‚ software ᴬᴳ

Adabas Cluster Services

Installation

Version 8.5.4

April 2020

ADABAS & NATURAL

Table of Contents

Installation

This documentation describes the installation process for Adabas Cluster Services.

It is important to review the entire installation procedure before starting the physical installation.

Depending on which components you actually install, it may be useful to group certain installation activities together, even though they may not be in the same installation section in this section. For example, if you are installing the Entire Net-Work XCF line driver, modifications to the CFRM policy for both the line driver and for parallel sysplex cache and lock structures can be made in a single step rather than in two separate steps.

Refer to *Adabas Cluster Services Release Notes* for specific information about late changes to this section.

Using System Maintenance Aid

If you use Software AG's System Maintenance Aid (SMA), refer to the *System Maintenance Aid* documentation for information about the installation process.

The installation procedures described in this section correspond to the jobs that SMA creates to install the product.

If you do not use SMA, you can modify and use the sample JCL provided to unload the Adabas Cluster Services and Entire Net-Work libraries from the installation tape.

The Adabas Cluster Services Installation documentation is organized in the following topics:

System Requirements
What to Install
Overview of the Installation Procedure
Modifying the z/OS Common Storage (CSA) Key

1 Conventions

Notation vrs or vr: When used in this documentation, the notation vrs or vr stands for the relevant version, release, and system maintenance level numbers. For further information on product versions, see *version* in the *Glossary*.

2 About this Documentation

Document Conventions

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

Online Information and Support

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at **https://documenta-tion.softwareag.com**.

Software AG Empower Product Support Website

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

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

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

To submit feature/enhancement requests, get information about product availability, and download products, go to **Products**.

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the **Knowledge Center**.

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

Software AG Tech Community

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

- Access product documentation, if you have Tech Community credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.

- Access articles, code samples, demos, and tutorials.

- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.

- Link to external websites that discuss open standards and web technology.

Data Protection

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

3 System Requirements

Adabas Cluster Services requires a parallel sysplex environment.

> **Note:** Fujitsu Technology Systems' HIPLEX is not supported.

This chapter describes the operating environment and the other Software AG software required to install and run Adabas Cluster Services successfully.

For information regarding Software AG product compatibility with IBM platforms and any IBM requirements for Software AG products, please review the *Product Compatibility for IBM Platforms* web page.

Supported Operating System Platforms

Software AG generally provides support for the operating system platform versions supported by their respective manufacturers; when an operating system platform provider stops supporting a version of an operating system, Software AG will stop supporting that version.

For information regarding Software AG product compatibility with IBM platforms and any IBM requirements for Software AG products, please review the *Product Compatibility for IBM Platforms* web page.

Before attempting to install this product, ensure that your host operating system is at the minimum required level. For information on the operating system platform versions supported by Software AG products, complete the following steps.

1. Access Software AG's Empower web site at *https://empower.softwareag.com*.

2. Log into Empower. Once you have logged in, you can expand **Products & Documentation** in the left menu of the web page and select **Product Version Availability** to access the Product Version Availability screen.

3. Use the fields on the top of this screen to filter its results for your Software AG product. When you click the **Search** button, the supported Software AG products that meet the filter criteria are listed in the table below the filter criteria.

   This list provides, by supported operating system platform:

   - the Software AG general availability (GA) date of the Software AG product;
   - the date the operating system platform is scheduled for retirement (OS Retirement);
   - the Software AG end-of-maintenance (EOM) date for the product; and
   - the Software AG end-of-sustained-support (EOSS) date for the product.

> **Note:** Although it may be technically possible to run a new version of your Software AG product on an older operating system, Software AG cannot continue to support operating system versions that are no longer supported by the system's provider. If you have questions

about support, or if you plan to install this product on a release, version, or type of operating system other than one listed on the Product Version Availability screen described above, consult Software AG technical support to determine whether support is possible, and under what circumstances.

Required Operating Environment

This section describes the operating environment required to install and run Adabas Cluster Services successfully.

- Operating System Requirements
- Sysplex System Requirements
- Coupling Facility
- XCF and XES Services

Operating System Requirements

Adabas Cluster Services can be installed on z/OS operating systems. For information regarding Software AG product compatibility with IBM platforms and any IBM requirements for Software AG products, please review the *Product Compatibility for IBM Platforms* web page.

One or more z/OS operating system images defined as members of a parallel sysplex are required that are:

- connected with an IBM Sysplex Timer™, and

- connected to a coupling facility, and

- running z/OS (any version supported by IBM)

Before attempting to install Adabas Cluster Services, ensure that the host operating system is at the minimum required level. For information on the platform versions supported by Software AG products, access the Software AG web site at *http://www.softwareag.com/corporate/products/bis/platforms/default.asp*.

Software AG provides support for the operating system versions supported by their respective manufacturers. Generally, when an operating system provider stops supporting a version of an operating system, Software AG will stop supporting that operating system version.

To determine the particular platforms supported by this release of Adabas Cluster Services, review the product version availability information, available on Software AG's Empower web site: *https://empower.softwareag.com*. Log into Empower. Once you have logged in, you can expand **Products** in the left menu of the web page and select **Product Version Availability** to access the Product Version Availability application. This application allows you to review support information for specific products and releases.

Although it may be technically possible to run a new version of Adabas or Adabas Cluster Services on an old operating system, Software AG cannot continue to support operating system versions that are no longer supported by the system's provider.

If you have questions about support, or if you plan to install Adabas or Adabas Cluster Services on a release, version, or type of operating system other than those included at the link above, consult Software AG technical support to determine whether support is possible, and under what circumstances.

