

# ADAX\* - Adabas Cluster Nucleus Messages

## Overview of Messages

ADAX01	ADAX02	ADAX03	ADAX04	ADAX05	ADAX09	ADAX11	ADAX12
ADAX13	ADAX14	ADAX15	ADAX16	ADAX20	ADAX21	ADAX22	ADAX23
ADAX24	ADAX26	ADAX27	ADAX28	ADAX29	ADAX2A	ADAX2B	ADAX2C
ADAX2D	ADAX2E	ADAX2F	ADAX2G	ADAX2H	ADAX2I	ADAX31	ADAX32
ADAX33	ADAX40	ADAX41	ADAX42	ADAX43	ADAX44	ADAX45	ADAX46
ADAX47	ADAX48	ADAX49	ADAX50	ADAX51	ADAX52	ADAX53	ADAX54
ADAX55	ADAX56	ADAX57	ADAX58	ADAX59	ADAX5A	ADAX5B	ADAX5C
ADAX5D	ADAX60	ADAX61	ADAX62	ADAX63	ADAX64	ADAX65	ADAX66
ADAX67	ADAX68	ADAX69	ADAX6B	ADAX70	ADAX71	ADAX72	ADAX73
ADAX74	ADAX75	ADAX76	ADAX77	ADAX78	ADAX79	ADAX80	ADAX81
ADAX82	ADAX83	ADAX84	ADAX85	ADAX86	ADAX87	ADAX88	ADAX89
ADAX90	ADAX91	ADAX92	ADAX93	ADAX94	ADAX95	ADAX96	ADAX97
ADAX98	ADAX99	ADAX9A	ADAX9B	ADAX9C	ADAX9D	ADAX9E	ADAX9F
ADAX9G	ADAX9H	ADAX9J	ADAX9K	ADAX9L	ADAXS4		

### **ADAX01**     *dbid NUCID nucid on system system status*

**Explanation**     A nucleus entered or left the Adabas cluster.

Adabas nucleus cluster members that were identified as active at initialization are indicated with status 'is present'. Post-initialization changes to Adabas cluster membership are indicated with status 'has joined'; 'has withdrawn'; or 'has failed'.

**Action**             No action is required for this informational message.

### **ADAX02**     *dbid System sysn status monitor update missing*

**Explanation**     XCF reported that the named system (*sysn*) is not responding. This affects all XCF users on that system, not just Adabas. Until communications are restored, all nuclei on that system may not be able to respond to critical intracluster synchronization messages before the message times out.

**Action**             Attempt to restore XCF communications with the named system.

### **ADAX03**     *dbid System sysn status monitor update resumed*

**Explanation**     After missing one or more status monitoring intervals, the named system (*sysn*) is now responding to XCF communications.

**Action**             No action is required for this informational message.

**ADAX04**     *dbid* NUCID *nucid* on system *sysn* status monitor update missing

**Explanation**    XCF reported that the identified nucleus (*nucid*) on the identified system (*sysn*) is not updating its heartbeat monitor. It is likely that the nucleus is not able to process commands, including critical intracluster synchronization messages.

**Action**            Determine why the nucleus is not able to update its heartbeat monitor. The nucleus may be hung or looping, or there may be other processes in the system preventing the nucleus from getting sufficient CPU allocations.

**ADAX05**     *dbid* NUCID *nucid* on system *sysn* status monitor update resumed

**Explanation**    The identified nucleus (*nucid*) on the identified system (*sysn*) has resumed updating its heartbeat monitor.

**Action**            No action is required for this informational message.

**ADAX09**     *dbid* Post NUC *nucid* failed - RET *return-code* RSN *reason-code*

**Explanation**    This message may appear when another nucleus terminates.

**Action**            No action is necessary if the other nucleus is terminating abnormally. If the message occurs in an otherwise normal nucleus session, contact your Software AG technical support representative.

**ADAX11**     *dbid* Adabas cluster messaging initialization failed

**Explanation**    An error described in preceding messages prevented successful initialization of Adabas cluster messaging services. Nucleus initialization fails with PARM-ERROR 092.

**Action**            Correct the problem identified in the preceding messages.

**ADAX12**     *dbid* Unable to obtain {AXMVT | ADAMCB} storage

**Explanation**    A request to obtain storage from the work pool for Adabas cluster messaging service control structures failed. Nucleus initialization fails with PARM-ERROR 092.

**Action**            Increase the amount of virtual storage available. Alternatively, adjust ADARUN parameters to allow for a larger work pool or decrease ADARUN parameters NT and NU to require fewer AXMCBs.

**ADAX13**     *dbid Error subcode replying to nucid message msgnum*

**Explanation** An error occurred while replying to an intracluster command originating from the nucleus whose nucleus ID is given in the message (*nucid*). The reply could not be sent. The error code given in the message (*subcode*) is one of the possible subcodes for response code 123 (ADARSP123). The originating nucleus message sequence number is also given in the message (*msgnum*).

The originating nucleus will continue to wait for a reply until the message times out or the nonresponding nucleus ends.

**Action** If the reason for the error cannot be determined and corrected, notify your Software AG technical support representative.

**ADAX14**     *dbid Statistics for type-type messages*  
*dbid Messages sent nn replies sent nn*  
*dbid Messages arrived nn messages accepted nn*

**Explanation** Produced during normal nucleus termination, this message provides Adabas Cluster Services messaging service statistics:

messages sent	reflects the number of internucleus messages initiated from this nucleus
messages arrived	is the count of asynchronous incoming messages queued for the nucleus (normally, the same as 'messages accepted')
messages accepted	is the count of those messages the nucleus processed (normally, the same as 'messages arrived')
replies sent	is the count of nucleus responses to accepted messages that required a response.

**Action** No action is required for this informational message.

**ADAX15**     *dbid AXMCB allocated nn used nn total requests nn*

**Explanation** Produced during normal nucleus termination, this message provides Adabas cluster messaging service statistics:

AXMCB allocated	number of internucleus message control blocks allocated.
AXMCB used	number of internucleus message control blocks used.
total requests	total number of requests to use the allocated internucleus message control blocks.

**Action** No action is required for this informational message.

**ADAX16**     *dbid date time statistic*

**Explanation**     This message is used to display the output from the DXMSG operator command on the console.

**Action**             No action is required for this informational message.

**ADAX20**             ***dbid XCF transport initialization complete***

**Explanation**             The z/OS XCF transport service successfully initialized.

**Action**                 No action is required for this informational message.

ADAX21     *dbid error-text*

**Explanation** The error specified by one of the following error texts occurred during the z/OS XCF transport service initialization check.

