Oracle Identity Sync

Resolving Identity Column Drift after Data Pump Operations

Technical Deep Dive: "START WITH LIMIT VALUE"

The Identity Gap

Understanding how metadata and actual data drift apart during large-scale database migrations.

What are Identity Columns?

Internal Sequences

Introduced in 12c, Identity Columns use hidden internal sequences to manage PK values. **GENERATED ALWAYS** prevents manual insertion, ensuring data integrity.

Metadata Storage

The `LAST_NUMBER` is stored in metadata. During Data Pump exports, this definition is captured as it existed at the start of the job.

The Data Pump Conflict

- **Export Phase:** Captures sequence metadata and table data.
- **Import Phase:** Recreates the sequence and appends rows.
- **The Risk:** If target tables already have rows, or if metadata lags, the sequence resets to a lower value.

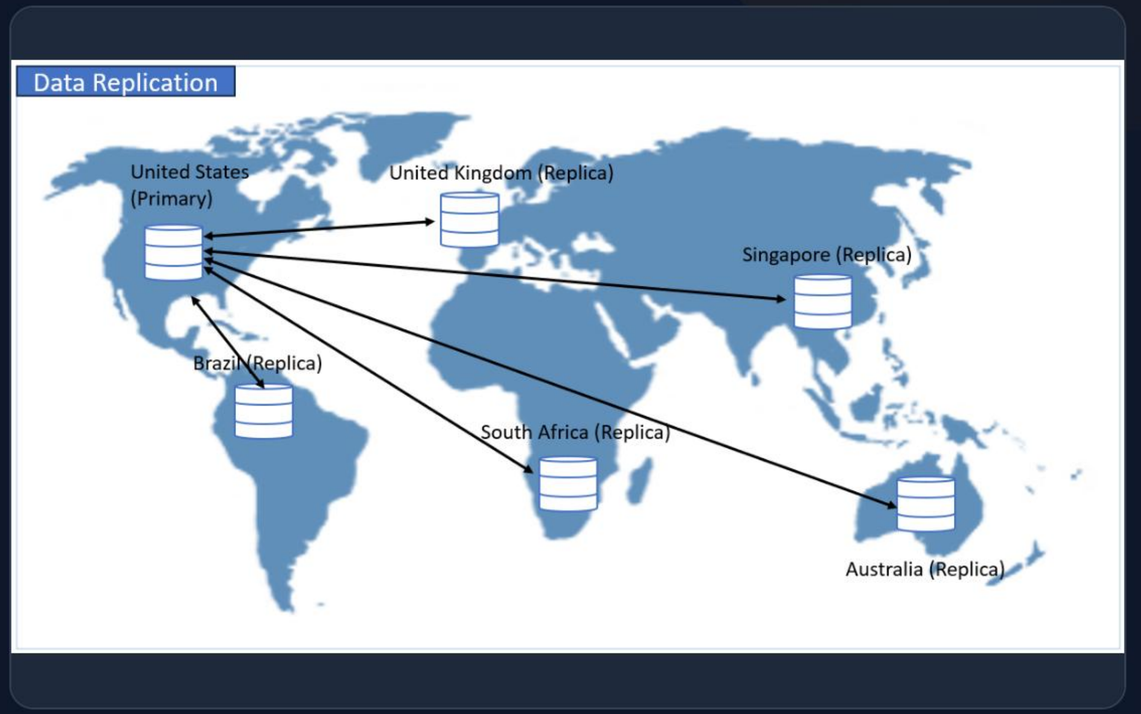

Metadata Drift

During parallel Data Pump operations, the sequence value captured in metadata may not reflect the highest value actually loaded into the table block.

This results in a **Drift**: the "Next Val" is lower than the "Current Max," leading to immediate primary key collisions upon application startup.