Sysplex System Requirements

An operational IBM parallel sysplex environment is required, if Adabas Cluster Services is to provide 24-by-7 availability for an indefinite period of time. This includes the necessary hardware to support the configuration, the necessary links and channels for connectivity, and the power to provide the required performance during peak production periods. The knowledge and experience required to implement and maintain such an environment should be acquired from IBM.

Adabas Cluster Services uses:

- a coupling facility cache structure
- a coupling facility lock structure
- multiple protection logs, command logs, and Work data sets on shared DASD

For testing purposes only, the hardware, operating system, and software environment may be totally different. You may simulate the sysplex cluster using a series of LPARs, including the coupling facility itself. However, under these circumstances, system performance cannot be accurately measured.

For general information regarding Software AG product compatibility with other platforms and their requirements for Software AG products, visit Software AG's *Hardware Supported* web page.

Coupling Facility

The coupling facility configuration must have sufficient resources to support the cache and lock structures used by Adabas Cluster Services. You can duplex these structures, if at least two coupling facilities are available. Adabas Cluster Services can coexist without restriction with other coupling facility users such as GRS Star.

Each Adabas Cluster Services database uses one cache structure and one lock structure. For each, you must define all of the following in the coupling facility resource management (CFRM) policy :

- the name of the structure (NAME)
- the structure size (SIZE or INITSIZE)

- the preference list of one or more coupling facilities where the structure should be allocated (PREFLIST).

You can also allow duplexing (DUPLEX) or dynamic reallocation (ALLOWREALLOCATE) of either structure.

XCF and XES Services

Adabas Cluster Services uses parallel sysplex XCF (Cross-System Coupling Facility) services for intracluster communication and XES (Cross-System Extended Services) for intracluster data sharing. Entire Net-Work may also use XCF services. XCF and XES must be available and functioning at all times for Adabas Cluster Services to operate smoothly.

Performance Recommendations

The inherent necessity to synchronize and share data across system boundaries makes it a challenge for Adabas Cluster Services to perform as well or better than the base Adabas, especially when measured by CPU consumption. Adabas Cluster Services performance is best when the following recommendations are applied:

- The coupling facility processors (CFCPs) should be at least as fast as the central processors running Adabas Cluster Services and they should have sufficient capacity to avoid delays in the execution of coupling facility requests.

- The coupling facility link speed should match the coupling facility processor speed.

- The buffer efficiency in a cluster nucleus when functioning as a non-cluster nucleus should be equal to or greater than 50.

- The buffer pool size (LBP) used for a cluster nucleus should not be lower than when the nucleus is functioning as a non-cluster nucleus.

- The rate of throwbacks due to ISN contention should not be greater than two percent (2%).

- The size of the cache and lock structures specified in the coupling facility resource management (CFRM) policy should be appropriate. For more information, read *Optimizing Lock and Cache Structures in the Coupling Facility* , in *Adabas Cluster Services Operations*.

  You may use the cache and lock structure size calculators provided for Adabas Cluster Services in Adabas Online System to estimate adequate structure sizes. For more information about these calculators, read *Estimating Sizes for the Cache Structure in a Cluster Environment* and *Estimating Sizes for the Lock Structure in a Cluster Environment*, in *Adabas Cluster Services Operations*.

Performance measurements of Adabas Cluster Services in a test environment may yield inconclusive results if these recommendations are not followed.

Prerequisite Software AG Software

This section describes the Software AG product and SVC (router) requirements for Adabas Cluster Services.

- Adabas Requirements
- Adabas Online System (AOS) Requirements
- Entire Net-Work Requirements
- Router (SVC) Requirements and Support

Adabas Requirements

Adabas Cluster Services version 8.5 SP1 requires Adabas version 8.5 SP1, with appropriate Adabas 8.5 SP1 zaps applied.

Please be sure to check the Knowledge Center in Software AG's Empower (*https://empower.softwareag.com*) web site for additional cluster-related Adabas zaps when installing Adabas Cluster Services. The complete Adabas zaps may not be supplied with your Adabas Cluster Services product.

The appropriate version of Adabas must be installed and operational on your systems before you install Adabas Cluster Services.

- Use the ADADEF utility to define a new Adabas database.
- Use the ADACNV utility to convert an existing database from an earlier version.

For more information, refer to your *Adabas Utilities Manual* documentation.

All Adabas nuclei in the same cluster must run with the same version, release, and modification levels of Adabas and of Adabas Cluster Services. For example, if one nucleus in a cluster runs with Adabas Cluster Services 8.5 SP1 and Adabas 8.5 SP1, the other nuclei in the cluster must also run with Adabas Cluster Services 8.5 SP1 and Adabas 8.5 SP1; they cannot run with different versions.

The following table identifies the Software AG cluster products and the Adabas product combinations with which they are compatible:

Software AG Cluster Product		Compatible Adabas	
Name	Library Version	Library Version	SVC Version
Adabas Cluster Services	8.5 SP1	8.5 SP1	8.5 SP1
	8.4 SP1	8.4 SP1	8.4 SP1

Adabas Online System (AOS) Requirements

If you install Adabas Online System (AOS), either the demo version delivered with Adabas 8.5 SP1 or the version 8.5 SP1 selectable unit is required.

Entire Net-Work Requirements

Entire Net-Work for mainframes is required for communication between users and non-local Adabas Cluster Services nuclei. The essential Entire Net-Work components, including the XCF line driver, are part of the Adabas Cluster Services package.

Adabas Cluster Services version 8.5 SP1 requires Entire Net-Work version 6.3 or above. The necessary components of Entire Net-Work are delivered with Adabas Cluster Services. Please check **Empower** for Entire Net-Work zaps that must be applied to Entire Net-Work 6.3.

