ORACLE OPEN WORLD

experience
OPENWORLD

November 11–15, 2007

ORACLE

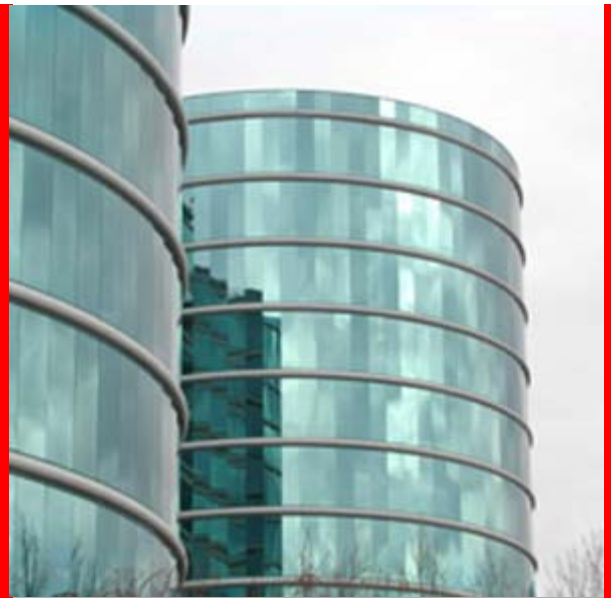

ORACLE

Best Practices for Implementing Replication with Oracle Streams in Oracle Database 10g and 11g

Patricia McElroy
Principal Product Manager,
Oracle Streams

Dirk Duellmann
Database Project Leader
CERN IT

What is Oracle Streams?

- Simple solution for information sharing
- Provides
  - replication
  - message queuing
  - data warehouse loading and simple ETL
  - database migration
  - application upgrade
  - event management and notification

Streams: Functional Components

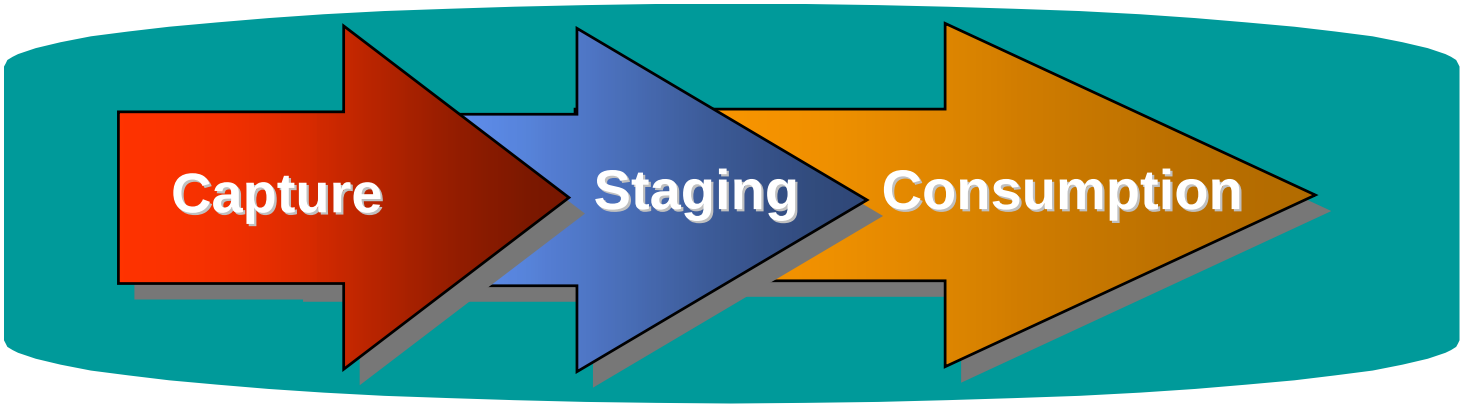

Asynchronous Information Sharing Infrastructure

Capture

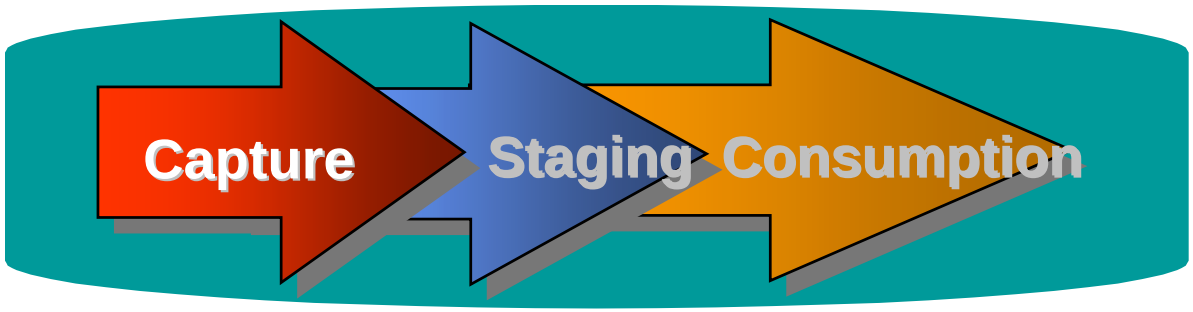

- Low overhead, low latency change capture
  - Multiple modes of capture
    - Asynchronously capture locally or at a downstream database with Redo-based Capture
      - Oracle Streams extracts changes from the redo as it is written
      - Log Buffer, Online Redo, or Archived log files
    - Sync Capture
      - Capture as part of executing transaction and persist to disk
  - Changes are formatted as a Logical Change Record (LCR), a representation of the change

Logical Change Record (LCR)

- Database change = LCR
  - DML
    - Row change = LCR
    - Object name, owner, Type of DML, SCN,…
    - OLD, NEW values
    - Optional attributes: username, session, thread,…
  - DDL
    - Object name, owner, Type of DDL, SCN,…
    - DDL text
    - Optional attributes: username, session, thread,…
  - LOB and LONG columns
    - Multiple LCRs per LOB or LONG column
    - Piecewise chunks

