FILE EXAMPLE.ADAyyyyy.CLOGR2, LINK=DDCLOGR2, SHARUPD=YES
/FILE EXAMPLE.ADAyyyyy.RLOGR1, LINK=DDRLOGR1, SHARUPD=YES
/FILE ADAvrs.TZ00
/EXEC (ADARUN, ADAvrs.MOD)
ADARUN PROG=ADANUC, DB=yyyyy
ADARUN LBP=600000
ADARUN LWP=320000
ADARUN LS=80000
ADARUN LP=400
ADARUN NAB=24
ADARUN NC=50
ADARUN NH=2000
ADARUN NU=100
ADARUN TNAE=180, TNA=180, TNAX=600, TT=90
ADARUN NPLOG=2, PLOGSIZE=1800, PLOGDEV=ddd
ADARUN NCLOG=2, CLOGSIZE=1800, CLOGDEV=ddd
/REMARK
/LOGOFF NOSPOOL

```

JCL Required for UES Support (BS2000)

If you are using universal encoding support (UES), the following additional JCL is required for BS2000 environments:

In SDF Format:

```

/SET-FILE-LINK BLSLIB00, ADAvrs.MOD
/SET-FILE-LINK BLSLIB01, ADAvrs.LIB
/CLP (APSVrs.LIB, LMSLINKLIB), (LNK-NAME=BLSLIB02)
.
.
/SET-FILE-LINK DDEC SOJ, ADAvrs.ALLECSO
.
/START-PROGRAM *M(ADAvrs.MOD, ADARUN), PR-M0=ANY, R-M=*A(ALT=Y, LOAD-INF=*REF, -
/ UNRES=*DELAY, MESSAGE=*ERROR)
.

```

In ISP Format:

```

/FILE ADAvrs.ALLECSO, LINK=DDEC SOJ

```

BS2000 Job Notes

The following notes apply to Adabas startup jobs in various platforms (as described by the notes).

1. This data set is used to provide the Adabas session parameters.
2. This data set is used to print messages produced by the control module ADARUN or the Adabas nucleus.
3. The Adabas Associator and Data Storage data sets. These data sets are mandatory.

n and *nn* represent the number of the Associator and Data Storage data set, respectively.

If more than one data set exists for Associator or Data Storage, a separate statement is required for each. For example, if the Associator consists of two data sets in a z/VSE environment, DD statements for ASSOR1 and ASSOR2 are required.

If *less* than 10 data sets exist for each, you must use the ASSOR* and DATAR* DD names in the JCL. For example, if, in a z/OS environment, the Associator consists of two data sets and Data Storage consists of three data sets, the following names would be used in the JCL: DDASSOR1, DDASSOR2, DDDATAR1, DDDATAR2, and DDDATAR3. If 10 or more data sets exist, the first nine must use the ASSOR* and DATAR* DD names in the JCL and the remainder must use the ASSO* and DATA* DD names in the JCL (dropping the "R" in the DD names). For example, the tenth Associator data set in a z/OS environment would be identified in the JCL using the name DDASSO10, while the third Associator data set in the same JCL would be identified using the name DDASSOR3.

A maximum of 99 physical extents is now set for Associator and Data Storage data sets. However, your actual real maximum could be less because the extent descriptions of all Associator, Data Storage, and Data Storage Space Table (DSST) extents must fit into the general control blocks (GCBs). For example, on a standard 3390 device type, there could be more than 75 Associator, Data Storage, and DSST extents each (or there could be more of one extent type if there are less for another).

So, the range of Associator DD names that can be used in JCL is DDASSOR1 (ASSOR1 in z/VSE) to DDASSOR9 (ASSOR9 in z/VSE) and DDASSO10 (ASSO10 in z/VSE) to DDASSO99 (ASSO99 in z/VSE). And the range of Data Storage DD names that can be used in JCL is DDDATAR1 (DATAR1 in z/VSE) to DDDATAR9 (DATAR9 in z/VSE) and DDDATA10 (DATA10 in z/VSE) to DDDATA99 (DATA99 in z/VSE).

4. The Adabas Work data sets. The WORKR1 data set is mandatory. If you have Adabas Transaction Manager version 7.5 or later installed, an additional work data set, WORKR4 is also mandatory.

Software AG recommends running the nucleus with DISP=OLD (under z/OS; share not specified for BS2000 and z/VSE) for the WORKR1 data set as a way of preventing two nuclei from writing to the same WORK data set and corrupting the database. This could otherwise happen if the ADARUN parameters FORCE and IGNDIB are improperly used.

Work part 4 of WORKR1 is no longer supported if you have Adabas Transaction Manager Version 7.5 or later installed. Instead, you should use the WORKR4 data set. WORKR4 is used for the same purpose as Work part 4, but it can be used in parallel by all members in a cluster. It is used to store the PET (preliminary end-of transaction) overflow transactions (those that cause a work overflow) of a database or of all members in a multiplex/SMP cluster.

