

COMMERCIAL IN CONFIDENCE

This document is copyright. No part of this document may be reproduced or copied, except as permitted under the Copyright Act 1968 (Commonwealth), by any means or process whether electronic, photocopying or otherwise, without the prior written consent of MGA

MGA DUPLICA 3.01

TECHNICAL WHITE PAPER



1. Contents

- 1. CONTENTS2
- 2. FIGURES2
- 3. DOCUMENT STATUS3
 - Document Status3
 - Summary of Changes3
- 4. INTRODUCTION4
- 5. COMMON BUSINESS PROBLEMS4
 - The Performance Advantage4
 - The Tailoring Advantage4
 - The Security Advantage5
 - High Availability - Backup/Recovery5
 - Disaster Recovery Planning5
 - Geographical Dispersion and Accumulation of Data5
 - Embedding the technology into your custom application5
- 6. ARCHITECTURE & DESIGN6
 - Overview6
 - Further Technical Knowledge6
 - 6.1.1 Replication Manager6
 - 6.1.2 Replication Monitor7
 - 6.1.3 Data Deletion – a special case7
 - Data Flagging8
 - Oracle to SQL Server Replication9
- 7. ESTIMATED CAPACITY OVERHEAD9
- 8. CAPABILITY OF REPLICATING ON LINE CHANGES10
- 9. NON-SUPPORTED REPLICATED CHANGES10
- 10. HARDWARE / SOFTWARE REQUIREMENTS10
- 11. CAPABILITY OF REPLICATING PART OF TABLES11
- 12. SUPPORT FOR ORACLE PARALLEL SERVER11
- 13. PARALLEL SESSIONS PROCESSING11
- 14. BENCHMARKS12
- 15. CONTACTING MGA12

2. Figures

- Figure 1 Architecture I of MGA Duplica8

3. Document Status

Document Status

Document Id.	Technical Info.doc
Document Version	3.01.03
Status	Final
Prepared by	Mark Rance
Last Changed Date	22 nd May 2003
Printed	12/11/03 02:16 PM

Summary of Changes

Version No.	Release Date	Author	Approval	Pages Changed	Change Request No.	Description
0.8	11 Dec 2001	Hiep Nguyen	Mark Gurry	All	1	Draft
0.9	20 Dec 2001	Robert Cox	Mark Gurry	All	2	Draft
1.0	21 Dec 2001	Mark Gurry	Mark Gurry	All	3	Final
2.0	03 Jun 2002	Mark Rance	Mark Gurry	All	4	Amended
3.01.01	20 Mar 2003	Mark Rance	Mark Gurry	All	5	Amended for Version 3.01 of MGA Duplica
3.01.02	27 Mar 2003	Mark Rance	Mark Gurry	All	6	Amended
3.01.03	28 Mar 2003	Robert Cox	Mark Gurry	All	7	Amended
3.01.04	20 May 2003	Robert Cox		All	8	Updated to include SQL Server and Detail further technical information

4. Introduction

MGA Duplica (Duplica) is a software tool developed by Mark Gurry and Associates (MGA) for replicating data from one database to another.

The software consists of two major components:

- The GUI interface, which assists users in monitoring the replication activities; and
- The database component (or informally, the backend), which carries out and monitors the actual data replication and gathers statistics.

The GUI interface is designed to run on Windows platforms and the database component resides on the destination Oracle database.

Duplica supports the following databases:

- Oracle to Oracle, on versions of Oracle 7.3, 8.0, 8.1 and 9.2; and
- Oracle to SQLServer, on versions of Oracle 8.1 and 9.2 and SQL Server 2000 and above.

Aside: Whilst the GUI enables full visual control of the replication process, some sites choose to administer the application via scripts and log files.

5. Common Business Problems

The software was developed to solve business problems that require replication. The following is a summary of the uses clients have identified and implemented. MGA would be interested to learn more about other specific requirements.

