Oracle Infrastructure Systems Development
Cross-Platform Database Migration using Oracle Streams

"Oracle Streams let us complete our platform migration from single instance Solaris SPARC to Linux RAC with virtually no down time."

Clark Elms
Senior Director, Infrastructure Tools Development,
Oracle Corporation

Corporate Profile

- Internal Oracle development group

Oracle Technology Used

- Oracle Database 10*g* Release 2 (10.2.0.2 and 10.2.0.3)
- Oracle Streams replication
- Solaris Single Instance database migration to Oracle Enterprise Linux RAC

OVERVIEW

Infrastructure Systems Development (ISD) is one of Oracle's internal groups that cater to internal as well as external applications. The ISD/ARU database is used by several, critical applications requiring 24x7 availability, including Oracle E-Delivery and the Patches and Updates interface in Metalink, as well as the following internal applications:

- Autobuild
- Automated Release Updates (ARU)
- Environment Management System (EMS)
- Patch Tracking System (PTS)
- UI Review System (URS)

These applications have nearly 1.5 million registered users. Each week, approximately 700,000 of these users initiate over six million requests and download nearly 100,000 patches. As they watched these numbers continue to grow, ISD, like many organizations, began looking for a way to reduce overall costs while managing their ever-increasing database load.

This case study describes the successful implementation of Oracle Streams [1] replication technology by the ISD team to achieve their goal of migrating a 230 GByte single instance Oracle Database 10g Release 2 (10.2.0.2) database running on Solaris SPARC to Oracle Database 10g Release 2 (10.2.0.3) using Real Application Clusters (RAC) running on Oracle Enterprise Linux. ISD used a phased migration approach to allow for real time testing using the production load and minimal downtime during the migration.

ISD REQUIREMENTS

The ISD/ARU database was a single instance database running on Solaris SPARC. ISD determined that they could achieve significant cost savings, as well as support a rapidly growing database load by migrating to a RAC database on Linux.

Because of the critical nature of both the external, as well as the internal, applications supported by this database, ISD preferred to take the following phased approach:

> Phase 1 – ISD created the Linux target RAC database and set up Streams unidirectional replication between the Solaris primary database and the Linux target RAC database. All applications continued to connect to the primary Solaris database. Changes made to the Solaris primary database were replicated to the Linux target RAC database using Streams. The target database was available during planned, as well as unplanned downtimes, which allowed the ISD team to experience the real world performance of the new system prior to completing the switchover. Having this cross-platform database available during outages increased the confidence of the ISD team in the stability of the proposed configuration, which met their defined success criteria.
>
> Phase 2 – The team moved a limited number of critical applications to connect to the target database instead of the primary database.
>
> Phase 3 – After ISD successfully completed all testing of the new environment; they moved the remaining applications to the target RAC database and were able to repurpose the original Solaris SPARC platform.

ORACLE STREAMS REPLICATION

Oracle Streams captures DML and DDL changes made to database objects and replicates those changes to one or more destination databases. The destination databases allow DML and DDL changes to the same database objects, and these changes can also be propagated to other databases in the environment should the user desire. Streams can be configured to propagate changes between one or more databases bi-directionally. The tables for which data is shared do not need to be identical copies at all databases; both the structure and the contents of these tables can differ at different databases.

There are three steps to the Streams replication process – Capture, Propagation, and Apply. While the basics of replication are summarized below, please note that the Streams documentation includes many ways to customize how and where these processes are executed in order to address a wide range of replication requirements. Please refer to Oracle Streams Concepts and Administration [2] for a complete discussion of capabilities.

- **Capture:** A Streams capture process mines the redo log to create one or more logical change records (LCRs) and queues them to a Capture Queue. An LCR is a message with a specific format that describes a database change.

- **Propagation:** Streams propagation propagates the staged LCR to another queue residing in the destination database where apply will occur.
- **Apply:** Once the LCR has reached the destination database, a Streams apply process consumes the change by applying the LCR to the shared database object.

ISD Streams Configuration

ISD needed a way to support their rapid growth while reducing costs. They chose Oracle Streams for its ability to support a phased migration from a Solaris single instance database to a Linux RAC database.

ISD needed their applications to remain available throughout the migration. They needed to implement a solution where they could replicate between the two databases and the applications could be moved over a few at a time. This gradual transition allowed them to fine tune the new environment as necessary, and to address any needs as they arose. It was not feasible to have the database unavailable for the estimated 20 hours that a traditional export/import of this 230 Gbyte database would take. Instead, by using Oracle Streams, ISD took only 15 – 20 minutes of downtime to verify their new environment before completing the switchover.

Downstream Capture

ISD did not want any load from Oracle Streams replication to affect their production environment, so they chose to take advantage of downstream capture. Downstream capture means that a capture process runs on a database other than the source database. Oracle Streams supports both real-time and archived-log downstream capture. With downstream capture, the alternative database must be on a similar platform type and operating system.

The following diagram illustrates the Oracle Streams replication environment that was in place during the migration, involving the primary ARU database, downstream ARUPRIM database and target ARURAC database.

