

Adabas Cluster Services

Glossary

Version 8.5.4

April 2020

This document applies to Adabas Cluster Services Version 8.5.4 and all subsequent releases.

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

Copyright © 2020 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

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

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third-Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Use, reproduction, transfer, publication or disclosure is prohibited except as specifically provided for in your License Agreement with Software AG.

Document ID: ALS-GLOSSARY-854-20210929

Table of Contents

Glossary	1
A	1
C	1
D	2
E	2
L	2
N	2
P	2
S	3
T	4
V	4
X	5

Glossary

A

ADACOM

Either an ADACOM initialization module or the first nucleus to start, whichever starts first, sets up the sysplex environment on an operating system image. Parameters of the ADACOM control program define the characteristics of any number of specified SVC/DBID sets in the sysplex.

The ADACOM module *may* be installed on each operating system image that has a nucleus that participates in a sysplex cluster and *must* be installed on each operating system image that has users that access a cluster database but no cluster nuclei.

After initialization, an ADACOM task may either be quiesced or remain in operation as a command manager to monitor and control cluster nuclei running under the specified SVC/DBID sets.

C

cache structure

Each Adabas cluster uses a sysplex cache structure to hold ASSO/DATA blocks that have been updated during the session. The cache structure provides the cluster nuclei with up-to-date block images to ensure data integrity. It is defined in the sysplex coupling facility resource management (CFRM) policy and specified using the ADARUN parameter `CLUCACHENAME`. The same name must be provided to each nucleus in the cluster and the name may not be used outside the cluster.

coupling facility

The coupling facility in a parallel sysplex (SYStems comPLEX) environment is located in a separate machine or a special logical partition on the same machine. It makes it possible for program instances in different operating system images to efficiently share data.

The operating systems are directly connected to the coupling facilities through high-bandwidth, high-speed channels. These channels

provide a fast and efficient means for systems in the sysplex to access data in the coupling facility.

Storage in a coupling facility is divided into objects called "structures". Structure types are cache, list, and lock. The storage is managed by coupling facility resource management (CFRM) "policies". Policies contain information about the structures you plan to use. Once created, policies resides in the CFRM COUPLE_{xx} data sets.

D

database, single physical

One set of Associator and Data Storage data sets identified by a single database ID number (DBID).

Using Adabas Cluster Services, a single physical database is accessed simultaneously by up to 32 Adabas nuclei running under multiple operating system images.

E

ECSA

Extended Common Service Area (ECSA) is a virtual storage area above the 16MB line that can be accessed from all address spaces in the operating system image.

L

lock structure

Each ADABAS cluster uses a sysplex lock structure to manage the setting, status, and release of various locks imposed during multiple update nucleus processing. The lock structure synchronizes the nuclei, users, and transaction processing to ensure data integrity. It is defined in the sysplex coupling facility resource management (CFRM) policy and specified using the ADARUN parameter CLULOCKNAME. The same name must be provided to each nucleus in the cluster and the name may not be used outside the cluster.

N

NUCID

An ADARUN parameter that identifies a particular Adabas nucleus as a cluster nucleus.

P

PPT

The parallel participant table (PPT) is an area of 32 blocks in the database's Associator. It is used to track the Work, PLOG, and CLOG

data sets/files used by the nuclei running in a cluster (as well as by a noncluster nucleus).

Each cluster nucleus reserves one PPT block for itself and records in this block the names of the Work, PLOG, and CLOG data sets/files that it uses, together with pertinent status indicators.

Individual cluster nuclei use the PPT information to access the Work data sets/files of the peer nuclei when performing recovery after a nucleus failure.

The ADARES utility uses the PPT information to access the PLOG or CLOG data sets/files of the cluster nuclei when performing the automated PLOG or CLOG merge, respectively.

Furthermore, cluster and noncluster nuclei use the PPT information to warn about changes in the Work or PLOG JCL, if the previously used Work or PLOG data sets/files still contain data.

policies

See the entry *coupling facility* for information about coupling facility resource management (CFRM) policies.

S

structures

See the entry *coupling facility* for information about coupling facility structures.

SVCCLU

As a component of the Adabas SVC/Router, SVCCLU routes commands to the cluster nuclei. It uses the nucleus and user tables in common storage, which are updated based on nucleus or ADACOM information, to decide to which nucleus to assign a new user and route the user's commands.

There is no predefined limit on the number of Adabas clusters, each servicing a separate database, that can run under a single SVC/Router.

sysplex environment

A parallel sysplex (SYStems comPLEX) environment is a set of software services and hardware components that allow multiple instances of the OS/390 and z/OS operating systems to communicate and cooperate with each other to process customer work. For certain kinds of work, the sysplex provides parallel processing and improved data sharing.

High-speed caching, list processing, and locking functions are provided by the sysplex hardware components.

T

target IDs

Target IDs identify servers that Entire Net-Work can pass messages to. Such servers include Adabas nuclei, Natural global buffer pools, Entire Broker, among others. In the realm of Entire Net-Work, target IDs must be uniquely assigned to these servers as numbers in the range 1-65000.

The NUCIDs assigned to Adabas nuclei running on cluster databases are also such target IDs and must therefore be unique across all target IDs. Furthermore, Adabas Cluster Services assigns target IDs in the range 65001-65481 (a restricted range) to identify operating system images in the sysplex to the network. Software AG recommends that you assign unique values in the range 65482-65535 to identify Entire Net-Work nodes themselves.

V

version

A product version is identified by the first two digits of the versioning number. Software AG distinguishes between major and minor versions according to the amount of functionality or technology added to the product. All other digits indicate correction levels.

In the product documentation, the notation *vrs*, *vr*, or simply *v* is often used as a placeholder for the current product version, for example, in data set or module names.

Placeholder	Meaning	Definition
<i>v</i>	version	Major Version The first digit of the product version number indicates major architecture and functionality implementation or enhancement that adds value to the product.
<i>r</i>	release	Minor Version The second digit of the version number indicates functionality addition or enhancement that adds value to the product.
<i>s</i>	system maintenance level	Correction Level Correction levels contain error corrections only, without new functionality, including documentation of all modifications and repairs.

Placeholder	Meaning	Definition
		<p>In case it is necessary to include functional changes into a correction level, an exception handling process ensures that corresponding quality assurance activities are triggered. These functional changes are documented. The main target is to avoid impacts when you install such a correction level.</p> <p>The third number of an Adabas Cluster Services version denotes the system maintenance level.</p> <p>On certain platforms supported by Adabas, additional levels may exist, such as update package, patch level, service pack and hot fix.</p>

X

XCF IBM's cross-system coupling facility (XCF) services allow authorized applications on one system to communicate with applications on the same system or on other systems. XCF transfers data and status information between members of a group that reside on one or more operating systems in a sysplex.

XCF Group Each Adabas cluster in a sysplex environment uses a cross-system coupling facility (XCF) group to control communication among the cluster nuclei. The operating system automatically creates an XCF group for a cluster when the first connection is made to the cluster's lock structure.

When the sysplex COUPLE_{xx} data sets are first formatted, the maximum number of XCF groups to be used for the entire sysplex environment is defined. However, XCF group names are *not* pre-defined in the CFRM policy.

A name for a cluster's XCF group must be provided to each cluster nucleus using the ADARUN parameter CLUGROUPNAME. The same name must be provided to each nucleus in the cluster and the name may not be used outside the cluster.
