




SouthEastern Michigan Oracle Professionals

# GoldenGate for Oracle DBAs

November 18, 2015

**METAT** Solutions for the Red Stack

Daniel A. Morgan  
email: dmorgan@forsythe.com  
mobile: +1 206-669-2949  
skype: damorgan11g



# Introduction



# Dan Morgan

- Principal Adviser: Forsythe Meta7
-  Oracle ACE Director
- More than 45 years technology experience
  - First computer was an IBM 360/40 mainframe in 1970
  - Fortran IV and Punch Cards
-  Curriculum author and primary Oracle instructor at University of Washington
-  Guest lecturer on Oracle at Harvard University
- Decades of hands-on SQL, PL/SQL, and DBA experience
- The "Morgan" behind Morgan's Library on the web  
[www.morganslibrary.org](http://www.morganslibrary.org)
- 10g, 11g, and 12c Beta tester
- Co-Founder Intl. GoldenGate Oracle Users Group
- Contact email: [dmorgan@forsythe.com](mailto:dmorgan@forsythe.com)



# My Websites: Morgan's Library

 **Morgan's Library**   [www.library](http://www.library)

**International Oracle Events 2015-2016 Calendar**

Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

**The Library**

The library is a spam-free on-line resource with code demos for DBAs and Developers. If you would like to see new Oracle database functionality added to the library ... just email us. Oracle 12.1.0.2.0 has been released and new features will be showing up for many weeks. The first updates have already been made.

**Home**

**Resources**

- [Library](#)
- [How Can I?](#)
- [Code Samples](#)
- [Presentations](#)
- [Links](#)
- [Book Reviews](#)
- [Downloads](#)
- [User Groups](#)
- [Blog](#)
- [Humor](#)

**General**

[Contact](#)

[About](#)

[Services](#)

[Legal Notice & Terms of Use](#)

[Privacy Statement](#)

**Presentations Map**



**MadDog Morgan**



**Training Events and Travels**

- [IOUG, Chicago, Illinois - Mar 10](#)
- [UTOUG, Salt Lake City, Utah - Mar 11-12](#)
- [OUGN, Oslo, Norway - Mar 12-14](#)
- [Collaborate, Las Vegas, Nevada - Apr 12-16](#)
- [IIYOUG, New York, NY - May 19](#)
- [GLOC, Cleveland, Ohio - May 19-20](#)

**Next Event: 27 January, Redwood Shores, CA**

**Oracle Events**



Click on the map to find an event near you

**Morgan**

  
aboard USA-71



**Library News**

- [Morgan's Blog](#)
- [Join the Western Washington OUG](#)
- [Morgan's Oracle Podcast](#)
- [US Govt. Mil. STIGS \(Security Checklists\)](#)
- [Bryn Llewellyn's PL/SQL White Paper](#)
- [Bryn Llewellyn's Editioning White Paper](#)
- [Explain Plan White Paper](#)





**ACE News**

Would you like to become an Oracle ACE? 


Learn more about becoming an ACE



- [ACE Directory](#)
- [ACE Google Map](#)
- [ACE Program](#)
- [Stanley's Blog](#)

Congratulations to our newest ACE Director Jim Czuprynski

# My Websites: International GoldenGate Oracle Users Group




# Travel Log: 2010



# Travel Log: 2013





## Travel Log: 2014



## Why Meta7

- We are the Forsythe division dedicated exclusively to the Oracle Red Stack
  - Subject matter expertise, focused on Oracle's customers, with extensive experience across multiple industries
  - Internationally recognized experts in
    - Oracle Database
    - Oracle GoldenGate
    - Oracle Engineered Systems
    - Full support for Zero Downtime Migrations
    - Database Performance Tuning and Health Checks
    - Database Security Audits
  - Reliable on-time and on-budget delivery
  - A professional and agile customer-focused team
  - New, State-of-the-Art Technology Evaluation Center
  - Secure hosting and Managed Services
  - Flexible financial support






Manitoba Oracle Users Group

# GoldenGate for Oracle DBAs

## Zero Downtime Migrations

Daniel A. Morgan  
email: [dmorgan@forsythe.com](mailto:dmorgan@forsythe.com)  
mobile: +1 206-669-2949  
skype: [damorgan11g](skype:damorgan11g)




Daniel A. Morgan  
email: dmorgan@forsythe.com  
mobile: +1 206-669-2949  
skype: damorgan11g



# Technical Briefing: How Do You Safeguard the Database Against Today's Cyber Threats?


October 21, 2015





# Database Partitioning


Daniel A. Morgan  
email: [dmorgan@forsythe.com](mailto:dmorgan@forsythe.com)  
mobile: +1 206-669-2949  
skype: damorgan11g



A dramatic image of a large fire with intense orange and yellow flames. In the foreground, two firefighters in dark uniforms and helmets are silhouetted against the fire, working together to extinguish it. The background is filled with the chaotic and powerful nature of a large-scale fire.

# IT Fire Fighting

Daniel A. Morgan  
email: [dmorgan@forsythe.com](mailto:dmorgan@forsythe.com)  
mobile: +1 206-669-2949  
skype: [damorgan11g](#)



## Rhetorical Question

- If an operating room is not the place to experiment ...