Entire Net-Work must be installed on each operating system image in the parallel sysplex where either Adabas cluster nuclei or users are located. Each system with Entire Net-Work installed becomes a *node* in the network. Each node's adjacent links to other nodes are defined by name and driver type.

All nodes in your network must be running Entire Net-Work 6.3.

Router (SVC) Requirements and Support

Adabas Cluster Services requires a compatible version of the Adabas router (ADASVC), which contains the Adabas Cluster Services component SVCCLU. The router must be installed on each operating system image in the parallel sysplex where either Adabas cluster nuclei or users are located.

The Adabas router (SVC) provided with Adabas 8.5 SP1 supports clusters running Adabas Cluster Services 8.4 SP1 as well as clusters running Adabas Cluster Services 8.5.

For more information about Adabas router compatibility, review the table showing valid Adabas and cluster product combinations in *Adabas Requirements*, elsewhere in this section.

4 What to Install

Adabas Cluster Services requires Entire Net-Work to provide communication across operating system images between users (including utilities and Adabas Online System) and the nuclei in the Adabas cluster as well as among ADACOM jobs. Communication among the cluster nuclei is performed via XCF.

> **Note:** ADACOM and several Adabas cluster nuclei could run within a single operating system image without Entire Net-Work, but there could then be no application programs for that database in any other image.

Entire Net-Work Limited (WDXvrs) Library

If you are *not* currently running Entire Net-Work, you must install the Entire Net-Work limited (WDX*vrs*) library supplied on the Adabas Cluster Services installation tape. This library includes Entire Net-Work's XCF line driver (also known as the *Entire Net-Work XCF Option*).

Adabas Limited (WALvrs) Libraries

The Adabas limited (WAL*vrs*) libraries provide the Adabas components necessary to support the Entire Net-Work connection on a node that participates in an Adabas sysplex cluster (that is, has cluster users) but does not have Adabas installed.

Using an Existing Entire Net-Work System

If you are currently running Entire Net-Work, you can use your existing system as long as the following conditions are met:

- An Entire Net-Work node is running on each operating system image that will also be running a cluster nucleus and/or ADACOM.

- The Entire Net-Work nodes supporting the cluster nuclei are all interconnected using one of the currently supported Entire Net-Work line drivers: XCF, VTAM, CTC, or TCP/IP.

- The Entire Net-Work nodes use the same SVCCLU linked with the ADASVC that is used for the Adabas Cluster Services components they support.

If these conditions are not met, you cannot use your existing Entire Net-Work configuration. You must then use the Entire Net-Work limited library supplied on the Adabas Cluster Services installation tape.

Using the XCF Line Driver with an Existing Entire Net-Work

You may incorporate the XCF line driver (also known as the *Entire Net-Work XCF Option*) into your existing Entire Net-Work configuration. The XCF line driver provides performance comparable to the CTC line driver.

When using the XCF line driver, the physical connection mechanism between Entire Net-Work nodes is transparent to Entire Net-Work. The connections may be over CTC links, XCF links, or both.

The connection mechanism/methods are defined by the system programmer using the COUPLE*xx* SYS1.PARMLIB member and the coupling facility resource management (CFRM) policy.

Order the STEPLIB Concatenation

If you are installing the XCF line driver into an existing set of Entire Net-Work nodes, Software AG recommends that you concatenate the Entire Net-Work limited libraries from the Adabas Cluster Services tape in front of the existing Entire Net-Work library in the STEPLIB sequence.

Concatenate the Adabas Cluster Services library ahead of the Adabas load library.

5 Overview of the Installation Procedure

This section provides information about installing Adabas Cluster Services in the operating system environment.

The installation procedure must be repeated for *each* operating system image (that is, LPAR) that will participate in the system.

> ⚠️ **Important:** Before proceeding with the installation, ensure that the prerequisite environment has been established. Review the earlier sections of this documentation for detailed information.

The procedure corresponds to the jobs that SMA creates to install the product. If you do not use SMA, you can modify and use the sample JCL described in *Unload the Installation Libraries to Disk* to unload the libraries from the installation tape.

Data Sets Delivered

The following data sets are delivered with Adabas Cluster Services on z/OS platforms:

Data Set Name	Contains
ADA*vrs*.LOAD	Adabas Cluster Services load modules. Packaged in the Adabas load library containing the load modules ADAXCF, ADAXEC, ADAXEL; are now executable with the distributed ALS license module.

Step 1: Unload the Installation Libraries to Disk

- Allocate DASD
- Copy the Contents of the Tape to Disk
- Copy the Adabas Limited (WALvrs) Libraries, If Required

Allocate DASD

If you need to install them, the minimum 3390 disk space requirements for the Adabas limited (WAL) load library is five cylinders; the WAL source library requires two cylinders.

The WDX limited load library for Entire Net-Work requires one cylinder; the WCP source library for Entire Net-Work requires one cylinder.

Copy the Contents of the Tape to Disk

Copy the data sets from the supplied installation medium to your disk before you perform the individual installation procedure for each component to be installed.

The way you copy the data sets depends on the installation method and the medium used:

- If you use System Maintenance Aid (SMA), refer to the copy job instructions provided in the *System Maintenance Aid* documentation.

- If you are not using SMA and want to copy the data sets from CD-ROM, refer to the README.TXT file on the CD-ROM.

- If you are not using SMA and want to copy the data sets from tape, follow the instructions in this section.

This section explains how to copy all data sets from tape to disk.