The WORKR4 data set is a container data set that should be allocated and formatted in the normal way (use ADAFRM WORKFRM), using a block size greater than or equal to WORKR1. WORKR4 can have the same or a different device type than WORKR1. It should be at least as large as the cluster's LP parameter of the database or cluster. The smaller WORKR1 Work part 1 is, the larger WORKR4 should be. This is because the nucleus must prevent a work overflow due to incomplete DTP transactions, but the nucleus must keep all PET transactions; they cannot be backed out.

5. If the Adabas Recovery Aid is being used, this logging data set is required.
6. The data protection log data set. This data set is required if the database will be updated during the session and logging of protection information is desired. This data set is not applicable if multiple protection logging is used.

The data protection log may be assigned to tape or disk. A new data set must be used for each Adabas session (DISP=MOD may not be used). See *Adabas Restart and Recovery* for additional information.

7. Multiple (two to eight) data protection log data sets. These data sets are required only if multiple data protection logging is to be in effect for the session.

Multiple data protection logging is activated by the ADARUN NPLOG and PLOGSIZE parameters. The device type of the multiple protection logs is specified with the ADARUN PLOGDEV parameter.

Whenever one of multiple protection log data sets is full, Adabas switches automatically to another data set and notifies the operator through a console message that the log which is full should be copied using the PLCOPY function of the ADARES utility. This copy procedure may also be implemented using the user exit 12 facility as described in the User Exits documentation.

If no command logging is to be performed, this data set may be omitted.

8. The command log data set. This data set is required if command logging is to be performed during the session. Command logging is activated by the ADARUN LOGGING parameter.
9. Multiple (two to eight) command log data sets. These data sets are required only if multiple command logging is to be in effect for the session.

Multiple command logging is activated by the ADARUN NCLOG and CLOGSIZE parameters. The device type of the multiple command log data sets is specified with the ADARUN CLOGDEV parameter.

Whenever one of multiple command log data sets is full, Adabas switches automatically to another data set and notifies the operator through a console message that the log which is full should be copied using the CLCOPY function of the ADARES utility. This copy procedure may also be implemented using the user exit 12 facility as described in the Adabas User Exits documentation.

10. The Entire Conversion Services (ECS) objects data set is required for universal encoding support (UES) in all operating environments.
11. The TZINFO data set is required when defining fields with the TZ option and using the user session open parameter TZ. This data set is a library or partitioned data set containing information on local time offsets and daylight savings time offsets and its transition times.

z/OS Session

The following data sets are required when executing an Adabas session under z/OS.

Data Set	DD Name	Storage Medium	Additional Information
ADARUN parameters	DDCARD	card image	note 1
ADARUN / Adabas messages	DDPRINT	printer	note 2
Associator	DDASSOR <i>n</i> or DDASSO <i>nn</i>	disk	note 3
Data Storage	DDDATAR <i>n</i> or DDDATA <i>nn</i>	disk	note 3
Work	DDWORKR1 DDWORKR4	disk	note 4
Recovery Aid log	DDRLOGR1	disk	note 5
Protection log multiple log 1 multiple log 2	DDSIBA DDPLOGR1 DDPLOGR2	tape/disk disk disk	note 6 note 7 note 7
Command log multiple log 1 multiple log 2	DDLOG DDCLOGR1 DDCLOGR2	tape/disk disk disk	note 8 note 9 note 9
ECS encoding objects	DDECSOJ	tape/disk	note 10
Abnormal termination	MPMDUMP	printer	note 11
Abnormal termination (if large buffer pools are used)	SVCDUMP	printer	note 12
SMGT dump and snap dump	ADASNAP	printer	note 13
SMGT print file	DDTRACE1	printer	note 14
Time zone files	TZINFO	disk	note 15

- Example of Adabas Session Job Control (z/OS)
- JCL Required for UES Support (z/OS)
- JCL Required for UES and TCP/IP Support (z/OS)
- Using the z/OS SVC Dump Facility
- z/OS JCL Notes

Example of Adabas Session Job Control (z/OS)

This job includes multiple protection logging, multiple command logging, and Recovery Aid logging:

```
//NUC099 EXEC PGM=ADARUN
//STEPLIB DD DISP=SHR,DSN=ADABAS.ADAvrs.LOAD
//DDASSOR1 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.ASSOR1
//DDDATAR1 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.DATAR1
//DDWORKR1 DD DISP=OLD,DSN=EXAMPLE.ADAyyyyy.WORKR1
//DDPLOGR1 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.PLOGR1
//DDPLOGR2 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.PLOGR2
//DDCLOGR1 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.CLOGR1
//DDCLOGR2 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.CLOGR2
//DDRLOGR1 DD DISP=SHR,DSN=EXAMPLE.ADAyyyyy.RLOGR1
//TZINFO DD DISP=SHR,DSN=ADABAS.Vvrs.TZ00
//DDPRINT DD SYSOUT=X
//DDTRACE1 DD SYSOUT=X
//SYSUDUMP DD SYSOUT=X
//MPMDUMP DD SYSOUT=X
//ADASNAP DD SYSOUT=X
//DDCARD DD *
ADARUN PROG=ADANUC,DB=yyyyy
ADARUN LBP=600000
ADARUN LWP=320000
ADARUN LS=80000
ADARUN LP=400
ADARUN NAB=24
ADARUN NC=1000
ADARUN NH=2000
ADARUN NU=100
ADARUN TNAE=180,TNAA=180,TNAX=600,TT=90
ADARUN NPLOG=2,PLOGSIZE=1800,PLOGDEV=dddd
ADARUN NCLOG=2,CLOGSIZE=1800,CLOGDEV=dddd
//
```

where:

dddd is a valid device type.

