

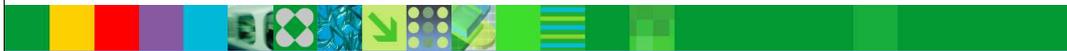


IBM Software Group

Informix Dynamic Server

Enterprise replication - Changes and new features in V10.00.xC9

IBM Information Management software



@business on demand.

© 2009 IBM Corporation
Updated June 15, 2009

This presentation will provide an overview of the changes and new features within Enterprise Replication in the 10.00.xC9 version.

Major changes for V10.00.xC9

- **Memory changes**
 - ▶ Allocation during cdr sync
 - ▶ New cdr sync options
- **Sync and check changes**
 - ▶ Sync and validate
 - ▶ Reduce false failures of in-flight data
- **Error message clean up**
 - ▶ Removal of -1 errors in all ER code
 - ▶ More descriptive messages for certain functions
- **Recompress rows in grouper paging file**
- **Miscellaneous changes**

The goal of this presentation is to provide details regarding the changes and new features introduced with 10.00.xC8W3 and 10.00.xC9. These features were added based on feedback from business partners after using ER within an embedded environment.

The major changes in this release include memory, sync and check changes, monitoring, Error Message clean up, recompress rows in grouper paging file, and miscellaneous other changes.

The first area of topic is Memory changes made to the 10.00.xC9 release.

Memory changes – Overview

- There are two main areas of memory changes:
 - ▶ cdr sync
 - Previous releases
 - Running cdr sync, ER would dynamically allocate memory by increasing the CDR_QUEUEMEM internally. In embedded systems, this has caused memory limitation issues for both ER and the engine.
 - ER can behave erratically in certain low memory situations, causing Assert Failures, component failure, or engine abends
 - V10.00.xC9
 - Controlling memory usage during a sync gives customers flexibility with memory resources
 - Running cdr sync doesn't interfere with other engine and ER activities

There are two main areas where Memory Changes were focused.

During a **cdr sync**, ER would dynamically allocate memory by increasing CDR_QUEUEMEM internally, which caused problems for embedded systems with memory limitations.

ER can behave erratically when encountering memory limitations – causing AF's, or worse. Customers needed more control over resources used during a **cdr sync**. Running a **cdr sync** should not interfere with other engine or ER operations.

Memory changes – -memadjust (-J) option

- As of V10.00.xC9, a new option is available for controlling memory usage with several ‘sync’ features
 - ▶ Long Form –memadjust={size K|M}
 - ▶ Short Form -J {size K|M}
 - ‘–memadjust=’ size must be a positive number, but not larger than the available memory - bufferpool size
 - ▶ Other available options that can be used with cdr:
 - cdr start replicate
 - cdr realize template
 - cdr sync replicate
 - cdr sync replicateset

A new option named ‘--memadjust=’ or ‘-J’ was added to control memory allocation during a sync operation.

You can also use this option with **cdr start replicate**, **cdr realize template**, and with **cdr sync** for either a single replicate or a replicate set.

The ‘–memadjust=’ size must be a positive number, not exceeding available memory.

Memory changes – -memadjust details

- Current documentation states:
 - ▶ The --memadjust= size will not
 - Limit the amount of memory that the send queue can use during synchronization
 - Specify the size of the send queue
 - ▶ The size specified is temporally added to the current value of CDR_QUEUEMEM
 - The duration of the change remains until the sync completes
 - If no size is provided, the default value is based on the calculation:
 - $512 * SHMVIRT SIZE$

5

Enterprise replication Changes and new features in V10.00.xC9

© 2009 IBM Corporation

The information found in the documentation regarding the --memadjust= option is a bit misleading. In reality, the size specified with this option is temporally added to the value of CDR_QUEUEMEM. This value is removed after the sync operation completes.

If this option is not used, a default value for the sync process is calculated as $(512 * SHMVIRT SIZE)$.