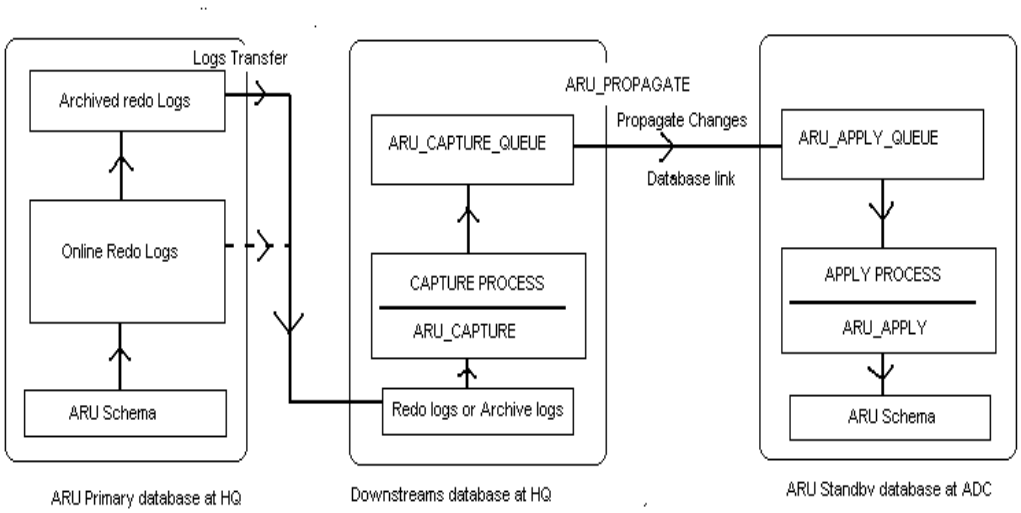

System Configuration

Primary

- Single node Solaris Enterprise 6800 server

  Sun Fire 6800
  20 CPU's
  40 GB Memory
  SunOS 5.9 Generic_117171-17 (64-bit)

- Oracle Database 10g Release 2 (10.2.0.2 Solaris SPARC 64-bit)

Downstream

- Single node Solaris Enterprise 3800 server

  Sun Fire 3800
  6 CPU's
  10 GB Memory
  SunOS 5.9 Generic_117171-17 (64-bit)

- Oracle Database 10g Release 2 (10.2.0.2 Solaris SPARC 64-bit)

Target

- 5 node Dell PowerEdge 1950q

  16 GB Memory
  2 Quad Core CPU's (Intel(R) Xeon(R) L5320 1.86 GHz)
  Oracle Enterprise Linux 4 Upgrade 5 (32-bit)
  Kernel - 2.6.9-55.0.0.0.2.ELhugemem

- Oracle Database 10g Release 2 (10.2.0.3 Linux x86 32-bit)

The archive logs (average redo generation of 5MB/minute to a maximum of 100MB/minute during peak load in the database) from the primary database were transferred to the downstream database. A capture process at the downstream database captures all the changes in the form of logical change records by mining the archive logs received from the primary database and enqueues them to a queue, which the propagation job propagates to the secondary site using a database link. These LCRs are then applied on the target site by an apply process.

The target site was always writable throughout the apply process. Access to the target database was controlled by routing application connections to the proper database. For more information about configuring replication, consult the Oracle Streams Replication Administrators Guide [3].

Rule-based Replication

Oracle Streams uses rule-based replication. A rule is a database object that enables a client to perform an action when an event occurs and a condition is satisfied. Each rule is specified as a condition that is similar to the condition in the `WHERE` clause of a SQL query. Users can group related rules together into rule sets. Rule sets are evaluated by a rules engine, which is a built-in part of Oracle.

Rules determine what Oracle Streams clients replicate. Users can configure rules for each Oracle Streams client independently, and the rules for different Oracle Streams clients do not need to match.

Oracle Streams supports both positive as well as negative rule sets. As a part of their application design, ISD chose not to replicate selected application related tables, as well as queues and materialized views.

MONITORING STREAMS REPLICATION

ISD used a "heartbeat table", to enable the administrator to determine the status of the complete configuration in a single view. Inside the heartbeat table, each row represents information about a location, using ID, Database Name (or Global Name) and a timestamp. The target database periodically updates the heartbeat table with timestamps for the most recent data replicated. The table also includes other attributes for apply, propagation, capture, queuing, and apply error information. The Streams Replication Best Practices section of the Oracle Streams Replication Administrators Guide [3] contains details on creating a heartbeat table as well as other recommended administrative practices in a Streams environment.

CONCLUSION

Oracle's Infrastructure Systems Development (ISD) Group needed a method of keeping up with their ever-increasing database load, as well as controlling costs. They needed a way of incrementally migrating from their Solaris single instance database to a more performant Linux RAC database. By using Oracle Streams replication, they were able to perform a phased migration between the two platforms with minimal downtime.

REFERENCES

1. Oracle Streams
   http://www.oracle.com/technology/products/dataint/index.html
2. Oracle Streams Concepts and Administration
   http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14229/strms_over.htm
3. Oracle Streams Replication Administrators Guide
   http://download-west.oracle.com/docs/cd/B19306_01/server.102/b14228/toc.htm

ISD and Oracle Streams
March 2009

Authors: Jaideep Paidimarri, Senior Manager, ISD Development, Oracle Corporation
Rahul Saraf, Principal Database Administrator, ISD Development, Oracle India
Oracle HA Product Management

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