vrs is the *version* of the product.

yyyyy is the physical database ID.

JCL Required for UES Support (z/OS)

A sample startup job for a UES-enabled database is provided in member ADANUCU in the ADAvrs.JOBS data set. If you are using universal encoding support (UES), you must

- include the following additional libraries for internal products in the STEPLIB:

```
//STEPLIB DD DISP=SHR,DSN=ADABAS.APSvrs.LDnn
// DD DISP=SHR,DSN=ADABAS.APSvrs.LD00
```

where *nn* is the load library level. If the library with a higher level number is not a full replacement for the lower level library(s), the library with the higher level must precede those with lower numbers in the STEPLIB concatenation.

Note:

If you are using an Adabas load library prior to version 7.2.2, it contains internal product libraries with an earlier version number and must be ordered below the current internal product libraries in the STEPLIB concatenation.

- add the following additional JCL related to internal product libraries:

```
//DDECSOJ DD DISP=SHR,DSN=ADABAS.ADAvrs.EC00
//SYSPARM DD *
SYSTEM_ID=ADAAPS
ABEND_RECOVERY=NO
THREAD_ABEND_RECOVERY=NO
```

JCL Required for UES and TCP/IP Support (z/OS)

If you are connecting your UES-enabled database directly through a TCP/IP link, you must also

- include the ADATCP library in the STEPLIB:

```
//STEPLIB DD ....
//          DD DISP=SHR,DSN=WATvrs.LOAD
```

Note:

This library is distributed with Entire Net-Work.

- identify the TCP/IP stack you intend to use with the CDI_DRIVER parameter of the SYSPARM statement:

```
//DDECSOJ DD DISP=SHR,DSN=ADABAS.ADAvrs.EC0n
//SYSPARM DD *
SYSTEM_ID=ADAAPS
ABEND_RECOVERY=NO
THREAD_ABEND_RECOVERY=NO
* User must choose one of the following depending on the TCP/IP stack used:
*CDI_DRIVER=('tcpip,PAALSOCK,SUBSYS=ACSS') <--Interlink TCP/IP stack
*CDI_DRIVER=('tcpip,PAALSOCK,ADDRSPCE=STACKNAME') <--IBM TCP/IP stack for HPS
*CDI_DRIVER=('tcpip,PAAOSOCK,ADDRSPCE=STACKNAME') <--IBM TCP/IP stack for OE
```

Using the z/OS SVC Dump Facility

The dump produced by MPMDUMP may be too slow for users with very large buffer pools. You may instead elect to use the z/OS SVC dump facility to speed up nucleus dump processing. An SVC dump is triggered by the presence of an //SVCDUMP DD statement in the nucleus startup JCL.

If //SVCDUMP DD DUMMY is specified, and the job is running with APF authorization, a z/OS SVC dump is produced on the system dump data set, normally SYS1.DUMPxx. If //SVCDUMP DD DUMMY is specified and the job is *not* running with APF authorization, message ADAM77 is issued and dump processing continues as if the SVCDUMP DD statement had *not* been specified.

If //SVCDUMP DD DSN=*dsm* is specified (with an appropriate data set name), a z/OS SVC dump is produced on the specified data set. Note that the SVCDUMP data set needs to be allocated with DCB attributes RECFM=FB, LRECL=4160, BLKSIZE=4160. Note also that, for APF-authorized jobs, secondary extents are ignored.

For APF-authorized jobs, the SVC dump title consists of "ADABAS System Dump" plus additional job name, DBID, and timestamp information in the form: Job *jjjjjjjj* DBID *nnnnn* *yyyy/mm/dd hh:mm:ss.hht tmm*. For non-authorized jobs, the title consists of "ADABAS Sys Tx Dump" plus the additional job name, DBID, and timestamp information. Note that the timestamp reflects the time at which the SVC dump request was passed to z/OS -- not the time at which any preceding abend occurred.

The shortest dump processing time occurs when you specify `//SVCDUMP DD DUMMY`. This is because the nucleus only needs to wait for the dump to be captured, not written out to a dump data set. A specification of `//SVCDUMP DD DSN=dsn` will give you a shorter processing time for an APF-authorized job than for a non-APF-authorized job. In both cases, the time taken for dump processing may actually be longer than with `MPMDUMP`.

When an `SVCDUMP DD` statement is included in your JCL, any `MPMDUMP DD` statement is ignored, unless Adabas detects that it is unable to proceed with SVC dump processing.

If an error is encountered while writing the SVC dump, message ADAM78 appears. If dump writing completes successfully, message ADAM79 appears. For more information on these messages, refer to your Adabas messages and codes documentation.

No error message is produced when a dump can only be partially written. You should therefore ensure that sufficient space is available on the dump data set to accommodate the dump.