- Step 1: Copy Data Set COPY.JOB from Tape to Disk
- Step 2: Modify hilev.COPY.JOB on Your Disk
- Step 3: Submit COPY.JOB

Step 1: Copy Data Set COPY.JOB from Tape to Disk

- Modify the following sample job according to your requirements:

```
//SAGTAPE JOB SAG,CLASS=1,MSGCLASS=X
//* ------------------------------
//COPY EXEC PGM=IEBGENER
//SYSUT1 DD DSN=COPY.JOB,
// DISP=(OLD,PASS),
// UNIT=(CASS,,DEFER),
// VOL=(,RETAIN,SER=tape-volser),
// LABEL=(2,SL)
//SYSUT2 DD DSN=hilev.COPY.JOB,
// DISP=(NEW,CATLG,DELETE),
// UNIT=3390,VOL=SER=disk-volser,
// SPACE=(TRK,(1,1),RLSE),
// DCB=*.SYSUT1
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//
```

where:

tape-volser is the VOLSER of the tape, for example: T12345,
hilev is a valid high-level qualifier, and
disk-volser is the VOLSER of the disk.

- Execute the job to copy the data set COPY.JOB to your disk.

Step 2: Modify hilev.COPY.JOB on Your Disk

■ Modify *hilev*.COPY.JOB according to your requirements:

Set EXPDT to a valid expiration date, for example, 99365.

Set HILEV to a valid high-level qualifier, for example, USERLIB.

Set LOCATION to a storage location, for example, STORCLAS=ABC or UNIT=3390,VOL=SER=USR123.

Step 3: Submit COPY.JOB

■ Execute *hilev*.COPY.JOB to copy single, multiple, or all data sets to your disk.

Copy the Adabas Limited (WALvrs) Libraries, If Required

Entire Net-Work requires some Adabas modules. For systems where Adabas is not installed, these modules are supplied on the installation tape in the Adabas limited libraries (WAL*vrs*). See the *Software AG Product Delivery Report* for specific information regarding these limited libraries.

Software AG recommends that you concatenate the Adabas limited (WAL) libraries with the Entire Net-Work (WDX) load library first in the sequence.

Step 2: APF-Authorize All Load Libraries

Ensure that Adabas Cluster Services runs authorized.

To run authorized, the Adabas, and limited Entire Net-Work load libraries and all other load libraries in the STEPLIB concatenation must be APF-authorized.

Ensure that all load libraries referenced in the STEPLIB concatenation for your Entire Net-Work, ADACOM, and Adabas cluster nuclei startup procedures are defined to the operating system as authorized libraries. If this is not done, the Entire Net-Work, ADACOM, or the Adabas cluster nuclei will not initialize and may abnormally terminate, usually with an ABENDS047 or ABENDS306.

Either copy the base Adabas modules to an existing APF-authorized library or APF-authorize the Adabas load library unloaded from the base Adabas installation tape.

Step 3: Install the WALvrs Libraries, if Needed

The Adabas router (ADASVC), which contains the SVCCLU component used by Adabas Cluster Services, can and should also be used for Entire Net-Work. It must be installed on any node where Entire Net-Work is to be run, even if no servers are running on that node. Use the WAL*vrs* libraries for such nodes.

Link Routines

Adabas client application programs need an Adabas link routine module suitable for the environment in which they run. For example, batch programs should be linked with ADAUSER, which loads current versions of ADALNK and supporting modules from the Adabas load library. TP environments usually require specialized link routine modules that may require an additional license.

Software AG recommends that link routines have the same version level as Adabas in a sysplex cluster environment.

Step 4: Modify COUPLExx and the CFRM Policy

Make any needed additions and modifications to the COUPLE*xx* SYS1.PARMLIB member and the sysplex CFRM policy. See the *Performance and Tuning* documentation for sizing recommendations.

For XCF Groups

A separate XCF group is activated for each

■ Adabas sysplex cluster;

■ Entire Net-Work instance using an XCF line driver; and

■ active lock structure; each Adabas cluster uses one lock structure.

Ensure that the COUPLE*xx* data sets are formatted with sufficient capacity for the necessary XCF groups.

No changes to the COUPLE*xx* member are required. If values are not specified, the system defaults are used. However, rather than use the default values for CLASS, GROUP, and STRUCTURE, Software AG recommends that you add specific values to the COUPLE*xx* member in SYS1.PARMLIB for use by the Entire Net-Work XCF line driver, as follows:

1. Add a CLASSDEF statement with CLASS(ENTIRE) and GROUP(NETWORK) specified.

2. Add a PATHIN statement with STRNAME(IXC_NETWORK) specified.

3. Add a PATHOUT statement with STRNAME(IXC_NETWORK) and CLASS(ENTIRE) specified.

A sample section from the COUPLE*xx* member in SYS1.PARMLIB is as follows:

```
COUPLE SYSPLEX(USRSPLX1)
    PCOUPLE(SYS1.SYSPLEX.USRSPLX1.CDS01)
    ACOUPLE(SYS1.SYSPLEX.USRSPLX1.CDS02)
    INTERVAL(120)
CLASSDEF
    CLASS(ENTIRE)
    CLASSLEN(8192)
        GROUP(NETWORK)
    MAXMSG(4096)
LOCALMSG
    MAXMSG(4096)
    CLASS(ENTIRE)
PATHIN STRNAME(IXC_DEFAULT ,IXC_NETWORK)
PATHOUT STRNAME(IXC_DEFAULT)
PATHOUT STRNAME(IXC_NETWORK) CLASS(ENTIRE)
```

