

Log Stream Staging

The Why and The How

Author: Gwyn Jones
Senior Solution Specialist

Date: 15th October 2019

Version: 2.1

Introduction **The Why**

Log Stream Staging (LSS) has been developed by Attunity to provide for scenarios where there is a requirement to write to multiple targets WITHOUT the necessity of these multiple tasks all reading the same source CDC transaction log.

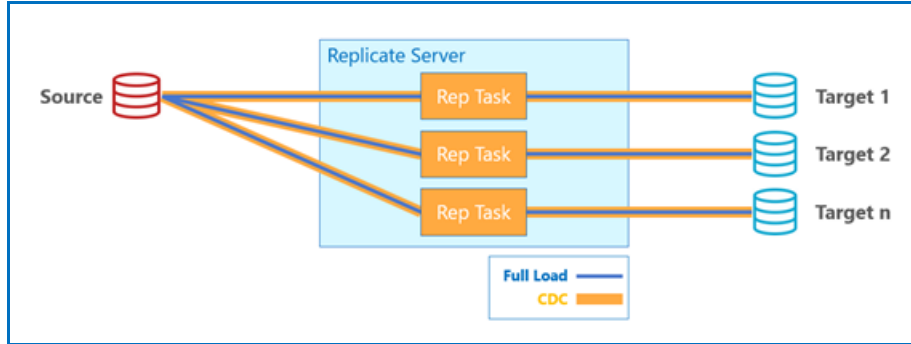
This document is intended to itemise and demonstrate the steps needed to utilise this solution.

Contents

Introduction The Why	2
Contents.....	2
High-Level Configuration	3
Standard Singular Source to Multiple Target Scenario.....	3
Log Stream Staging Scenario – Single On-Going CDC Task	3
Detailed Configuration The How	4
Steps.....	4
Step 1 Standard Source Endpoint	5
Step 2 Log Stream Staging (LSS) Target Endpoint	6
Step 3 CDC Only Task – Source to LSS for all required tables from this endpoint.....	7
Step 4 Run CDC Only Task	8
Step 5 Duplicate Original Source.....	9
Step 6 Standard Target Endpoint	10
Step 7 Standard Task – LSS to Target	11
Step 8 Repeat Standard Task <i>as many as required</i>	14
Step 9 How the Storage looks	15
In Summary	16

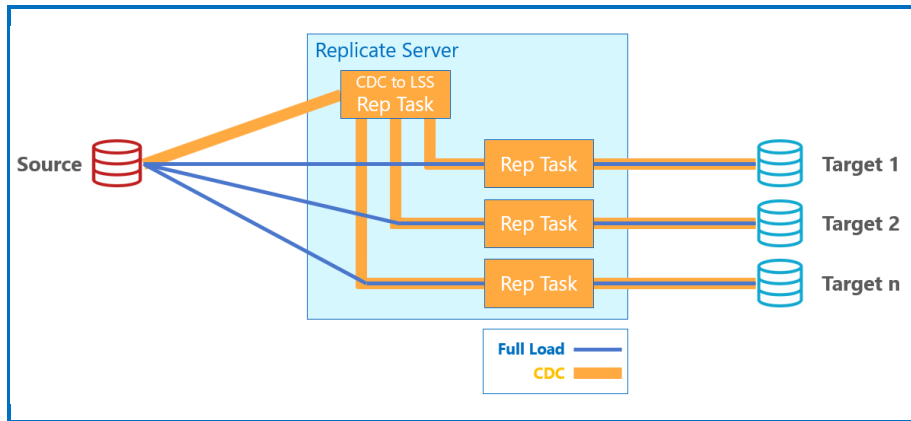
High-Level Configuration

Standard Singular Source to Multiple Target Scenario



Rep Task 1	Full Load from Source <i>when required</i>	CDC from Source
Rep Task 2	Full Load from Source <i>when required</i>	CDC from Source
Rep Task 3	Full Load from Source <i>when required</i>	CDC from Source
On Source Server	n x Full Load <i>when required</i>	n x CDC

Log Stream Staging Scenario – Single On-Going CDC Task



CDC LSS Task	<i>No Full Load</i>	CDC from Source
Rep Task 1	Full Load from Source <i>when required</i>	CDC from LSS
Rep Task 2	Full Load from Source <i>when required</i>	CDC from LSS
Rep Task 3	Full Load from Source <i>when required</i>	CDC from LSS
On Source Server	n x Full Load <i>when required</i>	1 x CDC <i>just the one</i>