When an `SVCDUMP DD` statement is included in the JCL, but the SVC dump is unable to complete successfully, dump processing reverts to the standard dump options as specified in the JCL via the `MPMDUMP`, `SYSUDUMP`, `SYSABEND` or `SYSMDUMP DD` statements.

Note:

SVC dump processing might be suppressed due to installation `SLIP` or `DAE` options. If dump processing is still required in this case, the relevant `MPMDUMP`, `SYSUDUMP`, `SYSABEND` or `SYSMDUMP DD` statement should be specified in the JCL in addition to the `SVCDUMP DD` statement.

When the `SVCDUMP DD` statement is omitted from the JCL, existing dump options, specified via the `MPMDUMP`, `SYSUDUMP`, `SYSABEND` or `SYSMDUMP DD` statements, continue to operate as normal.

z/OS JCL Notes

The following notes apply to Adabas startup jobs in various platforms (as described by the notes).

1. This data set is used to provide the Adabas session parameters.
2. This data set is used to print messages produced by the control module ADARUN or the Adabas nucleus.
3. The Adabas Associator and Data Storage data sets. These data sets are mandatory.

n and *nn* represent the number of the Associator and Data Storage data set, respectively.

If more than one data set exists for Associator or Data Storage, a separate statement is required for each. For example, if the Associator consists of two data sets in a z/VSE environment, `DD` statements for `ASSOR1` and `ASSOR2` are required.

If *less* than 10 data sets exist for each, you must use the `ASSOR*` and `DATAR*` `DD` names in the JCL. For example, if, in a z/OS environment, the Associator consists of two data sets and Data Storage consists of three data sets, the following names would be used in the JCL: `DDASSOR1`, `DDASSOR2`, `DDDATAR1`, `DDDATAR2`, and `DDDATAR3`. If 10 or more data sets exist, the first nine must use the `ASSOR*` and `DATAR*` `DD` names in the JCL and the remainder must use the `ASSO*` and `DATA*` `DD` names in the JCL (dropping the "R" in the `DD` names). For example, the tenth Associator data set in a z/OS environment would be identified in the JCL using the name `DDASSO10`, while the third Associator data set in the same JCL would be identified using the name

DDASSOR3.

A maximum of 99 physical extents is now set for Associator and Data Storage data sets. However, your actual real maximum could be less because the extent descriptions of all Associator, Data Storage, and Data Storage Space Table (DSST) extents must fit into the general control blocks (GCBs). For example, on a standard 3390 device type, there could be more than 75 Associator, Data Storage, and DSST extents each (or there could be more of one extent type if there are less for another).

So, the range of Associator DD names that can be used in JCL is DDASSOR1 (ASSOR1 in z/VSE) to DDASSOR9 (ASSOR9 in z/VSE) and DDASSO10 (ASSO10 in z/VSE) to DDASSO99 (ASSO99 in z/VSE). And the range of Data Storage DD names that can be used in JCL is DDDATAR1 (DATAR1 in z/VSE) to DDDATAR9 (DATAR9 in z/VSE) and DDDATA10 (DATA10 in z/VSE) to DDDATA99 (DATA99 in z/VSE).

4. The Adabas Work data sets. The WORKR1 data set is mandatory. If you have Adabas Transaction Manager version 7.5 or later installed, an additional work data set, WORKR4 is also mandatory.

Software AG recommends running the nucleus with DISP=OLD (under z/OS; share not specified for BS2000 and z/VSE) for the WORKR1 data set as a way of preventing two nuclei from writing to the same WORK data set and corrupting the database. This could otherwise happen if the ADARUN parameters FORCE and IGNDIB are improperly used.

Work part 4 of WORKR1 is no longer supported if you have Adabas Transaction Manager Version 7.5 or later installed. Instead, you should use the WORKR4 data set. WORKR4 is used for the same purpose as Work part 4, but it can be used in parallel by all members in a cluster. It is used to store the PET (preliminary end-of transaction) overflow transactions (those that cause a work overflow) of a database or of all members in a multiplex/SMP cluster.

The WORKR4 data set is a container data set that should be allocated and formatted in the normal way (use ADAFRM WORKFRM), using a block size greater than or equal to WORKR1. WORKR4 can have the same or a different device type than WORKR1. It should be at least as large as the cluster's LP parameter of the database or cluster. The smaller WORKR1 Work part 1 is, the larger WORKR4 should be. This is because the nucleus must prevent a work overflow due to incomplete DTP transactions, but the nucleus must keep all PET transactions; they cannot be backed out.

5. If the Adabas Recovery Aid is being used, this logging data set is required.
6. The data protection log data set. This data set is required if the database will be updated during the session and logging of protection information is desired. This data set is not applicable if multiple protection logging is used.

The data protection log may be assigned to tape or disk. A new data set must be used for each Adabas session (DISP=MOD may not be used). See *Adabas Restart and Recovery* for additional information.

7. Multiple (two to eight) data protection log data sets. These data sets are required only if multiple data protection logging is to be in effect for the session.

Multiple data protection logging is activated by the ADARUN NPLOG and PLOGSIZE parameters. The device type of the multiple protection logs is specified with the ADARUN PLOGDEV parameter.