Memory changes – -memadjust details

- ▶ Contrary to the documentation, --memadjust cannot be used with these commands:
 - cdr check replicate
 - cdr check replicateset

--memadjust= cannot be used with **cdr check replicate** or **cdr check replicateset**.

Memory changes – (-F) option

- The foreground option specifies that the synchronization operation is performed as a foreground process
 - ▶ Previous releases; sync operations were ran as background processes by default
 - Unless 'cdr sync' was ran manually
 - Long Form --foreground
 - The long form --foreground is not available in the 10.x family
 - Short Form -F
 - ▶ The --foreground option can be used with:
 - cdr start replicate
 - cdr start replicateset
 - cdr realize template

The --foreground option specifies the synchronization operation as a foreground process. Before this option, these commands can only be ran as background processes, unless you ran the **cdr sync** command manually. This command has these parameters available: **cdr start replicate**, **cdr start replicateset**, and **cdr realize template**.

Memory changes - Low resources

- Memory changes made in this version, relate to how ER behaves in low memory situations
 - ▶ Previous releases
 - ER behaves erratically when memory resource limits are reached
 - If cdr sync or cdr check threads cannot allocate memory - they can fail and provide a useless error message such as '-1'
 - ER can shutdown the Snooper or the Grouper, or both, but leave the DS threads active for receiving

The second type of Memory Changes made in this version, relate to how ER behaves in a low memory situation.

Previous releases to this version, if a memory issue occurred - ER might behave erratically, return an odd error message, or shutdown some ER components.

Memory changes - Low resources

▶ V10.00.xC9

- ER maintain consistent behavior when memory resource limits are reached
 - If cdr sync or cdr check processes fail, they will return a meaningful error message
 - If memory resources are low, ER will attempt:
 - First – free unused ER memory
 - Second – free unused IDS memory
 - If memory cannot be freed to meet the request:
 - ER will shutdown individual components, such as Snoopy or Grouper
 - ER will send a new alarm 'ALRM_CDR_SHUT_DOWN_ERR_COM' indicating the shutdown
 - Class ID #47 in 11.50.latest
 - Class ID #41 in 10.00 0 in alarmprog.h

In version 10.00.xC9, ER behaves consistently when memory limits occur. When this happens, ER will attempt to free its own memory first, and then ER will try to free IDS memory.

If enough memory cannot be freed to meet the request, ER will shutdown individual components and send a new alarm ALRM_CDR_SHUT_DOWN_ERR_COM.

Section

Sync and check changes

The next section covers sync and check changes.

Sync and check changes - Overview

- Behavior in previous releases
 - ▶ Running **cdr sync** or **cdr check**, ER only sends the data to the target once
 - ▶ Data already in the replication stream, or values modified after the row has already been checked, might still be 'out-of-sync'
- Reported problem
 - ▶ On Busy ER systems, running **cdr sync** or **cdr check** might not be enough to insure data is in sync on all nodes
 - ▶ Multiple sync's or check's were required to confirm data across all nodes were in sync

Currently, when running a **cdr sync** or **cdr check**, ER will only send the data to the target server once. If the table or tables being checked were in the process being replicated, or a row in the table was modified after it was checked, the table might still be 'out-of-sync'.

Customers running on busy systems were forced to run multiple checks or syncs to confirm data on all nodes were in sync.

Sync and check changes - Overview

- Improved behavior in 10.00.xC9
 - ▶ The -R (--repair) option repairs and validates repaired rows
 - Only initial out of sync rows are re-checked
 - ▶ The -v (--verbose) option prints information about re-checked rows
 - #rows out of sync initially versus #rows corrected (processed)

12

Enterprise replication Changes and new features in V10.00.xC9

© 2009 IBM Corporation

The positive changes in **cdr sync** and **cdr check** behavior in 10.00.xC9 includes a **repair (-R)** option to repair and validate rows.