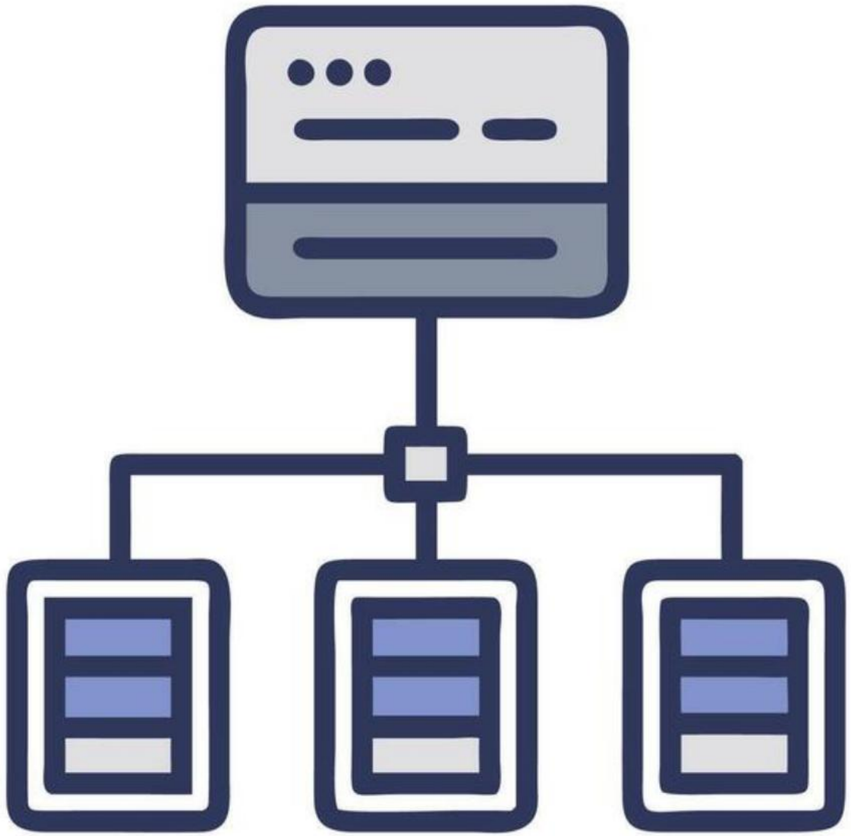

Critical Impact: ORA-00001

00001

UNIQUE CONSTRAINT VIOLATED

Broken Application Workflows

When the app attempts an INSERT, the identity column generates a duplicate value. This halts all transactional processing for that entity until the sequence is manually reset.

Diagnostic Workflow

- **Check Max Value:** Always run `SELECT max(id) FROM table;` to identify the high-water mark of the data.
- **Verify Metadata:** Inspect `user_tab_identity_cols` and `user_sequences` to find the `LAST_NUMBER`.
- **Auto-Resolution:** Use Oracle's built-in `LIMIT VALUE` clause to sync them instantly.

Manual vs. Automated Reset

METHOD	WORKFLOW COMPLEXITY	HUMAN ERROR RISK
Manual Alter	High (Calculate Max + 1)	Significant
Recreate Sequence	Extreme (Drop/Create/Grant)	Very High
LIMIT VALUE	Single SQL Command	Negligible

Why LIMIT VALUE Works

Atomic Scanning

Oracle performs an internal scan to find the current maximum value, ensuring the sequence starts exactly where the data ends.

Preserves Definitions

It modifies the existing identity metadata without changing grants, caching, or cycling properties.

Recovery Efficiency (Seconds)

Recovery time is reduced by ~98% using the native automated command.

Implementation Roadmap

Step 1
Complete Impdp Operation

Step 2
Run Audit Queries

Step 3
Apply LIMIT VALUE Fix

Step 4
Verify Insertion

```
ALTER TABLE employees MODIFY employee_id GENERATED ALWAYS AS IDENTITY (START WITH LIMIT VALUE);
```

Questions?

Thank you for your attention.

DBA Resources: docs.oracle.com

Topic: Identity Columns & Data Pump

Image Sources

https://skyvia.com/learn/images/data-replication.png

Source: skyvia.com

https://static.vecteezy.com/system/resources/previews/065/410/834/non_2x/database-architecture-icon-with-server-and-client-nodes-illustration-vector.jpg

Source: www.vecteezy.com