

Oracle Database 11g: Oracle Streams

An Oracle White Paper
July 2007

Oracle Database 11g: Oracle Streams

EXECUTIVE OVERVIEW

Businesses have a need to share information in a variety of ways. This can include disseminating price lists to field offices, consolidating sales data from multiple stores, provisioning data in a grid environment, or notifying a payroll application that a new employee has been hired. By sharing information in near real time, businesses can synchronize multiple copies of data to provide improved availability and scalability, for example, by synchronizing multiple call center locations.

Oracle Database 11g provides a unified solution for information sharing—Oracle Streams. Oracle Streams captures and distributes database updates, events, and application messages. It can automatically apply updates to destination databases, or pass events and messages to custom procedures and applications. Combining these capabilities provides an extremely flexible solution for replication, message queuing, and event notification solutions. Unlike traditional solutions, you can use Streams to create powerful information sharing solutions that incorporate some or all of the capabilities of classic information sharing solutions, all through the use of a single feature.

This paper discusses how the three basic elements of the Oracle Streams technology (capture, stage, consumption) are used to share information in the Oracle 11g database. It further describes how these features enable not only efficient information sharing and verification, but efficient data provisioning for grid computing as well.

ORACLE STREAMS OVERVIEW

As a company grows and expands, it becomes increasingly important for that company to be able to share information among multiple databases and applications. Traditionally, a company may have selected from among a variety of information sharing technologies, each aimed at solving a specific business problem. While these targeted solutions may have initially appeared easier to use and implement, they fail to deliver once the needs of the company grow beyond the capabilities of the simple targeted solution. Suddenly, developers must implement multiple tools to build a solution, and complexity grows exponentially.

Oracle Database 11g provides a single, unified solution for information sharing—Oracle Streams. Oracle Streams provides a set of elements designed to facilitate the capture, staging, and consumption of events within the Oracle database. These

events include both messages enqueued into a database queue by applications, as well as, modifications to database objects, such as tables or procedures, using DML or DDL.

Each element of streams can be configured explicitly or implicitly. Explicit capture allows user applications to explicitly enqueue messages into a staging area on the source database. Using the implicit capture mechanism of Oracle Streams, changes made to a database can be efficiently captured and replicated to one or more remote systems with little impact to the originating system. This capture mechanism extracts both data changes (DML) and structure changes (DDL) from the redo log and publishes these updates to the staging area.

Both explicitly and implicitly captured changes can be propagated to one or more remote staging areas where they can be applied at the destination site. Changes in a staging area are consumed by the apply engine, where the changes they represent are applied to a database, or they are consumed by an application. Oracle Streams includes a flexible apply engine, that allows use of a standard or custom apply function. This enables data to be transformed when necessary.

Existing applications can use Oracle Streams with minimal modification or special handling. Oracle Streams can specify rules at multiple levels of granularity: database, schema, and table. Each element of Oracle Streams can use these rules to easily configure, for example, the capture of changes for an entire database, or a set of schemas, or individual tables.

Using these basic elements, users control: how information is put into a stream, how the stream flows or is routed from node to node, what happens to events in the stream as they flow into each node, and how the stream terminates. By specifying the configuration of the elements acting on the stream, a user can address specific requirements. To simplify the deployment of Oracle Streams, Oracle provides applications specially configured for specific markets.

Unique Databases

Databases using Oracle Streams technology need not be identical. Participating databases can maintain different data structures using Streams to transform the data into the appropriate format. Streams provides the ability to transform the stream at multiple points: during capture, while propagating to another destination, or during application at the destination site. These transformations can be Oracle-supplied or user-defined functions registered within the Oracle Streams framework.

Transformations can be used to change the data type representation of a particular column in a table, change the name of a column in a table, or change a table name.

The data at each site can be subsetted based on content as well. For example, you could implement a rule that only events (or changes) related to the employees of a particular division, based on the department identifier column, be applied to a particular table. Oracle Streams automatically manages the changes to ensure that the data applied to the table matches the subset rule criteria.

Directed Networks

Although locally captured or enqueued events can be consumed locally, they can also be propagated to one or more remote locations. The directed network capability of streams allows changes to be directed through intermediate databases as a pass-through. Changes at any database can be published and propagated to or through other databases anywhere on the network. By using the rules-based publish and subscribe capabilities of the staging area queues, database administrators can choose which changes are propagated to each destination database, and can specify the route messages will traverse on their way to a destination. This directed network approach is also friendly to Wide Area Networks (WAN), enabling changes to subsequent destinations to traverse the network once to a single site for later fan-out to other destinations, rather than sending to each destination directly.