Error Text	Explanation	Action
Existing XCF group member <i>scf-member</i> uses different AXMCB version	The cluster member listed in the message is not using the same version of Adabas as the other nuclei in the cluster. All Adabas sysplex cluster nuclei generate <i>scf-member</i> names in the format: <i>db4444pppppnn</i> -where <i>db4444</i> is the database ID, <i>ppppp</i> is the nonzero nucleus ID, and <i>nn</i> is an internal ordinal identifier.	Check that all nuclei in the cluster use the same Adabas version and the same Adabas Cluster versions.
Existing XCF group member <i>scf-member</i> uses different DBID	An Adabas sysplex cluster nucleus that is already active in the same XCF group is using a different DBID. All Adabas sysplex cluster nuclei generate <i>scf-member</i> names in the format: <i>db4444pppppnn</i> -where <i>db4444</i> is the database ID, <i>ppppp</i> is the nonzero nucleus ID, and <i>nn</i> is an internal ordinal identifier.	Verify that ADARUN parameters DBID and CLUGROUPNAME are correct in all nuclei participating in the sysplex cluster. Contact your Software AG technical support representative if you are unable to resolve the problem.
Incompatible AXMCB version	The Adabas Cluster Services modules you are using are not compatible with the Adabas running.	Contact your Software AG technical support representative for assistance.
Incompatible AXMVT version	The Adabas Cluster Services modules you are using are not compatible with the Adabas running.	Contact your Software AG technical support representative for assistance.
Invalid group name	The ADARUN parameter CLUGROUPNAME was omitted or invalid.	Contact the ADARUN parameter CLUGROUPNAME must begin with an alphabetic character, may not begin with SYS, and may not be UNDESIG.
Invalid user state data from existing member <i>scf-member</i>	The control information presented for a member already connected to the XCF group was not formatted as a proper Adabas sysplex cluster nucleus or had a different DBID. The preexisting member may be an Adabas sysplex cluster nucleus associated with a different DBID, or it may be another process using the same XCF group name. All Adabas sysplex cluster nuclei generate <i>scf-member</i> names in the format: <i>db4444pppppnn</i> -where <i>db4444</i> is the database ID, <i>ppppp</i> is the nonzero nucleus ID, and <i>nn</i> is an internal ordinal identifier.	Identify the source of <i>scf-member</i> . If it is an Adabas sysplex cluster nucleus, make sure the parameters NUCID, CLUGROUPNAME, and DBID are correct. There may be additional information in messages generated by the other nucleus. If it is not an Adabas sysplex cluster nucleus, contact your systems programmer or support representative. If you are unable to resolve the problem, contact your Software AG technical support representative.
IXCJOIN failed, duplicate NUCID and XCF member name	XCF service IXCJOIN reported the member name requested by this nucleus was already active in the XCF group. The member name is derived from the ADARUN parameters DBID and NUCID and an internal number assigned during nucleus initialization. Any return and reason codes included with this message are defined in the IBM documentation entitled <i>MVS Programming: Sysplex Services Reference</i> .	Verify that ADARUN parameter DBID is correct and NUCID is unique among all nuclei participating in the Adabas sysplex cluster. Contact your Software AG technical support representative if you are unable to resolve the problem.
IXCJOIN failed or retry count exhausted	An error was reported by XCF service IXCJOIN. Message ADAX28 provides the return and reason code from IXCJOIN. These return and reason codes are defined in the IBM documentation entitled <i>MVS Programming: Sysplex Services Reference</i> . An error may result from XCF options specified for your installation by your systems programmer.	Contact your Adabas technical support representative if you are unable to resolve the problem.
IXCQUERY failed	An error was reported by IBM XCF service IXCQUERY. Message ADAX28 provides the return and reason code from IXCQUERY. These are defined in the IBM documentation.	Contact your Adabas technical support representative for assistance.
NUCID already active	XCF initialization found an active Adabas sysplex cluster nucleus with the same NUCID.	Verify that ADARUN parameter NUCID is unique among all nuclei participating in the Adabas sysplex cluster.
Too many members exist in XCF group	IXCQUERY identified an unexpected number of preexisting members in the XCF group.	Contact your Software AG technical support representative for assistance.
XCF latch set creation failed	An error was reported by the IBM latch set creation routine ISGLCRT.	Contact your Software AG technical support representative for assistance.
XCF transport initialization failed	The initialization of the Adabas sysplex cluster's messaging service failed and nucleus initialization failed with PARM error 092. The reason for the failure is indicated in a previous message.	Review the previous messages to resolve the error. Contact your Software AG technical support representative for assistance.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX22**     *dbid* Status monitor update missed

**Explanation** This nucleus was notified by XCF that it failed to update its heartbeat monitor within the interval specified by ADARUN parameter MXSTATUS. XCF will notify all other nuclei registered to the same XCF group, each of which may issue message ADAX04.

**Action** Determine why this nucleus is not able to update its heartbeat monitor. The nucleus may be hung or looping, or there may be other processes in the system preventing the nucleus from getting sufficient CPU allocations.

**ADAX23**     *dbid* Status monitor update resumed after *nn* seconds

**Explanation** After missing one or more status monitoring intervals, this nucleus has resumed updating its heartbeat monitor. XCF will notify all other nuclei registered to the same XCF group, each of which may issue message ADAX05. This message reports the number of seconds (*nn*) since the missing status update was first detected and the related ADAX22 message was printed.

**Action** No action is required for this informational message.

ADAX24     *dbid error-text*

**Explanation** An error occurred while processing an incoming asynchronous message. This message appears only in the z/OS Adabas nucleus JESLOG listing, SYSLOG or operator's console. One of the following error texts is given in this message:

Error Text	Explanation	Action
AXMCB allocation failed in XCF message exit	A message control block to describe the incoming message could not be obtained from the pool.	Adjust ADARUN parameters NT and NU to increase the number of AXMCBs created at initialization. Contact your Software AG technical support representative if you are unable to resolve the problem.
Buffer allocation failed in XCF message exit	A buffer for the incoming message could not be obtained.	Provide more storage by increasing the REGION JCL parameter. Contact your Software AG technical support representative if you are unable to resolve the problem.
Invalid incoming MSGCNTL header in XCF message exit	The control information presented for the incoming message was not formatted as a proper Adabas cluster nucleus or had a different DBID. This may result if a message was sent from an XCF group member previously cited in message ADAX27.	See message ADAX27. Contact your Software AG technical support representative if you are unable to resolve the problem.
Out of sequence or missing segments in XCF message exit	The segments of a message whose length required it to be sent in multiple segments did not arrive in the expected order.	Contact your Software AG technical support representative for assistance.
Segmented message timed out in XCF message exit	A message whose length required it to be sent in multiple segments was not complete at the expiration of the timeout interval. This may be the result of an error on the sending nucleus, an XCF error, or contention for system resources.	Contact your Software AG technical support representative if you are unable to resolve the problem.
Unable to receive segment in XCF message exit	An error was reported from the XCF IXCMMSGI service when attempting to receive the message. Message ADAX28 provides the IXCMMSGI return and reason codes. These return and reason codes are defined in the IBM documentation entitled <i>MVS Programming: Sysplex Services Reference</i> .	Contact your Software AG technical support representative if you are unable to resolve the problem.
Unable to save message segment in XCF message exit	An error was reported from the XCF IXCMMSGC service when attempting to save the message. Message ADAX28 provides the IXCMMSGC return and reason codes. These return and reason codes are defined in the IBM documentation entitled <i>MVS Programming: Sysplex Services Reference</i> . There may not be sufficient resources allocated in your installation to save the message.	Contact your systems programmer or technical representative to determine if XCF is experiencing a shortage of resources. Contact your Software AG technical support representative if you are unable to resolve the problem.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX26**     *dbid Invalid user state data from xcf-member*

**Explanation** An error occurred when processing a member state change event generated by the z/OS XCF messaging transport service as a member attempted to join the XCF group. The control information presented for the XCF group member attempting connection was not formatted as a proper Adabas sysplex cluster nucleus or had a different DBID. The joining member may be an Adabas sysplex cluster nucleus associated with a different DBID or it may be another process using the same XCF group name. All Adabas sysplex cluster nuclei generate *xcf-member* names in the format

DB`ddddpppppNnn`

where:

<code>dddd</code>	is the database ID
<code>ppppp</code>	is the nonzero NUCID
<code>nn</code>	is an internal ordinal identifier

The member state change event is discarded.

**Action** Identify the source of *xcf-member*. If it is an Adabas sysplex cluster nucleus, make sure the parameters NUCID, CLUGROUPNAME, and DBID are correct and notify your Adabas technical support representative if this does not correct the problem. There may be additional information in the messages for the nucleus attempting to join. If it is not an Adabas sysplex cluster nucleus, notify your system programmer or support representative.

**ADAX27**     *dbid No room in AXCFVT table for xcf-member*

**Explanation** An error occurred when processing a member state change event generated by the z/OS XCF messaging transport service.

**Action** Contact your Adabas technical support representative. The member state change event is discarded.

**ADAX28**     *dbid xcf-service-routine RET return-code RSN reason-code*

**Explanation** This message appears only in the z/OS Adabas nucleus JESLOG listing, SYSLOG, or operator's console. It is issued for certain XCF message transport service requests during initialization and termination, and whenever a request fails. Refer to IBM documentation entitled *MVS Programming: Sysplex Services Reference* for descriptions of the various return and reason codes for each XCF service.

**Action** This message may be associated with an error condition reflected in other messages, or may have caused an Adabas cluster message request to fail. If an associated error is identified, include this message when contacting your Software AG technical support representative.



**ADAX29**     *dbid* Adabas abend in XCF exit DBID *dbid* NUCID *nucid* abend routine *Snnn*  
Unnnn reason *rsn* PSW *psw* REG 0-3 *reg 0 reg 1 reg 2 reg 3* REG 4-7 *reg 4 reg 5*  
*reg 6 reg 7* REG 8-B *reg 8 reg 8 reg 10 reg 11* REG C-F *reg 12 reg 13 reg 14 reg 15*

**Explanation**   A program check or system abend was intercepted by the z/OS XCF messaging transport service FRR or ESTAE routine. The error may have occurred under an SRB.

The nucleus should terminate. A dump may be produced in SDUMP format, either in one of the nucleus-allocated files SYSUDUMP, SYSMDUMP, or SYSABEND; or in a system-allocated file such as SYS1.DUMPn.