Steps

- ❖ Standard Source Endpoint
- ❖ Log Stream Staging (LSS) Target Endpoint
- ❖ CDC Only Task – Source to LSS for all required tables from this endpoint
- ❖ Run CDC Only Task
- ❖ Duplicate Original Source
 - Full Load details for the original source
 - CDC details for the previous LSS Task
- ❖ Standard Target Endpoint
- ❖ Standard Task – LSS Source to Target
 - Full Load reads from the Original Source
 - CDC reads from the LSS Task

- ❖ Repeat Standard Task *as many as required*

The steps below are using the example of "*Oracle*" to "*SQL Server*"

with screen shots to demonstrate the steps

Step 1 Standard Source Endpoint

No difference to normal

The screenshot shows the 'Manage Endpoint Connections' dialog box. The 'Sources' tab is active, and a list on the left contains 'demo' and 'demo ORA Source'. The 'demo' entry is selected. The 'General' tab is active, showing the following configuration:

- Name: demo ORA Source
- Description: Standard Oracle Source
- Role: Source Target
- Type: Oracle
- Read changes from log stream staging folder
- Log stream staging task: (empty dropdown)
- Connection string: localhost:1521/xe
- Username: system
- Password: (masked with dots)

Buttons at the bottom include 'Test Connection' (with a green checkmark), 'Save', 'Cancel', and 'Close'.

Step 2 Log Stream Staging (LSS) Target Endpoint

Main points here are:

- Storage Path for the events on the Replicate server
 - There will be a task specific folder and sub-folders under this location
- Batching of events
- Retention of events
- Rollover of events

The screenshot shows the 'Manage Endpoint Connections' dialog box with the 'General' tab selected. The 'Storage path' field is highlighted with a red box and contains the text 'C:\AttunityLogStream'. Other fields include 'Name' (demo LSS Target), 'Description' (empty), 'Role' (Target selected), and 'Type' (Log Stream). There are also buttons for 'Batch', 'Retention', and 'Rollover'.

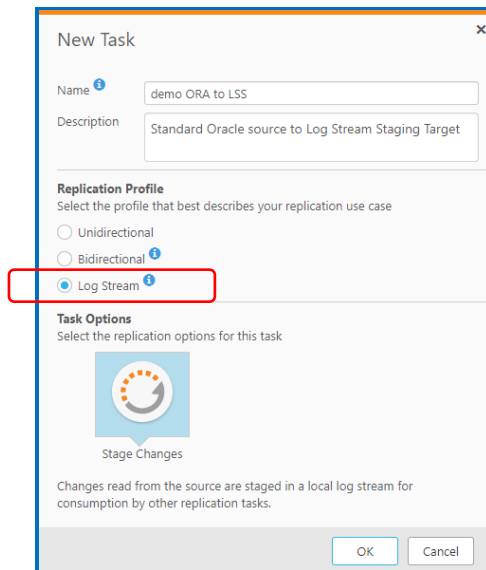
The screenshot shows the 'Batch' configuration section. It has a checked checkbox for 'Batch'. Below it are two settings: 'Apply batch after (seconds)' with a value of 5, and 'Apply batch when number of changes reaches' with a value of 10000.

The screenshot shows the 'Retention' configuration section. It has a checked checkbox for 'Retention'. Below it are two settings: 'Delete files after (hours)' with a checked checkbox and a value of 48, and 'Delete oldest files when the total size of all files exceeds (MB)' with a value of 100000.

The screenshot shows the 'Rollover' configuration section. It has a checked checkbox for 'Rollover'. Below it are two settings: 'Roll over file after (minutes)' with a value of 120, and 'Roll over files larger than (MB)' with a value of 500.