Table Data Comparison and Convergence

An important adjunct to sharing data in multiple locations is the ability to verify data consistency between databases. The Oracle Database 11g includes a package to compare table data between databases. Complete data, data ranges, or data subsets may be checked on a periodic or as needed basis, without interference to running applications. Data consistency can be confirmed at the table or row level and differences can be re-examined in case of transient differences due to in-flight transactions. If necessary, divergent data can be converged so that both compared databases reflect the same data.

STREAMS ARCHITECTURE

There are three basic elements of Oracle Streams: capture, staging, and consumption.



Capture

Oracle Streams supports capture of events into the staging area. These events are captured in three ways. Explicit capture allows applications to explicitly generate events and place them in the staging area. Implicit capture automates the capture of changes within the database. Oracle Database 11g provides two techniques for implicitly capturing changes: log-based capture and synchronous capture. With log-based capture, the server captures DML and DDL events for a source database by mining the source redo logs and archive logs locally. Alternatively, the server can mine the redo and/or archived logs of the source database at an alternate database,

assuming the alternative database is on a similar platform type and operating system. With synchronous capture, DML changes are captured via an efficient internal mechanism during the user's transaction activity.

Log-Based Capture

The log-based capture mechanism leverages the fact that changes made to tables are logged in the redo log to guarantee recoverability in the event of a crash or media failure. Capturing changes directly from the redo log files minimizes the overhead on the system. Oracle can read, analyze, and interpret redo information, which contains information about the history of activity on a database. Oracle Streams mines the information and delivers change data to the capture process. Tables identified for replication with Oracle Streams automatically log supplemental information into the redo stream, such as primary key columns, to facilitate the delivery of this information. Additional information, such as columns required for conflict resolution, must also be logged.

The capture process consists of several subcomponents. A reader server reads the redo log and divides the redo log into regions. Preparer servers scan the regions defined by the reader server in parallel and perform prefiltering of changes found in the redo log, based upon user-defined rules. Thus, only changes to desired objects are captured.

A builder server merges redo records from the preparer servers and passes the merged redo records to the capture process. The capture process then formats each change into a Logical Change Record (LCR). If the LCR is found to satisfy the defined rules it then enqueues the LCR into the staging area for further processing. The capture process is an Oracle background process.

Logical Change Records

Logical change records describe changes made to a single row of a table modified with a single DML statement. A single DML statement that operates on multiple rows within a table will generate multiple LCRs, and a single transaction can consist of multiple DMLs. Each logical change record includes the name of the changed table, the old and new values for any changed columns, and the values for the key columns. Armed with this information, changes can be applied to the correct rows at the destination sites and conflicts can be detected.

It can be important to differentiate between changes that originate at different sites, so each LCR is tagged with identification information from the source database about the origin of the change. These tags are typically used during the consumption phase of Streams. In certain cases this information is used to filter out changes that originated at another site and were made by the apply process, to prevent their being placed in the staging area and cycling back to the originating site. This is most notably the case in an N-way replication configuration, where each replicated site communicates its changes directly to every other site in the configuration. In other cases, for example a hub and spoke configuration, it is

important to capture all changes, including those that originate at other sites. In these cases, it may be desirable to have changes from an individual spoke reflected through the hub to other spokes in the configuration

Local vs. Downstream Capture

Oracle Streams supports mining from the in-memory log buffer, hot mining of the active redo log, as well as mining archived log files. In the case of hot mining, the redo stream is mined for change data at the same time it is written to the active redo log, reducing the latency of capture. Streams seamlessly traverses the log buffer, redo log, or archived log as necessary without any additional configuration. Only LCRs that satisfy the selection criteria for capture are placed in the staging area.

Streams also has the ability to capture changes for a source database at a different server. Using Oracle's Log Transport Services, the archived log files of the source database are written to the remote server as they are written locally, using the same data protection modes as Oracle Data Guard. The remote server, capturing changes into the staging area on this remote server, then mines these remote logs. Should the remote server be a subscriber to the changes, they can also be applied. This has two important benefits. First, it allows for complete offloading of Streams activities from the production database server—often a critical requirement for high volume OLTP databases. Secondly, since the log files are written remotely using Oracle Data Guard protection modes, you can choose the mode appropriate to your environment. A single remote server can be used to capture changes for multiple source databases.