**Action**        Contact your Software AG technical support representative for assistance.

**ADAX2A**     *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

Message Text	Explanation	Action
Incompatible AXMVT version	The Adabas Parallel Services modules you are using are not compatible with the Adabas running.	Check that your Adabas Parallel Services library is compatible with your Adabas library. If the problem persists, contact your Software AG technical support representative for assistance.
TI-0, initialized, RC <i>return-code</i>	ADASMM initialized successfully.	No action is required for this informational message.
TI-1, cannot get work memory	The attempt to acquire memory for the ADASMM work area failed. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
TI-2, cannot get PLXCB	The attempt to acquire an Adabas Parallel Services control block (MPM 76 call) failed. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
TI-3, cannot load PLXDEP	The attempt to load the operating system interface module PLXDEP failed. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
TI-4, memstate call NUCID: <i>nucid</i> , RC <i>return-code</i>	The specified hexadecimal return code was received from the member state table interface call for the specified nucleus ID.	Analyze the return code and correct the error.
TI-5, Error in post NUCID: <i>nucid</i> , RC <i>return-code</i>	The specified hexadecimal return code was received from the cross-memory post (MPM 80) routine to the specified nucleus ID.	Analyze the return code and correct the error.
TI-6, NUCID: <i>nucid</i> reported active - inconsistent PLXCB	A fatal error occurred during initialization. The cluster control block PLXCB reported incorrectly that the nucleus (NUCID) was active. The PLXCB is therefore inconsistent and initialization fails with response code 8 (ADARSP008).	Restart the cluster cleanly.
<i>dbid</i> TI-9, error set process token: <i>xx</i>	A fatal error occurred in obtaining the process token value <i>xx</i> as returned from the Adabas Operating System interface ADAIOR, where <i>dbid</i> is the database ID of the SMP cluster.	Note the response code delivered and contact your Software AG technical support representative for assistance.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX2B** *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

Message Text	Explanation	Action
TT-0, SMM not yet initialized	The termination call was made without a previous successful initialization call. The SMM facility terminates.	Contact your Software AG technical support representative for assistance.
TT-1, SMM terminating	The SMM facility is terminating.	Contact your Software AG technical support representative for assistance.
TT-2, memstate call, NUCID: <i>nucid</i> , RC <i>return-code</i>	The specified hexadecimal return code was received from the member state table interface call to the specified nucleus ID.	Analyze the return code and correct the error.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX2C** *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

<b>Message Text</b>	<b>Explanation</b>	<b>Action</b>
SM-0, SMM not initialized yet	A send call was made without a previous successful initialization. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
SM-1, timer call, RC <i>response-code</i>	The specified hexadecimal response code was returned from a call to set up a timeout interval.	Analyze the response code and correct the error.
SM-2, No UB available, RC <i>response-code</i>	The specified hexadecimal response code was returned by the call to acquire a user buffer.	Analyze the response code and correct the error.
SM-3, Cannot find ECB element	An event control block is required in order to send a message. This control block cannot be acquired because the table is full.	The size of the table is based on the ADARUN NC parameter. Increase the value of the ADARUN NC parameter to increase the table size.
SM-4, Reply error, NUCID: <i>nucid</i> , RC <i>response-code</i>	The specified hexadecimal response code was returned by the specified (external) NUCID.	Analyze the response code and correct the error.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX2D** *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

Message Text	Explanation	Action
RM-0, SMM not yet initialized	A receive call was made without a previous successful initialization. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
RM-1, reply ECB not found, RC <i>response-code</i>	Every ADASMM message needs an event control block, which is held in a table. The required receive messages does not have an equivalent event control block entry set by a send.	Analyze the response code and correct the error.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX2E** *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

Message Text	Explanation	Action
QU-0, SMM not yet initialized	A query member call was made without a previous successful initialization. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
QU-1, bad function code <i>code</i>	The calls to ADASMMQU have a function code. The function code "code" is out of range. This is an internal error. ADASMM terminates.	Contact your Software AG technical support representative for assistance.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX2F**     *dbid* TM-0, SMM not initialized yet

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

A call was made to terminate the cluster session without a previous successful initialization call.

**Action** ADASMM terminates.

**ADAX2G**     *dbid* CM-0, SMM not yet initialized

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

A cancel call was made without a previous successful initialization.

**Action** ADASMM terminates.

**ADAX2H** *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

Message Text	Explanation	Action
CME-0, SMM not yet initialized	A receive exit call was made without a previous successful initialization. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
CME-1, cannot find ECB element	An event control block was received in the database for ADASMM. The equivalent event control block for the message that was sent cannot be found.	Ensure that the ECB for the message sent is available.
TIMEX-0, SMM not yet initialized	A timer exit call was made without a previous successful initialization. ADASMM terminates.	Contact your Software AG technical support representative for assistance.
<i>dbid</i> TIME-1, msg to <i>cccc</i> timed out	An attempt to send a message to cluster <i>cccc</i> timed out, with no response.	This is a warning. Contact your Software AG technical support representative for assistance.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.

**ADAX2I** *dbid message-text*

**Explanation** This message is returned by the Adabas Parallel Services messaging module ADASMM, also called the SMM facility.

One of the following message texts can appear in this message:

Message Text	Explanation	Action
SS-1, target gone, ID <i>nucid</i>	The cluster nucleus with the specified (external) NUCID left the cluster.	No action is required for this informational message.
SS-2, memstate, ID <i>int-nucid</i> IND <i>idx-num</i> EXT <i>nucid</i> state <i>nn</i>	This message provides information about calls to the member state table manager from the SMM facility where: <ul style="list-style-type: none"> <li>● <i>int-nucid</i> is the nucleus indicator entry for the cluster nucleus in the member state table (internal).</li> <li>● <i>idx-num</i> is the index number of the internal nucleus indicator entry.</li> <li>● <i>nucid</i> is the user-specified (external) NUCID number, or zero (0) for a noncluster nucleus.</li> <li>● <i>nn</i> is the status of the nucleus: 03 to activate or 00 to release.</li> </ul>	No action is required for this informational message.
SS-3, ACB to ID <i>nucid</i> , RC <i>rsp-code</i> , AD2 <i>value</i>	An error has occurred in cluster communication. The cluster nucleus that returned the message is identified by the (external) NUCID. The response code is provided as well as the contents of the command ACB's or ACBX's Additions 2 field.	This is a warning. Contact your Software AG technical support representative for assistance.

**Action** The user actions are given in the table above. If the problem persists, contact your Software AG technical support representative for assistance.



**ADAX31      Opening work data set for NUCID=*nucid***

**Explanation** While recovering from the failure of one or more Adabas cluster nuclei, this nucleus is about to open the Work data set of the nucleus with the specified (external) NUCID, which terminated abnormally.

**Action** No action is required for this informational message.

**ADAX32      Open failed. IOR code=X'*rc*', system code=X'*ssss*'**

**Explanation** While recovering from the failure of one or more Adabas cluster nuclei, this nucleus tried to open the Work data set of the nucleus with the NUCID '*nucid*' (message ADAX31), but the open failed. ADAIOR issued response code *rc* (in hexadecimal), and the system's return code is '*ssss*' (in hexadecimal). The nucleus job protocol (DD/PRINT) contains an ADAI63 message detailing the name of the Work data set that could not be opened.

**Action** Determine why opening the Work data set failed. Possible causes include the following:

- the Work data set or its catalog entry is inaccessible or has been damaged; or
- the PPT block containing the name of the Work data set has been damaged.
- If possible, correct the error and restart the nucleus. Otherwise, the database must be restored and regenerated.

If in doubt, contact your Software AG technical support representative.

**ADAX33      Bad work block for NUCID=*nucid* timestamp mismatch -- RABN=*rrrr***

**Explanation** While recovering from the failure of one or more Adabas cluster nuclei, this nucleus encountered a Work block in which the timestamp at the beginning of the block did not match the control timestamp at the end of the block. The last write of the block was incomplete, or the block has been damaged for another reason. It is inconsistent and cannot be used for recovery. The block was read from the Work data set of the nucleus with the NUCID '*nucid*'. Its RABN is '*rrrr*'.

**Action** Restore and regenerate the database. If in doubt, contact your Software AG technical support representative.

**ADAX40**      *dbid* Adabas abend in cache exit DBID *dbid* NUCID *nucid* abend routine Snnn  
 Unnnn reason *rsn* PSW *psw* REG 0-3 *reg 0 reg 1 reg 2 reg 3* REG 4-7 *reg 4 reg 5*  
*reg 6 reg 7* REG 8-B *reg 8 reg 8 reg 10 reg 11* REG C-F *reg 12 reg 13 reg 14 reg 15*

**Explanation**      A program check or system abend was intercepted by the z/OS XES parallel sysplex cache service FRR routine. The error occurred under an SRB.

The nucleus should terminate. A dump may be produced in SDUMP format, either in one of the nucleus-allocated files SYSUDUMP, SYSMDUMP, or SYSABEND; or in a system-allocated file such as SYS1.DUMPn.

**Action**              Contact your Adabas technical support representative.

**ADAX41**      *dbid* ADANCX GETMAIN failed

**Explanation**      ADANCX is the nucleus extension module for Adabas cluster environments. The attempt to allocate space for this module failed.