For List, Cache, and Lock Structures

To define the optional list structure for Entire Net-Work and the required cache and lock structures for each ADABAS nucleus cluster, you will need to run a job similar to the following to update the CFRM administrative policy data in the COUPLE data set:

```
//STEP20   EXEC PGM=IXCMIAPU
//SYSPRINT DD   SYSOUT=*
//SYSABEND DD   SYSOUT=*
//SYSIN    DD   *
    DATA TYPE(CFRM) REPORT(YES)
    DEFINE POLICY NAME(POLICY1) REPLACE(YES)
    CF NAME(CF)
        TYPE(009674)
            MFG(IBM)
            PLANT(02)
            SEQUENCE (000000040101)
            PARTITION(1)
            CPCID(00)
            DUMPSPACE(2000)

            .
            .
        STRUCTURE NAME(IXC_DEFAULT)
            SIZE(1000)
            PREFLIST(CF)
        STRUCTURE NAME(ADA_CACHE1)
            SIZE(48000)
            PREFLIST(CF)

            .
```

```
                .
      STRUCTURE NAME(ADA_LOCK1)
          SIZE(16000)
          PREFLIST(CF)
          .
          .
      STRUCTURE NAME(IXC_NETWORK)
          SIZE(10000)
          PREFLIST(CF)
          .
          .
```

The CACHE structure must be large enough to retain

■ all changed blocks between buffer flushes; and

■ directory elements for all blocks in all buffer pools.

For information on cache structure sizes, read *Optimizing Lock and Cache Structures in the Coupling Facility*, in *Adabas Cluster Services Operations*.

System-Managed Rebuild and Duplexing Rebuild

Adabas Cluster Services supports rebuild processes and duplexing rebuild processes for the cache and lock structures that it uses. These processes are managed by the operating system.

■ *System-managed rebuild* dynamically rebuilds a coupling facility structure in the same or another coupling facility, while an active application (in this case, Adabas Cluster Services) may be using the structure. It can be used to reduce *planned* outage time for applications using the coupling facility.

The option of the SETXCF operator command initiating this system-managed process is SETXCF START,REBUILD.

■ *System-managed duplexing rebuild* dynamically duplicates a coupling facility structure into another coupling facility, while an application may be using the structure. It can be used to reduce *unplanned* outage time for applications using the coupling facility.

The option of the SETXCF operator command initiating this system-managed processes is SETXCF START,REBUILD,DUPLEX.

See the following IBM documentation for additional information:

■ MVS Setting Up a Sysplex

■ MVS Programming: Sysplex Services Guide

■ MVS System Commands (the `SETXCF` operator command)

Step 5: Customize the Entire Net-Work Startup

Modify the Startup Procedure

Customize the Entire Net-Work startup procedure and copy it to a valid system procedure library.

The following example execution job can be used to start Entire Net-Work in a z/OS environment. More JCL may be needed, depending on the node configuration (types of line drivers, number of links, and so on). All load libraries in the STEPLIB concatenation must be APF-authorized.

```
//NETWORK  JOB . . . .
//*-------------------------------------------------------------*
//*               ENTIRE NET-WORK Vvrs STARTUP               *
//*-------------------------------------------------------------*
//XCFNET01 EXEC PGM=ADARUN,REGION=4M,TIME=1440  <--- see Note 1
//STEPLIB  DD DISP=SHR,DSN=NETWRK.Vvrs.LOAD
//         DD DISP=SHR,DSN=ADABAS.Vvrs.LOAD      <--- see Note 2
//DDPRINT  DD SYSOUT=*                           <--- see Note 3
//NETPRNT  DD DISP=SHR,DSN=WDXvrs.NETPRNT        <--- see Note 4
//MPMDUMP  DD SYSOUT=*                           <--- see Note 5
//SYSUDUMP DD SYSOUT=*
//DDCARD   DD *                                  <--- see Note 6
*-------------------------------------------------------------*
*             SAMPLE ADARUN CARDS FOR ENTIRE NET-WORK
      *
*-------------------------------------------------------------*
ADARUN PROG=NETWRK
ADARUN TARGETID=65490       /* MUST BE UNIQUE WITHIN SYSTEM  */
ADARUN SVC=209              /* SUPPLY ADABAS SVC NUMBER      */
ADARUN NAB=16,LU=65535,NC=20
ADARUN FORCE=N
*
//DDKARTE  DD *                                  <--- see Note 7
*-------------------------------------------------------------*
*
```

📄 **Notes:**

1. The region size required varies with the number and type of links, as well as other operating parameters.

2. The STEPLIB should always designate the most current Adabas load library unless you have been specifically instructed otherwise by Software AG. Adabas Cluster Services requires the corresponding Adabas limited library (WAL*vrs*) provided on the Adabas installation tape.

3. All Entire Net-Work print output is written to DDPRINT.

4. All diagnostic information from tracing, logging, and ABENDs is written to the NETPRNT file if it is open (otherwise it is written to the DDPRINT file). NETPRNT can be allocated to a large data set that can be copied when closed. The data set should be created with the DCB attributes RECFM=FBA and LRECL=121. To do this, allocate the file SHR. This causes the data set to be erased at the time the file is opened. Be aware that the diagnostic information is very large and will fill a data set quickly. When this happens, the file is closed and all additional output is sent to DDPRINT. This diagnostic information is created by Entire Net-Work and does not include the operating system dump information written to SYSUDUMP.

5. If MPMDUMP is defined, a snap dump is produced during any abnormal termination. In most error situations, the MPMDUMP dump may contain more pertinent information than the SYSUDUMP dump.

6. DDCARD contains the ADARUN control statements defining the interregion communication parameters for Entire Net-Work. Some ADARUN parameters are the same as for Adabas; others are similar, but are designed specifically for Entire Net-Work. Read *Entire Net-Work (NETWRK) Initialization Parameters* , in *Adabas Cluster Services Reference*, for information about coding ADARUN statements for Entire Net-Work.