Step 3 CDC Only Task – Source to LSS for all required tables from this endpoint

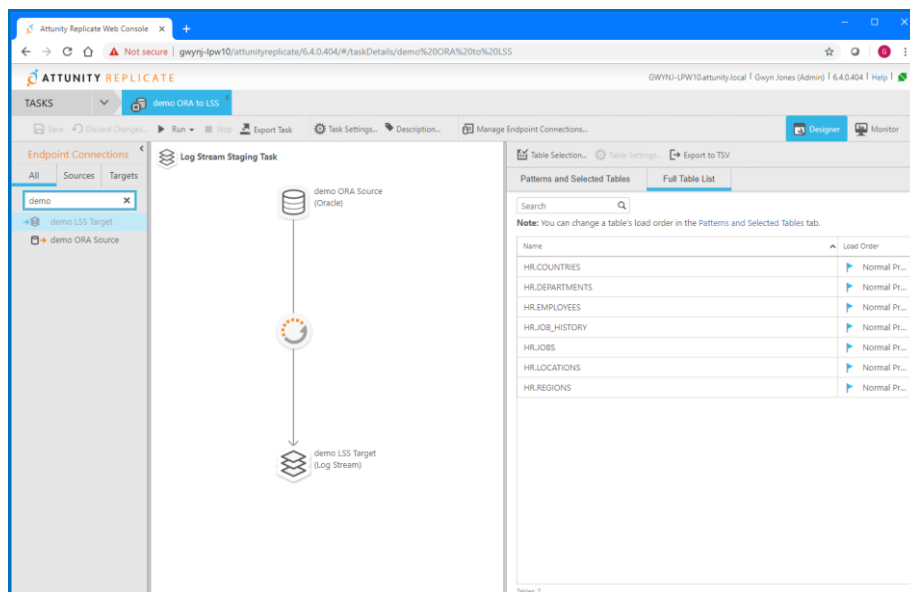
Main point here is to select Log Stream as the Replication Profile



Select the standard source endpoint

Select the new Log Stream Staging target endpoint

Select ALL the required tables required from this source



❖ No Filters, No Table Transformations, No Global Transformations

Step 4 Run CDC Only Task

Important:

This step does need to be performed before starting the rest of the solution.

This task not only contains/manages the CDC events but also the necessary table metadata ... a further saving of source resources

The screenshot displays the Attunity Replicate Web Console interface for a task named 'demo ORA to LSS'. The dashboard shows the following metrics:

- Incoming Changes:** 0 Transactions
- Applied Changes:** A pie chart showing the distribution of Inserts, Updates, and Deletes.
- Apply Throughput:** rec/sec
- Apply Latency:** Now, 00:00:00

The 'Applied Changes Details' table is as follows:

Table Name	INSERT	UPDATE	DELETE	DDL	Total Applied	Data Errors	Last Modified
HRREGIONS	100	85	35	0	220	0	5:32 PM
HRLOCATIONS	0	0	0	0	0	0	5:32 PM
HRJOBS	0	0	0	0	0	0	5:32 PM
HRJOB_HISTORY	0	0	0	0	0	0	5:32 PM
HREMPLOYEES	0	0	0	0	0	0	5:32 PM
HRDEPARTMENTS	0	0	0	0	0	0	5:32 PM
HRCOUNTRIES	0	0	0	0	0	0	5:32 PM

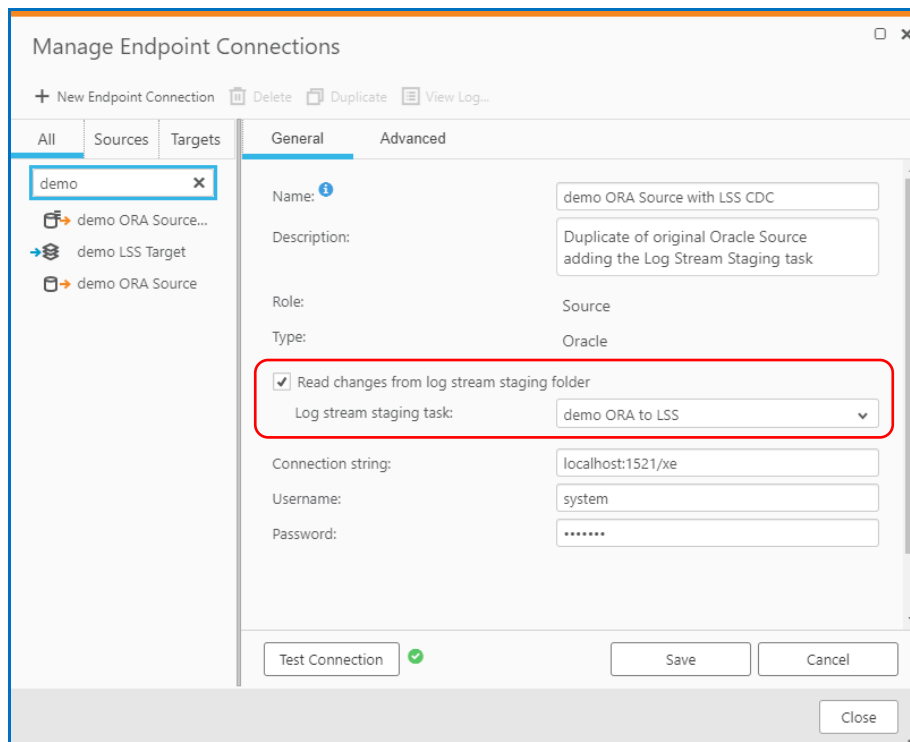
The 'Messages' section on the right includes a search bar and a table for log messages, with columns for Date and Time, Level, and Message.