**Action**              Reduce memory requirements or expand the amount of memory available.

**ADAX42**      *dbid* GETMAIN failed

**Explanation**      An attempt to allocate space failed.

**Action**              Reduce memory requirements or expand the amount of memory available.

**ADAX43**      *dbid* {ADAXEC | ADASMC} initialization error - xxx

**Explanation**      Internal error.

**Action**              Contact your Software AG technical support representative.

**ADAX44**      *dbid* ADANCX cache-related GETMAIN failed

**Explanation**      ADANCX is the nucleus extension module for Adabas cluster environments. A second attempt to allocate space for this module failed.

**Action**              Reduce memory requirements or expand the amount of memory available.

**ADAX45**      *dbid* Unexpected cache connection error - xxx

**Explanation**      An error occurred while connecting to the sysplex cache structure. See the following messages for the reason.

**Action**              Correct the problem. If the action to take is not obvious, contact your Software AG technical support representative.

**ADAX46**     *dbid* GETMAIN failed after connect to cache

**Explanation** An attempt to allocate space failed after Adabas was connected to the external cache structure.

**Action** Reduce memory requirements or expand the amount of memory available.

**ADAX47**     *dbid* Cache connect problem RC X'xxxxxxxx' reason X'yyyyyyyy'

**Explanation** An error occurred while connecting to the sysplex cache structure. The 'x's identify cache-related return codes; the 'y's identify reason codes to explain the error.

**Action** If you are using Adabas Parallel Services, contact your Software AG technical support representative for assistance.

If you are using Adabas Cluster Services, see the IBM manual *MVS Programming: Sysplex Services Reference* for an explanation of the codes. Relevant information can be found in the Return and Reason Codes section of the macro IXLCONN. Common reason codes that occur due to configuration errors include the following:

RC	Reason	Explanation
X'08'	X'xxxx084C'	Improper SAF authorization. Adabas is not authorized to connect to the structure.
X'0C'	X'xxxx0C05'	Structure not defined in CFRM policy, possibly due to a bad structure name parameter.
X'0C'	X'xxxx0C08'	No suitable coupling facility found for structure allocation.
X'0C'	X'xxxx0C29'	The CFRM function is not active or not available.

**ADAX48**     *dbid* Cache disconnect RC *rrr* CRC X'xxxxxxxx'X reason X'yyyyyyyy'

**Explanation**     The message text of this message varies. The following table explains and describes the user action that should occur as a result of the different message texts.

Message Text	Explanation	Action
Cache disconnect RC <i>rrr</i> CRC X'xxxxxxxx'X reason X'yyyyyyyy'	An error occurred while disconnecting from the sysplex cache structure. The return codes from the ADAXEC module ( <i>rrr</i> ); the return codes from the cache structure ( <i>xxxxxxxx</i> ); and the reason codes ( <i>yyyyyyyy</i> ) are provided in the message to explain the error.	If you are using Adabas Parallel Services, contact your Software AG technical support representative for assistance.  If you are using Adabas Cluster Services, see the IBM manual <i>MVS Programming: Sysplex Services Reference</i> for an explanation of the codes. Relevant information can be found in the Return and Reason Codes section of the macro IXLDISC.
Block <i>blk-name</i> cast-out locked at disconnect <i>nnnn</i> cast-out locks released at disconnect	A Parallel Services nucleus held a cast-out lock on one or more cache blocks when disconnecting from the global cache. The cast-out locks have been released. This may happen during an online recovery process.  The first message is repeated for every block affected; the second one summarizes how many blocks were affected.	If these messages occur during online recovery, no action is required. However, if these messages appear during normal session termination, contact your Software AG technical support representative.

**Action**     Perform the action described in the table above.

**ADAX49**      *dbid Unexpected cache return code encountered*  
*dbid Function X'ff' xxxxxxxxxxxxxxxxxxxxxxxxxx*  
*dbid CRC X'yyyyyyyy' reason X'zzzzzzz'*

**Explanation**    An unexpected return code was returned during execution of a macro related to the cache structure.

**Action**            If you are using Adabas Parallel Services, contact your Software AG technical support representative for assistance.

If you are using Adabas Cluster Services, see the IBM manual *MVS Programming Sysplex Services Reference* for an explanation of the codes. Relevant information can be found in the Return and Reason Codes section of the macro IXLCACHE, in the chapter corresponding to the function displayed in the message.

**ADAX50**            *dbid GETMAIN failed for cast-out directory buffer*  
*dbid Size requested Xnnnnnnnn*

**Explanation**    An attempt to allocate space for the cast-out directory buffer failed.

**Action**            Reduce memory requirements or expand the amount of memory available.

**ADAX51**            *dbid Cache structure allocation unacceptable*  
*dbid Requested cache allocation values*  
*dbid Storage classes X'ee'*  
*dbid Cast-out classes X'ffff'*  
*dbid ADJUNCT=YES*  
*dbid Maximum data elements per entry X'gg'*  
*dbid Data element characteristic X'hh'*  
*dbid Actual cache allocation values*  
*dbid Storage classes X'ii'*  
*dbid Cast-out classes X'jjj'*  
*dbid ADJUNCT={YES | NO}*  
*dbid Maximum data elements per entry X'kk'*  
*dbid Data element characteristic X'mm'*

**Explanation**    The Adabas nucleus connected successfully to the cache structure; however, the attributes of the cache structure are inappropriate for the Adabas sysplex cluster. Most likely, the cache structure is already being used by another program.

**Action**            Determine which other program is using the cache structure. Either terminate this other program or specify a different cache structure for use by Adabas Cluster Services.

**ADAX52**      ***Idbid Incompatible existing user(s) of the  
dbid cache structure cache-structure-name***

**Explanation**    The cache structure with the name indicated is already being used by another Adabas cluster. Cache structure names may only be used for a single Adabas cluster of nuclei.

**Action**            Use the cache structure name identified for use by your particular cluster.

**ADAX53**      ***dbid Internal error - no available XQRB***

**Explanation**    An internal error has occurred.

**Action**            Contact your Software AG technical support representative for assistance.

**ADAX54**      ***dbid Insufficient cache data elements***

**Explanation**    The cache structure connected successfully; however, the number of cache data elements in the external cache structure (or global cache area) is not sufficient. There must be enough data elements to hold 80,000 or more bytes of information. The number of data elements allocated is indicated earlier in message ADAX57.

**Action**            Increase the size of the external cache structure (or global cache area). Alternatively, modify the ADARUN parameters DIRRATIO and/or ELEMENTRATIO to ensure that enough cache data elements are allocated.

**ADAX55**      ***dbid This job will now terminate***

**Explanation**    An internal error occurred that caused the nucleus to terminate abnormally. A message issued prior to this one provides more information related to the error.

**Action**            Contact your Software AG technical support representative for assistance.

**ADAX56**     *dbid message-text*

**Explanation**     Various message texts can be associated with the message number, as described in the following table:

Message Text	Explanation
Unable to obtain PLXCB address	The nucleus is unable to locate the PLXCB in common storage.
Unable to obtain ALET of XQRB data space	The nucleus could not connect to the messaging data space.
More than 31 XQRB areas in data space	The nucleus did not find free space for its own use in the messaging data space.

**Action**     For all possible message texts, contact your Software AG technical support representative for assistance.

**ADAX57**     *dbid Connected to cache structure cache-structure-name*  
*dbid Directory elements xxxxxx*  
*dbid Data elements yyyyyy*  
*dbid Data element size zzzzzz*

**Explanation**     An Adabas cluster nucleus successfully connected to the specified cache structure (or global cache area) in a cluster environment. This message provides a count of the cache's directory entries and data elements, along with the data element size.

**Action**     No action is required for this informational message.

**ADAX58**     *dbid Time expired waiting for notification of*  
*dbid existing connectors to the cache structure*

**Explanation**     The attempt to connect an Adabas cluster nucleus to the cluster cache structure in a sysplex environment timed out waiting for information about existing connections to the cache structure.

**Action**     Try again to start the Adabas sysplex cluster nucleus. If the error continues to occur, contact your Software AG technical support representative.

**ADAX59**     *dbid Unexpected return code from {ADAXEC | ADASMC}*  
*dbid function X'ff' xxxxxxxxxxxxxxxxxxxxxxxxxx*  
*dbid RC rrr*

**Explanation**     An unexpected return code was received during a call to the referenced module. The message includes the function code and return code.

**Action**     Contact your Software AG technical support representative.

**ADAX5A**     *dbid* DSST update omitted because cache or buffer pool full

**Explanation** After data storage updates, the nucleus is unable to read a DSST block into the buffer pool or global cache or to write the updated DSST block to the cache, because the buffer pool or cache was full. It is possible that the DSST bytes for one or more of the updated data storage blocks are incorrect.

This message is printed at most once in 10 minutes.

Incorrect DSST bytes may have slightly adverse effects on the effectiveness of space reuse in data storage. If the associated data storage blocks are updated again, incorrect DSST bytes are implicitly corrected.