Whenever one of multiple protection log data sets is full, Adabas switches automatically to another data set and notifies the operator through a console message that the log which is full should be copied using the PLCOPY function of the ADARES utility. This copy procedure may also be implemented using the user exit 12 facility as described in the User Exits documentation.

If no command logging is to be performed, this data set may be omitted.

8. The command log data set. This data set is required if command logging is to be performed during the session. Command logging is activated by the ADARUN LOGGING parameter.
9. Multiple (two to eight) command log data sets. These data sets are required only if multiple command logging is to be in effect for the session.

Multiple command logging is activated by the ADARUN NCLOG and CLOGSIZE parameters. The device type of the multiple command log data sets is specified with the ADARUN CLOGDEV parameter.

Whenever one of multiple command log data sets is full, Adabas switches automatically to another data set and notifies the operator through a console message that the log which is full should be copied using the CLCOPY function of the ADARES utility. This copy procedure may also be implemented using the user exit 12 facility as described in the Adabas User Exits documentation.

10. The Entire Conversion Services (ECS) objects data set is required for universal encoding support (UES) in all operating environments.
11. This data set is used to take an Adabas dump including SVC, ID-TABLE and allocated CSA in the event that an abnormal termination occurs.

The line count in the JCL must be set appropriately; otherwise, the dump cannot be printed in its entirety.

12. The z/OS SVC dump facility can be used to speed up nucleus dump processing if an MPMDUMP is too slow for users with very large buffer pools. We recommend that you specify both an MPMDUMP and an SVCDUMP DD statement in your JCL to ensure that one of the dumps is produced when needed. If an SVCDUMP DD statement is included, an SVC dump is created if possible and the MPMDUMP DD statement is ignored. Should problems arise during processing of the SVC dump, an MPMDUMP will be taken, but only if the MPMDUMP DD statement is also specified in the JCL. For more information, read *Using the z/OS SVC Dump Facility*.
13. This data set is used under z/OS to take an Adabas dump (SMGT,DUMP) or snap dump (SMGT,SNAP) when using the error handling and message buffering facility.
14. This data set is used when the PIN output is to be directed to DDTRACE1 rather than DDPRINT as specified by the user in ADASMXIT when using the error handling facility. If DDTRACE1 has not been specified in the JCL and PRINTDD is set to "NO" in ADASMXIT, output will be lost. The PIN output will not be written to DDPRINT unless "YES" is specified.
15. The TZINFO data set is required when defining fields with the TZ option and using the user session open parameter TZ. This data set is a library or partitioned data set containing information on local time offsets and daylight savings time offsets and its transition times.

z/VSE Session

The following table contains all data sets used when executing an Adabas session under z/VSE. SYSnnn means that any programmer logical unit may be used.

Data Set	File Name	Logical Unit	Storage Medium	Additional Information
ADARUN parameters	none CARD CARD	SYSRDR SYS000 SYSnnn	reader tape disk	note 1
ADARUN/Adabas messages	none	SYSLST	printer	note 2
Associator	ASSOR <i>n</i> or ASSO <i>nn</i>	SYSnnn	disk	note 3
Data Storage	DATAR <i>n</i> or DATA <i>nn</i>	SYSnnn	disk	note 3
Work	WORKR1	SYSnnn	disk	note 4
Recovery Aid log	RLOGR1	SYSnnn	disk	note 5
Protection log multiple log 1 multiple log 2	SIBA PLOGR1 PLOGR2	SYS014 SYSnnn SYSnnn SYSnnn	tape disk disk disk	note 6 note 7 note 7
Command log multiple log 1 multiple log 2	LOG CLOGR1 CLOGR2	SYS012 SYSnnn SYSnnn SYSnnn	tape disk disk disk	note 8 note 9 note 9
ECS encoding objects	DDECSOJ	SYS020 *	tape disk	note 10
Time zone files	TZINFO	not required	disk	note 11

- Example of Adabas Session Job Control (z/VSE)
- JCL Required for UES Support (z/VSE)
- JCL Required for Time Zone Support (z/VSE)
- z/VSE Job Notes

Example of Adabas Session Job Control (z/VSE)

This job includes multiple protection logging, multiple command logging, and Recovery Aid logging:

```
// ASSGN SYS031,ddd,VOL=ADA001,SHR
// ASSGN SYS032,ddd,VOL=ADA002,SHR
// ASSGN SYS033,ddd,DISK,VOL=ADA003,SHR
// ASSGN SYS034,ddd,VOL=ADA004,SHR
// DLBL ASSOR1,'EXAMPLE.ADAyyyyy.ASSOR1',2099/365,DA
// EXTENT SYS031,ADA001,,15,1500
// DLBL DATAR1,'EXAMPLE.ADAyyyyy.DATAR1',2099/365,DA