- Should you trying things for the first time on your operational systems?

# Content Density Warning



Take Notes ... Ask Questions

# Business Case



# GoldenGate Business Cases (1:3)

- Load Distribution
  - The current server infrastructure is not sufficient to handle the demands placed upon it; for example it would save the organization both time and money if analytic queries were run against a separate database performance optimized for those queries
- Upgrading
  - The organization wishes to upgrade from one version of a product to a different version of the same product; for example 11.2.0.3 to 12.1.0.2 with near-zero downtime
- Homogeneous Migration
  - The organization wishes to migrate a database from one hardware platform to another; for example Oracle Database on Solaris to Oracle Database on Linux
- Heterogeneous Migration
  - The organization wishes to migrate from one database product to another; for example from stand-alone SQL Server to Oracle RAC on an ODA

# GoldenGate Business Cases (2:3)

- Data Distribution
  - The organization wishes to make data available at different physical locations without dependency upon the internet; for example the corporate head quarters are in London England but the manufacturing facilities are in Latin America
- Data Consolidation
  - The organization has multiple locations where data is collected and wishes to have a single source that combines them all for management reporting and analysis, for example there are 8,200 retail stores in North America, each with its own on-site database and the corporate headquarters needs to be able to run a single query across data from one region's stores or all stores in all regions
- Data Access Sharing
  - The organization wishes to have two separate locations with equal access to data with the ability to perform DML and DDL, for example the organization considers each office a DR site for the other


## GoldenGate Business Cases (3:3)

- Security
  - The organization stores both public and confidential information some of which it does not wish to make vulnerable by making it visible through-out the organization; for example data collection from an e-commerce site that processes both orders and credit cards
- Auditing
  - The organization wants a database of suspicious transactions for fraud investigators where the source might be multiple databases, from multiple vendors; for example the organization has both Oracle and DB2 databases and wants a single repository for internal auditors
- Data Transformation
  - The organization wants to load data from OLTP (On-Line Transaction Processing) systems into one or more data warehouses, Operational Data Stores (ODS) or Decision Support Systems (DSS); for example the data warehouse collects data from only specific tables and columns and must load them into cubes

# What Is GoldenGate



# Part of a Family of Data Integration Technologies



# Other Oracle Database Replication Technologies


- Data Guard
  - Physical
  - Logical
  - Usable only when the source and the target are identical
    - Operating System and Database Version
- Streams
  - While still supported and used, Streams had a number of issues (which is why Oracle acquired GoldenGate) and is no longer Oracle's strategic replication product, it's on a continue and converge path and will continue to be supported but not extended to support multi-tenant container databases
- Change Data Capture (CDC)
  - A subset of Streams capabilities: Support will continue but capabilities will not be enhanced for multi-tenant container databases
- Materialized View, Database Links, 3<sup>rd</sup> Party Products

## What Makes An Ideal Replication Solution?


- Replication apply time configurable from near-zero to a user defined pre-set delay
- Scalable to handle high transaction volumes
- Minimal footprint and impact to existing systems
- Modular architecture supporting multiple data sources and targets
- Fault tolerant
- Maintains transactional integrity
- Reliable
- Secure
- Stable

## What Is GoldenGate?

- A software tool that creates a tunnel between two, or more physically separate systems such that they can behave as a single logical system
- Middleware that provides a data communications channel between a source from which transactions will be read and a target to which the transactions will be written
- GoldenGate connections can be homogeneous, for example Oracle Database to Oracle Database, or heterogeneous, for example Teradata to Oracle Database
- Data transformation and enrichment support
- Can replicate DML
- Can replicate DDL
- Unlike Oracle Service Bus ... understands transactions



# Logical Replication Architectures



## Query Offloading

- Using GoldenGate for offloading queries and reporting tasks reduces the load on the primary system
  - Create a database with either a full set or subset of the data from the original, primary, system and use it as a second target to reduce the connection and I/O load on the primary system
  - The second system can be a single database (replication) or multiple secondary databases (broadcast replication)
  - The secondary databases(s) can be co-located in the same data center as the primary database or be geographically separated
  - A secondary database can be optimized in a number of ways
    - It can be smaller and more efficient with a subset of schemas and objects
    - It can contain indexes not present in the primary system
    - It can eliminate unnecessary indexes and constraints
    - Custom statistics can be created that optimize querying and reporting
    - It can utilize partitioning not present on the source system

# Extraction Transformation and Loading (ETL)

- ETL is the name we give to a three-step process
  - Extract (capture) data from a source
  - Transform the data if necessary
  - Load the data into a target system
- Numerous transformation options are available
  - Value substitutions can be performed ("Yes" changed to "T" true)
  - Character sets can be changed
  - Columns and/or rows can be dropped
  - Nulls can be replaced with default values
  - Strings can be concatenated or parsed
  - Currency exchanges calculated
  - Time zones changed
  - Functions can be applied to perform mathematical operations
  - Denormalized data can be normalized