Synchronous Capture

Synchronous capture uses an internal mechanism to capture DML changes to specified tables. When a DML change is made to a table, it can result in changes to one or more rows in the table. Synchronous capture captures each row change and converts it to a **Logical Change Record (LCR)**. After capturing an LCR, synchronous capture writes a message containing the LCR to disk into a **queue**. Synchronous capture can be especially useful in situations where you want to replicate only a small subset of the total number of tables in your database.

Synchronous capture offers the same tagging capabilities as log-based capture. Each LCR is tagged with identification information from the source database about the origin of the change. These tags are typically used during the consumption phase of Streams. In certain cases this information is used to filter out changes that originated at another site and were made by the apply process, to prevent their being placed in the staging area and cycling back to the originating site. This is most notably the case in an N-way replication configuration, where each replicated site communicates its changes directly to every other site in the configuration. In other cases, for example a hub and spoke configuration, it is important to capture all changes, including those that originate at other sites. In these cases, it may be

desirable to have changes from an individual spoke reflected through the hub to other spokes in the configuration

Explicit Capture

Explicit capture allows applications to explicitly generate events and place them in the staging area. Applications can use strongly typed queues, as has been characteristic of the queues in previous releases of the database, or can use the Streams AnyData queues. AnyData queues can hold messages of different types, so it is possible to enqueue different types of messages into a single staging area. If desired, these messages can be formatted as LCRs, and subsequently consumed by an apply engine.

Staging

Once captured, events are placed in a staging area. The staging area is a queue designed to store and manage captured events. Database changes, already formatted as LCRs by the capture process, are stored in an in-memory queue until consumed by all subscribers. LCRs captured via synchronous capture are stored in a disk staging area. Events explicitly queued to the staging area by applications for future processing are also maintained in the staging area. The queue provides a holding area with security, as well as auditing and tracking of event data.

Subscribers examine the contents of the staging area and determine whether or not they have an interest in the message representing that event. A subscriber can either be a user application, another staging area (usually on another system), or the default apply process. The subscriber can optionally evaluate a set of rules to determine whether the message meets the criteria set forth in the subscription. If so, then the message will be consumed by the subscriber.

If the subscriber is a user application, the application will explicitly dequeue the message from the staging area in order to consume the message. If the subscriber is another staging area, the message will be propagated and enqueued to that staging area. If the subscriber is an apply process, the messages will be dequeued and consumed by the apply process.

Propagation

Events in the staging area can be propagated to staging areas in other databases. To simplify network routing and reduce WAN traffic, events need not be transmitted to all databases and applications. Rather, they can be directed through queues on one or more systems until they reach the subscribing system. For example, consider a network between three sites A, B, and C where sites A and B can communicate directly, as can B and C, but sites C and A do not have direct communication. Streams is configurable to allow the changes at site A to be reflected at site C via site B. Site B can pass the events through to site C without requiring site B to apply the change locally. Of course, if the change is also desired at this intermediate site

as well, the change can be applied at site B without impacting the pass-through to site C.

Administrators have a great deal of flexibility to use in specifying the routing of the Streams. By using the rules-based publish and subscribe capabilities of the staging area queues, they can choose which changes are propagated to each destination database, and can specify the route messages will traverse on their way to a destination. Throughout the system, Oracle Streams maintains the source database information for each event, as it can be important to differentiate between changes that originate at different sites. The administrator controls whether to select all changes, including those that originate at other sites, or only those changes made at the local site and not by the apply engine

Consumption

Messages in a staging area are consumed by the apply engine, where the changes they represent are applied to a database, or they are consumed by an application. The Oracle Streams apply engine applies DML changes, DDL changes, user-supplied LCRs, and user-enqueued messages. When the destination database is an Oracle database, the apply engine runs locally on the system hosting the Oracle database. For greater flexibility, the Oracle Streams apply engine supports both default and custom apply procedures. Support for explicit dequeue allows application developers to use Oracle Streams to notify applications of changes to data, while still leveraging the change capture and propagation features of Oracle Streams.

Default Apply

The default apply engine can apply automatically captured DML and DDL changes, as well as user-supplied LCRs. The default apply engine detects conflicts where the destination row has been changed and does not contain the expected original values. If a conflict is detected, a conflict resolution routine may be invoked, if desired.

The apply engine is a background process consisting of a set of server processes: an apply coordinator, a reader server, and one or more apply servers. The reader server assembles transactions for capture-generated LCRs. The apply coordinator performs transaction dependency and DML level dependency scheduling to maximize concurrency. The apply servers actually apply the changes. Since the apply coordinator and apply servers are typically in the same Oracle instance, there is no network roundtrip involved in dependency scheduling. The apply engine can invoke user-supplied functions to transform changes to the correct format or name for each LCR.