Staging

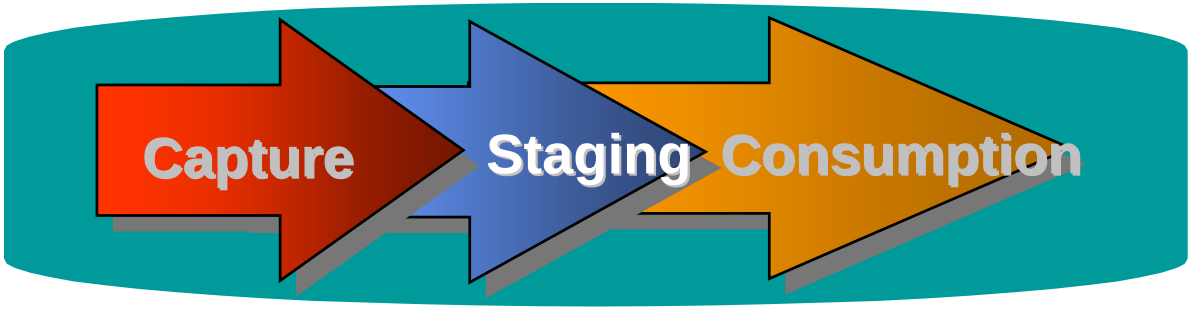

- Redo-based capture publishes messages into an in-memory staging area
  - Streams Pool of SGA
- Sync Capture stores messages to disk queue
- Subscribers: other staging areas or Streams processes
- Messages remain in staging area until consumed by all subscribers
- Propagation between staging areas is scheduled through a propagation job
  - LCRs propagated from a queue in the local database to a queue in the target database

Apply

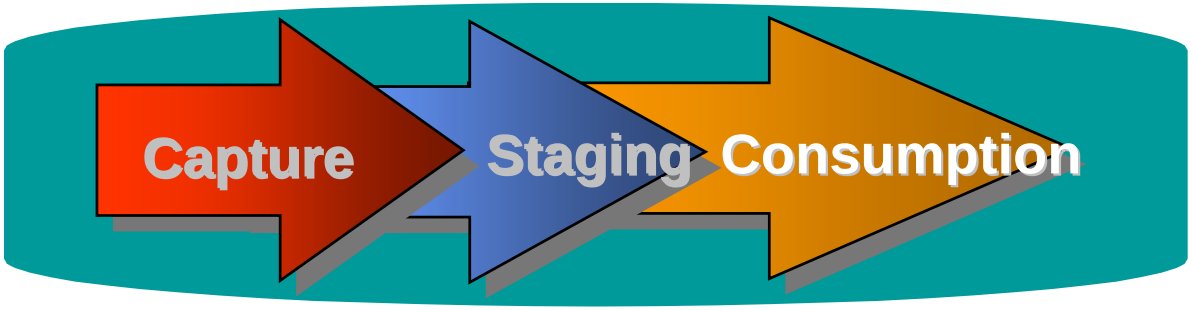

- The default apply engine will directly apply the DML or DDL represented in the LCR
  - Apply to local Oracle table or to non-Oracle table via dbLink
- Automatic conflict detection with optional resolution
  - Unresolved conflicts placed in error queue
  - Transactions can be re-applied or deleted from error queue
- Parallel apply maximizes concurrency
- Customizable apply processing
  - DML, DDL, Pre-commit, and Error handlers modify Apply processing

Rule-based Configuration

- Rule is expressed as SQL WHERE clause

```
dbms_rule_adm.create_rule(
   rule_name=>'scott.rule1',
condition=>':dml.get_object_owner() = ''SCOTT'' AND
            :dml.get_object_name()=''EMP''');
```

- Rule sets govern capture, staging, and apply
  - Inclusion
  - Negative
- Dynamic rule maintenance
- Tailored Replication API

  DBMS_STREAMS_ADM.MAINTAIN_*
  {Tables | Schemas | Global | TTS}

  DBMS_STREAMS_ADM.ADD_*_RULES
  {Table | Schema | Global |Subset}

Simple Configuration
Schema Replication

```
BEGIN dbms_streams_adm.MAINTAIN_SCHEMAS(
    SCHEMA_NAMES        => 'HR,SCOTT',
    SOURCE_DIRECTORY_OBJECT => null,
    DESTINATION_DIRECTORY_OBJECT  => null,
    SOURCE_DATABASE  => NULL,
    DESTINATION_DATABASE    => 'TARGET_global_name',
    BI_DIRECTIONAL => TRUE,
    INSTANTIATION  => DBMS_STREAMS_ADM.INSTANTIATION_SCHEMA_NETWORK,
   PERFORM_ACTIONS  => TRUE,
   SCRIPT_DIRECTORY_OBJECT => 'SCRIPT_DIR',
   SCRIPT_NAME   =>  'generated_schemas_script.sql',
);
END;
/
```