## Load Distribution

- Load Distribution allows the load of database operations to be distributed among several systems
  - Active Data Guard might be a better option if no transformation is being done
  - GoldenGate is ideal where transformation is required
  - The load can be distributed both locally within the data center or to remote locations
  - If using an Oracle 12c database connectivity can be configured using Global Data Services (GDS)

## Auditing

- Auditing is an important part of database operations
- Often custom audit tables are created
- These audit tables can be replicated via GoldenGate to local or remote audit tables
- Data can be subsetted
- Data can be masked
- Depending on your needs, initial replication may not be required
  - For example: Monthly, Quarterly or Yearly Audit tables

## Data Warehouse or Decision Support System

- GoldenGate can be used to loading Data Warehouse, Decision Support Systems and schemaless Big Data from OLTP
- GoldenGate has transformed data loading
  - Batch, trickle and real-time loading are supported
  - GoldenGate can be integrated with Oracle Data Integrator (ODI) in a configuration where ODI's Journaling Knowledge Module (JKM) is used to generate GoldenGate parameter file mappings saving manual work and reducing the errors


# Logical Architecture



# GoldenGate Components

- Source & Source Processes
  - Manager
    - Monitors and manages running GoldenGate processes
    - Writes the alert log
    - Performs file maintenance
  - Extract
  - Pump
- Parameter Files
- Trails
  - also referred to as Extract Files, Trail Files, Local and Remote Trail Files
- Target & Target Processes
  - Manager
  - Collector
  - Replicat

## Logical Flow (1:4)




# Supported Environments

Transaction Sources	Transaction Targets	Operating Systems
<ul style="list-style-type: none"><li>▪ HP SQL/MP</li><li>▪ HP SQL/MX</li><li>▪ IBM DB2</li><li>▪ IBM DB2 for I (AS/400)</li><li>▪ IBM Informix</li><li>▪ JMS message queues</li><li>▪ Microsoft SQL Server</li><li>▪ Oracle Database</li><li>▪ Oracle MySQL</li><li>▪ Oracle TimesTen</li><li>▪ Sybase ASE</li><li>▪ Tandem</li><li>▪ Teradata</li></ul> <p><b>Coming in 12.2 support for:</b> <b>ALO Framework</b> <b>AVRO files across the Big Data stack</b> <b>Kafka</b></p>	<ul style="list-style-type: none"><li>▪ Apache Hadoop (Big Data)</li><li>▪ EMC Greenplum</li><li>▪ HP SQL/MP</li><li>▪ HP SQL/MX</li><li>▪ IBM DB2 (LUW)</li><li>▪ IBM Informix</li><li>▪ IBM Netezza</li><li>▪ IBM System I (AS/400)</li><li>▪ IBM System z (mainframe)</li><li>▪ JMS Message Queue</li><li>▪ Microsoft SQL Server</li><li>▪ ODBC Databases</li><li>▪ Oracle Database</li><li>▪ Oracle MySQL</li><li>▪ Oracle TimesTen</li><li>▪ Sybase ASE</li><li>▪ Tandem</li><li>▪ Teradata</li><li>▪ Flat Files</li><li>▪ XML Files</li></ul>	<ul style="list-style-type: none"><li>▪ HPUX IA64</li><li>▪ HP NonStop Itanium</li><li>▪ IBM AIX</li><li>▪ IBM z/OS</li><li>▪ IBM iSeries</li><li>▪ Linux x86-64</li><li>▪ Oracle Solaris Sparc</li><li>▪ Oracle Solaris x86-64</li><li>▪ Windows x86-64</li><li>▪ z/Linux (IBM mainframe)</li></ul>




## Logical Flow (2:4)




The Extract process captures transactions and writes them to a Local Trail File

## Logical Flow (3:4)




The Pump process routes and copies the trail file from the Local directory to the Remote location(s)

## Logical Flow (4:4)




The Replicat process reads the Remote Trail File and Applies transactions

# Complete Logical Flow with Managers



The Pump process communicates with the Target's Manager and requests a Collector process to transfer the trail file


## Logical Flow: Bi-directional



# Physical Architecture and Installation




# eDelivery.oracle.com



The screenshot shows a web browser window for the Oracle Software Delivery Cloud. The URL in the address bar is [https://edelivery.oracle.com/osdc/faces/SearchSoftware?\\_afrLoop=30609748004624528\\_afrWindowMode=0&\\_adf.ctrl-state](https://edelivery.oracle.com/osdc/faces/SearchSoftware?_afrLoop=30609748004624528_afrWindowMode=0&_adf.ctrl-state). The page title is "SearchSoftware". The Oracle Cloud logo is at the top, along with a sign-in link for "damorgan12c@gmail.com". A message at the top of the content area says: "To select products for download, enter the Oracle Product into the type-ahead field below, then select from the list of available platforms. Once you have made your selection, the title will be displayed in the 'Selected Products' section below. Repeat this step for all titles you wish to download. Once complete, click on the 'Continue' button." Below this, a bold message states: "You must agree to Oracle's trial license terms before downloading products that you do not have a current valid license to use." A "Selected Products" table is shown with one row: "Product" (Oracle GoldenGate) and "Platform" (empty). A dropdown menu for "Product" is open, showing a list of Oracle GoldenGate components. A "Select Platform" button is also visible.