**Action** Either the buffer pool or the global cache or both are too small to support the number of parallel commands running in the nucleus or the whole cluster. Increase the buffer pool or global cache size.

**ADAX5B**     *dbid* Connecting to S64 cache at address  
*dbid* Connect to S64 cache return code  
 ADAIOR-return-code/zOS-return-code/zOS-reason-code  
*dbid* Disconnecting from S64 cache  
*dbid* Disconnect from S64 cache return code  
 ADAIOR-return-code/zOS-return-code/zOS-reason-code

**Explanation** ADABAS Parallel Services is using z/OS shared 64-bit addressable storage as part of its cache configuration. Connecting to the S64 object establishes a local affinity and makes the object addressable. Disconnecting deletes the affinity, after which the object is no longer addressable. Any non-zero return code received from z/OS IARV64 is formatted.

**Action** If the attempt fails, examine the IARV64 return and reason code description in the IBM documentation entitled *MVS Programming: Authorized Assembler Services Reference, Volume 2 (EDTINFO-IXGWRITE)*. If the cause is not clear, notify your Software AG technical support representative.

**ADAX5C**     *dbid* Unchanged blocks will [NOT] be written to cache

**Explanation** This message reports the setting of the ADARUN CLUCACHEUNCHANGED parameter for this Adabas Parallel Services nucleus. Unchanged blocks either will or will not be written to the global cache.

**Action** No action is required for this informational message.



**ADAX5D**     *dbid Invalid cache token tt...tt at disc*

**Explanation**     A logic error occurred. An invalid token was passed to the cache disconnect function. If the disconnect is already identified as abnormal, termination continues. If not, the nucleus ends with abend S0C6.

**Action**             Contact your Software AG technical support representative for assistance.

**ADAX60**             **{Peer nucleus | Unknown connector} *connection-name***  
**{is already | has} connected to**  
**{cache | lock} structure *structure-name***

**Explanation**     An event occurred related to another connector of a cache or lock structure.

**Action**             No action is required for this informational message.

ADAX61     *date time statistic-text statistic-value*

**Explanation** Whenever a nucleus disconnects from a cache, this message is issued with cache activity statistics. These may be useful in diagnosing problems and selecting a suitable cache configuration. Bear in mind that the statistics show only one nucleus's activity. Full understanding may require examining the statistics of other nuclei.

This message is also used to display the output from the DXSTAT, DXCACHE, DXLOCK, DXFILE, and DMEMTB operator commands on the console. For examples of this message used with these commands, refer to the documentation for Adabas Cluster Services operator commands.

Here is an example of the statistics provided in this message:

```
ADAX61 Statistics at disconnect for cache C0007231
Cache Directory Management Activity
33,545 Read Located active
5,230 Read obtained from free pool
0 Read reclaimed, first choice criteria
0 Read reclaimed, second choice criteria
0 Read reclaimed, third choice criteria
0 Read reclaimed, fourth choice criteria
0 Read unable to obtain (cache full)
0 Total number of directory reclaim attempts
0 Number of directories examined for reclaim
0.000 Average number of directories examined per attempt
680 Write obtained from free pool
496 Directory high-water mark (this nucleus)
1,134 Directory high-water mark (cluster-wide)

Cache Write Requests
309,756 Sufficient preallocated space
235 Free space allocated
5 Space reclaimed, first choice criteria
6,322 Space reclaimed, second choice criteria
885 Unable to obtain space (cache full)
4,252 Space search examined part of space chain
3,155 Space search examined entire space chain
7,447 Total number of space allocation attempts
462,840 Number of space chain descriptors examined
62.151 Average number of descriptors examined per attempt

Cache Space Element Reclamation
48 Elements reclaimed, first choice criteria
28,431 Elements reclaimed, second choice criteria

Latch Attempts
Exclusive Shared
Cache Space Chain
Get 7,456
WaitFor 8
Release 7,447
Cache Directory Index
Get 5,351 249,009
WaitFor 5,228 Upgrade 0
Release 10,579 Upgrade 243,781
Cache Directory
Get 454,179 398,747
WaitFor 5,335 Upgrade 53
Release 459,297 393,339
Cache Cast-Out Class
Get 44,865 3,510
WaitFor 0 0
Release 44,865 3,510
```

These statistics are described in the following table:

Message Statistic Type	Description
Cache Directory Management Activity	<p>Cache directory elements describe ASSO and DATA blocks that have been referenced. Directory elements are also used to describe unallocated space in the cache data storage pool. The number of directory elements created during cache initialization is shown in message ADAX57.</p> <p>The first reference to an ASSO or DATA block is a cache read request. If the block is already known in the cache, the existing directory element is located. Otherwise, a directory element must be allocated. If there are none in the free pool, the nucleus must search for an existing element that may be reclaimed. Directory elements for blocks that have been modified ("dirty") and not yet written to external storage ("hardened") may not be reclaimed. The remaining directory elements are examined to select one that describes a block that is least likely to be needed in the near future. There are four sets of criteria in order of decreasing desirability. In the best case, there is less likelihood of additional I/O from a cache miss, and in the worst case we have reclaimed an element that describes a block likely to be reused, and thus there's a higher probability of additional I/O. Factors such as the number of nuclei that have referenced the block, whether the block has been referenced recently, and the status of the block's contents are part of the selection criteria.</p> <p>If all directory elements describe dirty blocks, the cache is considered full and the directory allocation fails. The nucleus then executes a buffer flush to harden the dirty blocks, after which the request is tried again.</p> <p>A write request may also attempt to allocate a directory element to describe fragmented space. This is not a frequent occurrence.</p> <p>If reclaimed directory elements are a significant portion of the total activity, especially reclaims from the higher criteria, consider increasing the number of directory elements for best performance.</p> <p>The cache control structures are rebuilt as part of online recovery after a nucleus fails. This will reset the cluster-wide high water mark but will not affect the individual member high water mark.</p>
Cache Write Requests	<p>If the directory element does not already describe sufficient space to contain the block, additional space is first obtained from the free pool. If a suitable block cannot be found, space from one or more existing blocks will have to be reclaimed to assemble a contiguous area. If no suitable area can be assembled without including space containing dirty blocks, the cache is considered full. The nucleus executes a buffer flush and retries the request. As with directory elements, the selection criteria are used to reduce the chance that additional I/O will be needed because an active block was removed.</p>
Cache Space Element Reclamation	<p>The cache data element is shown in message ADAX57. A cached ASSO or DATA block requires one or more data elements.</p>
Latch Attempts	<p>Access to data structures shared among multiple nuclei is serialized by high-performance latches. These are administered entirely by the nuclei—they are not z/OS-administered latches. A latch may be acquired either exclusive or shared, and a shared latch may be upgraded to exclusive. If the latch cannot be acquired immediately, the nucleus may elect to wait for it.</p> <ul style="list-style-type: none"> <li>• There is one cache space chain latch to serialize space allocations.</li> <li>• There is a directory element index latch for each index element. It is held when searching for an existing directory element. It is also held when allocating or releasing a directory element.</li> <li>• There is a separate latch to serialize updates for each directory element.</li> <li>• There is a separate latch for each cast-out class to serialize access to directory elements during buffer flushes.</li> </ul>

**Action** No action is required for this informational message.

**ADAX62 Unexpected lock return code encountered  
function X'xx'  
LRC X'ccccccc' reason X'rrrrrrrr'**

**Explanation** An attempt to lock or unlock a logical resource failed. The lock manager function number is 'xx'; its response code is ccccccc; and its reason code is rrrrrrrr. All variables are in hexadecimal.

The nucleus terminates abnormally.

**Action** This is an unexpected error. Contact your Software AG technical support representative for assistance.

If you are using Adabas Cluster Services, see the IBM manual *MVS Programming: Sysplex Services Reference* for an explanation of the codes, . Relevant information can be found in the Return and Reason Codes section of the macro IXLLOCK.

**ADAX63 Lock connect problem RC X'ccccccc' reason X'rrrrrrrr'**

**Explanation** An error occurred while connecting to the lock structure. The lock manager's response code is ccccccc, its reason code is rrrrrrrr. Both variables are in hexadecimal.

The nucleus terminates abnormally.

**Action** This is an unexpected error. Contact your Software AG technical support representative for assistance.

If you are using Adabas Cluster Services, see the IBM manual *MVS Programming: Sysplex Services Reference* for an explanation of the codes. Relevant information can be found in the Return and Reason Codes section of the macros IXCQUERY and IXLCONN. Common reason codes that occur due to configuration errors include the following:

RC	REASON	Explanation
X'08'	X'00000024'	Structure not defined in CFRM policy.
X'08'	X'xxxx084C'	Improper SAF authorization. Adabas is not authorized to connect to the structure.
X'0C'	X'00000154'	No CFRM policy active.
X'0C'	X'xxxx0C05'	Structure not defined in CFRM policy, possibly due to a bad structure name parameter.
X'0C'	X'xxxx0C08'	No suitable coupling facility found for structure allocation.
X'0C'	X'xxxx0C29'	The CFRM function is not active or not available.