Example Configurations using MAINTAIN_SCHEMAS

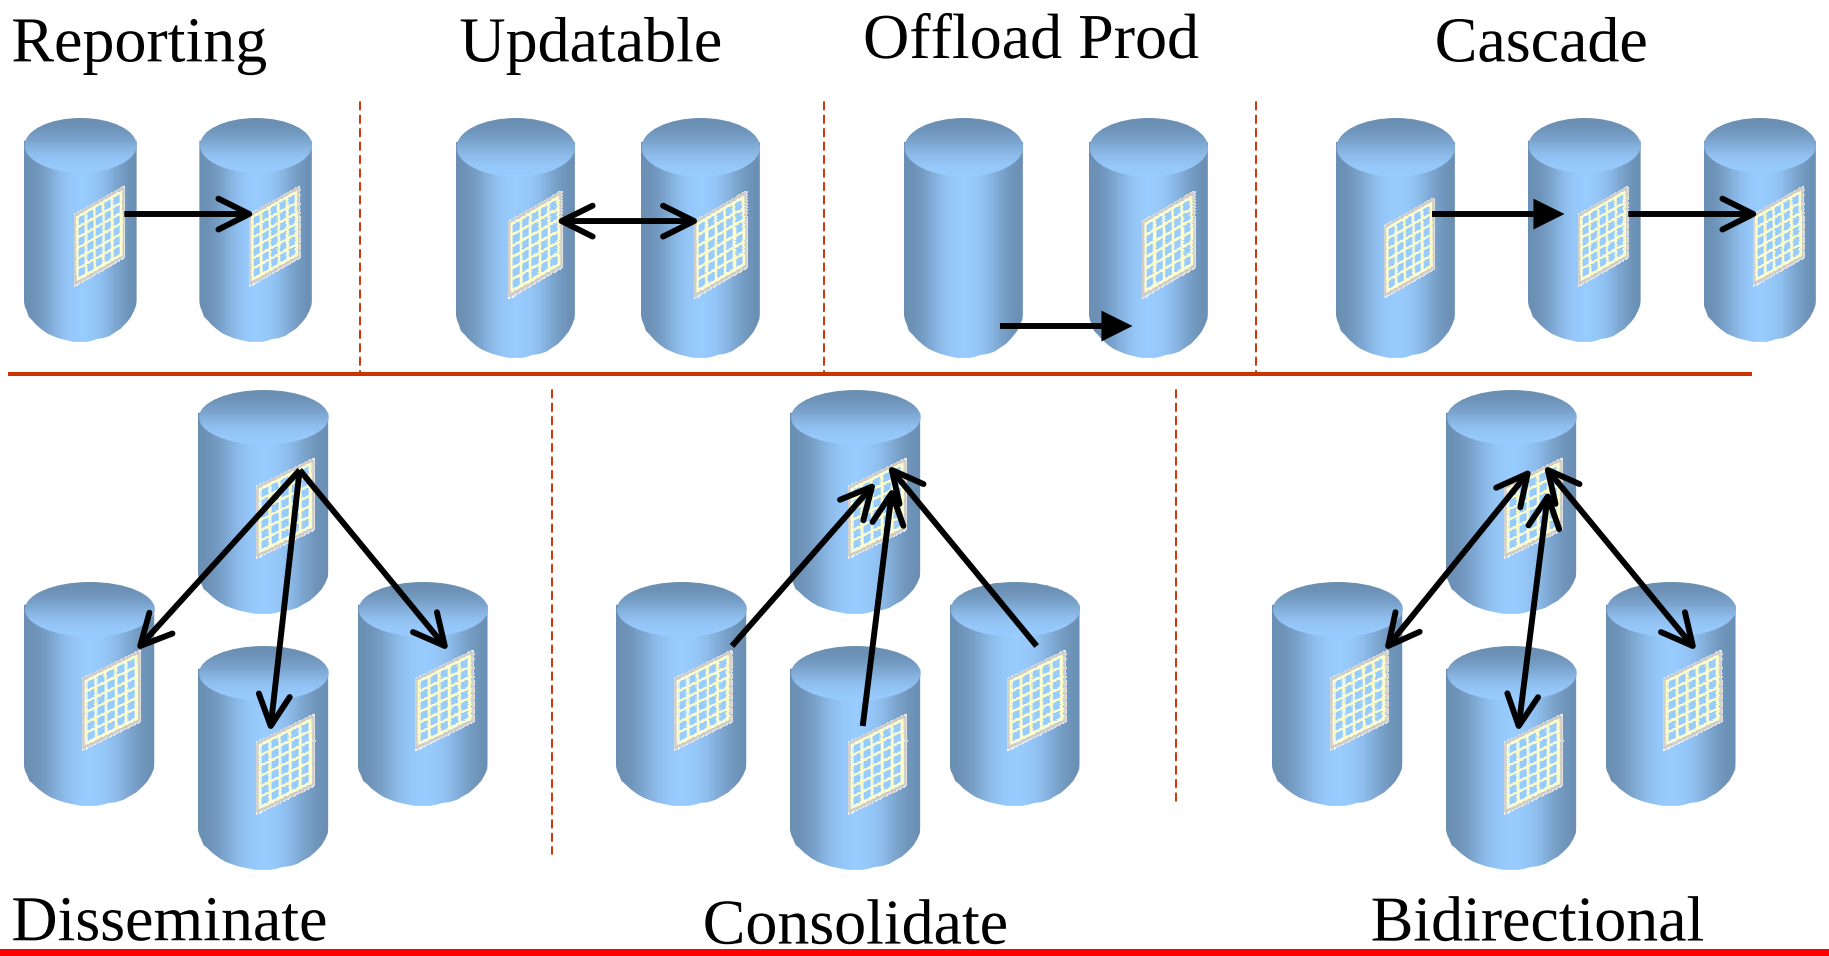

N-way Multimaster Replication

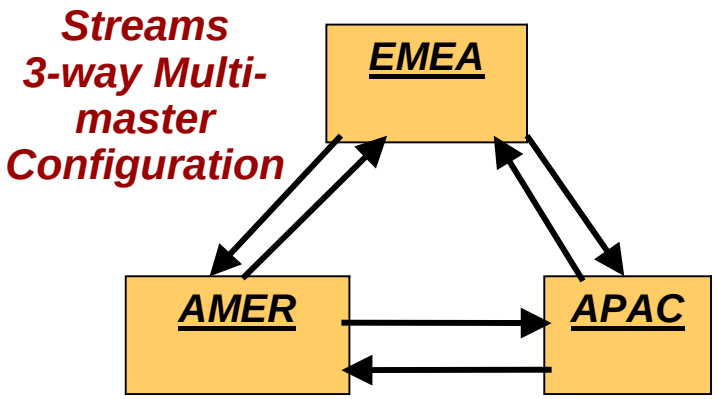

- Real-time data replication useful to provide site autonomy and balancing usage
- Geographically distinct locations, frequently across WAN
- Consistent data available at each database
- Updates at each database shared directly with each of the other databases
- Automatic conflict detection, customizable conflict resolution

Each Database

3 queues: 1 for each process

1 capture (local changes)