# Download

- Oracle Technology Network: Middleware



The screenshot shows the Oracle Downloads website. At the top, there is a banner with a red 'Download Database 12c' button and a video thumbnail for 'Experience Oracle Database 12c'. Below the banner is a navigation bar with categories: Applications, Database, Enterprise Management, Java, Middleware, Server and Storage Systems, Developer Tools, and More Downloads. A red arrow points to the 'Middleware' category. The main content area lists various Oracle Middleware products in a grid format. The products listed under 'Middleware' include: AIA Foundation Pack, BPEL Process Manager, Business Activity Monitoring, Business Intelligence Suite EE, Business Intelligence Applications, Business Intelligence Mobile Security Toolkit, Business Intelligence Publisher, Business Process Analysis Suite, Business Process Management, Coherence, Crystal Ball, Data Integrator and Suite, Discoverer, Endeca Information Discovery, Endeca Server, Enterprise Repository, Entitlements Server, Event-Driven Architecture Suite, Forms & Reports Services, GlassFish Server, GoldenGate, Hyperion Performance Management and BI, Identity Management, Internet Application Server 10g and Components, JRockit, Outside In Technology, Portal, Real-Time Decisions, Service Bus, Service Bus for Financial Services, Service Registry, SOA Suite, Tuxedo, Web Services Manager, Web Tier, WebCenter Content, WebCenter Portal, WebCenter Sites, WebLogic Integration, and WebLogic Server 12c.

Middleware	Endeca Information Discovery	Real-Time Decisions
AIA Foundation Pack	Endeca Server	Service Bus
BPEL Process Manager	Enterprise Repository	Service Bus for Financial Services
Business Activity Monitoring	Entitlements Server	Service Registry
Business Intelligence Suite EE	Event-Driven Architecture Suite	SOA Suite
Business Intelligence Applications	Forms & Reports Services	Tuxedo
Business Intelligence Mobile Security Toolkit	GlassFish Server	Web Services Manager
Business Intelligence Publisher	GoldenGate	Web Tier
Business Process Analysis Suite	Hyperion Performance Management and BI	WebCenter Content
Business Process Management	Identity Management	WebCenter Portal
Coherence	Internet Application Server 10g and Components	WebCenter Sites
Crystal Ball	JRockit	WebLogic Integration
Data Integrator and Suite	Outside In Technology	WebLogic Server 12c
Discoverer	Portal	


## Installation (1:2)

- Uses the Oracle Universal Installer (OUI) used by the database


```
[oracle@gg00a Disk1]$ ls
install response runInstaller stage
[oracle@gg00a Disk1]$ ./runInstaller
Starting Oracle Universal Installer...

Checking Temp space: must be greater than 120 MB      Actual 4028 MB      Passed
Checking swap space: must be greater than 150 MB      Actual 2047 MB      Passed
Checking monitor: must be configured to display at least 256 colors.      Actual 65536      Passed
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2015-01-25_01-54-05PM. Please
wait ...[oracle@gg00a Disk1]$ ./runInstaller
Starting Oracle Universal Installer...
```

# Installation (2:2)



# Directory Structure from the GGHOME directory



```
oracle@DB-OINT-LIN01:~/u04/app/gg
[oracle@DB-OINT-LIN01 gg]$ ls
121200_fbo_ggs_Linux_x64_shiphome.zip
bcfmt.tpl
bcrypt.txt
BR/
cachefiledump*
cfg/
cfgtoollogs/
chkpt_ora_create.sql
convchk*
convprm*
db2ctl.tpl
ddl_cleartrace.sql
ddl_create.sql
ddl_ddl2file.sql
ddl_disable.sql
ddl_enable.sql
ddl_filter.sql
ddl_orai0.sql
ddl_oral0upCommon.sql
ddl_ora11.sql
ddl_ora9.sql
ddl_pin.sql
ddl_remove.sql
ddl_session1.sql
ddl_session.sql
ddl_setup.sql
ddl_status.sql
ddl_staymetadata_off.sql
ddl_staymetadata_on.sql
ddl_tracelevel.sql
defgen*
deinstall/
demo_more_ora_create.sql
demo_more_ora_insert.sql
demo_ora_create.sql
demo_ora_insert.sql
demo_ora_lob_create.sql
demo_ora_misc.sql
demo_ora_pk_befores_create.sql
demo_ora_pk_befores_insert.sql
demo_ora_pk_befores_updates.sql
diagnostics/
dirchk/
dircrd/
dirdat/
dirdef/
dirdmp/
dirjar/
dirobey/
dirout/
dirpcs/
dirprm/
dirpt/
dirsrl/
dirtmp/
dirwt/
dirwww/
discard/
emsclnt*
enable_gg_replication.out
extract*
fbo_ggs_Linux_x64_shiphome/
freeBSD.txt
ggcmd*
ggMessage.dat
ggsci*
ggser.log
help.txt
install/
inventory/
jagent.sh*
jdk/
keygen*
label.sql
libantlr3c.so*
libbdb-5.2.so*
libglog.so*
libggnnzitp.so*
libggperf.so*
libggrep.so*
libicudata.so.48*
libicudata.so.48.1*
libicu18n.so.48*
libicu18n.so.48.1*
libicuuc.so.48*
libicuuc.so.48.1*
libxml2.txt
logdump*
marker_setup.sql
marker_status.sql
mgr*
notices.txt
ogg_12c_silent_install_response.rsp
oggerr*
OGG_WinUnix_Rel_Notes_12.1.2.0.0.doc
OGG_WinUnix_Rel_Notes_12.1.2.0.0.pdf
OPatch/
Oracle-GoldenGate-12.1.2.0-README.doc
Oracle-GoldenGate-12.1.2.0-README.txt
oraInst.loc
oui/
params.sql
prvtclkm.plb
pw_agent_util.sh*
remove_seq.sql
replicat*
retrace*
reverse*
role_setup.sql
sequence.sql
server*
sqlldr.tpl
srvm/
toperrs
ucharset.h
ulg.sql
UserExitExamples/
usrdecs.h
zlib.txt
```