Changes from multiple databases may be sent to a single destination staging area. Multiple apply engines apply the changes from each source database concurrently. For example, queue1 from database A and queue2 from database B can propagate to queue3 at database C. Database C will use two apply engines to apply these

changes. The administrator configures the two apply engines using rule sets to ensure each apply engine applies changes from each source site. Throughout the system, Oracle Streams maintains the source database information, as it can be important to differentiate between changes that originate at different sites. The administrator controls whether to capture all changes, including those that originate at other sites, or only those changes made at the local site and not by the apply engine.

User-Defined Custom Apply

User-defined custom apply procedures enable total control over the events to be applied. With customized apply, the apply engine passes the LCR or user-enqueued message to a user-defined procedure, or apply handler. This provides greater flexibility in processing the message. To further maximize flexibility in apply handling, users can define multiple apply handlers for non-LCR messages, as well as separate apply handlers for each type of DML operation (inserts, updates, or deletes) performed on a particular table.

For example, using this custom apply capability, a user could write a procedure to skip the apply of all deletes for the Employee table. Only the behavior of LCRs that perform deletes on the table Employee are affected in this scenario. Inserts and updates to the Employee table continue to be applied using the default apply engine. Thus no delete LCRs from other sites will be performed on the Streams site replica and all of the rows in that Employee table will still exist at the replica site. These custom DML handlers allow the user to exercise complete control over the behavior of the apply engine, enabling users to implement custom Streams sites to satisfy any business requirement.

Explicit Dequeue

Similar to the explicit capture feature where applications explicitly enqueue messages into the Streams queues, user applications can explicitly dequeue LCRs or other messages from the staging area. Messages directly enqueued into the source staging area can be directly dequeued by a subscribing application without any intervention from the apply engine. Subscribing applications can directly dequeue messages from the staging area. If the application is remote, a staging area may be created in a remote database that will subscribe to messages published in the source staging area. The destination application can then dequeue from the remote staging area. Alternatively, the destination application can directly dequeue from the source staging area using a variety of standard protocols.

CUSTOMIZING ORACLE STREAMS

It is possible to control which events, or changes, are captured, propagated and applied using a combination of Oracle Streams rules and transformations.

Rules

Each element of Oracle Streams uses a set of rules to distinguish the events to be managed. These user-specified rules are enforced by capture to selectively enqueue change events into the staging area. Similarly, the apply engine will selectively apply changes based on the rules defined for apply. Rules are also used to selectively propagate events between staging areas. Rules can specify the selection of the entire database, or schemas, as well as individual tables. Rules are SQL-like expressions that evaluate to either TRUE or FALSE. Rules can also be used to limit events based on the content of the event itself. An example of a rule limited to any untagged DML changes to the HR schema is:

```
:dml.get_object_owner() = 'HR' AND  
:dml.is_null_tag() = 'Y'
```

Note the use of the tag in this rule. In most cases, changes local to a database have no tag (null tag,) while changes applied from other databases are tagged in the redo log with a non-null value.

For simplicity, rules are organized into named collections referred to as rule sets. Rules can easily be added to or subtracted from the rule set. Multiple rules can belong to a rule set and multiple rule sets can contain the same individual rules. The DBMS_RULES_ADM package is provided to facilitate the management of these rules.

Each Streams process uses two rule sets, a positive rule set and a negative rule set, and rule set evaluation is optimized to efficiently identify the events that meet the specified criteria for each process. The positive rule set identifies the rules for objects to be included for Streams processing. The negative rule set lists the rules for objects to be discarded from Streams processing. Each Streams process evaluates the negative rule set first, discarding any message that satisfies a rule in this negative rule set. The process then evaluates rules using the positive rule set. Only messages that evaluate to TRUE as a result of a rule in the positive rule set are passed on to the process.

Horizontal and Vertical Subsetting

Each of the Oracle Streams elements supports subsetting of objects, so a destination need only subscribe to a subset of the data at the source. Consider a typical distributed scenario involving headquarter databases and branch office databases. The branch office database only needs the data relevant for that branch. A branch office would therefore only subscribe to events affecting that branch. Other branches would receive different subsets.