Step 5 Duplicate Original Source

Select your previous source and DUPLICATE

Add the configuration for "Read changes from Log Stream Staging task"

Note: the Password requires to be re-entered

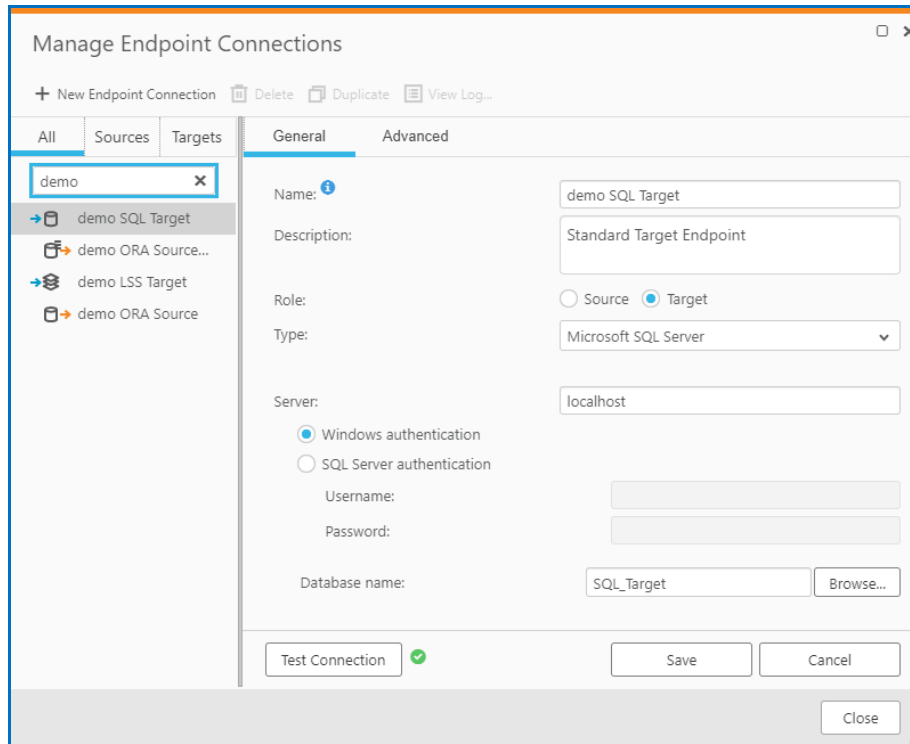


This source endpoint now has configuration to:

- Full Load details for the original source
- CDC details for the previous LSS Task