# GGSCI

- GoldenGate Software Command Interface

```
GGSCI (gg00a) 1> info mgr

Manager is running (IP port gg19a.7809, Process ID 14259).

GGSCI (gg19a) 2> info all

Program      Status      Group      Lag at Chkpt  Time
Since Chkpt

MANAGER      RUNNING
EXTRACT      RUNNING      E1SH      00:00:12      00:00:03
EXTRACT      RUNNING      P1SH      00:00:01      00:00:05

GGSCI (gg00a) 4> info extract e1sh

EXTRACT      E1SH      Last Started 2015-08-14 15:21
Status RUNNING
Checkpoint Lag 00:00:08 (updated 00:00:11 ago)
Process ID 21465
Log Read Checkpoint Oracle Integrated Redo Logs
2015-08-14 12:17:54
SCN 0.14674261 (14674261)

GGSCI (gg00a) 5> info extract p1sh

EXTRACT      P1SH      Last Started 2015-08-14 26:15
Status RUNNING
Checkpoint Lag 00:00:00 (updated 00:00:00 ago)
Process ID 21496
Log Read Checkpoint File dirdat/la000001
First Record RBA 2307
```

# Configuration



## GoldenGate Configuration

- Configuration is performed using parameter files very similar in concept and editing to the database initSID.ora
- GG parameter files are ASCII text ... they are not compiled
- The major parameter files are:
  - GLOBALS
  - manager (mgr.prm)
  - extract
  - pump
  - replicat
- Some of the more important minor parameter files are
  - CMDSEC (command security)
  - ENCKEYS (encryption keys)
  - TCPERRS (Posix-based TCP/IP network error handling)
  - The Collector process does not have an editable parameter file

# GLOBALS Parameter File

```
-----  
-- Sample GLOBALS File - Author: Daniel Morgan Date: 14-AUG-2015  
-----  
CredentialStore /home/oracle/ggcredentials  
AllowInvisibleIndexKeys  
CharSet UTF-8  
GGSchema ggadmin  
MaxGroups 256  
OutputFileUMask 022  
Syslog ALL
```

# Manager Parameter File

```
-----  
-- Sample Manager - Author: Daniel Morgan Date: 14-AUG-2015  
-----  
UserIDAlias ggadm  
AutoStart ER *  
AutoRestart Extract *, WaitMinutes 5, Retries 10  
CheckMinutes 15  
DownCritical  
DownReportMinutes 2  
DynamicPortList 9500, 9800-9899  
LagCriticalSeconds 60  
LagInfoMinutes 3  
LagReportMinutes 30  
Port 7809  
PurgeDDLHistory MinKeepDays 2, MaxKeepDays 7, FrequencyMinutes 60  
PurgeMarkerHistory MinKeepDays 2, MaxKeepDays 7, FrequencyMinutes 60  
PurgeOldExtracts ./dirdat/*, UseCheckPoints, MinKeepDays 7,  
    FrequencyMinutes 60  
PurgeOldTasks EXTRACT *, After 3 DAYS  
StartupValidationDelayCSecs 100  
SysLog ALL  
UpReportMinutes 60
```

# Capture/Extract Parameter File

```
-----  
-- Sample Extract - Author: Daniel Morgan Date: 14-AUG-2015  
-----  
EXTRACT E1SH  
SETENV (NLS_LANG = AMERICAN_AMERICA.AL32UTF8)  
SETENV (ORACLE_HOME=/app/oracle/product/11.2.0.3)  
USERIDALIAS ggadm  
  
TranLogOptions IntegratedParams (max_sga_size 2048, parallelism 4)  
TranLogOptions User ggadmin  
EXTTRAIL dirdat/la  
LogAllSupCols  
UpdateRecordFormat COMPACT  
  
StatOptions ReportFetch  
WarnLongTrans 1h, CheckInterval 2m  
ReportCount Every 30 Minutes, Rate  
Report At 01:01  
ReportRollover At 01:02 On SUNDAY  
DiscardFile dir rpt/E1SH.dsc, Append  
DiscardRollover at 01:03 On Sunday  
  
TABLE SH.*;
```

