

Informix Enterprise Replication Overview, Configuration, Monitoring

James Edmiston
Database Consultant
Quest Information Systems, Inc.
www.questinfosys.com
james@questinfosys.com

January 28, 2008



Purpose

Identify key items within the Informix engine to configure, monitor and trouble shoot Informix Dynamic Server Enterprise Replication (ER).

January 28, 2008



Outline

- Overview of Enterprise Replication
- Server Configuration
- Replication Configuration
- Monitoring
- Trouble shooting
- Documentation

January 28, 2008



Washington Area
Informix Users Group

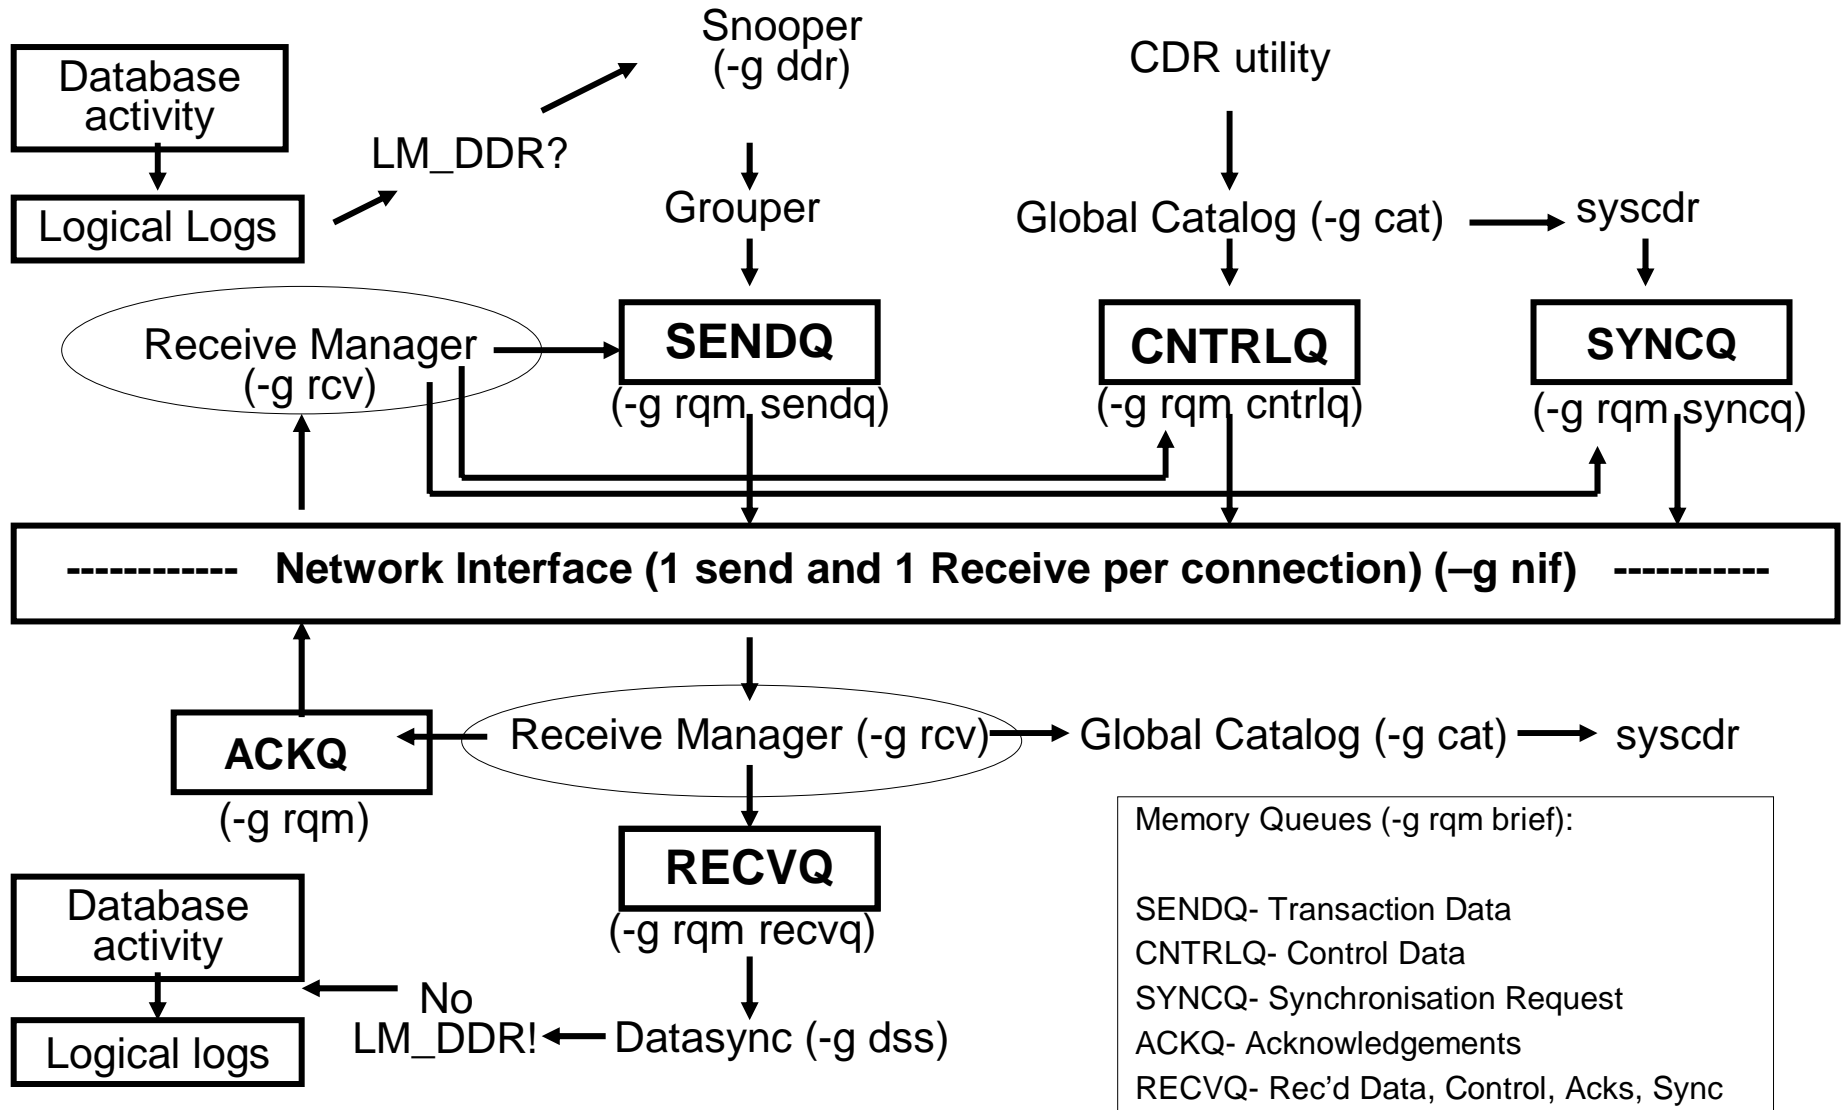


Enterprise Replication Overview

From IBM: Enterprise Replication (ER) or Continuous Data Replication (CDR) is a built-in, logical log based, asynchronous mechanism to distribute changes to specific tables and rows across an arbitrary number of participating nodes on a transactional level. Uses:

- Data Dissemination - Distributing data to other sites
- Data Consolidation - Distributing data to other sites
- Workload Partitioning - Distributing data to "load balance"
- Workflow Replication - Moving data through a workflow process
- Failover - Providing an alternate environment in case of failure

Enterprise Replication Component Diagram



Considerations

- Selecting Replication model
 - Primary/Target
 - Update Anywhere
- Conflict resolution?
- Topology
 - Fully connected
 - Hierarchical

January 28, 2008



Server Configuration

- Server to Server communications configured about all/for all servers involved in ER
 - /etc/hosts
 - /etc/services
 - Trusted environment
 - Hosts.equiv
 - .rhosts
 - Sqlhosts
 - Servers be members of server groups
 - Unique

January 28, 2008



Server Configuration (disk space considerations)

- Logical logs
- Extra database space for CRCOLS and delete tables
- Dbspaces for send and receive queues
- Dbspace for grouper paging file
- Disk space/directories for ATS and RIS files

January 28, 2008



Server Configuration (sqlhosts)

| dbservername | nettype | hostname | servicename | options |
|--------------|----------|----------|-------------|------------|
| g_group1 | group | - | - | i=1 |
| ifx_db1 | onsoctcp | host1 | ifx_db1 | g=g_group1 |
| g_group2 | group | - | - | i=2 |
| ifx_db2 | onsoctcp | host2 | ifx_db2 | g=g_group2 |