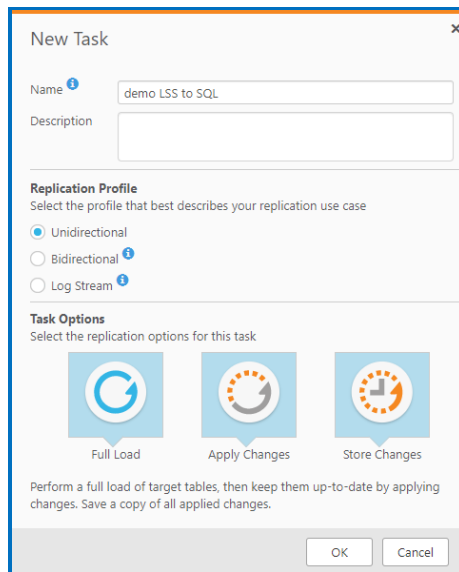
Step 6 Standard Target Endpoint

No difference to normal



Step 7 Standard Task – LSS to Target

No difference to normal



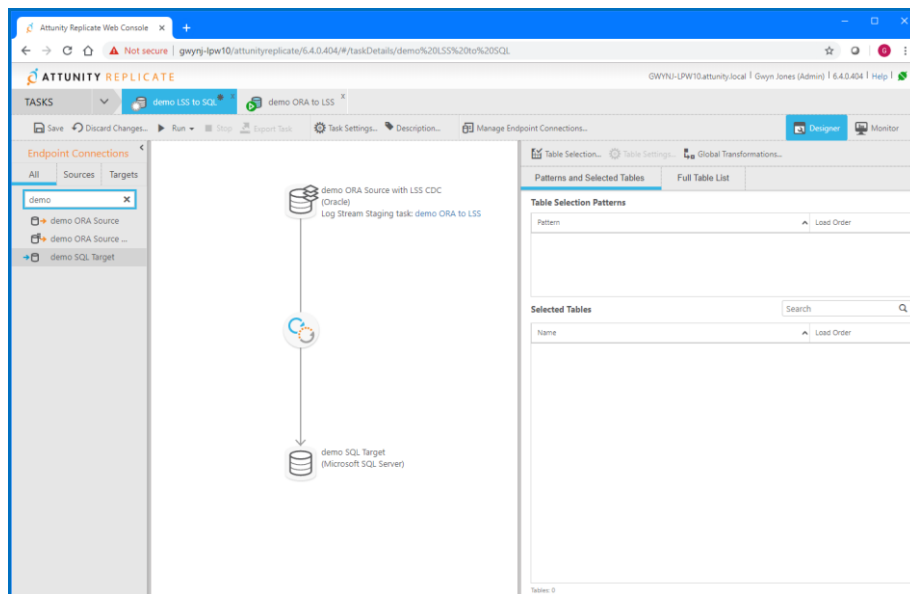
The 'New Task' dialog box is shown with the following fields and options:

- Name:** demo LSS to SQL
- Description:** (empty)
- Replication Profile:** Unidirectional (selected), Bidirectional, Log Stream
- Task Options:** Full Load (selected), Apply Changes, Store Changes

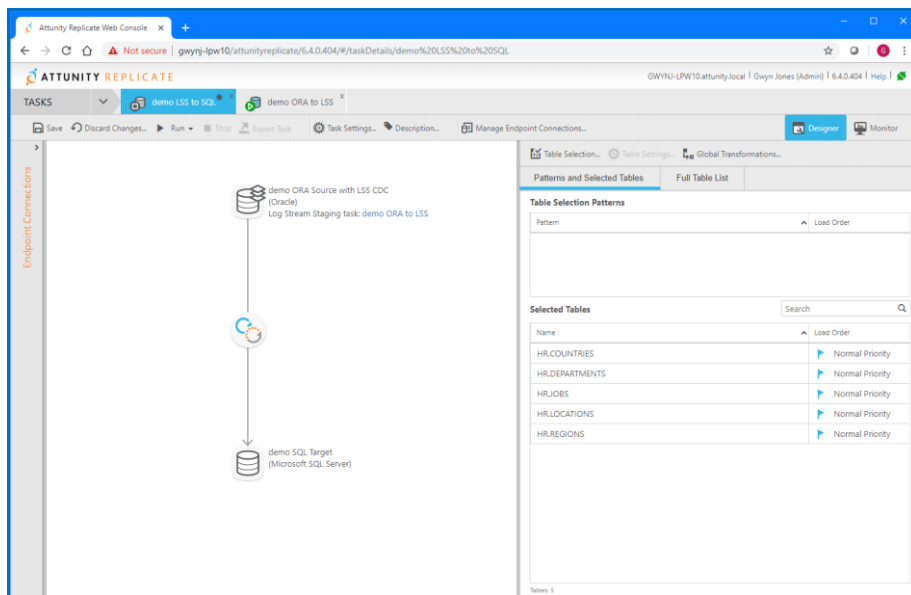
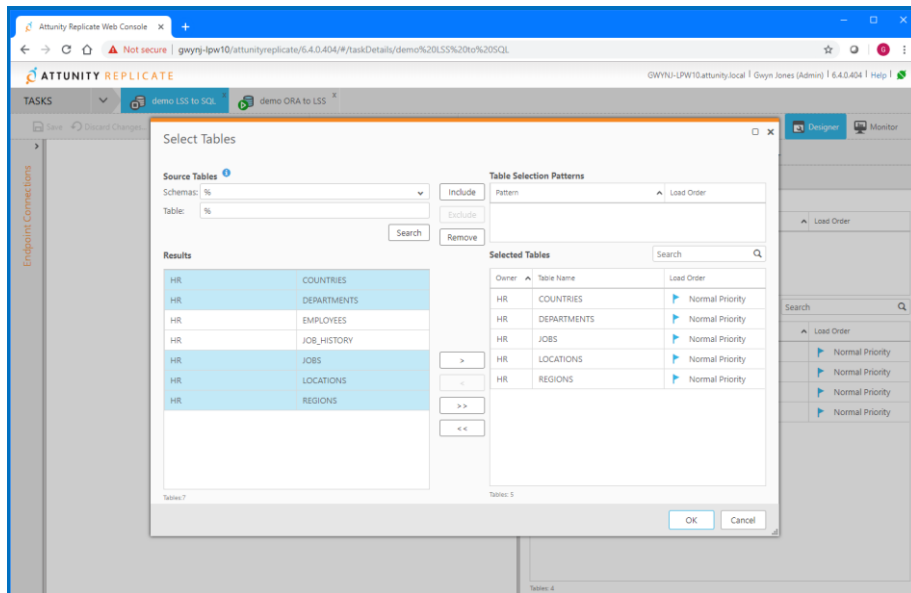
Below the options, a note states: "Perform a full load of target tables, then keep them up-to-date by applying changes. Save a copy of all applied changes." Buttons for 'OK' and 'Cancel' are at the bottom right.