Oracle Streams supports migration of data from one subset to another. Subscriptions are content-based, and should that content change, the subscriptions may change. This may require data to be deleted from one subset database, and inserted into another. Streams automatically manages the changes and row migration issues to ensure that the data applied to the replica reflects the subset

criteria. By configuring subset rules for propagation to a particular destination, Insert and Delete LCRs that do not match the subset criteria are not sent to the replica site. Update LCRs are examined to determine if the change affects the replica and, if so, converts the update into the appropriate command (insert or delete) to apply the change correctly. The following example shows how to configure the rules for an apply site to maintain data only for the customers in California (CA). The table Customers is in the HR schema and this subset of data is replicated to the MYREP.WORLD database into the Streams queue:

```
DBMS_STREAMS_ADM.ADD_SUBSET_PROPAGATION_RULES (
  table_name           => 'HR.CUSTOMER' ,
  dml_condition        => 'STATE=' 'CA' ' ' ,
  streams_type         => 'PROPAGATION' ,
  queue_name           => 'strmadmin.streams_queue' ,
  destination_queue_name =>
    'strmadmin.streams_queue@MYREP.WORLD'
  source_database      => 'MYDB.WORLD'
);
```

In this example, the CUSTOMER table contains a column STATE, and each replica site is limited to customers for a particular state. In this instance, state = CA. If an LCR with an insert or delete for a customer has the state column equal to CA, the LCR is sent to the target site. If the insert or delete is for any other state, the LCR is not sent. However, suppose an update to the customer table occurs showing that customer Joe has moved from California to New York. Streams will automatically remove the customer record for Joe, so that Joe will no longer appear in the replica as a California customer. If the update to the customer table shows that customer Jim moved into California from Pennsylvania, then the customer record for Jim will be automatically inserted into the database.

Vertical subsetting of data can be implemented using transformations, as described in the next section.

Transformations

A transformation is a change in the form of an object being captured, propagated or applied, or a change in the data it holds. Transformations can include changing the data type representation of a particular column in a table at a particular site, or renaming a column to a table at a particular site, or even changing the data values. Streams supplies built-in transformation codes for handling the most common customer modifications. Changing the owner of a table, renaming a table, or renaming a column within a table is easily configured. Other custom transformations can be represented by a user-supplied PL/SQL function that takes the source data type as input and returns an object of the target data type.

When a transformation is used, every event for the specified object is manipulated using the transformation. Transformations can be performed as events are placed into or removed from the staging area. A transformation during capture modifies the message before inserting it into the staging area. For example, a company may

want to replicate a table that contains a column with confidential information that should not exist at any other site. Using a transformation during capture, the column can be eliminated from the event or LCR. As a result, no other site will see this restricted information. Transformations can also be configured for propagation, which may be useful for formatting the message before it is sent to subsequent remote sites. Finally, a transformation can be specified as the event is consumed with the apply process, which can be useful for formatting a message in a manner appropriate for a specific destination.


USING ORACLE STREAMS

Oracle Streams provides a flexible architecture that supports a wide variety of information sharing and data provisioning requirements. A few of the most common uses include the following:

- Replication
- Advanced Queuing
- Integration with non-Oracle data stores
- Provisioning data in a Grid environment
- Availability during database or application upgrade

Using Oracle Streams for Replication

The following figure demonstrates the streams architectural elements as used in replication. Oracle Streams replication captures changes from a source database, stages and propagates those changes to one or more remote databases, and then consumes or applies the changes at each target database.




The figure shows an Update statement that changes the state for a row in the EMP table representing an employee whose employee id (empid) is 100. As usual, the change is recorded in the redo log. The capture process, previously configured to

collect changes made to the EMP table, retrieves the changes from the redo log, formats the information into logical change records, and places the LCR into the staging area at the local database. This LCR is propagated from the staging area (or queue) at this source database and delivered to another database that has subscribed to the changes on the EMP table. In addition, the apply engine at this second database is configured to receive the changes on the EMP table. Once the change has been propagated to the new site, the local default apply process at this second database automatically applies the change to the database.

Configuring a Replication Environment

Oracle Streams provides an easy way to instantiate replicated objects at the target sites, using Oracle export and import utilities or the RMAN duplicate database feature. The instantiation is performed at the target database without affecting any existing sites in the Streams configuration. All sites continue to process their own work while instantiation to new sites is in progress. After a new object is prepared at the source site to capture changes for replication, the Streams replication metadata for the object is marked with the appropriate SCN, to ensure no changes are lost. As copies of the objects are created at the target site using export and import, the system will note the source SCN of the copied data, and will begin applying any changes that committed after the object was copied.