January 28, 2008



Server Configuration (onconfig)

onconfig variables:

```
CDR_EVALTHREADS      1,2      # evaluator threads (per-cpu-vp,additional)
CDR_DSLOCKWAIT       5        # DS lockwait timeout (seconds)
CDR_QUEUEMEM         4096     # Maximum amount of memory for any CDR queue (Kbytes)
CDR_NIFCOMPRESS      0        # Link level compression (-1 never, 0 none, 9 max)
CDR_SERIAL            0        # Serial Column Sequence
CDR_DBSPACE           <dbspace> # dbspace for syscdr database
CDR_QHDR_DBSPACE     <dbspace> # CDR queue dbspace (default same as catalog)
CDR_QDATA_SBSPACE    <sbpace> # List of CDR queue smart blob spaces
CDR_MAX_DYNAMIC_LOGS 0        # Dynamic log addition disabled by default
CDR_SUPPRESS_ATSRISWARN (see documentation for list of suppressions)
SBSPACETEMP          <temp sbpace> # Auto
                        (Grouper paging for evaluating large transactions)
```

January 28, 2008



ER Configuration

- Define servers
- Define replicates*
- Define participants*
*or use template (define and realize)
- Start Replication

January 28, 2008



ER Configuration

- cdr utility
 - See Quick Reference Guide
 - See ER Guide – Appendix A
- Keep it simple
 1. Define Servers
 2. Define Template
 3. Realize Template

January 28, 2008



cdr Example - server

- Define Server (first server)

```
cdr define server --connect=ifx_db1 --idle=0 --init g_group1  
--ats=/informix/ats --ris=/informix/ris
```

OR

```
cdr def serv -c ifx_db1 -i 0 -I g_group1  
-A /informix/ats -R /informix/ris
```

cdr Example - server (cont)

- Define server

```
cdr define server --connect=ifx_db2 --idle=0 --init  
--sync=g_group1 g_group2 --ats=/informix/ats --ris=/informix/ris
```

OR

```
cdr def serv -c ifx_db2 -i 0 -I -S g_group1 g_group2  
-A /informix/ats -R /informix/ris
```

ER Configuration (cont)

- Replicates
 - Participants
 - Master replicate
 - Conflict resolution rules and scope
- Replicate set
 - Grouping replicates so they have the same characteristics

ER Configuration (cont)

- Templates
 - Easier set up of replication with large numbers of tables to replicate
 - Defines a group of master replicates and a replicate set

cdr Example (template)

```
cdr define template dbtemplate --conflict=always --scope=row  
  --ats --ris --master=g_group1  
  --database=stores_demo  
  --file=/informix/dbschema.txt
```

- *dbtemplate* – the name of the template
- *file* - contains list of table names

```
cdr realize template dbtemplate g_group1
```

```
cdr realize template dbtemplate g_group2
```

Monitoring

- cdr utility
 - cdr list
 - cdr view
 - cdr error
- onstat commands
- Message log (online log)

January 28, 2008



Monitoring

- cdr list serv

| SERVER | ID | STATE | STATUS | QUEUE | CONNECTION | CHANGED |
|----------|----|--------|-----------|-------|------------|----------|
| g_group1 | 1 | Active | Local | 0 | | |
| g_group2 | 2 | Active | Connected | 0 | Jan 9 | 14:28:06 |

- cdr list ...

- Replicate
- Replicate set
- Template

- **cdr view (new in 10.00.UC9)**

You must first execute \$INFORMIXDIR/etc/syscdrsmiadd.sql on all nodes
(no connections allowed to syscdr or sysmaster)

```
cdr view [-c server] [-r interval] object(s) [options]
```

```
-c server --connect=server connect to server
```

```
-r interval --repeat=repeat interval in seconds
```

```
objects: list of objects separated by space
```

List of supported objects and their suboptions are:

```
ddr
```

```
servers
```

```
sendq
```

```
nif
```

```
apply
```

```
rcv
```

```
ris
```

```
ats
```

```
profile
```

```
state
```

```
atsdir [-R | -C | -v | -d | -q]
```

```
risdir [-R | -C | -v | -d | -q]
```

```
-R repair
```

```
-C check
```

```
-v verbose
```

```
-d delete option for repair
```

```
-q quiet option for repair
```

January 28, 2008



onstat -g options

cat [scope | replname]

Print Enterprise Replication global catalog information

cdr Print Enterprise Replication statistics

dtc Print statistics for the Enterprise Replication delete table cleaner

dss [UDR | UDRx]