2 apply processes

- changes from each of the other source databases

Rule-based Transformations

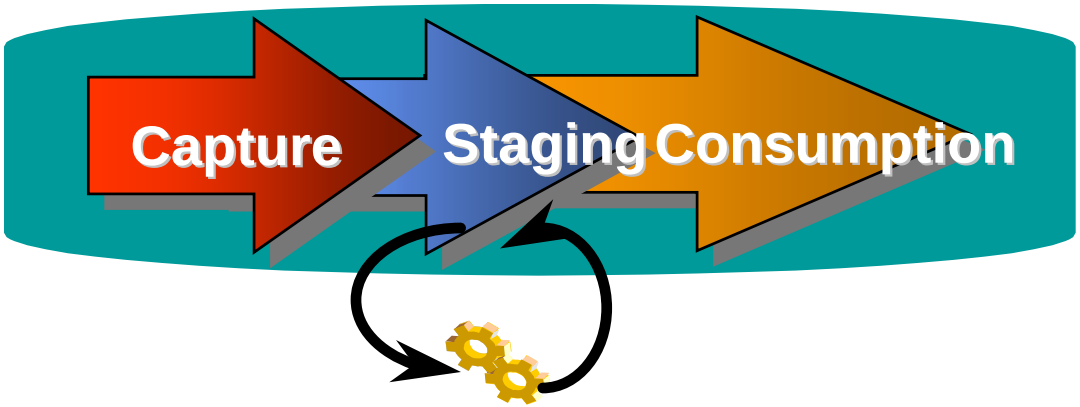

- Transformations can be performed:
  - As events enter the staging area
  - As events leave the staging area
  - As events propagate between staging areas
- Declarative Transformations
  - Rename schema, table, column
  - Add or delete column
- Custom Transformations
  - User-supplied PL/SQL function

Declarative Transformations

```
BEGIN
   DBMS_STREAMS_ADM.RENAME_SCHEMA(
   rule_name => 'STRMADMIN.HR51',
   from_schema_name => 'HR',
   to_schema_name => 'HR_REPL',
   operation => 'ADD');
END;
/
```

```
SELECT rule_owner||'.'||rule_name
  rule,transform_type,from_schema_name,to_schema_name from
  DBA_STREAMS_TRANSFORMATIONS;
```