# Pump Parameter File

```
-----  
-- Sample Pump - Author: Daniel Morgan Date: 14-AUG-2015  
-----
```

```
EXTRACT P1SH
```

```
SETENV (NLS_LANG = AMERICAN_AMERICA.AL32UTF8)  
SETENV (ORACLE_HOME=/app/oracle/product/11.2.0.3)  
USERIDALIAS ggadm
```

```
PASSTHRU
```

```
RMTHOST GG00B, MGRPORT 7809  
RMTTRAIL /dirdat/ra
```

```
StatOptions ReportFetch  
WarnLongTrans 1h, CheckInterval 2m  
ReportCount Every 30 Minutes, Rate  
Report at 01:00  
ReportRollover at 01:15 on SUNDAY
```

```
TABLE SH.*;
```

# Coordinated Deliver/Replicat Parameter File

```
-----  
-- Sample Delivery - Author: Daniel Morgan Date: 14-AUG-2015  
-----  
REPLICAT R1SH  
SETENV (NLS_LANG = AMERICAN_AMERICA.AL32UTF8)  
SETENV (ORACLE_HOME=/app/oracle/product/12.1.0.2)  
USERIDALIAS ggadm  
  
ASSUMETARGETDEFS  
map sh.sales, target sh.sales, ThreadRange (1-2);  
map sh.sales_arch, target sh.sales_arch, Thread (3);  
map sh.events, target sh.events, Coordinated, ThreadRange (4-6);  
map sh.transfers, target sh.transfers;  
  
StatOptions ReportFetch  
ReportCount Every 20 Minutes, Rate  
Report At 01:01  
ReportRollover at 01:02 On SUNDAY  
  
DiscardFile dir rpt/r1sh.dsc, Purge  
DiscardRollover At 01:03 On Sunday
```

# Integrated Deliver/Replicat Parameter File

```
-----  
-- Sample Delivery - Author: Daniel Morgan Date: 14-AUG-2015  
-----  
REPLICAT R1SH  
SETENV (NLS_LANG = AMERICAN_AMERICA.AL32UTF8)  
SETENV (ORACLE_HOME=/app/oracle/product/12.1.0.2)  
USERIDALIAS ggadm  
  
ASSUMETARGETDEFS  
DBOPTIONS INTEGRATEDPARAMS (parallelism 6)  
REPERROR (1403,DISCARD)  
  
StatOptions ReportFetch  
ReportCount Every 20 Minutes, Rate  
Report At 01:01  
ReportRollover at 01:02 On SUNDAY  
  
DiscardFile ./dir rpt/REP1HR.dsc, Append, MEGABYTES 100  
DiscardRollover At 01:03 On Sunday  
  
MAP sh.sales, Target sh.sales_archive WHERE (CUST_ID > 101000)
```

# Database Preparation



## Database Preparation Steps

- GoldenGate Home Directory
- Archivelog Mode
- Force Logging
- Supplemental Logging
- GG Admin Schema
- Roles and System Privileges

## GoldenGate Home Directory

- A separate GoldenGate instance must be created for every Oracle instance on a server, or in a VM, you wish to replicate
- One GoldenGate instance connects to only one Oracle Database instance
- Create a directory under ORACLE\_BASE/product to hold the GoldenGate application files
- If one ORACLE\_BASE services multiple Oracle homes then perform multiple GG installs and be sure you name the directories so that you know which directory corresponds with which Oracle home
- A GoldenGate installation requires less than 250MB of disk
- The trail files can consume substantial space and should be on a separate mount point(s) with very fast disk

## Archive Log Mode and Force Logging

- Simply put ... transactions that do not get written into the redo logs do not get replicated
- A danger to data integrity is any object created or modified using NOLOGGING

```
SQL> conn / as sysdba

SQL> SHUTDOWN IMMEDIATE;

SQL> STARTUP NOMOUNT;

SQL> ALTER DATABASE ARCHIVELOG MODE;

SQL> ALTER DATABASE MOUNT;

SQL> ALTER DATABASE FORCE LOGGING;

SQL> ALTER DATABASE OPEN;
```

## Supplemental Logging

- Supplemental logging is essential for the following situations
  - Tables without primary keys
  - Tables without unique constraints
  - Tables without unique indexes
  - Transactions that may alter one or more rows but the SQL does not alter the columns corresponding to the unique identifier or use it to define the columns to be updated

```
SQL> conn / as sysdba

SQL> ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;

SQL> ALTER SYSTEM SWITCH LOGFILE;
```

- The GoldenGate administrator will issue commands from within GoldenGate using the syntax SCHEMATRANDATA or TRANDATA that will enhance supplemental logging