Note: Only initial out of sync rows are re-checked.

Additionally, a **-verbose (-v)** mode option will check and print information about the re-checked rows.

Sync and check changes - Overview

- ▶ As mismatches are found, they are added to a 'recheck' list
 - List is checked again during the processing of other rows
 - Maximum of 120 times per row
 - If the row cannot be synced, it is listed on the report
 - Once the row is in sync, it is removed from the recheck list
 - A recheck can occur every second for the same row
 - This new behavior has been altered in 11.50xC4
 - `--inprogress` option can be configured to control how long to recheck for failed rows

13

Enterprise replication Changes and new features in V10.00.xC9

© 2009 IBM Corporation

As mismatches are found, they are added to a re-check list. This list is checked while processing other rows. A mismatched row might be checked a maximum of 120 times. Once the row is in sync, it's removed from the re-check list. If the row cannot be synced, it is listed on the report as a mismatch. A row can be rechecked multiple times during the same second.

NOTE: This behavior will change in 11.50.xC4 which will allow you to control the length of time a failed row will be re-checked, using the **--inprogress** option.

Sync and check changes - Reporting

- Example of new reporting options for rechecked rows:

```

cdr check repl -m g_srv1 -r repl1 -a -v -R
----- Statistics for repl1 -----
Creating Shadow Repl sync_20104_1310721_1219952381
Node          Rows   Extra  Missing Mismatch Processed
-----
g_srv1        424    0      0       0      11
g_srv2        416    3      11      0       3

The repair operation completed. Validating the repaired rows ...
Validation failed for the following rows:

row missing on <g_srv2>
key:cl:424
-----
row missing on <g_srv2>
key:cl:425
-----
row missing on <g_srv2>
key:cl:426
-----
marking completed on g_srv1 status 0

```

Here is an example of the new reporting options with **cdr check**. At the bottom of this output is a listing of all the mismatched rows that can not be re-checked.

Section

Error message clean up

The next section will discuss Error Message clean up.

Error message clean up - Overview

- Behavior in previous releases
 - ▶ In certain Low Memory situations, ER will fail and report the error:
General Error (-1) (-1)
 - ▶ The default error '*-100 fatal server error*' was used too often within ER code
- Reported problem
 - ▶ Customers had reported several instances of the '*General Error (-1) (-1)*' and '*-100 fatal server error*'
 - The error is not very descriptive, and typically leads to a Support Call to resolve

Previously, during certain low memory situations, ER will fail and report the error: *General Error (-1) (-1)*. Also, a default error *-100 fatal server error* was used too much with ER code. Because these errors were not descriptive, customers receiving these errors do not know what to do to correct the problem.

Error message clean up - Overview

- Improvements in 10.00.xC9
 - ▶ Removal of '-1' errors in all ER code
 - ▶ More descriptive messages added for specific functions

To correct the issue, all ER code was scanned, and all -1 errors were removed and replaced with more descriptive messages.

Section

Recompress rows in grouper paging file

This section will discuss the recompressing of rows within the grouper paging file.

Recompress rows in grouper paging file – Overview

- Behavior in previous releases
 - ▶ Before the grouper evaluates a row, it must decompress it (from Compressed Log Format)
 - ▶ If the grouper later needs to send the row to the Grouper Paging File, it does not recompress the row. Instead the grouper writes the row to the Grouper Paging File in an 'uncompressed' format.
- Reported problem
 - ▶ A small transaction in the logical log can fill the Grouper Paging File dbspace. Eventually this might cause a DDRBLOCK, if the request for more paging file space is not addressed.
 - A 1 or 2 MB transaction can consume so much of the paging file, that it runs out of space in dbspace. This, in turn, causing ER to potentially enter DDRBLOCK mode

19

Enterprise replication Changes and new features in V10.00.xC9

© 2009 IBM Corporation

In previous releases, if the grouper needed to evaluate a row, the row from the logical log would need to be decompressed (from the Compressed Log Format). If ER needed to spool this row to the Grouper Paging File – then the Grouper would spool the row in an uncompressed form.