RULE	TYPE	FROM	TO
STRMADMIN.HR51	DECLARATIVE TRANSFORMATION	HR	HR_REPL

Streams Replication

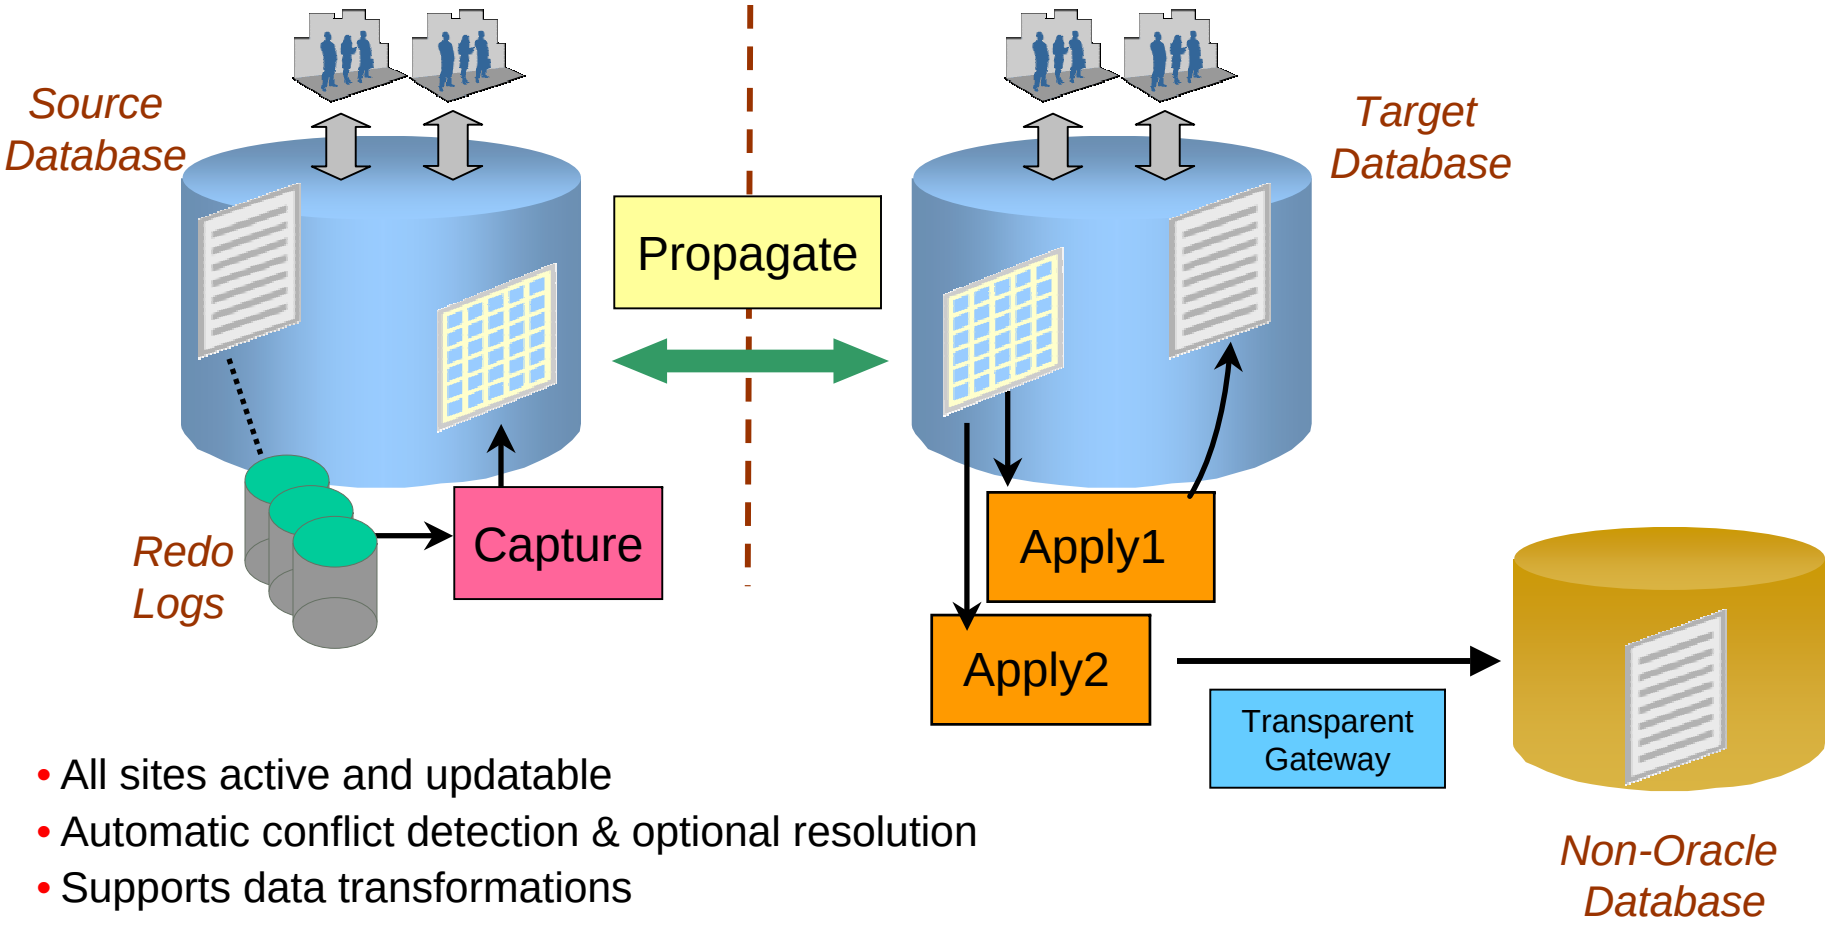

- All sites active and updatable
- Automatic conflict detection & optional resolution
- Supports data transformations
- Flexible configurations – n-way, hub & spoke, …
- Database platform / release / schema structure can differ
- Provides HA for custom apps where update conflicts can be avoided or managed

Downstream Capture

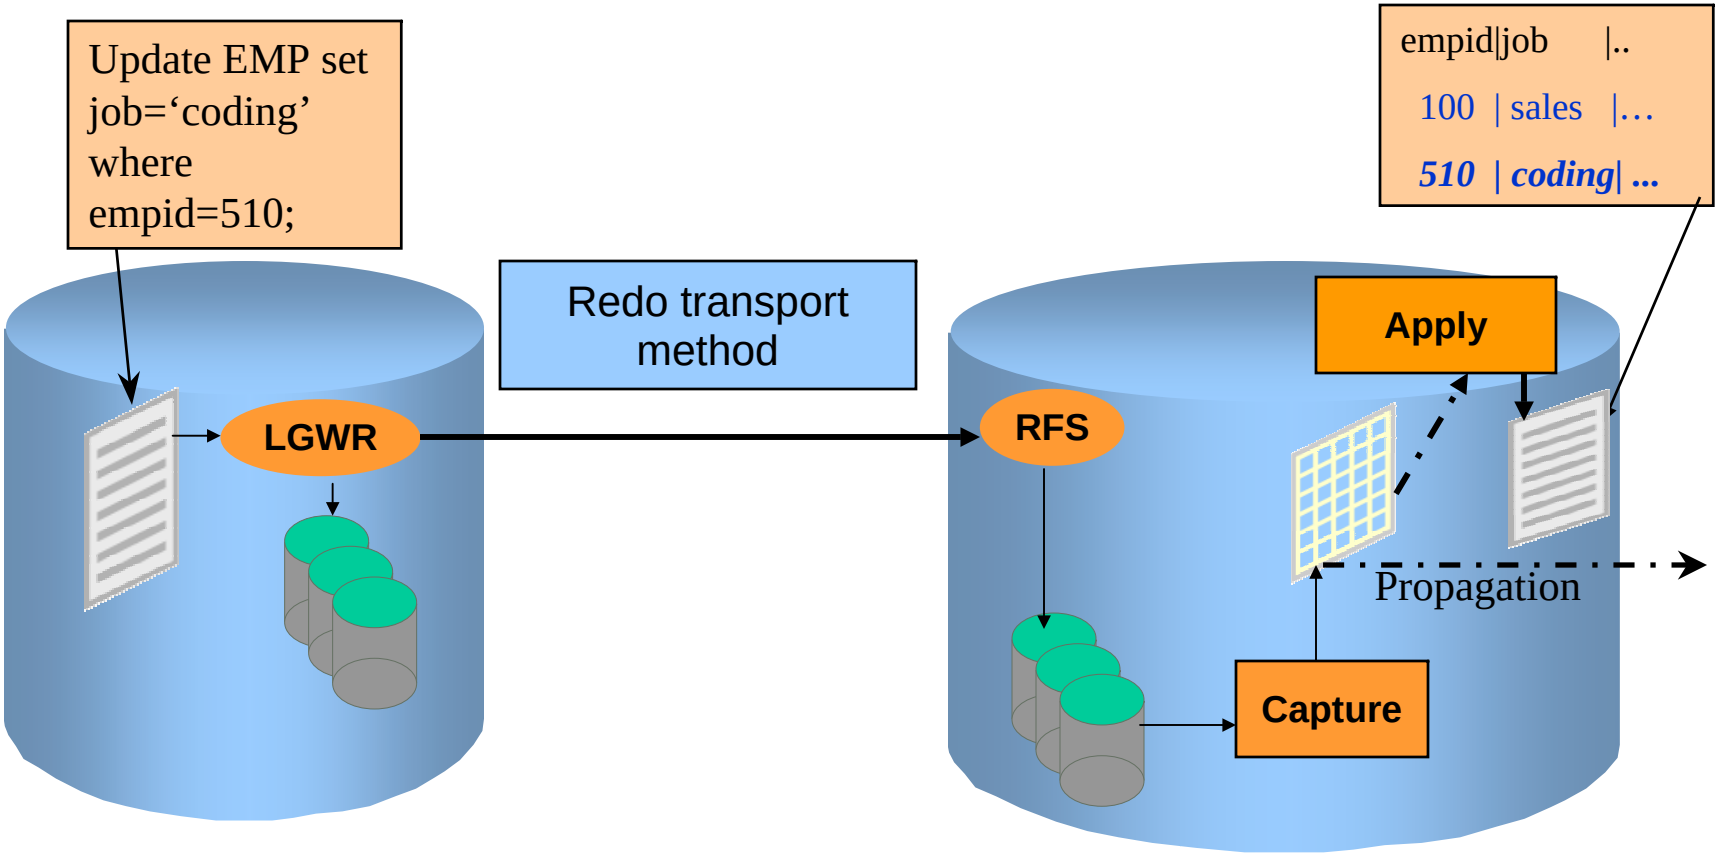

Streams Enhancements

ORACLE DATABASE 11g

- Documentation: 2Day+ Data Replication & Integration
- Source and Target data compare & converge
- Streams Performance Advisor
- Split/Merge of Streams for Hub & Spoke replication
  - Maintains high performance for all replicas
  - Automated, fast “catch-up” for unavailable replica
- Cross-database LCR tracking
  - Trace Streams messages from start to finish in single view
- Performance optimizations
- Streams Synchronous Capture
  - Available in all Editions of Oracle Database 11*g*
  - Efficient internal mechanism to immediately capture change

Streams Replication Best Practices

Ref. Chap 15 - "Best Practices for Streams Replication Databases", Oracle Streams Replication Administrator's Guide 11g .