7. DDKARTE contains Entire Net-Work parameter statements describing the local environment and the network connections for this node. Read *Entire Net-Work Parameter Statements*, in *Adabas Cluster Services Reference*.

Specify the Entire Net-Work Control Statements

As shown in the job example in section *Modify the Startup Procedure*, there are four types of Entire Net-Work control statements:

- ADARUN control statements define the interregion communication environment.

- NODE statements define the node's name and operating characteristics. Read *XCF NODE Statement* , in *Adabas Cluster Services Reference*.

- DRIVER statements define the line driver type(s) to be loaded. Add one or more DRIVER statements that specify your requirements for this Entire Net-Work node. At least one DRIVER statement is required. For the XCF line driver, read *XCF DRIVER Statement*, in *Adabas Cluster Services Reference*.

- LINK statements define the links to other nodes and are currently not required. Optionally, add one or more LINK statements. For the XCF line driver, read *XCF LINK Statement*, in *Adabas Cluster Services Reference*.

Customize the Entire Net-Work ADARUN and NODE statement parameters.

Add the XCF DRIVER statement that specifies your requirements to the Entire Net-Work parameter data set. Optionally, add one or more XCF LINK statements to describe Entire Net-Work partner nodes.

Step 6: Apply Entire Net-Work Maintenance

Apply corrective maintenance.

Refer to the *Software AG Product Delivery Report* to determine whether any files containing corrective maintenance (i.e., data sets named *pppvrs*.ZAPS) are supplied on the installation tape. If so, restore the data sets using IEBCOPY, and then follow the corrective maintenance instructions in the $READMVS member.

Step 7: Start Entire Net-Work

When execution JCL, including the ADARUN and Entire Net-Work parameter statements, is created for each node and all access method-specific definitions and installation steps are completed, Entire Net-Work can be started.

Start Entire Net-Work on multiple nodes and establish connectivity with the partner nodes.

The execution job or procedure for Entire Net-Work must be started on each of the participating nodes. Once started, Entire Net-Work determines the interregion communications environment on its node, connects to neighboring nodes through the defined links, and exchanges all relevant information with the other nodes during the "handshaking" process following link connection.

While the Entire Net-Work tasks are active and connected, user programs on any node can access all active targets on any other node, regardless of the targets' locations.

Run verification tests. For example:

■ Test your applications running across Entire Net-Work. Run one application at a time and verify the results.

■ Run a load test through the network (that is, multiple users on each node accessing data on the partner node).

If a link fails or cannot be connected properly, Entire Net-Work searches for possible alternate routes to complete the call. Once a disconnected link becomes available again, it is automatically considered for traffic.

Entire Net-Work can also be installed as a started task; no special considerations apply.

Entire Net-Work uses cross-memory services similar to Adabas in z/OS systems. As a result, the operating system may terminate the ASCB/initiator when Entire Net-Work terminates operation unless the system is configured with reusable address spaces. This is normal and is not an error.

During execution, a number of operator commands may be issued to Entire Net-Work to display or modify the system status. Read *Entire Net-Work Operator Commands*, in *Adabas Cluster Services Reference*.

Step 8: Create a Startup Procedure for ADACOM

An ADACOM initialization task is provided. This task must execute on each operating system image that has users who access the cluster database but no cluster nuclei; Software AG recommends that you also execute it on each operating system image that has a nucleus that participates in the cluster so that it is available as a command manager. All load libraries in the STEPLIB concatenation of ADACOM must be APF-authorized.

Either a cluster nucleus or the ADACOM program can allocate the nucleus table for monitoring the active nuclei and the user table for monitoring users in the extended CSA (ECSA) above the 16MB line.

The ADACOM task can be set up to manage any number of SVC/DBID sets in a sysplex. Read *ADACOM Initialization Parameters* in *Adabas Cluster Services Reference* for more information about specifying values for ADACOM parameters.

The following is a sample job for running ADACOM:

```
//ADACOM JOB . . . .
//*------------------------------------------------------------*
//*         Adabas Cluster Services ADACOM Vv.r.s STARTUP       *
//*------------------------------------------------------------*
//ALSvrs EXEC PGM=ADARUN,REGION=4M,TIME=nnnnn
//STEPLIB           DD DISP=SHR,DSN=ADABAS.Vvrs.LOAD
//*
//COMPRINT DD SYSOUT=*
//DDPRINT  DD SYSOUT=X
//DDCARD DD *
ADARUN PROGRAM=ADACOM
/*
//DDKARTE DD *
ADACOM SVC=sv1,DBID=dbid1,NU=150
ADACOM SVC=sv1,DBID=dbid2,NU=150,C=N,F=Y
ADACOM SVC=sv2,DBID=dbid1,NU=500
ADACOM SVC=sv2,DBID=dbid3,NU=500
/*
```

The COMPRINT DD statement must be specified when running ADACOM. It defines an output data set for all general messages printed by ADACOM. For each SVC/DBID set specified in the ADACOM task, two subtasks are attached and SYSOUT data sets are dynamically allocated to receive all messages specific to that combination. The DD-name of the SYSOUT data sets are

"P*ssddddd*" and "D*ssddddd*", where *ss* represents the last two digits of the SVC number and *ddddd* are the five digits of the DBID.

ADACOM can also be installed as a started task; no special considerations apply.

You may also want to add the ADACOM task's start command to member COMMND*xx* of SYS1.PARMLIB; this enables the ADACOM task to begin automatically at IPL.

Step 9: Create a Startup Procedure for Each Cluster Nucleus

1. Customize a startup procedure to execute ADARUN.

   For each Adabas cluster nucleus, customize the appropriate startup parameters and execute ADARUN from the Adabas load library.

2. Verify you have prepared and installed the product license file, as described in the Adabas for Mainframes installation..