**ADAX64      ADANCX lock-related GETMAIN failed**

**Explanation** The lock manager failed to acquire main storage.

**Action** Restart the nucleus with a larger REGION parameter or make the nucleus parameters NH, NU, or LDEUQP smaller.

**ADAX65      Parameter taken over: *parameter-name*  
old: *old-value* new: *new-value***

**Explanation** This message is printed in two instances. It occurs when:

1. The parameter value of a newly started cluster nucleus is taken over (replaced) by the global parameter value in use by the already active cluster nuclei. The old value (*old-value*) is the parameter value read from the ADARUN statements of the starting nucleus; the new value (*new-value*) is the global parameter value in use by the already active cluster nuclei.
2. A global parameter value is changed on one nucleus in the cluster. This new parameter value is propagated to all of the other active cluster nuclei and is taken over (used) by them.

**Action** No action is required for this informational message.

**ADAX66      Incompatible global parameter *parameter-name*  
specified: *value-specified* in effect: *value-in-effect***

**Explanation** An incompatible parameter value has been specified for a second or subsequent cluster nucleus when it starts. When this occurs, the starting nucleus:

- Does not take over (use) the global parameter value.
- Does not come up.

The parameter value specified (*value-specified*) by the second or subsequent cluster nucleus and the parameter value currently in effect (*value-in-effect*) for all cluster nuclei are given in the message.

**Action** Resolve the incompatibility of the parameter values so that the newly started cluster nucleus and the other active cluster nuclei can use the same parameter value.

**ADAX67      Incompatible existing user(s) of the  
lock structure *lock-structure-name***

**Explanation**    The lock structure with the name specified in the message is already being used by another Adabas sysplex cluster or by other software. Lock structure names must be unique in the sysplex environment and for each Adabas nucleus cluster.

**Action**            Use the lock structure name that has been identified for use by your particular Adabas sysplex cluster.

**ADAX68      Time expired waiting for notification of  
existing connectors to the lock structure**

**Explanation**    An attempt to connect an Adabas sysplex cluster nucleus to the lock structure timed out waiting for information about existing connections to the lock structure.

**Action**            Determine whether any conditions exist in the coupling facility or the system itself that could interrupt the flow of information or make the flow extremely slow.

**ADAX69      Lock structure too small  
expected min number of records *nnn,nnn,nnn***

**Explanation**    The lock structure defined in the CFRM policy is too small to handle the current settings of the ADARUN parameters. The minimum number of records expected by the processes is indicated.

**Action**            Either decrease the value of the ADARUN parameters NU, NH, or LDEUQP; or increase the size of the lock structure.

**ADAX6B** **IXLEERSP request got RSP *rsp-code* RSN *reason-code***

**Explanation** After a peer nucleus disconnected abnormally from the lock structure, this nucleus issued an IXLEERSP request to XES to acknowledge the peer failure, but the IXLEERSP request got the response code and reason code shown. These are documented in *MVS Programming: Sysplex Services Reference* IBM manual in the *Return and Reason Codes* section of the IXLEERSP macro.

This nucleus will continue performing the online recovery process after the peer failure. It will disconnect from the lock structure as part of the recovery process, which is another way of acknowledging the peer failure to XES. However, due to the unsuccessful IXLEERSP request, there is a very small chance of a cluster-wide deadlock involving XES.

**Action** In the case of a hangup of the online recovery process, cancel a nucleus that has not yet printed an ADAX89 message. This may unravel the deadlock. If there is no such nucleus, cancel the nucleus that will do the session autostart, as per the ADAX89 messages. Restart the cluster after all nuclei have terminated.

In either case, report the occurrence of the ADAX6B message to your Software AG technical support representative.

**ADAX70** **Connected to lock structure *lock-structure-name*  
number of lock entries *nnn,nnn*  
max number of record elements *nn,nnn***

**Explanation** An Adabas cluster nucleus successfully connected to the specified lock structure in a z/OS parallel sysplex environment. This message provides

- a count of lock entries; and
- the maximum number of records elements.

**Action** No action is required for this informational message.

**ADAX71** **Retrying cache write for RABN X'*rrrrrrrr*'**

**Explanation** This message identifies the RABN value in hexadecimal for which a cache write is being retried. The error leading to the write retry is identified in previous messages.

**ADAX72** **GETMAIN failed for lock element table  
size requested X*nnnnnn***

**Explanation** An attempt to allocate space for the lock element table failed.

**Action** Reduce the memory requirements for the table or expand the amount of memory available.

**ADAX73      Lock structure size error**

**Explanation** An error was detected in the lock structure policy: both the SIZE and INITSIZE values are zero.

**Action** Review the lock structure policy and make the necessary changes.

**ADAX74      *dbid* Warning: Now it is too late to copy DDPLOGRn**

**Explanation** Corresponds to the ADAN05 message at startup, but occurs during online recovery.

Adabas has begun to write data protection log data to the data set identified by DD/PLOGRn. This means that the data set can no longer be copied to tape for subsequent use as input to the REGENERATE or BACKOUT functions of the ADARES utility. A user exit 2 (dual log processing) or user exit 12 (multiple log processing) call either was not made or did not successfully copy the DD/PLOGRn data set with the ADARES utility.

**Action** If the database is running without user exit 2, overwriting the PLOG data is normal and this message can be ignored.

If the database is running with user exit 2, this message occurs only when the user exit asks the nucleus to proceed even though the PLOG has not been copied. Whether this is an error or not depends on the logic the user has implemented in the user exit.

**ADAX75      *dbid* Protection log PLOGRn started**

**Explanation** Corresponds to the ADAN21 message at startup, but occurs during online recovery.

Adabas is now ready to begin writing data protection information to the dual or multiple data protection log identified by DD/PLOGRn.

**Action** Execute the PLCOPY function of the ADARES utility at this time to reinitialize the PLOGs.

**ADAX76      *dbid* Nucleus run with protection log nnnnn**

**Explanation** Corresponds to the ADAN02 message at startup, but occurs during online recovery.

The Adabas nucleus session has been initiated and database protection logging has been specified. Subsequent execution of the REGENERATE and BACKOUT functions of the ADARES utility for any updates applied during the session is possible.

**ADAX77**     *dbid* IDTH prefix problem

**Explanation**   A query request to ADAMPM to get the address of the IDTH failed. The nucleus terminates abnormally.

**Action**        This is an unexpected error. Contact your Software AG technical support representative.

**ADAX78**     *dbid* ADACOM is not running or CLULOCKSIZE is not specified

**Explanation**   While trying to connect to the global lock area, ADASML detects that either ADACOM is not running or the CLULOCKSIZE parameter is not specified. The nucleus terminates abnormally.

**Action**        Start ADACOM before starting the Adabas Parallel Services cluster nucleus and specify a nonzero value for the CLULOCKSIZE parameter.

**ADAX79**     *dbid* Global resource lock on this system is invisible to NUCID *nucid* on the system *sysname*

**Explanation**   Cluster Services nuclei working on the same database synchronize some of their actions using resource locks via Global Resource Serialization (GRS).

During session start one nucleus detected that a resource lock it obtained was not effective against the peer nucleus with the NUCID shown, which was active on the system shown.

The starting nucleus terminates with parm-error 105.

**Action**        Contact your system programmer to ensure that GRS is configured in a way that GRS resource locks are mutually effective against one another on all systems on which you intend to run Cluster Services nuclei.

**ADAX80**     *dbid* Online recovery initiated

**Explanation**   An Adabas cluster nucleus initiated an online recovery process after it detected that a peer nucleus in the same cluster terminated abnormally. (Each surviving nucleus initiates its own online recovery process.) The online recovery process stops all ongoing work in the nucleus, performs a session autorestart (including the backout of all open transactions), or waits until a peer nucleus performs the session autorestart, and then resumes normal processing.



**ADAX80**     *dbid* **Online recovery in progress**

**Explanation**   A nucleus started while other nuclei that were already active in the same cluster were performing online recovery in response to a nucleus failure. The starting nucleus waits until the online recovery process completes and then continues with its startup sequence.

**ADAX80**     *dbid* **{Online save | Trans suspend | ADAEND/halt} process canceled**

**Explanation**   In order to recover from the failure of a peer nucleus (online recovery), the nucleus canceled

- a running online save operation in which case the save operation fails;
- a running transaction suspension operation; or
- an ADAEND or HALT shutdown request, in which case the nucleus does not shut down after the recovery process has finished.