Ref. Chap 14 - "Best Practices for Streams Replication Databases", Oracle Streams Replication Administrator's Guide 10g Release 2 (10.2) .

General Configuration Tips

- Separate queue for each capture and apply, also for each source database

- init.ora:
  - **`_job_queue_interval=1`** (set at the source site, indicates scan rate interval (seconds) of job queue) (10g recommendation)
  - **`streams_pool_size=200M`** (set at source and destination, ref. to **`V$STREAMS_POOL_ADVICE`**)

- For WANs: (SQL*NET parameters)
  - Increase SDU (sqlnet.ora, tnsnames.ora, listener.ora)
  - Increase send_buf_size, recv_buf_size
  - For downstream capture, refer to MAA best practice paper: *Oracle Database 10g Release 2 Best Practices: Data Guard Redo Transport & Network Configuration*

Rule Tips

- Spelling counts!
- source_database_name = GLOBAL_NAME of source
  - Must be correctly specified for each process
- Eliminate duplicate or overlapping rules in a rule set
- Use the negative rule set (available with10*g* and above)
  - Ensure that rules do not allow objects with unsupported data types
- Avoid complex rules
  - LIKE
  - Functions
  - NOT (9.2)

Streams Process Parameters

Capture (DBMS_CAPTURE_ADM package)

- Set retention time for capture checkpoints as needed

  Alter_capture(‘captureName’, checkpoint_retention_time=>7)

  (checkpoint_retention_time available in 10.2 and above)

Propagation (DBMS_PROPAGATION_ADM package)

- Use queue_to_queue parameter (set to TRUE)
  - Source and target must be 10.2 or above

Apply (DBMS_APPLY_ADM package)

- Set_parameter(‘applyName’,’parallelism’,’4’) (as needed)
- Set_parameter(‘applyName’,’disable_on_error’,’N’)

10g Additional Recommendations

Capture (DBMS_CAPTURE_ADM package)

- Reduce the capture checkpoint frequency parameter

  Set_parameter(‘captureName’,’_checkpoint_frequency’,’1000’)

Apply (DBMS_APPLY_ADM package)

- Set_parameter(‘applyName’,’_hash_table_size’,’1000000’)

Apply Performance Tips

- **`ALTER TABLE SYS.STREAMS$_APPLY_PROGRESS INITRANS 16 PCTFREE 10;`**
  - **`INITRANS`** should match the apply parallelism

- Batch Processing
  - Frequent commits (transaction size < 1000 LCRs)
  - Consider procedural replication – replication of a PL/SQL procedure call, rather than the DML generated by the PL/SQL procedure – sample code available

Handling Errors

- Error queue stores information about transactions that could not be successfully applied
  - Created automatically when queue created
  - Only contains information about local errors
  - Contains all LCRs for each error transaction
- **DBMS_APPLY_ADM.EXECUTE_ERROR()** has optional error handler capability
- Apply Error Management
  - **MESSAGE_NUMBER** column in **DBA_APPLY_ERROR**
  - Ref. scripts in Manual "*Oracle Streams Concepts and Administration*", "*Monitoring Streams Apply Processes*", and Section – "*Displaying Detailed Information About Apply Errors*" [Chap 22 in 10.2, Chap 27 in 11g]

Operational Tips

- Streams processes restart automatically, if enabled before database shutdown
- Automatic flow control between capture process and targets
  - Rate controlled by fastest subscriber (target)

- Replicating DDL?
  - Avoid replicating system-generated constraint/index names
  - Modify manual hotbackup scripts to set an apply tag
    - Eg: **`dbms_streams.set_tag('99')`**

- Removing archived log files from disk ?
  - Do not remove log files that may be needed for capture restart
    - Minimum **`REQUIRED_CHECKPOINT_SCN`**
    - Configure the source database to store archived redo log files in a location other than the flash recovery area – applies to downstream database also

Active/Active Considerations

- Potential for updates to same row at same time in multiple locations
  - Avoid conflicts, as much as possible
  - Implement conflict resolution methods to automate error handling
- If replicating DDL, perform DDL from single database
- Manage keys generated from sequences

Streams Monitoring

- Views
- STRMMON – MetaLink Note:290605.1 – *Oracle Streams STRMMON Monitoring Utility*
- UTLSPADV – 11g STRMMON replacement
  - Package in *$ORACLE_HOME/rdbms/admin* directory
- Healthcheck – MetaLink Note:273674.1 – *Configuration Report and Health Check Script*
- Alert Log –
  - Look out for messages from the Capture process, logged when Capture sees at least 1 large or long-running transaction in its active transaction workload