3. Allocate and format a Work data set for each nucleus.

   All nuclei in an Adabas cluster share a common database resource; i.e., the same ASSO and DATA data sets. Each nucleus in the cluster must have its own Work data set; and all Work data sets within a cluster must have the same size and device type as defined in the general control block (GCB).

   Use DISP=SHR on the DD card for the Work data set (DDWORKR1). During an offline or online restart/recovery, a nucleus may access the Work data sets belonging to other nuclei in the cluster.

4. Specify the ADARUN `CLUSTER`, `NUCID`, `CLUGROUPNAME`, `CLUCACHENAME`, and `CLULOCKNAME` parameters for each nucleus.

   Although each nucleus of an Adabas cluster shares the same database resource (DBID), each nucleus must have a unique `NUCID` value:

   ▪ a single (noncluster) nucleus: NUCID=0 (default)

   ▪ a cluster nucleus: NUCID=1-65000

   Each cluster has cache and lock structure names that are unique in the sysplex system. ADARUN parameters are used to identify these names to each cluster nucleus.

   Use current values for all other ADARUN parameters, then reevaluate the values after monitoring the result. Ensure that each nucleus in the cluster is prepared to handle the entire workload for the common database, if necessary.

5. If protection logs are used, they must be dual or multiple logs and each nucleus must have its own. If one nucleus in the cluster runs with PLOGs, all nuclei in the cluster must run with PLOGs. The ADARUN `PLOGRQ` parameter must be the same for all nuclei (global parameter).

If user exit 2 or user exit 12 is supplied for one nucleus, the same user exit must be supplied for all nuclei in the cluster. User exit 12 must be used instead of user exit 2 if NCLOG/NPLOG is specified.

6. If command logs are used, each nucleus must have its own. If command logs are to be merged, they must be dual or multiple command logs and each nucleus in the cluster must have the same CLOG definition. `CLOGMRG=YES` must be specified in the ADARUN parameters to invoke automatic CLOG merging.

The following sample nucleus JCL executes the Adabas ADARUN program to implement session parameters for an Adabas cluster nucleus.

```
//PLX240N1 JOB . . . .
//*----------------------------------------------------------------*
//*                  STARTUP FOR ADABAS CLUSTER NUC #1              *
//*----------------------------------------------------------------*
//NUC240   EXEC PGM=ADARUN,REGION=8M,TIME=nnnnn
//STEPLIB          DD DSN=ADABAS.Vvrs.LOAD,DISP=SHR
//*
//DDASSOR1 DD DSN=MPM240.ASSOR1,DISP=SHR
//DDDATAR1 DD DSN=MPM240.DATAR1,DISP=SHR
//DDWORKR1 DD DSN=MPM240.WORKR1.CLU01,DISP=SHR
//*
//DDPRINT  DD SYSOUT=*
//DDDRUCK  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//*
//DDCARD   DD *
ADARUN PROG=ADANUC
ADARUN CLUSTER=SYSPLEX              ADABAS CLUSTER SERVICES
ADARUN NUCID=1
                  ADABAS CLUSTER NUCLEUS ID
ADARUN CLUGROUPNAME=ADA240          XCF GROUP NAME
ADARUN CLUCACHENAME=ADA240_CACHE    CACHE STRUCTURE NAME
ADARUN CLULOCKNAME=ADA240_LOCK      LOCK STRUCTURE NAME
ADARUN SVC=209                      SVC NUMBER
ADARUN DBID=240                     DATABASE ID
ADARUN PLOGRQ=NO                    PLOG REQUIRED
ADARUN MODE=MULTI                   MULTIUSER MODE
ADARUN LOCAL=NO
      .
      .
      .
//*
```

Step 10: Starting Adabas Cluster Services

Software AG recommends the following first start sequence:

1. Start Entire Net-Work on each operating system image that is participating in the Adabas nucleus cluster. Ensure that all nodes required for the cluster are connected.

2. Start the ADACOM initialization task on each operating system image that is participating in the Adabas cluster.

3. Start the Adabas cluster nuclei in any order.

The Adabas cluster is now ready to process user requests.

Rules for subsequent starts of Adabas Cluster Services are described in *Adabas Cluster Services Operations*.

When the Adabas cluster nuclei are active, issue the Entire Net-Work operator command `DISPLAY TARGETS` to ensure that the required targets have been established. For more information, read *DISPLAY - Display Information About a Network Component*, in *Adabas Cluster Services Reference*.

6 Modifying the z/OS Common Storage (CSA) Key

With z/OS version 1.9, the default setting of the ALLOWUSERKEYCSA parameter has been changed to NO, prohibiting the allocation of CSA storage in a non-privileged storage key. In past releases of Adabas Cluster Services and Adabas Parallel Services, its data structures in common storage (PLXCB and subordinate control blocks) were stored in user key storage. Due to the change of the ALLOWUSERKEYCSA default, you must either:

- Change the ALLOWUSERKEYCSA setting in z/OS 1.9 to YES. In this case, you need do nothing more, but you incur the security risk that unauthorized programs might modify storage owned by Adabas Cluster Services and Adabas Parallel Services.

- Accept the ALLOWUSERKEYCSA default setting in z/OS 1.9 (NO), thus eliminating the security risk. If you choose this option, however, you will need to modify your Adabas Cluster Services or Adabas Parallel Services installation and the Adabas Cluster Services, Adabas Parallel Services, and ADACOM startup procedures to use a system key. This chapter provides information on modifying the Adabas Cluster Services and the startup procedures to use a system key and on the prerequisites for doing so.