Select the new LSS source

Select the standard target

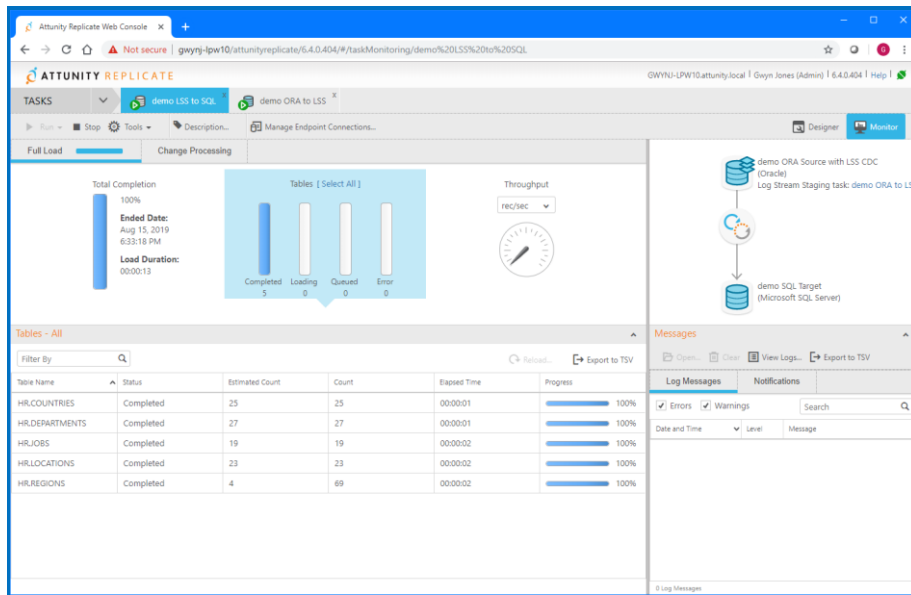


- ❖ ONLY tables included within the Source to LSS task will be available

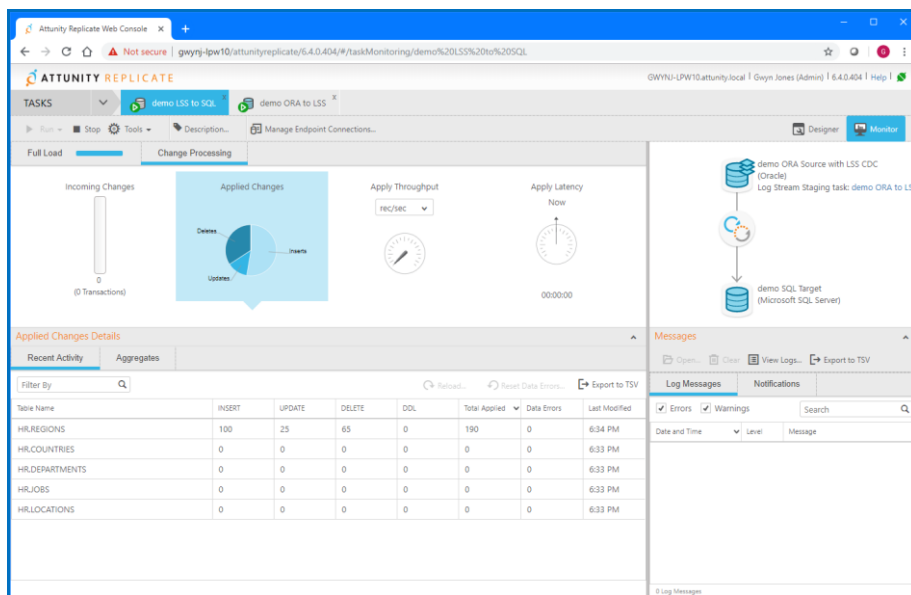


- ❖ Filters, Table Transformations, Global Transformations can be applied

❖ Full Load as normal from Source



❖ CDC as normal but from the Log Stream Staging



Step 8 Repeat Standard Task *as many as required*

Now repeat the task build for as many targets as required

Step 9 How the Storage looks

```
C:\AttunityLogStream> dir /s

Volume in drive C is OS
Volume Serial Number is 72E9-7E91

Dir of C:\AttunityLogStream

16/08/2019  09:56    <DIR>          .
16/08/2019  09:56    <DIR>          ..
15/08/2019  17:26    <DIR>          demo ORA to LSS
                0 File(s)            0 bytes

Dir of C:\AttunityLogStream\demo ORA to LSS

15/08/2019  17:26    <DIR>          .
15/08/2019  17:26    <DIR>          ..
15/08/2019  17:26    <DIR>          LOG_STREAM
                0 File(s)            0 bytes

Dir of C:\AttunityLogStream\demo ORA to LSS\LOG_STREAM

15/08/2019  17:26    <DIR>          .
15/08/2019  17:26    <DIR>          ..
15/08/2019  17:32    <DIR>          audit_service
                0 File(s)            0 bytes

Dir of C:\AttunityLogStream\demo ORA to LSS\LOG_STREAM\audit_service

15/08/2019  17:32    <DIR>          .
15/08/2019  17:32    <DIR>          ..
15/08/2019  17:26    <DIR>          20190815172613781936
16/08/2019  08:27    <DIR>          20190815173225782012