Enterprise Manager 11g Database Control

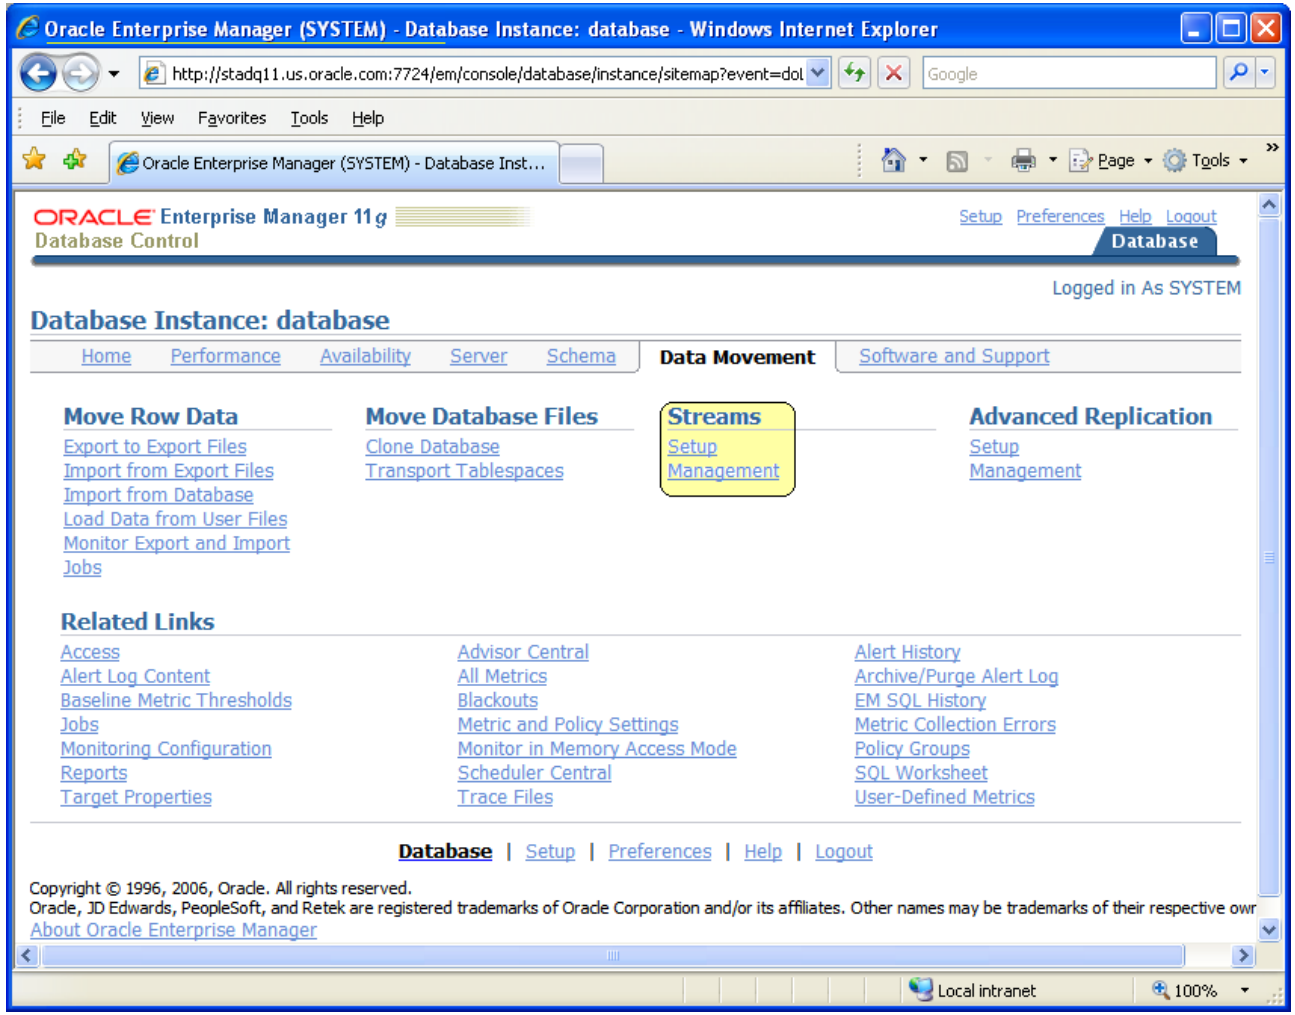

Enterprise Manager 10g Grid Control

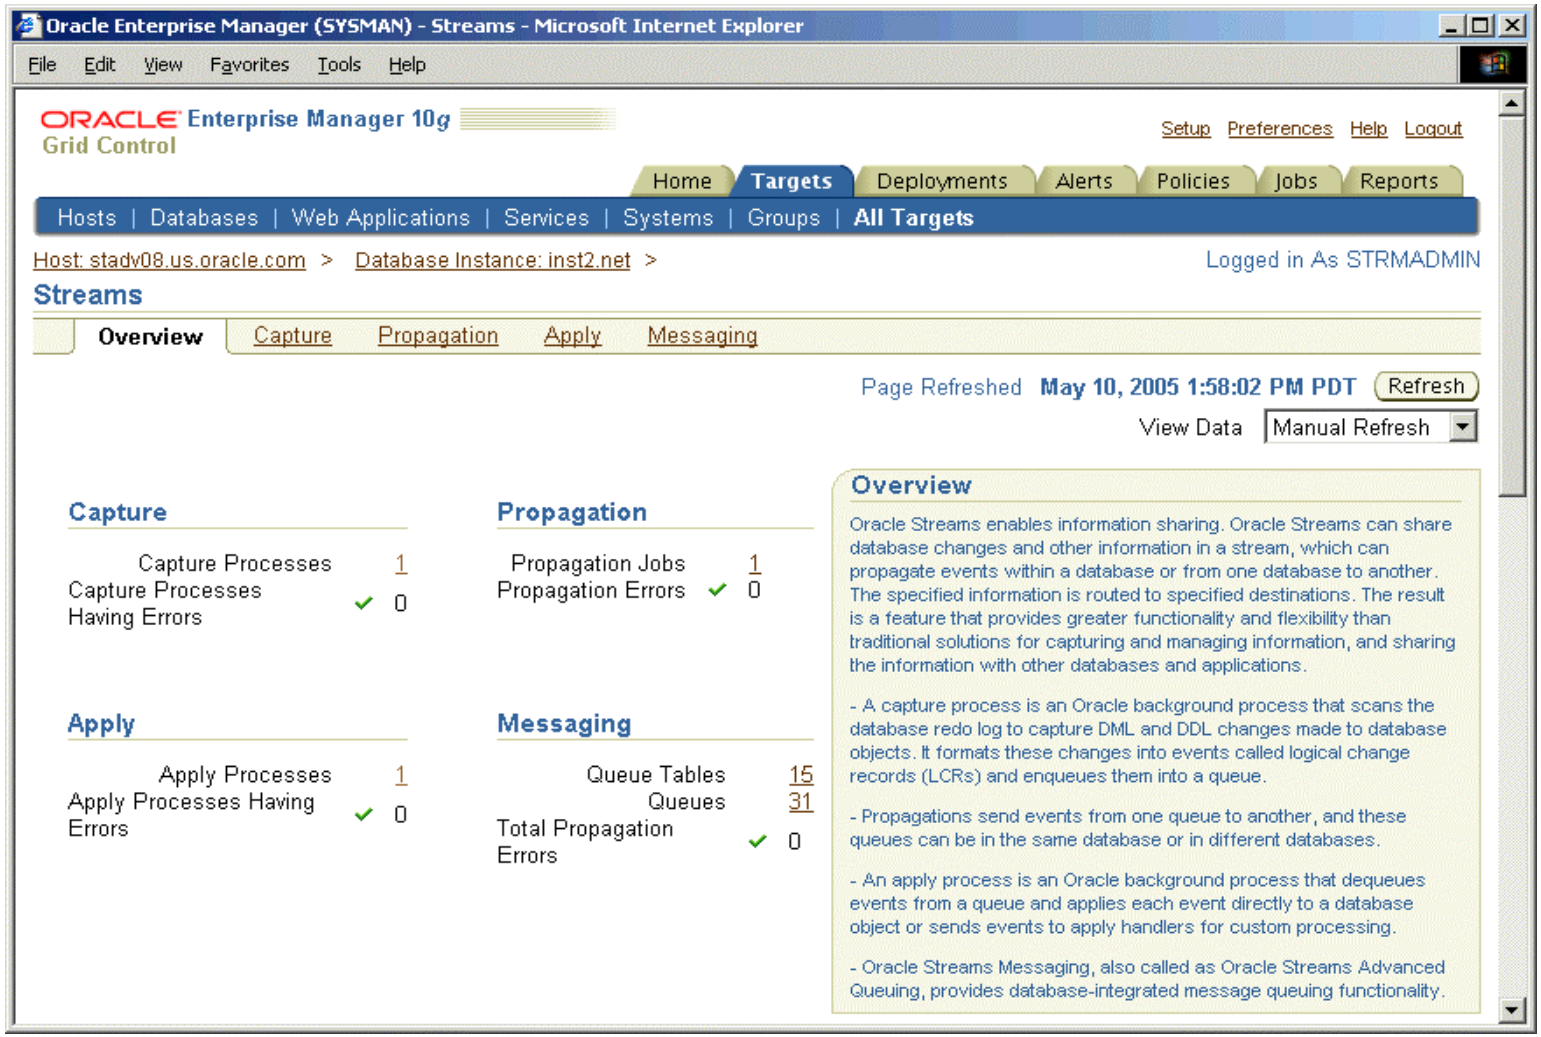

Source Site: Heartbeat

- Implement periodic Job to update “heartbeat” table
  - 2 columns: Dbname, timestamp
  - Configure Streams capture/apply for “heartbeat” table
  - Configure Conflict Resolution OVERWRITE or Latest Timestamp
- Generates activity in Database
- Quick status for DBA

Periodic Maintenance

- **DBMS_CAPTURE_ADM.BUILD()**
- **DBMS_CAPTURE_ADM.PREPARE_GLOBAL_INSTANTIATION()**
  - Default prepare enables supplemental logging
  - To avoid supplemental logging configuration, use ‘NONE’ instead of default ‘KEYS’

Streams Summary

- Features:
  - Redo-based Change Capture
  - Customizable Apply Engine
  - Schema Evolution
  - Transformations
  - Heterogeneous Support
- Versatile:
  - Replicate Data
  - Consolidate Information
  - Provide High Availability during database migration, upgrade

CERN Case Study

Resources

- Collateral, best practices, sample code: http://otn.oracle.com/products/dataint/
- MetaLink Note 418755.1 – 10.2 Streams Recommendations

Database HA Sessions From Oracle Development

Monday, Nov 12

- S291483 - The Fastest and the Most Cost-Effective Backup for Oracle Database: What's New in Oracle Secure Backup 10.2, 11:00 am - 12:00 pm, Moscone South 304
- S291492 - Oracle Database 11*g*: Next-Generation High Availability, 12:30 - 1:30 pm, Moscone South 103
- S291923 - Implementing Oracle Maximum Availability Architecture (MAA) at Allstate Insurance Using Oracle 10*g* RAC, ASM, Oracle Data Guard and Oracle Grid Control, 3:15 - 4:15 pm, Moscone South 304
- S291484 - Oracle Database 11*g* Data Repair Technologies: Comprehensive, Intelligent Recovery, 4:45 - 5:45 pm, Moscone South 304

Tuesday, Nov 13