## GoldenGate Administration Schema

- On the both source and target databases a GoldenGate administration schema must be created
- The most common name for this schema is GGADMIN
- The Oracle docs recommend giving the admin user privileges far in excess of what is actually required
  - If operational security is important to you do not follow Oracle's recommendations
  - For example Oracle recommends giving the GG administrator the privilege of flashing back the database
  - Giving anyone other than a senior DBA the flashback database privilege is **totally irresponsible** (and you can tell Oracle I said so)

## Database Roles and System Privileges (1:3)

- GoldenGate capabilities, as of GoldenGate version 12.1 are integrated into the database and require the following

```
SQL> conn / as sysdba

SQL> ALTER SYSTEM SET enable_goldengate_replication=TRUE SID=* SCOPE=spfile;
-- restart the database

-- on the source only
SQL> exec dbms_goldengate_auth.grant_admin_privilege ('GGADMIN', 'CAPTURE', TRUE);

-- on the target only
SQL> exec dbms_goldengate_auth.grant_admin_privilege ('GGADMIN', 'APPLY', TRUE);

-- if Transparent Data Encryption is in use in database version 12.1 or above
SQL> GRANT execute ON dbms_internal_clkm TO ggadmin;
```

- Some versions of GoldenGate with an Oracle Database may require running scripts that create tables, sequences, triggers, and roles that are granted to the GoldenGate administrator schema

## Database Roles and System Privileges (2:3)

- Source System Privileges
  - CREATE SESSION
  - CREATE TABLE
  - SELECT ANY DICTIONARY
  - SELECT ANY TABLE
  - SELECT ANY TRANSACTION

## Database Roles and System Privileges (3:3)

- Target System Privileges
  - CREATE SESSION
  - CREATE TABLE
  - SELECT ANY DICTIONARY
  - INSERT ANY TABLE (if DML inserts will be replicated)
  - UPDATE ANY TABLE (if DML updates will be replicated)
  - DELETE ANY TABLE (if DML deletes will be replicated)
  - ALTER ANY <object\_type> (for DDL replication is enabled)
  - CREATE ANY <object\_type> (if DDL replication is enabled)
  - DROP ANY <object\_type> (if DDL replication is enabled)

# Subsetting & Mapping

## Subsetting

- Subsetting consists of starting with a full data set, the primary host server database and replicating a selected set of qualified transactions
- For example the subset might consist of
  - Only insert and update statements: Not deletes
  - Only invoice data (header and detail) but not look-up tables
  - Only columns that do not contain PII (Personally Identifiable Information)
  - Only rows that are modified by a person whose title is District Manager or Assistant District Manager

## @RANGE Filtering

- You can use the `@RANGE` function in the filter clause to divide the processing workload among multiple processes, using separate `MAP` statements, to improve performance
  - Syntax:  
`@RANGE (range, total_ranges [, column] [, column]  
[, ...])`
- For example, the following splits the replication workload into two ranges (between two Replicat processes or two threads of a coordinated Replicat) based on the ID column of the source `EMPLOYEES` table
  - One Replicat parameter file will include the first `MAP` statement

```
MAP hr.employee, TARGET emp.staff, FILTER (@RANGE (1, 2, PID));
```

- Second Replicat parameter file will include the second `MAP` statement

```
MAP hr.employee, TARGET emp.staff, FILTER (@RANGE (2, 2, PID));
```

## WHERE Clause Filtering (1:2)

- The syntax for **WHERE** is identical in the **TABLE** and **MAP** statements
- Each **WHERE** clause must be enclosed within parentheses
- Literals must be enclosed within single quotes
- Permissible **WHERE** operators
  - Column names (*i.e.* LAST\_NAME, SALARY, ID)
  - Numeric values (*i.e.* 4096, -545, 3187.775)
  - Literal strings (*i.e.* 'FL', 'Fred', 'Auto')
  - Built-in column tests (*i.e.* @NULL, @PRESENT, @ABSENT)
  - Boolean Comparison operators (*i.e.* =, <>, >, <, >=, <=)
  - Conjunctive operators (*i.e.* **AND**, **OR**)
  - Grouping parentheses (Use open and close parentheses ( ) for logical grouping of multiple elements)

# WHERE Clause Filtering (2:2)

- TABLE Syntax

```
TABLE <table_name>, WHERE (<where_clause>);
```

- TABLE Example

```
TABLE hr.employee, WHERE first_name = @PRESENT;  
TABLE hr.employee, WHERE (employee_id < 100);
```

- MAP Syntax

```
MAP <source_table>, TARGET <target_table>, WHERE (<where_clause>);
```

- MAP example

```
MAP hr.employee, TARGET hr.employee_data,  
WHERE salary > 50000 AND active_flag = 'Y';
```

# Schema, Table, and Column Mapping

- COLMAP - Table-level Column Mapping
  - Valid for both TABLE and MAP
  - Map individual source columns to target columns that have different names
  - Specify default column mapping when an explicit column mapping is not needed
  - Provide instructions for selecting, mapping, translating, and moving data from a source column into a target column
  - Syntax  
`COLMAP ([USEDEFAULTS, ]  
target_column = source_expression)`
  - Examples

```
MAP sh.sales, TARGET sh.sales, COLMAP (UseDefaults);
```

```
MAP sh.sales, TARGET sh.sales_remap, COLMAP (prod_id = prod#, amount_sold = invoiced_quantity);
```