Print statistics about data sync threads and user-defined data types

grp [A|E|Ex|G|L|Lx|M|Mz|P|pager|R|S|SI|Sx|T|UDR|UDRx]

Print statistics about the Enterprise Replication grouper

nif [all | sites | serverid | sum]

Print statistics about the Enterprise Replication network interface

que Print statistics for the Enterprise Replication high-level queues

rcv [serverid]

Print statistics about the Enterprise Replication receive manager

rep [replname]

Print events that are in the queue for the schedule manager

rqm [ACKQ | CNTRLQ | RECVQ | SENDQ | SYNCQ | FULL | BRIEF | VERBOSE]

Print statistics of the Enterprise Replication low-level queues

sync Print the Enterprise Replication synchronization status

January 28, 2008



Message log (online log)

```
14:25:20 CDR connection to server lost, id 2, name <g_group2>  
Reason: disconnect server  
14:25:20 CDR NIF shutdown waiting for 1 site(s) to disconnect  
14:25:22 CDR NIF shutdown waiting for 0 site(s) to disconnect  
14:25:22 CDR NIF Shutdown: connections all shutdown.  
14:25:22 CDR The NIF sub-component has shut down.  
14:25:27 DDR Log Snooping - Shutdown  
14:25:28 CDR shutdown complete
```

January 28, 2008



ATS and RIS files

- Aborted Transaction Spooling
 - Transactions that fail to be applied to the target database
- Row Information Spooling
 - replicate row data that fails conflict resolution or encounters replication order problems

January 28, 2008



Are you in sync? (courtesy IBM)

| Method | Description |
|--|--|
| Select count(*) | A simple way to check if the row counts are the same. This method will not tell you about rows that are different. Consider setting the isolation level to 'dirty read' before running this on a live table. |
| Select from (site A) where PK not in (site B) | This method will check the Primary Key between sites. However, it will not check for data outside the PK that may have been updated. |
| Select from (site A) where PK not in (site B) and A.crcols <> B.crcols | This method checks the PK and checks for updates to other row data by comparing the shadow columns. However, this method will only work with replicates with timestamp conflict resolution, and may not work well in ER environments with more than 2 servers. |
| Unload to files, then use diff to compare | This is a 'manual' operation and can be very time consuming. Consider only as a last resort. |
| cdr check | Uses the checksum() udr to compare full rows (10.00.xC5 and above) |

Getting Back in sync (courtesy IBM)

| Method | Description |
|---|--|
| High Performance Loader | This is the fastest option for large data sets, but requires a good deal of setup before running. Use 'Express' mode for data loads as it will load the CRCOLS data. |
| Dbimport/dbexport | Not a good option for ER. They require exclusive access to the database and cannot be used while ER is active. |
| Dbload | A good choice if you have a set of data that needs to be replicated to other sites. Data loaded into active replicates will be replicated. |
| Onunload/onload | Not a good fit with ER active. Both utilities require database logging to be off. The table or database loaded cannot exist, as the utility creates it. Lastly, when tables are loaded, they do not retain any information on constraints, triggers, and default values. |
| Unload/Load | Data loaded into active replicates will be replicated unless the transaction is started with the following syntax: BEGIN WORK WITHOUT REPLICATION. |
| Select from (site A) where PK not in (site B) | If you have a small set of data that needs to be synced, this may be a good option. There are issues with logging which you should be aware of before using this method. |
| Update col1 = col1 | If your replicate uses timestamp conflict resolution, this option will update existing any rows, and insert any missing rows. |
| Drop ER, Backup, Restore, Start ER | This option requires an ER outage, but if you have a large set of data and/or replicates – this option insures you will be in sync. |
| cdr sync | Synchronize a repl/replset across many nodes |
| cdr check -R | Uses 'cdr check' to "Repair" any replicates that are out of sync |

January 28, 2008



Washington Area
Informix Users Group



Documentation

- IBM Informix Dynamic Server Enterprise Replication Guide
- Administrator Guide
- Online
 - www.ibm.com/informix
 - IDS Information Centers (by version 10, 11.10, 11.50)

January 28, 2008



Summary

- Informix Dynamic Server Enterprise Replication
- Configuration (Server and ER)
- Monitor
 - cdr, onstat, logs
- Trouble shoot
 - cdr, onstat, logs
- Documentation

January 28, 2008



Questions?



James Edmiston
Database Consultant
Quest Information Systems, Inc.
www.questinfosys.com
james@questinfosys.com

January 28, 2008