- S290710 - Maximum Availability Architecture Best Practices: Oracle E-Business Suite 12, 12:15 - 1:15 pm, Marriott Salon 10 & 11

Wednesday, Nov 14

- S291915 - What's New in Oracle Data Guard 11*g*: Revolutionizing Data Protection and Availability, 9:45 - 10:45 am, Moscone South 304

Database HA Sessions From Oracle Development

Wednesday, Nov 14

- S291487 - Backup and Recovery Best Practices for Very Large Databases (VLDB), 11:15 am - 12:15 pm, Moscone South 304
- S291920 - Oracle Active Data Guard: How to Utilize Your Standby Databases for Production Workload - What They Didn't Print in the Manuals, 3:00 - 4:00 pm, Moscone South 304
- S291917 - Oracle Data Guard Tips and Tricks: Direct From Oracle Development, 4:30 - 5:30 pm, Moscone South 102

Thursday, Nov 15

- S291495 - Oracle Streams Replication and Advanced Queuing (AQ): What's New in Oracle Database 11*g*, 8:30 - 9:30 am, Moscone South 304
- S291499 - Best Practices for Implementing Replication with Oracle Streams in Oracle Database 10*g* and 11*g*, 10:00 - 11:00 am, Moscone South 304
- S291525 - Maximum Availability Architecture (MAA) Best Practices: Online Patching, Rolling Upgrades and Planned Maintenance with Minimal Downtime with Oracle Database, 11:30 am - 12:30 pm, Moscone South 104
- S290542 - Maximum Availability Architecture (MAA) Best Practices for Siebel 8.0, 2:30 pm - 3:30 pm, Marriott Salon 10 & 11

Database HA Demos From Oracle Development

Monday, Nov 12 – Thursday, Nov 15
Oracle DEMOgrounds, Moscone West

Oracle Active Data Guard

Oracle Streams: Replication and Advanced Queuing

Oracle Secure Backup

Recovery Manager (RMAN) and Flashback Technologies

Maximum Availability Architecture


Q&A
QUESTIONS
ANSWERS

For More Information

search.oracle.com

or

oracle.com


ORACLE®

ORACLE IS THE INFORMATION COMPANY