The problem with this method is that a seemingly small transaction can end up filling up the Grouper Paging File and causing a DDRBLOCK.

Recompress rows in grouper paging file – Overview

- Improvements in 10.00.xC9
 - ▶ If the grouper needs to page a row to the Grouper Paging File – it will first recompress the row to conserve space

20

The solution is simple, if the grouper needs to page a row to the Grouper Paging File – it will first recompress the row to conserve space.

Section

Miscellaneous changes

The final section will cover miscellaneous changes with this release.

Miscellaneous changes

- CDR_ENV variables:
 - ▶ CDR_DISABLE_SPOOL
 - Allows user to disable the generation of any and all spool files
 - To avoid spool files being written to /tmp, set the CDR_ENV variable to 1
 - 0 = Allow ATS and RIS file generation
 - 1 = Prevent ATS and RIS file generation
 - ▶ CDR_ATSRISNAME_DELIM
 - Allows user to customize the time delimiter within ATS and RIS file names
 - Uses the delimiter set between the Hour, Minute, and Second values instead of the default : (colon) on UNIX® and . (period) on Windows®

22

Enterprise replication Changes and new features in V10.00.xC9

© 2009 IBM Corporation

A new CDR_ENV variable named CDR_DISABLE_SPOOL – allows you to disable the generation of any and all spool files. In some cases spool files were written to /tmp by default. To avoid this, set CDR_DISABLE_SPOOL to 1 – to prevent ATS or RIS file generation.

Another new CDR_ENV parameter is CDR_ATSRISNAME_DELIM. This allows you to customize the delimiter within the ATS and RIS file names. The default delimiters on UNIX is a : (colon), and on Windows a . (period).

Miscellaneous changes

- dbinfo argument 'cdrsession'
 - ▶ The new 'cdrsession' argument to the DBINFO function allows you to detect:
 - Insert, update, or delete statement being performed as part of a replicated transaction
 - ▶ Returns '1' if the thread performing the database operation is an ER apply or sync thread
 - ▶ Returns '0' if not an ER thread

23

Enterprise replication Changes and new features in V10.00.xC9

© 2009 IBM Corporation

A new dbinfo argument named 'cdrsession' will now return a '1' if the thread performing the database operation is an ER apply or sync thread. Knowing if certain threads are related to ER activities can be useful in developing your own triggers, stored procedures, or user-defined routines.

Miscellaneous changes (cdrsession continued)

- ▶ Example shows an SPL function that uses the 'cdrsession' option to determine if a thread is performing an Enterprise Replication operation
 - CREATE FUNCTION iscdr ()
 - RETURNING int;
 - DEFINE iscdrthread int;
 - SELECT DBINFO('cdrsession') into iscdrthread
 - from systables where tabid = 1;
 - RETURN iscdrthread;
 - END FUNCTION

A sample SPL function is shown here, which uses the 'cdrsession' option.

Miscellaneous changes

- `cdr -V`
 - ▶ Option shows the version of the cdr utility
- Inclusion of a pre-compiled version of the checksum UDR
 - ▶ Register pre-compiled version: version:
 - `$INFORMIXDIR/etc/idschecksum.sql`
 - ▶ If upgrading from an earlier version, drop the existing checksum UDR before registering the pre-compiled version

25

A new `cdr -V` option was added to show the exact version of the cdr utility.

This version includes a pre-compiled version of the checksum UDR. Running the SQL script – `idschecksum.sql` will register the new pre-compiled version of the UDR.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_IDS_1000xC9_ERFeaturesExt1_part1.ppt

This module is also available in PDF format at: ..\IDS_1000xC9_ERFeaturesExt1_part1.pdf



You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:
Informix

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

Windows, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