To simplify configuration, administration and monitoring of Oracle Streams environments, Oracle provides a Streams tool within the Oracle Enterprise Manager Database Control. The Streams tool provides wizards that you can use to configure a Streams replication environment. You can also use this Streams tool to generate scripts, which you can then modify to meet your specific requirements.



package performs multiple configuration steps for the administrator depending on the type of Streams process specified, including the automatic generation of rules for the process. For example, the following command configures replication from a local database to the target database, NYC, identified by the database link NYC:

```
DBMS_STREAMS_ADM.MAINTAIN_SCHEMAS (
    schema_names=>'HR',
    source_directory_object=>null,
    destination_directory_object=>null,
    source_database=>null,
    destination_database=>'NYC',
    bi_directional=>FALSE,
    Include_ddl=>TRUE,
    instantiation=>
DBMS_STREAMS_ADM.INSTANTIATION_SCHEMA_NETWORK
);
```

Issuing this simple command results in the creation of a uni-directional replication environment, which will replicate DML and DDL changes from the local database to the target, NYC. The source and destination directory object parameters are set to null, because Oracle is using the network option of Data Pump to perform the initial instantiation of the replicated site. All of the configuration, including queue creation, rule generation, and initial data loading is performed automatically with this command. The administrator can further customize this configuration, if desired, using other packages and procedures related to rules and Streams processing.

Conflict Resolution


When Oracle Streams is used to support replication, both the source and destination databases are fully available to end-users for reading and writing during any replication activity. Because users can update different copies of the same table anywhere, it is possible for changes made at different database sites to update the same data element at the same time, resulting in an update conflict. The default apply mechanism will detect these conflicts as incoming replication changes are applied. Oracle provides built-in conflict resolution routines, such as "latest timestamp" or "overwrite", to automatically resolve potential conflicts. Users can choose different resolution methods for different tables. Users also have the option to create their own routines to employ resolution rules tailored to their particular business needs. Any unresolved conflicts are logged in the database for special handling or re-execution after the conflict is resolved manually.

Using Oracle Streams for Message Queuing

Streams provides a robust message queuing capability called Oracle Streams Advanced Queuing. This feature allows user applications to enqueue messages into the staging area, propagate messages to subscribing staging areas, notify user applications that messages are ready for consumption, and dequeue messages at the destination. It supports all the standard features of message queuing systems including multi-consumer queues, publish and subscribe, content-based routing,

Internet propagation, transformations, and gateways to other messaging subsystems.

Oracle Streams Advanced Queuing is useful for building message queuing applications. After an administrator creates the staging area, applications can explicitly enqueue messages into the source database. Applications can use strongly typed queues, as has been characteristic of the queues in previous releases of the database, or can use the Streams AnyData queues. AnyData queues can hold messages of different types, so it is possible to enqueue different types of messages into a single staging area. Enterprise Manager Database Control provides a simple management tool for Oracle Streams Advanced Queuing.



The DBMS_STREAMS_ADM package includes procedures that help users easily configure common messaging scenarios. Using these customized procedures, configuration of the rules and rule set that govern the queue contents is simplified. In addition, new procedures in the DBMS_STREAMS_MESSAGING package provide a simple interface for enqueueing and dequeuing messages from the staging area. These procedures simplify the application development process.

For example, assume that you have created a table called ORDER_EVENT that has the following format:

```
CREATE OR REPLACE TYPE order_event_type AS OBJECT (
```

```

        order_number    NUMBER(10),
        part_number     VARCHAR2(10),
        status          VARCHAR2(10),
        delivery_date   DATE
    );
/

```

The following command sets up a rule to apply only those orders that have been entered with appropriate part numbers:

```

BEGIN
  DBMS_STREAMS_ADM.ADD_MESSAGE_RULE(
    message_type    => 'DEMO.ORDER_EVENT_TYP',
    rule_condition  => ':msg.part_number is not null',
    streams_type    => 'APPLY',
    streams_name    => 'apply_w_msghdlr',
    queue_name      => 'strmadmin.streams_queue');
END;
/

```

Event Management with Oracle Streams

Business events are valuable communications between applications or organizations (internal and external). They require the highest degree of reliability, integrity, and security. Oracle Streams message queues are stored in the Oracle database, and are retained after consumption. This allows Oracle Streams Advanced Queuing to be used as a business event management system. Advanced Queuing stores all messages in the database in a transactional manner, where they can be automatically audited and tracked. This audit trail can be used to extract intelligence about the business operations.

The key to using Streams as a business event management system is retention in the staging area. As with message queuing, a staging area is created to stage business events that are explicitly enqueued by applications. Business events can be consumed locally via an application explicitly dequeuing the event, or they can be propagated to other staging areas and then consumed by applications. All staging areas will retain messages, even after they are consumed, providing an audit trail critical for managing business events.