The Performance Advantage

By architecting your system with a transaction database and a reporting database, optimal performance is achieved. Each database can be tuned for its specific purpose, for example setting up indexes specifically for reporting in the reporting database. Value is created, by giving the online users improved online performance and throughput as well as delivering reports in best possible timeframes. With Oracle to SQLServer replication, clients are able to quickly use reporting tools that rely on Microsoft Architecture.

The Tailoring Advantage

To ensure that upgrades of the application are not affected by changes to database data model clients are choosing a replication strategy. With the main database schema remaining as the supplier provided issues or complications arising from change are no longer present.

Within the replicated database the extent of the change programme and value derived is more easily measurable. Some clients have begun with creating data views to solve security and performance issues with whilst others have integrated data from other databases to increase the reporting capability

The Security Advantage

By allowing clients access to a system that only holds data they are permitted to view and removing all data updating capabilities, your security is greatly increased. This can be used when firms consider upgrading their web strategy to include an extranet capability. Duplica allows you to decode secure columns, only select certain columns and only selected rows. This is achieved by replicating from views to tables. One large client in the Electricity industry is currently downloading secure data from the national market database to 38 external clients using the MGA developed Duplica technology.

High Availability - Backup/Recovery

As a key component of a high availability plan, a backup made on a second database ensures the primary database is available 24 hours a day 7 days a week. This database can be geographically distant from the primary database. The US electricity industry is currently using the product to achieve this purpose.

Disaster Recovery Planning

Extending the High Availability plan to include fail over and fail back capability a disaster recovery plan can be built using replication. This has become a high priority for many firms. Companies are currently using the Duplica solution to achieve this. The advantage of the timestamp approach is that no data can be lost and failover and failback is improved dramatically.

Geographical Dispersion and Accumulation of Data

Sites are currently using the technology to obtain the data that is relevant to their section of a company. For example, a Californian Company would just want to extract the Californian data from the national database. The tool can also be used to accumulate data from many sources into a central Schema. Using the same example as above, the Californian Company is now able to extract relevant data from several databases. This approach has further advantages in reducing the amount of data being sent via the network.

Embedding the technology into your custom application

Some clients are also using the technology in their custom developed applications. By understanding the replication process and timing, applications can be designed with fail over and fail back capabilities.

6. Architecture & Design

Overview

The system is designed to operate at the data level; this enables a high level of data integrity and control in the replication process. MGA has designed the product with the common objectives of performance, robustness and integrity.

The product operates in two modes:

- Polling from the target against a data that is date stamped for transactions; or
- An architecture based on flagging which data has been replicated.

The decision on which mode the product is set up will depend on the nature of the system being implemented. Typically a system operating with large numbers of small transactions may be best suited to polling the data based on time stamp information, whilst a system with complex longer running transactions for example, a financial ERP may be suited to the data flagging architecture.

MGA consults regularly in this field with clients before implementing MGA Duplica as a replication solution.

Further Technical Knowledge

MGA Duplica is installed on the target database server and connects over:

- SQL*Net or Net 8 in the Oracle to Oracle product; and
- The Microsoft OLE DB Provider for Oracle with the Oracle to SQL Server product to manage replication from one schema to another.

To replicate data, Duplica creates a small number of database objects in the target database (MGA Duplica repository). A process is scheduled to pull data from the source database to the target database. Typically this is set to 5 minute intervals, although consideration for the amount of data and size of tables will vary this interval.

For site without time stamp information available on the source data, MGA Duplica is able to create it. This addition of timestamp information is done without impacting the applications currently accessing the database, via the use of a trigger.

The graphical user interface allows the user to control the replication of data via two main windows, Replication Manager and Replication Monitor.

6.1.1 Replication Manager

Replication Manager allows the user to group related tables into suites. This is an important concept that allows suites of tables to be replicated simultaneously. Suites of tables are able to be replicated frequently for data that is critical and less frequently for non-critical data.

The rate at which data is replicated is controlled by setting:

- Time intervals;
- Maximum rows that can be downloaded per run;
- Maximum query times; and
- The order in which tables within a suite are replicated (to allow for parent-child relationships).