// EXTENT SYS032,ADA002,,15,3000
// DLBL WORKR1,'EXAMPLE.ADAyyyyy.WORKR1',2099/365,DA
// EXTENT SYS033,ADA003,,15,600
// DLBL PLOGR1,'EXAMPLE.ADAyyyyy.PLOGR1',2099/365,DA
// EXTENT SYS034,ADA004,,15,600
// DLBL PLOGR2,'EXAMPLE.ADAyyyyy.PLOGR2',2099/365,DA
// EXTENT SYS034,ADA004,,615,600
// DLBL CLOGR1,'EXAMPLE.ADAyyyyy.CLOGR1',2099/365,DA
// EXTENT SYS034,ADA004,,1215,600
// DLBL CLOGR2,'EXAMPLE.ADAyyyyy.CLOGR2',2099/365,DA
// EXTENT SYS034,ADA004,,1815,600
// DLBL RLOGR1,'EXAMPLE.ADAyyyyy.RLOGR1',2099/365,DA
// EXTENT SYS034,ADA004,,1300,600
// DLBL ADALIB,'ADABAS.ADAvrs.LOADLIB',2099/365
// EXTENT ,ADADSK
ADARUN PROG=ADANUC,SVC=xxx,DEVICE=ddd,DB=yyyyy
ADARUN LBP=600000
ADARUN LWP=320000
ADARUN LS=80000
ADARUN LP=400
ADARUN NAB=24
ADARUN NC=1000
ADARUN NH=2000
ADARUN NU=100
ADARUN TNAE=180,TNAA=180,TNAX=600,TT=90
ADARUN NPLOG=2,PLOGSIZE=1800,PLOGDEV=ddd
ADARUN NCLOG=2,CLOGSIZE=1800,CLOGDEV=ddd
/*
/*
```

JCL Required for UES Support (z/VSE)

The following additional JCL is required for universal encoding support (UES):

```
// ASSGN SYS020,disk,VOL=volume,SHR
// DLBL DDECSOJ,'ADABAS.ADAvrs.ECSLIB'
// EXTENT SYS020
```

JCL Required for Time Zone Support (z/VSE)

The following additional JCL is required for universal encoding support (UES):

```
// DLBL TZINFO,'/ADALIB/ADAvrSTZ/.TIMEZONE'
```

See note 11 below for more information.

z/VSE Job Notes

The following notes apply to Adabas startup jobs in various platforms (as described by the notes).

1. This data set is used to provide the Adabas session parameters.
2. This data set is used to print messages produced by the control module ADARUN or the Adabas nucleus.
3. The Adabas Associator and Data Storage data sets. These data sets are mandatory.

n and *nn* represent the number of the Associator and Data Storage data set, respectively.

If more than one data set exists for Associator or Data Storage, a separate statement is required for each. For example, if the Associator consists of two data sets in a z/VSE environment, DD statements for ASSOR1 and ASSOR2 are required.

If *less* than 10 data sets exist for each, you must use the ASSOR* and DATAR* DD names in the JCL. For example, if, in a z/OS environment, the Associator consists of two data sets and Data Storage consists of three data sets, the following names would be used in the JCL: DDASSOR1, DDASSOR2, DDDATAR1, DDDATAR2, and DDDATAR3. If 10 or more data sets exist, the first nine must use the ASSOR* and DATAR* DD names in the JCL and the remainder must use the ASSO* and DATA* DD names in the JCL (dropping the "R" in the DD names). For example, the tenth Associator data set in a z/OS environment would be identified in the JCL using the name DDASSO10, while the third Associator data set in the same JCL would be identified using the name DDASSOR3.

A maximum of 99 physical extents is now set for Associator and Data Storage data sets. However, your actual real maximum could be less because the extent descriptions of all Associator, Data Storage, and Data Storage Space Table (DSST) extents must fit into the general control blocks (GCBs). For example, on a standard 3390 device type, there could be more than 75 Associator, Data Storage, and DSST extents each (or there could be more of one extent type if there are less for another).

So, the range of Associator DD names that can be used in JCL is DDASSOR1 (ASSOR1 in z/VSE) to DDASSOR9 (ASSOR9 in z/VSE) and DDASSO10 (ASSO10 in z/VSE) to DDASSO99 (ASSO99 in z/VSE). And the range of Data Storage DD names that can be used in JCL is DDDATAR1 (DATAR1 in z/VSE) to DDDATAR9 (DATAR9 in z/VSE) and DDDATA10 (DATA10 in z/VSE) to DDDATA99 (DATA99 in z/VSE).

4. The Adabas Work data sets. The WORKR1 data set is mandatory. If you have Adabas Transaction Manager version 7.5 or later installed, an additional work data set, WORKR4 is also mandatory.

Software AG recommends running the nucleus with DISP=OLD (under z/OS; share not specified for BS2000 and z/VSE) for the WORKR1 data set as a way of preventing two nuclei from writing to the same WORK data set and corrupting the database. This could otherwise happen if the ADARUN parameters FORCE and IGNDIB are improperly used.

Work part 4 of WORKR1 is no longer supported if you have Adabas Transaction Manager Version 7.5 or later installed. Instead, you should use the WORKR4 data set. WORKR4 is used for the same purpose as Work part 4, but it can be used in parallel by all members in a cluster. It is used to store the PET (preliminary end-of transaction) overflow transactions (those that cause a work overflow) of a database or of all members in a multiplex/SMP cluster.