# Data Transformation & Enrichment

# Data Transformation

- Basic Transforms

- A commonly seen transformation is converting <first\_name><space><last\_name> into the equivalent <last\_name><comma><space><first\_name> for example "Daniel Morgan" becomes "Morgan, Daniel"

- Data Masking Transforms

- Also referred to as "Data Masking"
  - A technology that transforms most often sensitive information by changing its form
  - A commonly seen example of a data masking transform is converting a social security number (678-91-2345) into a masked (\*\*\*\*\*-\*\*-2345)

```
MAP hr.employee, TARGET hr.staff, COLMAP (USEDDEFAULTS,  
wages = @COMPUTE(salary * 12),  
full_name = @STRCAT(last_name, ", ", first_name));
```

## Data Enrichment

- Data enrichment is the process by which GoldenGate can be used to add new columns to the data stream using either
  - Default values
  - Dynamic values

```
MAP hr.employee, TARGET hr.staff, COLMAP (USEDEFAULTS,  
wages = @COMPUTE(salary * 12),  
full_name = @STRCAT(last_name, " ", " ", first_name),  
source_system_id = 'EBS',  
target_commit_timestamp = SYSDATE);
```

# Troubleshooting and Debugging



## DBA Considerations

- The most common problems are network stability issues
- The second most common problems are caused by DBAs and Developers making schema changes that create collisions
  - For example inserting rows into the target that will later cause a failure when an identical key is replicated from the source
- The third most common issues relate to adding and altering tables creating supplemental logging failures
- Help your GoldenGate admin identify processing -intensive tables so they can be replicated using separate processes
- There are special considerations for ASM and RAC dependent upon GoldenGate and Database version ... do your research
- Bidirectional replication is hard ... not because of the database and not because of GoldenGate but because of application designs not capable of supporting it

Security



## Credential Store

- Deploy GoldenGate version 12.1 with a Credential Store to hold login userid and passwords to protect the database

```
[oracle@gg00a gghome_1]$ cd dircrd
[oracle@gg00a dircrd]$ ggsci

Oracle GoldenGate Command Interpreter for Oracle
Version 12.1.2.1.0 OGGCORE_12.1.2.1.0_PLATFORMS_140727.2135.1_FBO
Linux, x64, 64bit (optimized), Oracle 11g on Aug  7 2014 09:14:25
Operating system character set identified as UTF-8.

Copyright (C) 1995, 2014, Oracle and/or its affiliates. All rights reserved.

GGSCI (gg00a) 1> ADD CREDENTIALSTORE

Credential store created in ./dircrd/.

GGSCI (gg00a) 1> exit

[oracle@gg00a dircrd]$ ls -l
total 4
-rw-r----- 1 oracle oinstall 324 May  9 12:06 cwallet.sso


GGSCI (gg00a) 1> ALTER CREDENTIALSTORE ADD USER ggadmin PASSWORD ora123 ALIAS ggadm

Credential store in ./dircrd/ altered.

GGSCI (gg00a) 1> dblogin useridalias ggadm
Successfully logged into database.
```

# What's New In Twelve-Dot-Two

# New GoldenGate 12.2 Capabilities Coming Soon



## New Capabilities

- ✓ Metadata in Trail File
- ✓ Enhanced Monitoring of Lag
- ✓ Performance Improvements
- ✓ Delivery to Kafka
- ✓ GG Studio for easy design
- ✓ MySQL 5.7, SQL Server SE, ALO framework

## Key Benefits

- ✓ **Reliability** most proven CDC on the planet – don't lose data!
- ✓ **Performance** when low latency matters most – GG is the best
- ✓ **Breadth** of coverage on all your targets, Oracle, Non-Oracle, Big Data and NoSQL

Wrap-Up



## Conclusion

- In our opinion GoldenGate is
- The best homogeneous and heterogeneous replication tool on the market
- Robust capabilities for data enrichment, transformation
- Substantial toolset for troubleshooting and debugging
- Very small cpu, memory, and storage footprint
- Substantial capabilities for performance tuning
- Robust capabilities for subsetting and filtering
- Capable of solid security lockdown
- Easy to upgrade
- Easy to patch
- Secure
- Stable

# Conclusion

- In our opinion GoldenGate is
  - The best homogeneous and heterogeneous replication tool on the market
  - Robust capabilities for data enrichment, transformation
  - Substantial subset for troubleshooting and debugging
  - Very small footprint
  - Substantial capabilities for performance tuning
  - Robust capabilities for subsetting and filtering
  - Capable of solid security lockdown
  - Easy to upgrade
  - Easy to patch
  - Secure
  - Stable
- SELECT more\_information  
FROM experience  
WHERE topic = 'GoldenGate'  
\*ERROR at line 1:  
ORA-00028: your session has been killed*

# Thank You For Your Time

Daniel A. Morgan  
email: [dmorgan@forsythe.com](mailto:dmorgan@forsythe.com)  
mobile: +1 206-669-2949  
skype: damorgan11g