                0 File(s)            0 bytes

Dir of C:\AttunityLogStream\demo ORA to LSS\LOG_STREAM\audit_service\20190815172613781936

15/08/2019  17:26    <DIR>          .
15/08/2019  17:26    <DIR>          ..
15/08/2019  17:32    <DIR>          5,464 1
15/08/2019  17:26    <DIR>          61,440 metastore.sqlite
                2 File(s)            66,904 bytes

Dir of C:\AttunityLogStream\demo ORA to LSS\LOG_STREAM\audit_service\20190815173225782012

16/08/2019  08:27    <DIR>          .
16/08/2019  08:27    <DIR>          ..
15/08/2019  20:25    <DIR>          79,717 1
15/08/2019  22:26    <DIR>          49,952 2
16/08/2019  00:26    <DIR>          49,855 3
16/08/2019  02:26    <DIR>          49,926 4
16/08/2019  04:27    <DIR>          49,907 5
16/08/2019  06:27    <DIR>          50,014 6
16/08/2019  08:27    <DIR>          49,781 7
16/08/2019  09:57    <DIR>          37,216 8
15/08/2019  17:32    <DIR>          61,440 metastore.sqlite
                9 File(s)            477,808 bytes

Total Files Listed:
    11 File(s)            544,712 bytes

C:\AttunityLogStream>
```

In Summary

- Log Stream Staging provides for a single reader of the source transaction log irrespective of how many variations of targets are required
- Log Stream Staging is only storing and managing CDC events for the selected tables
- Log Stream Staging has its own Batching, Retention, Rollover parameters
 - ❖ Alleviates any short log retention policies on source databases

- TEST

- *Full documentation in the Attunity Replicate User Guide and Reference*