**Action**        Either restart the save operation after the online recovery process has completed successfully; or issue another ADAEND or HALT request if you still want to shut down the nucleus.

**ADAX81**     *dbid* **Waiting for active transactions to finish**

**Explanation**   When the online recovery process started, one or more transactions were active. The recovery process allows them to continue for a while in an attempt to bring them to normal completion.

**ADAX82**     *dbid* **All transactions finished**

**Explanation**   All transactions that were active when online recovery started have finished normally.

**ADAX82**     *dbid count* **active transaction(s) interrupted**

**Explanation**   A number of transactions indicated in the message were active when online recovery started but did not finish within the allotted time and were interrupted. They are backed out during online recovery. The affected users receive response code 9 (ADARSP009), subcode 18, for their next commands.

**ADAX83**     *dbid* **Waiting for active commands to finish**

**Explanation**   When the online recovery process was ready to interrupt all ongoing work, one or more commands were still active. The recovery process allows them to continue for a short time in an attempt to bring them to normal completion.

**ADAX84**     *dbid* All commands finished

**Explanation**     All active commands that the online recovery process was waiting for have finished normally.

**ADAX84**     *dbid count* Active command(s) interrupted

**Explanation**     A number of active commands indicated in the message did not finish within the allotted time and were interrupted. They are sent back to their respective users with response code 9 (ADARSP009), subcode 19. Their associated command IDs, if any, are deleted.

**ADAX85**     *dbid* Waiting for active I/Os to finish

**Explanation**     When the online recovery process interrupted all ongoing work, one or more I/Os were active. The recovery process waits for these I/Os to finish.

**ADAX86**     *dbid* All I/Os finished

**Explanation**     All I/Os that the online recovery process was waiting for have finished.

**ADAX87**     *dbid* Waiting for outstanding messages to be answered

**Explanation**     When the online recovery process interrupted all ongoing work, one or more internucleus commands were still due a response. The recovery process waits for the responses to arrive.

**ADAX88**     *dbid* All outstanding messages answered

**Explanation**     All outstanding responses for internucleus commands that the online recovery process was waiting for have arrived.

**ADAX89**     *dbid* Session autorestart will be done by {this | peer} nucleus

**Explanation**     The session autorestart that is part of the online recovery process is performed either by this nucleus or by a peer nucleus, as indicated in the message.

**ADAX90**     *dbid Recovery syncpoint syncpoint initiated*

**Explanation** If more than one nucleus remains active when a peer nucleus terminates abnormally, the surviving nuclei synchronize their online recovery processes using several syncpoints, which all nuclei must reach before recovery processing can continue.

This message indicates that the nucleus that performs the session autorestart is ready to proceed when all other nuclei have reached the respective syncpoint.

**ADAX91**     *dbid Waiting on recovery syncpoint syncpoint*

**Explanation** The online recovery process is waiting for the nucleus that performs the session autorestart to initiate the recovery syncpoint indicated.

**ADAX92**     *dbid Recovery syncpoint syncpoint reached*

**Explanation** All nuclei involved in the collaborative online recovery have reached the recovery syncpoint indicated. The recovery process proceeds.

**ADAX93**     *dbid Beginning session autorestart*  
*dbid Beginning WORK4 interpretation*  
*dbid WORK4 handling failed*

**Explanation** Various message texts can appear for this message number. The explanation and action for each is given in the following table.

Message Text	Explanation
Beginning session autorestart	One of the nuclei surviving a peer failure (this nucleus) begins the key step of online recovery-the session autorestart.
Beginning WORK4 interpretation	One of the peer nuclei failed, so one of the surviving nuclei begins autorestart processing, if DTP=RM. The WORK4 interpretation must occur before the autorestart is executed.
WORK4 handling failed	WORK4 interpretation was not successful. Refer to messages ADAN85 and ADAN86 for more information. All nuclei will go down.

**ADAX94**     *dbid* Session autorestart executed successfully  
*dbid* DTP=RM-USERS are copied  
*dbid* DTM=RM-USER-COPY failed  
*dbid* DTM=RM-USER-LOCKS not gotten

**Explanation**     Various message texts can appear for this message number. The explanation and action for each is given in the following table.

Message Text	Explanation
Session autorestart executed successfully	The session autorestart performed during online recovery was successful. No action is required for this informational message.
DTP=RM-USERS are copied	PET/HEURI users from other nuclei are copied to the user queue of the recovery nucleus (including transaction IDs and file lists). No action is required for this informational message.
DTM=RM-USER-COPY failed	The user copy failed, probably due to a logic error. All nuclei will go down.
DTM=RM-USER-LOCKS not gotten	If DTP=RM, the nucleus must get hold queue and UQDE locks for all users on PET status who are copied to DDWORKR4. All nuclei will go down.

**ADAX95**     *dbid* Session autorestart failed  
*dbid* Response code = *response-code*  
*dbid* File number = *file-number*  
*dbid* All active nuclei will go down

**Explanation**     The session autorestart performed during online recovery was not successful. It received the response code shown. If the response code was associated with a particular file, the file number is also shown.

This nucleus and all peer nuclei participating in the online recovery process will go down.

**Action**             The situation is now equivalent to that after failure of session autorestart during nucleus session start. Determine why the session autorestart failed. Consider contacting your Software AG technical support representative.

**ADAX96**     ***dbid* Peer nucleus failed during online recovery**  
***dbid* This nucleus goes down too**

**Explanation**   A second nucleus failure occurred while an online recovery process was in progress to handle the abnormal termination of a peer nucleus. All nuclei active in the Adabas cluster will go down.

**Action**        Restart the Adabas cluster. Determine the reasons for the first and the second failure. Consider contacting your Software AG technical support representative.

**ADAX96**     ***dbid* Utility with exclusive database control is active**  
***dbid* This nucleus goes down too**

**Explanation**   A nucleus failure occurred while a utility with exclusive database control was running. All nuclei active in the Adabas cluster will go down.

**Action**        Restart the Adabas cluster and perform appropriate recovery actions for the utility with exclusive database control.

**ADAX97**     ***dbid* Online recovery completed successfully**  
***dbid* Resuming normal operation**

**Explanation**   The online process set up to handle the abnormal termination of a peer nucleus finished successfully. The nucleus resumes normal operation.

**ADAX98**     ***dbid* Received response code *rsp-code* from peer nucleus**

**Explanation**   An online recovery process that was started to recover from the failure of one nucleus received a response code while communicating with another, still alive nucleus. All remaining active nuclei terminate.

**Action**        Restart the nuclei. The first starting nucleus performs offline recovery (that is, session autorestart).

**ADAX98**     ***dbid* V2/xxx command received *rsp-rr/ss* from NUCID *nnn***

**Explanation**   An internal command used for inter-nucleus communication encountered a messaging failure; it got the response code/subcode shown from the peer nucleus shown.

**ADAX99**     *dbid* **Uncorrectable intracluster communication failure**

**Explanation** This message may follow message ADAX9E or ADAX9H. After a messaging failure during intracluster communication the nucleus has decided to terminate itself due to an uncorrectable communication failure.

The nucleus terminates abnormally with user abend code 79.

**Action** Investigate the cause of the intracluster communication failure, starting with the response code and subcode reported in one or more preceding ADAX9E messages.

If some kind of timeout (MXMSGWARN, MXMSG, MXCANCELWARN, or MXCANCEL parameter) was involved in the failure, ensure that all cluster nuclei run with high enough priority to get sufficient resources for participating in cluster-wide business.

Consider contacting your Software AG technical support representative.

**ADAX9A**     *dbid* **Could not determine message status for V2/xxx command return code = nn**

**Explanation** When a cluster nucleus tried to determine the status of an internal intracluster command (given in the message), an error occurred. The internal return code is shown. A preceding message might possibly contain more information about the error.

The nucleus ignores the error and continues normally. Because it could not determine the status of the responses to the intracluster command cited, it does not warn early about outstanding responses (ADAX9B and ADAX9C messages)

**Action** Contact your Software AG technical support representative.

**ADAX9B**     *dbid* **Caution: Waiting for V2/xxx CMD being sent to NUCID (nn)**

**Explanation** A cluster nucleus issued an XCF send request for an intracluster command (shown) to a peer nucleus (shown), but the command has not yet been sent within the time period set by the MXMSGWARN parameter. The number nn (shown) distinguishes different internal intracluster commands that may be in progress at the same time.

No direct system action occurs. However, if the command is not sent or the peer nucleus does not respond within the time period set by the MXMSG parameter, either nucleus may terminate abnormally.

**Action** This message is for your information. It may be useful for analysis in the case of a subsequent error.

**ADAX9C**     *dbid* **Caution: NUCID *nucid* (*jobname*) on system *sysn* is slow to respond to internal V2/*xxx* command (*nn*)**

**Explanation**     A cluster nucleus sent an internal intracluster command (shown) to a peer nucleus (shown) in the cluster. The peer nucleus has not yet responded to the command within the time period set by the MXMSGWARN parameter. The number *nn* (shown) distinguishes different internal intracluster commands that may be in progress at the same time.