The WORKR4 data set is a container data set that should be allocated and formatted in the normal way (use ADAFRM WORKFRM), using a block size greater than or equal to WORKR1. WORKR4 can have the same or a different device type than WORKR1. It should be at least as large as the cluster's LP parameter of the database or cluster. The smaller WORKR1 Work part 1 is, the larger WORKR4 should be. This is because the nucleus must prevent a work overflow due to incomplete DTP transactions, but the nucleus must keep all PET transactions; they cannot be backed out.

5. If the Adabas Recovery Aid is being used, this logging data set is required.
6. The data protection log data set. This data set is required if the database will be updated during the session and logging of protection information is desired. This data set is not applicable if multiple protection logging is used.

The data protection log may be assigned to tape or disk. A new data set must be used for each Adabas session (DISP=MOD may not be used). See *Adabas Restart and Recovery* for additional information.

7. Multiple (two to eight) data protection log data sets. These data sets are required only if multiple data protection logging is to be in effect for the session.

Multiple data protection logging is activated by the ADARUN NPLOG and PLOGSIZE parameters. The device type of the multiple protection logs is specified with the ADARUN PLOGDEV parameter.

Whenever one of multiple protection log data sets is full, Adabas switches automatically to another data set and notifies the operator through a console message that the log which is full should be copied using the PLCOPY function of the ADARES utility. This copy procedure may also be implemented using the user exit 12 facility as described in the User Exits documentation.

If no command logging is to be performed, this data set may be omitted.

8. The command log data set. This data set is required if command logging is to be performed during the session. Command logging is activated by the ADARUN LOGGING parameter.
9. Multiple (two to eight) command log data sets. These data sets are required only if multiple command logging is to be in effect for the session.

Multiple command logging is activated by the ADARUN NCLOG and CLOGSIZE parameters. The device type of the multiple command log data sets is specified with the ADARUN CLOGDEV parameter.

Whenever one of multiple command log data sets is full, Adabas switches automatically to another data set and notifies the operator through a console message that the log which is full should be copied using the CLCOPY function of the ADARES utility. This copy procedure may also be implemented using the user exit 12 facility as described in the Adabas User Exits documentation.

10. The Entire Conversion Services (ECS) objects data set is required for universal encoding support (UES).
11. The TZINFO DLBL is required when defining fields with the TZ option and using the user session open parameter TZ. This data set is a library containing information on local time offsets and daylight savings time offsets and its transition times. The TZINFO DLBL defines the library and sublibrary to be used for accessing time zone data. The physical name associated with the DLBL must be in the form `'library/sublib/.TIMEZONE'`. In the example above, time zone data is accessed from the ADALIB sublibrary "ADAvrsTZ". In this case, a label for ADALIB must be available to the Adabas

session job.

Single-User Mode

Although the normal mode of operation is multiuser mode, it is also possible to execute Adabas together with a user program or Adabas utility in the same region.

For single-user mode, you must include the Adabas nucleus job control that you use along with the job control for the utility or user program.

The Adabas prefetch option cannot be used in single-user mode; however, single-user mode *must* be used when running a read-only nucleus and an update nucleus simultaneously.

User Profile

Some information within an Adabas database is user-related and must be retained from session to session. One such kind of information is ET data records; another is the priority value assigned to a user.

A set of user-related information can be stored in a profile table. The values stored in this table are read at OPEN time and assigned to the user. The direct call user must OPEN the Adabas session with the proper call; that is, as an ID user with an ETID in the Additions 1 field of the Adabas control block. For Natural users, the profile table is identified by the Natural ETID.

The associated fields are user-related timeout and threshold values, and the OWNERID for multiclient fields. One record per user is stored. The profile table is maintained using Adabas Online System.

