

The Oracle logo is displayed in a bold, red, sans-serif font against a solid black rectangular background. The word "ORACLE" is written in all caps, with a registered trademark symbol (®) positioned at the top right of the letter "E".

**Bjørn Engsig**  
Principal Member Technical Staff  
Server Technologies

A small, white Oracle logo is located in the bottom right corner of a red horizontal bar. The word "ORACLE" is written in all caps in a sans-serif font.

# Oracle Streams Overview

## Foundation for Replication in Oracle10g

ORACLE

## Oracle Streams

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

ORACLE

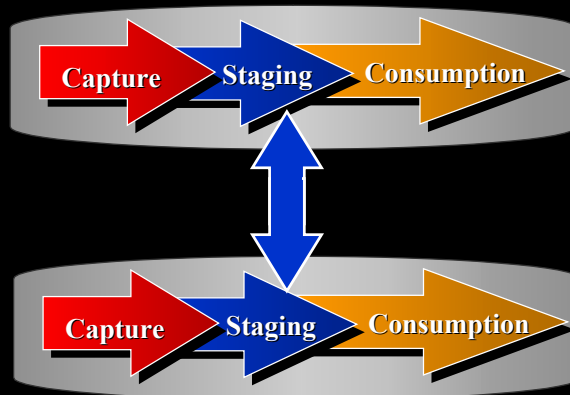
## Streams Basic Elements



- Capture
- Staging
- Consumption (apply)

ORACLE


## Multi-Database Streams



- A stream can contain multiple elements from multiple databases
- Events flow between staging areas

ORACLE


# Capture



- Streams captures events
  - Implicitly: log-based capture of DML and DDL
  - Explicitly: Direct enqueue of user messages
- Captured events are published in the staging area
- SQL and messaging APIs in multiple languages
  - JDBC, JMS, PL/SQL, C, SOAP

ORACLE

# Log-Based Change Capture



- Low overhead, low latency change capture
  - Changes to the database are written to the online redo log
  - Oracle Streams can extract changes from the log as it is written (hot mining)
  - Changes are formatted as a Logical Change Record (LCR), a representation of the change

ORACLE

# Logical Change Record (LCR)

- Database change = LCR
  - DML
    - Object name, owner, Type of DML, SCN
    - Row change = LCR
    - OLD, NEW values
  - DDL
    - Object name, owner, Type of DDL, SCN
    - DDL text
  - LOB
    - Multiple LCRs per LOB
    - Piecewise chunks

ORACLE

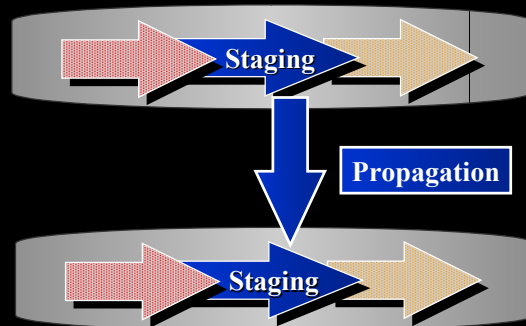
# Staging



- Streams publishes captured events into a staging area
  - Implemented as a queue
  - Supports for new type “any” datatype allows a single staging area to hold any type of data
  - All events, LCRs and user-messages, can be staged in the same queue
  - Messages remain in staging area until consumed by all subscribers

ORACLE


# Staging Area Propagation



- Other staging areas can subscribe to events
  - in same database
  - in a remote database
- Events can be routed through a series of staging areas

ORACLE


# Transformations



- Transformations can be performed
  - as events enter the staging area
  - as events leave the staging area
  - as events propagate between staging areas
- Transformation examples
  - change format, data type, column name, table name

ORACLE


# Consumption



- Staged events are consumed by subscribers
  - Implicitly: Apply Process
    - Default Apply
    - User-Defined Apply
  - Explicitly: Application dequeue via open interfaces
    - JMS, C, C++, PLSQL, SOAP (XML/HTTP)

ORACLE


# Default Apply



- The default apply engine will directly apply the DML or DDL represented in the LCR
  - apply to local Oracle table
  - apply via DB Link to non-Oracle table
- Automatic conflict detection with optional resolution
  - unresolved conflicts placed in exception queue
- Parallel apply maximizes concurrency

ORACLE

## User-defined Apply



- User-written custom apply procedures
- Written in PL/SQL, Java, C, C++
- Uses:
  - full control over apply
  - normalizing or denormalizing data
  - populating related fields or tables

ORACLE

## Rule-based Configuration

- Consumers subscribe to published events
- Content-based subscription
- Rule is expressed as SQL WHERE clause

```
dbms_rule_admin.create_rule(
  rule_name=>'scott.rule1',
  condition=>':dml.get_object_owner() = "SCOTT" AND
:dml.get_object_name()="EMP"');
```
- Rule sets for simplicity
- Rule sets govern capture, staging, and apply
  - Inclusion
  - ★ 10g – Negative
- Dynamic rule maintenance

ORACLE



# Replication Rules


- Rule-based
  - ★10g – DML, DDL, MSG
    - content
    - True/False
- Granularity
  - Table
  - ★10g – Subset of Table
  - Schema
  - Database
  - ★10g – Tablespace
- Tailored Replication API
  - DBMS\_STREAMS\_ADM
    - ADD\_TABLE\_RULES
    - ADD\_SCHEMA\_RULES
    - ADD\_GLOBAL\_RULES
    - ADD\_SUBSET\_RULES
    - ★10g – MAINTAIN\_TABLESPACES
    - ★10g – ADD\_MESSAGE\_RULE

ORACLE

# Directed Networks

- Propagation independent of Apply
- Rules-based subscription determine if event is locally applied
  - London applies UK only
- WAN Friendly
  - Send once, fan out
  - NY-->London, London-->Milan, London-->Paris

INSERT ...  
VALUES ('EUROPE',' ITALY')



ORACLE

## Automatic Conflict Detection

- Automatic conflict detection with user-selectable conflict resolution routines
  - latest timestamp, earliest timestamp, maximum or minimum value, overwrite, discard
  - User-definable resolution routines
- Conflict detection compares current row values at receiving site with “old” values of changed row from the originating site
  - if match, “new” values are applied to row
  - if not, conflict resolution method is used, if supplied
  - if still unresolved, place transaction in exception queue
- ★ Ability to disable conflict detection by column or table

ORACLE

QUESTIONS  
&  
ANSWERS

ORACLE