Event Notification with Oracle Streams

Information overload makes it difficult to find the data that is important. A new class of applications is used to monitor data and notify interested parties of relevant data. Oracle Streams Advanced Queuing supports event notification, where changes to an underlying database are captured and then sent to other applications or devices that subscribed to that change event. Streams includes a scalable and sophisticated rules engine that can evaluate captured changes and forward them on to the appropriate subscribers.

Configuring Streams for Event Notification requires setting up the capture process to capture the events of interest, be they DML, DDL, or explicitly enqueued messages. These events are published and subsequently consumed by subscribing

applications. Rules are used to restrict consumption to events of interest. The consuming application can then take action, which may include an email notification, or passing the message to a wireless gateway for transmission to a cell phone or pager.

Using Oracle Streams in Heterogeneous Environments


Oracle Streams is an open information sharing solution, supporting heterogeneous replication between Oracle and non-Oracle systems. Using a database gateway, changes initiated at Oracle databases can be applied to non-Oracle platforms. Only DML changes can be replicated to non-Oracle platforms.

In order to meet the needs of integrating with legacy applications, the messaging gateway is provided to automatically propagate messages between Oracle queues and IBM Websphere MQ and TIBCO Rendezvous queues.

Oracle to Non-Oracle Capture/Apply

To implement capture and apply of DML changes from an Oracle source to a non-Oracle destination, an Oracle system functions as a proxy and executes the apply engine that would normally be done at an Oracle destination site. The Oracle system then communicates with the non-Oracle system via a database gateway. The changes are dequeued in an Oracle database itself and a local Oracle apply process applies the changes to a non-Oracle system across a network connection through a database gateway specific to the non-Oracle database.

The following figure illustrates how heterogeneous propagation of DML statements work.



Non-Oracle to Oracle Capture/Apply

Users who want to propagate changes from a non-Oracle database to an Oracle database must write an application to capture the changes made to the non-Oracle database. The application can capture the changes by reading from transaction logs or by using triggers. The application is then responsible for assembling and ordering these changes into transactions, converting them into an LCR format and publishing them into the target Oracle database staging area.



Integrating with Legacy Applications

The message gateway functionality of Oracle integrates Oracle database applications with other message queuing systems such as IBM Websphere MQ (formerly called MQ Series) and TIBCO Rendezvous. Since many legacy applications on mainframes communicate with Websphere MQ, there is a need for integrating these applications into an Oracle environment. The message gateway makes non-Oracle message queues appear as if they were Oracle queues, and automatically propagates messages between Oracle queues and IBM Websphere MQ or TIBCO Rendezvous queues.

Using Oracle Streams for Provisioning Data in a Grid Environment

Oracle Real Application Clusters (RAC) is a key component of Oracle's grid solutions. RAC distributes a single database across multiple servers in a grid, enabling on-demand provisioning of CPU resources for a database workload. Oracle Database 11g enhances the Streams support for RAC environments with the ability to mine the active online log files for DML and DDL, realizing the advantages of low latency capture, propagation, and apply along with the many other benefits of RAC. Streams processes run from a single instance of a RAC cluster, identified as the primary instance for the Streams queue and processes. An alternative instance can be specified as the secondary instance in case the primary becomes unavailable. In the event the primary instance fails, the Streams processes will automatically be restarted on the secondary instance. Upon the recovery of the primary instance, the Streams processes will return to the primary instance. Thus,

Streams will exhibit the same unbreakable behavior as other database features when used in a RAC environment, without need for DBA intervention

Many of the same features of Oracle Streams that are useful for information sharing in a distributed environment are equally useful for data provisioning in a grid environment. For example, suppose you have a production database, and you need to perform some analysis. In a grid environment, you could simply choose to add nodes to your production RAC database to provide the additional CPU capacity needed for the analysis. However, it's possible that may not be feasible. Maybe there are departmental restrictions that prohibit running the analysis against the production database. Maybe you are not running RAC, or maybe you do not have any more suitable nodes on the proper network and SAN to allocate—but you do have nodes on another network or SAN. Using the data provisioning features of Oracle Database 11g, you can transport the tablespaces containing the information you plan to analyze, instantiate a Streams replica of the data in on another system, and perform the analysis on that system.