Requirements and Recommendations

The following requirements and recommendations apply when modifying the CSA storage key:

- The Adabas version 8 ADASVC is required. This procedure, when used with the Adabas 8 ADASVC, supports Adabas 7 and 8 used with Adabas Cluster Services versions 7.4 and 8, and with Adabas Parallel Services versions 7.4, 7.5 and 8. Older releases of Adabas Cluster Services or Adabas Parallel Services may require that some fixes be applied before attempting this, so be sure to check the Knowledge Center of Software AG's Empower (at *https://empower.soft-wareag.com*) web site for them.

- APF authorization is required to be assigned a system key. Adabas Parallel Services and Adabas Cluster Services nuclei must run authorized, but utilities and classic nuclei may bypass APF authorization by applying an optional fix to ADASVC. If APF authorization is not in effect and the execution JCL has been modified as described below, the following message will appear during initiation:

```
IEF188I PROBLEM PROGRAM ATTRIBUTES ASSIGNED
```

Steps to Modify the CSA Key

In general, you can modify the CSA key one system at a time for each database. On each system, you can apply the modifications to each DBID/SVC combination independently. One exception to this is when ADACOM controls more than one DBID/SVC pair. In this case, you must modify all DBID/SVC combinations defined to that ADACOM at the same time. All database member nuclei on the system using that ADASVC or ADACOM must be stopped and restarted.

The modification to the CSA storage key involves the use of program name ADARUNK, which is an alias for ADARUN. When your modifications are complete (see the steps below), your installation and all of the appropriate execution JCL will have been changed to use ADARUNK (`PGM=ADARUNK`) instead of ADARUN. Once this is done, z/OS will assign a system key to all tasks for the nuclei or ADACOM processes. Although only ADACOM tasks and nuclei associated with Adabas Cluster Services or Adabas Parallel Services obtain CSA storage and are required to perform these modification steps to avoid using a user key, the ADARUNK program can also be used for classic Adabas nuclei and utilities.

> **To modify the CSA storage key to a system key:**

1   Add the following entry for program ADARUNK to an active Program Properties Table description in SYS1.PARMLIB(SCHED*xx*):

```
PPT PGMNAME(ADARUNK) KEY(4)
```

2   Activate the new PPT entry. This will occur automatically at the next IPL, or you can use the z/OS operator command `SET SCH(xx)` to activate it without an IPL.

3   Optionally, create permissions for the FACILITY class resource IEAABD.DMPAKEY to allow SYSUDUMP to include system key storage.

When executing with a system key, SYSUDUMP will not contain the system key storage. The following message will appear:

```
IEA848I DUMP SUPPRESSED - ABDUMP MAY NOT DUMP STORAGE FOR KEY 0-7
```

Since cluster nuclei and ADACOM may obtain large storage objects, we suggest including a `//SVCDUMP` DD statement in the execution JCL instead of a `//SYSUDUMP` DD statement. Otherwise, enable SYSUDUMP for system key tasks by permitting the user to have read access to the IEAABD.DMPAKEY resource in the FACILITY class. For more information, refer to your *z/OS Security Server (RACF) Security Administrator's Guide*.

4   Examine the Adabas load library to see if an entry exists for ADARUNK. If it does not, create an alias ADARUNK for member ADARUN. You can do this using programs such as PDSTOOLS or the freeware PDS or use the Linker/Binder to relink ADARUN and create the alias:

```
//LINK    EXEC  PGM=IEWL,PARM='Map,Let,List(All),Ncal'
//SYSPRINT DD   SYSOUT=*
//SYSLMOD  DD   DSNAME=<Adabas load library>,DISP=SHR
//SYSLIN   DD   *
      Include  SYSLMOD(ADARUN)
        Order  RUNMVS,USRZAP,RUNIND
         Mode  AMode(31),RMode(24)
      Setcode  AC(1)
        Entry  RUNMVS
        Alias  ADARUNK
         Name  ADARUN(R)
/*
```

5   Stop all Adabas Parallel Services and Adabas Cluster Services member nuclei that use the same DBID/SVC pair.

6   Any existing user key PLXCB structures for a DBID/SVC combination must be deleted. This happens automatically at the next IPL or if ADASVC is reinstalled using ADASIP.

    ADACOM may be used to deallocate the PLXCB without an IPL by specifying NU=0 for every DBID/SVC pair, and then stop and restart ADACOM.

    The following sequence of operator commands may also be used to deallocate a PLXCB after stopping all nuclei but before stopping ADACOM:

```
F <ADACOM job>,DBID=<dbid>,SVC=<svc>,ADAEND
F <ADACOM job>,DBID=<dbid>,SVC=<svc>,NU=0
```

7   If ADACOM is running and controlling the database, it must be stopped.

   🔴   **Caution:** If ADACOM is controlling more than one DBID/SVC pair, all pairs must be processed at the same time.

8   Modify the startup JCL for each Adabas Parallel Services or Adabas Cluster Services member nucleus and ADACOM to specify program ADARUNK instead of ADARUN. For example, for an Adabas Parallel Services or Adabas Cluster Services nucleus you might specify:

```
//NUC240   EXEC PGM=ADARUNK
```

For an ADACOM startup procedure, you might specify:

```
//ALSvrs EXEC PGM=ADARUNK,REGION=4M,TIME=nnnnn
```

9   Modify the JCL for each Adabas Parallel Services or Adabas Cluster Services member nucleus and ADACOM to include a `//SVCDUMP DD DUMMY` statement.

10   Restart all ADACOM jobs and Adabas Parallel Services and Adabas Cluster Services nuclei that you stopped in Step 5.

11   Repeat steps 5 through 10 for each DBID/SVC pair.

> **Caution:** If ADACOM is controlling more than one DBID/SVC pair, all pairs must be processed at the same time.

Index