Replication on a suite level is able to be commenced or terminated at any time. You also have the ability to turn individual tables on and off as required. Conversely the entire replication control mechanism is able to be turned on or off for all suites if required. The Replication Manager Suites list is filtered to show specific suites and the suites tables and settings are printed for reporting purposes.

6.1.2 Replication Monitor

Replication Monitor allows administrators to monitor data whilst it is being replicated. Each time data is replicated an entry is saved to the audit log table. Any errors are recorded in the error log table. Both the audit log and error log reside in the Duplica repository and may be viewed from the Replication Monitor. These logs can be refreshed manually or automatically at intervals specified by the user. The logs are able to be printed, with the most recent logs also viewable as printable graphs. Target database system resources such as tablespace size are able to be viewed to ensure that it remains robust and is performing well.

Unless all of the data in the table is being replicated each time, Duplica requires a timestamp on the tables nominated for replication to determine if new records have been added or if existing records have changed. On many databases, the timestamp column is already implemented and maintained by the application. On other databases, timestamp column needs to be added where required. A pre-insert/update trigger is required to populate the timestamp column. Duplica can generate script for this activity.

6.1.3 Data Deletion – a special case

If deletions occur on the table, Duplica uses a delete control table, which keeps the primary key values of the deleted records and the time of deletion. A pre-delete trigger is required to maintain this table.

Figure 1 Architecture of MGA Duplica

Data Flagging

To ensure data integrity when timestamp information is not available in the source database a system of marking which data has been replicated, which data has not been replicated and which data is in the process of being replicated was introduced. This is used instead of the timestamp process.

When a transaction is created in the source database a trigger creates a record that contains 4 fields.

- The primary key of the record;
- The type of transaction (Insert, Update or Delete);
- The flag (where the record has been replication (N (No), Y (Yes), P (In Process)); and
- The time of the transaction.

When the replication process starts it changes the flag to (In Process), replicates the data to the target database and when completed sets the flag to Y for completed.

Oracle to SQL Server Replication

The Oracle server holding source tables for replication is added to SQL Server as a linked server using the Microsoft OLE DB Provider for Oracle.

A Suite is created and details for tables requiring replication are added. Each Oracle source table that requires replication must have a log table. When deletes, inserts and updates occur on the source table a trigger on the source table will insert the primary key, date, the type of update and a replicated flag into a source log table.

Upon starting replication, a job is added to SQL Server for each Suite requiring replication. Within a Suite job, a step is added for each table which calls a procedure to carry replication for that table. Within the replication procedure, a distributed transaction attempts to replicate modifications in the following order, deletes, updates and then inserts. If a step fails, a rollback for that table is performed and control is handed to the next step or table.

During replication of a table, firstly the source log table is queried for those records marked as deletes, but not replicated yet. Any records found are flagged as being in the process of being replicated, the marked records are then deleted from the target table and the source log table records are updated to indicate that they have been replicated. The same process is carried out for updates and then inserts.

7. Estimated Capacity Overhead

Due to a large number of variables such as the volume of data, the physical distance between the databases, the number of users, how data is used by various applications etc, accurate capacity overhead is difficult to estimate.

MGA technicians are available to identify the correct hardware capacity and architecture for each system. MGA's core competency is in database performance tuning; as such the product has been successfully implemented to assist sites with performance based problems.

Duplica has not been known as the cause of any noticeable performance degradation at any site where it is currently running.

Oracle to SQL Server replication is implemented so to 'pull data' from a production Oracle server to SQL Server. This minimises load on the production machine.

8. Capability of replicating on line changes

How long it takes for data to arrive at the target database largely depends on the frequency of replication, which could vary from seconds to days. Acceptable delay can also vary from table to table and needs to be determined by the user based on the intended use of the replicated changes. Replication frequency will then be set accordingly via the GUI tool.

DDL changes are made separately on each database. There is no automated DDL. The product dynamically builds the INSERT, UPDATE and DELETE statements from the columns existing at the target database. This allows columns to be added at the source and then added at a later time at the target database. Duplica will continue to replicate.