Leveraging the Oracle Transportable Tablespaces and Oracle Streams, the Oracle database offers an efficient way for migrating applications to the grid. With a single command, the database administrator can identify a set of tablespaces from one database, ship the tablespace set to another database even if the second database is on a different operating system or platform, and plug this set into the second database. During this time, both the source and destination databases are open and available for any user activity. Meanwhile, Oracle Streams has begun to capture any changes from the source database that occur during the tablespace copy to the replica database. After the tablespace set is available at the replica database, the replica database is synchronized by Oracle Streams with the changes from the source database. All of this is done with a single command with no downtime required.

Minimizing Downtime during Upgrades

Using the features of Oracle Streams, users can perform the following database maintenance operations with little or no downtime:

- Upgrading the version of the database from Oracle9i Release 2 or later to a later release, such as Oracle11g
- Migrating the database to a different operating system or character set
- Migrating to a new version of an application
- Applying an Oracle software patch

Ordinarily, these operations might require considerable database downtime to complete. Using Oracle Streams, however, users simply configure a replica of the source database, which will be used to perform maintenance operations, while the original database remains fully operational.


The following general steps outline the operations that must be completed in order to perform a database maintenance operation while remaining online:

1. Create a new, empty database.
2. Use Oracle Streams to configure a replication environment, where this new database is the destination database, and the original database is the source database. After instantiating (populating) the destination database (with Export/Import or RMAN), Oracle Streams will automatically ensure that DML and DDL changes occurring at the source database are ultimately applied at the destination database.
3. Perform the desired maintenance operations on the destination database. During this time, the original database remains available online with Streams configured, but not started.
4. After completing the maintenance operations, use Oracle Streams to apply changes made at the source database to the destination database.
5. Once these changes have been successfully applied at the destination, take the source database offline and make the destination database available for applications and users.

Example Streams Configuration

Oracle Streams is designed for maximum flexibility to satisfy customer information sharing requirements in a variety of markets. For example, customers can use Oracle Streams to create Messaging and Event Management, Event Notification, Replication, Data Warehouse Loading, and Data Provisioning solutions. Of course, all customers can utilize the full power of Oracle Streams, and create configurations that seemingly span multiple markets, enabling new classes of applications. In addition, all deployments and their associated meta-data are compatible. For example, a replication installation can easily be extended to load a data warehouse or enable bi-directional replication—a complete reconfiguration is not required.

The flexibility of Oracle Streams is demonstrated in the following example. A corporate IT organization is charged with maintaining data availability around the clock for a critical global application. Their strategy is to maintain three geographically distinct databases with all the requisite data for this high profile application. Additional requirements include a reporting database containing the most current information for the analysts in the company headquarters office in New York to perform ad-hoc querying as well as a disaster recovery database separately maintained from their New York office. A final requirement is to share data with existing applications that are hosted on a Sybase database.



Oracle Streams Usage Example

Oracle Streams is used to replicate data in an N-way configuration consisting of three regional sites: New York, London, and Tokyo. At each of these sites, Streams log-based capture will capture any changes that occur for subscribed tables in each local region, and will stage them locally in the queue. All changes captured in each region will be then forwarded to each of the other region's databases with the goal that all changes made at each site will be reflected at every other site, providing complete data for the subscribed objects throughout the world.

Since the updates will be automatically applied when received at each regional database, the Oracle Streams default apply engine is used to apply the changes. As updates are applied, Oracle Streams checks for conflicts, and resolves any conflicts that are detected. Streams can also be used to exchange data for particular tables with non-Oracle databases. Utilizing the Oracle Database Gateway for Sybase, the Streams apply engine will apply the changes to a Sybase database using the same mechanisms as it does for Oracle databases.

The databases for reporting and disaster recovery are hosted from the New York database site. The reporting database is a fully functional Oracle database that has a read-only copy of the relevant application tables. The reporting site will not be configured to capture changes on these application tables. Streams will impose no restrictions on the configuration or usage of this reporting database.

CONCLUSION

Oracle Streams simplifies sharing data between databases and database clusters, with its revolutionary concept of unifying message queuing and data replication capabilities. With Streams, the basic elements of capture, staging, and consumption in combination with user-defined transformations and heterogeneous interoperation produce more sophisticated replication, messaging, and notification solutions than ever before possible.

This resilient technology provides the configuration flexibility required to handle the real world situations encountered in modern global corporations. Streams technology is the premier feature for sharing and provisioning data in complex distributed database and grid computing environments.



Oracle Database 11g: Oracle Streams
July 2007
Author: Patricia McElroy, Maria Pratt
Contributing Authors:

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Copyright © 2007, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.