Profile Table Values

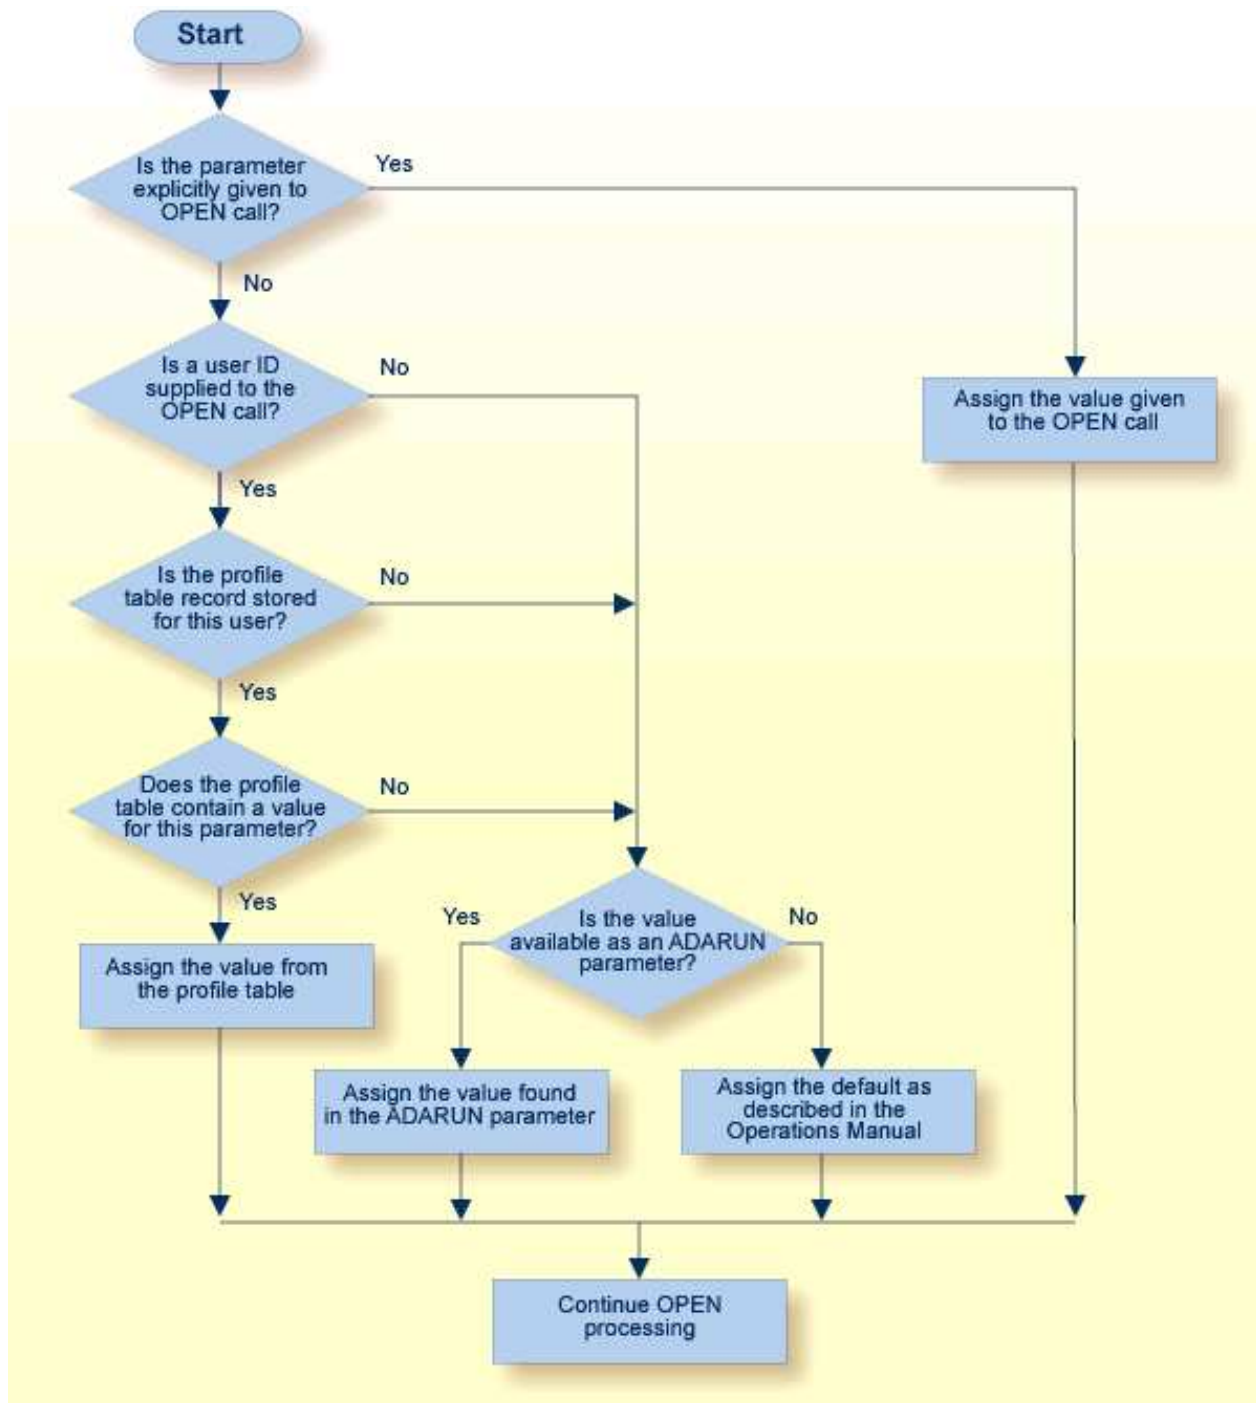
The user-related values shown below are currently stored in the profile table.

Value	Description
PRIORITY	User's priority (0-255)
TNAA*	Access user non-activity time
TNAE*	ET user non-activity time
TNAX*	EXU/EXF user non-activity time
TT*	Transaction time threshold
TLSCMD*	Sx command threshold
NSISN*	Maximum number of ISNs per TBI element
NSISNHQ*	Maximum number of records held by user
NQCID*	Maximum number of active command IDs per user
OWNERID	Owner ID for multiclient file access

* The decision sequence for determining the values for a user at the time of an open call is shown in *Managing the User Profile*.

Managing the User Profile

Adabas Online System (AOS) must be used to maintain the profile table. See the Adabas Online System documentation for detailed information about managing the profile table.



Profile Table Decision Flow