Duplica can also be turned off at the target database while the version of Oracle is updated or other changes made. When the product is re-started, Duplica will happily run without any loss of data or any downtime at the main database. This differs from other products that struggle if the target database is down for any considerable amount of time.

9. Non-supported replicated changes

Duplica supports replication of most data types. The exceptions are the following:

- Tables that have columns of type BFILE because they contain a locator to a large binary file outside of the database; and
- Tables with no primary keys unless “truncate and insert” option is selected (i.e. existing data will be truncated and the entire table repopulated).

To date MGA has not required these types to be copied due to technical constraints. MGA is interested to speak to clients who may require non supported types to be replicated.

10. Hardware / Software Requirements

The Oracle database versions supported are 7.3.4 through to 9.2. MGA recommends that clients work with Oracles supported versions of the database, using Oracles recommended version hardware and software requirements.

The Duplica GUI has been tested on the following workstations with the following minimum configuration outlined below:

- 300 MHz CPU
- 200 MB hard disk space
- 128 MB RAM
- Windows NT 4.0, 2000 Pro or XP Pro
- Oracle SQL*Net version 2.1 or Oracle Net8

For Oracle to SQLServer Replication:

- SQLServer 2000 Standard Edition
- MDAC 2.6 or greater
- You must setup the Oracle Server as a linked server within SQLServer.(using the Microsoft OLE DB Provider for Oracle).This is a standard feature within the SQLServer product.

- The database components are compatible with any version of Oracle above 8i and any hardware platform that supports the Windows operating system and Oracle RDBMS above.

11. Capability of replicating part of tables

Duplica allows the user to specify additional query predicate (in a form of a WHERE condition) to eliminate unwanted rows, thus further qualifying replicated data.

The product also allows you to replicate from views to tables.

Duplica does not store a list of columns to be replicated in its repository and it replicates all columns that exist in the target table. Columns that exist in the source table but not in the target table will be ignored. If there are more columns on the target table than on the source table, an error will occur. This eases the Change Management pressures.

The product supports complete and comprehensive horizontal and vertical replication of data.

12. Support for Oracle Parallel Server

Duplica works well in an Oracle Parallel Server environment. It relies on DBMS_JOBS to manage most of its database processing and therefore, is subject to the same limitation as those imposed on DBMS_JOBS.

Duplica currently does not specify the node when submitting its jobs, so the jobs can be run on any available participating server.

13. Support for Oracle 9i RAC

The MGA Duplica architecture does not rely on the redo logs, instead relying on the processing of transactions. This distinction allows Duplica to also operate in an Oracle 9i RAC environment.

14. Parallel sessions processing

Duplica uses the Oracle supplied package, DBMS_JOBS, to schedule its background jobs. Jobs are submitted for the control mechanism and for the suites scheduled to run. Once submitted, these jobs will be managed by Oracle RDBMS engine, which execute them at the specified interval, in parallel sessions that are independent from each other. This allows multi-threading of the replication. In addition to this, the product also utilises array processing.

Jobs submitted by Duplica are:

- Replication Monitor;
- Database Monitor (only if this option is selected);
- House Keeping job (executed only once a day); and
- One per each suite for which replication is scheduled.

Duplica also requires additional database connection to service its GUI interface.

15. Benchmarks

Duplica performs satisfactorily at all of our client's sites. Its performance varies vastly from site to site and from application to application due to a number of factors. Some of these are:

- Physical distance between the databases;
- The speed and the traffic of the communication link;
- The rate of changes at different times of the day when replication occurs;
- The database configuration;
- The size of the table being replicated; and
- The replication frequency.

As an example of the performance difference between Shareplex and Duplica, private timings were taken to replicate data between 2 sites 6 miles away. The typical timings were 14 seconds in average for Shareplex and 32 seconds for Duplica.

16. Contacting MGA

Information

Contact: Mark Gurry & Associates

Phone: Seattle (Washington – USA) (Toll Free): 888 495 1900
Melbourne (Australia) +61 2 9586 2780
Sydney (Australia) +61 3 9428 8454

Email: info@mga-it.com

WWW: www.mga-it.com