No direct system action occurs. However, if the peer nucleus does not respond within the time period set by the MXMSG parameter, it will be canceled.

**Action**     This message is for your information. You may want to use the provided information (NUCID, job name, system name) to make sure that the cited peer nucleus gets sufficient resources (CPU, storage, priority) for participating in cluster-wide business.

**ADAX9D**     *dbid* **Clear: received response to V2/*xxx* from NUCID (*nn*)**

**Explanation**     The peer nucleus listed in the message (*nucid*) that was slow to respond to an internal intracluster command (*xxx*) has finally responded. This message retracts the warning of a previous ADAX9B or ADAX9C message (shown with a matching *nn* number).

The peer nucleus is no longer in danger of being canceled (for this particular intracluster command).

**Action**     No action is required for this informational message.

**ADAX9E**     *dbid* **Error: V2/*xxx* cmd received RSP *rsp/sub* from NUCID *nucid***

**Explanation**     The internal command listed in the message (*xxx*) and used for intracluster communication encountered a messaging failure. The response code and subcode as well as the peer nucleus are given in the message. The response code and subcode are described in *Nucleus Response Codes*.

The reaction of the nucleus to this error varies, depending on the type of internal command and on the response code received. The nucleus may cancel the peer nucleus causing the error, wait for a failing peer nucleus to terminate, terminate itself due to an unhandled error condition, or pass the response code up the call chain.

**Action**     Investigate the cause of the response code and subcode to resolve the error.

If some kind of timeout (MXMSGWARN or MXMSG parameter settings) is involved in the error, ensure that all cluster nuclei run with high enough priority to get sufficient resources for participating in cluster-wide business.

If the problem persists, contact your Software AG technical support representative.

**ADAX9F**     *dbid* Canceling peer nucleus *nucid* (*jobname*) on system *sysn*

**Explanation** This message may follow message ADAX9E. A failure occurred when a nucleus issued an intracluster command to a target peer nucleus (listed in the message) because the target nucleus did not respond to the command in time. As a result, the sending nucleus is canceling the target peer nucleus.

The nucleus waits for the canceled peer nucleus to terminate and then performs an online recovery process.

**Action** Investigate the cause of the intracluster communication failure, starting with the response code and subcode reported in one or more preceding ADAX9E messages.

If some kind of timeout (MXMSGWARN or MXMSG parameter setting) was involved in the failure, ensure that all cluster nuclei run with high enough priority to get sufficient resources for participating in cluster-wide business.

If the problem persists, contact your Software AG technical support representative.

**ADAX9G**     *dbid* **Caution: NUCID *nucid* (*jobname*) on system *sysn* was canceled; Failure notification still outstanding**

**Explanation** The nucleus identified in the message by its nucleus ID (*nucid*), job name (*jobname*) and system name (*sysn*) was canceled, but the nucleus printing this message did not receive (within the time period set by the MXCANCELWARN parameter) the failure notifications reported in ADAX60 messages confirming that the canceled nucleus was terminated. As long as the canceled nucleus might still be active, the Adabas cluster cannot recover from the failure.

This message might also occur if the cited nucleus has not been canceled but is failing for another reason and is slow to terminate.

If the peer nucleus does not terminate within the time period set by the MXCANCEL parameter, this nucleus may ask for permission to terminate itself (if the MXWTOR parameter was set) or terminate itself without asking (if the MXWTOR parameter was not set).

**Action** Use the provided information (nucleus ID, job name, and system name) to make sure that the canceled peer nucleus gets sufficient resources (CPU, storage, priority) for terminating.



**ADAX9H**     *dbid* Error: Canceled NUCID *nucid* (*jobname*) on system *sysn* has not terminated; unable to perform cluster recovery

**Explanation**     The nucleus identified in the message by its nucleus ID (*nucid*), job name (*jobname*) and system name (*sysn*) was canceled, but the nucleus printing this message did not receive (within the time period set by the MXCANCEL parameter) the failure notifications confirming that the canceled nucleus has terminated. As long as the canceled nucleus might still be active, the Adabas cluster cannot recover from the failure.

This message might also occur if the cited nucleus was not canceled but failed for another reason and has not terminated.

The nucleus printing this message terminates itself with message ADAX99 and userabend 79.

**Action**             Check the status of the cluster nucleus cited in the message. Investigate why it did not terminate after being canceled or why the nucleus that printed this message did not receive the corresponding failure notifications (reported in ADAX60 messages).

Consider using the MXWTOR parameter to request that a cluster nucleus ask for permission before terminating itself after failing to cancel an unresponsive peer nucleus.

Ensure that all cluster nuclei run with high enough priority to get sufficient resources for participating in cluster-wide business.

If the problem persists, contact your Software AG technical support representative

**ADAX9J** *dbid* Error: Canceled NUCID *nucid* (*jobname*) on system *sysn* has not ended yet. Ensure that this nucleus ends to allow Adabas cluster recovery. Will terminate at *hh:mm:ss* (after *nnn* seconds). Reply 'W'ait, 'T'erminate, or 'R'eshow message

**Explanation** This message requests an operator response. The nucleus identified in the message by its nucleus ID (*nucid*), job name (*jobname*) and system name (*sysn*) was canceled, but the nucleus printing this message did not receive (within the time period set by the MXCANCEL parameter) the failure notifications reported in the ADAX60 messages confirming that the canceled nucleus has terminated. As long as the canceled nucleus might still be active, the Adabas cluster cannot recover from the failure.

This message might also occur if the cited nucleus was not canceled but failed for another reason and has not yet terminated.

The nucleus will wait for the time period set by the MXWTOR ADARUN parameter for either the failure notifications of the canceled peer nucleus or a response from the operator. If the nucleus receives the expected failure notifications of the canceled peer nucleus, it retracts the ADAX9J message and starts an online recovery process to recover from the failure.

**Action** Check the status of the other cluster nucleus cited in the message. If it terminates, this ADAX9J message will be retracted.

Respond to this message using one of the following responses:

Response	Causes the nucleus to:
R	Print the whole ADAX9J message again and continue to wait for resolution of this issue, but without setting a new time period.
T	Terminate itself with message ADAX99 and user abend 79.
W	Wait for another time period of length MXWTOR for resolution of this issue.

If you do not respond and the failure notifications of the canceled peer nucleus do not arrive by the MXCANCEL time, the nucleus terminates itself with messages ADAX9H and ADAX99 and user abend 79.

**ADAX9K** *dbid* Clear: Received failure notifications from NUCID *nucid*

**Explanation** The peer nucleus identified in the message (*nucid*) that was slow to terminate after being canceled has finally ended. This message retracts the warning of previous ADAX9G or ADAX9J messages about outstanding failure notifications. The nucleus is no longer in danger of terminating itself (for this particular incident).

**Action** No action is required for this informational message.

**ADAX9L**     *dbid All operator queries retracted*

**Explanation**     After the operator was asked in a previous ADAX9J message to check the status of a peer nucleus that had been canceled but did not terminate, the peer nucleus has now ended. The nucleus printing this message has retracted all outstanding ADAX9J operator queries.

**Action**             No action is required for this informational message.

**ADAXS4**     *Cluster nucleus cache-related storage above the 2gig baris [NOT] backed by large pages*

**Explanation**     Two different messages may appear using this message ID, depending on whether the word NOT appears in the message text. These messages appear because the LARGEPAGE ADARUN parameter was set to "YES" for your cluster nuclei and these messages indicate the ability of your operating system to support them.

If the message indicates that cluster nucleus cache-related storage above the 2 gigabyte bar is *NOT* backed by large pages, the LARGEPAGE ADARUN parameter was set to "YES", but either the system does not support large pages or insufficient 1-megabyte pages were available to support the request.

If the message indicates that cluster nucleus cache-related storage above the 2 gigabyte bar *is* backed by large pages, the LARGEPAGE ADARUN parameter was set to "YES", and is fully supported by the operating system.

**Action**             If the message indicates that cluster nucleus cache-related storage above the 2 gigabyte bar is *NOT* backed by large pages, consider the following actions:

- If your system does not support large pages, do not use the LARGEPAGE ADARUN parameter. Remove it from your cluster nuclei startup JCL and restart the nuclei
- If your system does support large pages, consider increasing the number of 1-megabyte pages specified for your operating system using the LFAREA parameter in PARMLIB member IEASYxx. Contact your system administrator for assistance. For more information, read your IBM *MVS Initialization and Tuning* documentation.

If the message indicates that cluster nucleus cache-related storage above the 2 gigabyte bar *is* backed by large pages, no action is required for this informational message.