The science behind the report:

Streamline heterogeneous database environment management with Toad Data Studio

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report *Streamline heterogeneous database environment management with Toad Data Studio*.

We concluded our hands-on testing on January 13, 2025. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on December 18, 2024 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to *http://facts.pt/calculating-and-highlighting-wins*. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Results of our testing.

Steps	Methodology	Screen descriptions
Scenario 1		
1	Next to the DB connection, click the drop-down menu, and select the PostgreSQL30 menu option	Main Toad® UI screen
2	In the Status bar, click Diagram	Main Toad UI screen
3	Hold the Shift key, click to highlight all of tables, and drag into the diagram pane	Diagram pane
Scenario 2		
1	Next to Compare, click the drop-down menu, and select the Schema Compare for PostgreSQL menu option	Main Toad UI screen
2	On the Wizard welcome screen, click Next	Schema compare for PostgreSQL Wizard welcome screen
3	Select the source and target database, and click Next	Schema compare Source & Target database page
4	Click Next	Schema compare Object filter page
5	Click Compare	Schema compare script options page

Steps	Methodology	Screen descriptions
Scenario 3		
1	Next to Compare, click the drop-down menu, and select the Data Compare menu option	Main Toad UI screen
2	Enter a Comparison name, select the source database, select the target server & database, and click Next	New Comparison Wizard screen
3	Confirm all checkboxes are checked, and click Next	Map Source Objects To Target screen
4	Click Compare	Select Comparison Options screen
Scenario 4		
1	Click Automate	Main Toad UI screen
2	Under Toolbox, click Execute a Toad Schema Compare for PostgreSQL project file, and publish output file (report)	Toad Automation Script screen
3	In the Schema compare file field, click the far right icon to create a new comparison	Toad Automation Script screen
4	Enter a name, and click Save	Save as pop up window
5	Select the source database, the target connection and database, and click Confirm	Comparison options Connection screen
6	In the Summary report field, click the three dots	Toad Automation Script screen
7	Enter a file name, and click Open	Select results file name window
8	In the Detail report field, click the three dots	Toad Automation Script screen
9	Enter a file name, and click Open	Select results file name window
10	Click Run	Toad Automation Script screen
11	Enter a file name for the Automation file, and click Save	Save as pop up window
Scenario 5		
1	Click Import Export, and select Import Export Data	Main Toad UI screen
2	Click Datasource Table (Existing or new table), click the drop-down under Put data in, select the DB, and click Next	Define Data Import Export screen
3	Select the categories table, and click Next	Select source data object screen
4	Select A single new table, enter a table name, select the schema, and click Next	Select target screen
5	Click Add table	Add Import Source screen
6	Select the customers table, and click Next	Select source data object screen
7	Select A single new table, enter a table name, select the schema, and click Next	Select target screen
8	Click Add table	Add Import Source screen
9	Select the cust_hist table, and click Next	Select source data object screen
10	Select A single new table, enter a table name, select the schema, and click Next	Select target screen
11	Click Add table	Add Import Source screen
12	Select the inventory table, and click Next	Select source data object screen
13	Select A single new table, enter a table name, select the schema, and click Next	Select target screen

Steps	Methodology	Screen descriptions
14	Click Add table	Add Import Source screen
15	Select the orderlines table, and click Next	Select source data object screen
16	Select A single new table, enter a table name, select the schema, and click Next	Select target screen
17	Click Add table	Add Import Source screen
18	Select the orders table, and click Next	Select source data object screen
19	Select A single new table, enter a table name, select the schema, and click Next.	Select target screen
20	Click Add table	Add Import Source screen
21	Select the products table, and click Next	Select source data object screen
22	Select A single new table, enter a table name, select the schema, and click Next	Select target screen
23	Click Add table	Add Import Source screen
24	Select the reorder table, and click Next	Select source data object screen
25	Select A single new table, enter a table name, select the schema, and click Next	Select target screen
26	Click Next	Add Import Source screen
27	Click Finish	Save Settings as Template screen
28	Right-click the DB → New SQL Editor	Main Toad UI screen
29	Copy the SQL code to create indexes, stored procedures, and foreign keys, and press F5 to execute the script	SQL script editor

System configuration information

We set up a virtual environment with 4 VMs for testing. On the Toad Data Studio VM, we installed Windows Server 2022 and Toad Data Studio software (version 1.1.0.27777). For PostgreSQL testing, we created a Red Hat® Enterprise Linux® (RHEL) 9.5 VM, installed PostgreSQL 16.6, and cloned a duplicate PostgreSQL VM. On our Oracle® VM, we installed Oracle Linux Enterprise 8.10 and Oracle Database 21c. We also set up a Snowflake trial (version 8.44) for our cloud-based testing. To create a simple 5GB schema for the manageability scenarios, we used DVD Store 2.1. Because there's no built-in schema build for Snowflake, we manually built the schema using the SQL scripts Snowflake provided.

How we tested

Setting up Toad Data Studio

For this testing, we used DVD Store 2.1 to generate a 5GB database and load that data into a production PostgreSQL 16 instance, a development PostgreSQL 16 instance, an Oracle Database 21c instance, and Snowflake. Once the data was loaded into each database, we installed Toad Data Studio (TDS) and connected each database. We did this by downloading, installing, and configuring the corresponding ODBC driver and then connecting the database in the TDS UI. Then, to demonstrate the management capabilities of TDS across multiple platforms, we ran through a series of use cases to demonstrate the management capabilities of TDS across multiple platforms.

Installing Toad Data Studio

1. In a browser, navigate to *https://www.quest.com/products/toad-data-studio/*.
2. Click Download Free Trial.
3. Enter a business email, accept the EULA, and click Download Free Trial.
4. Next to Toad Data Studio, under Latest Version, click download.
5. Double-click the downloaded installer.
6. Click Next.
7. Accept the EULA, and click Next.
8. Click Next three times.
9. Click Install.
10. Click Finish.

Installing the PostgreSQL ODBC driver

1. In a browser, navigate to the PostgreSQL ODBC driver download page: *https://www.postgresql.org/ftp/odbc/releases/*.
2. Select the driver directory for the driver version you wish to install. We used 17.00.00.02.
3. To download, click the setup .exe file.
4. Double-click the installer.
5. Accept the EULA, and click Install.

Installing the Oracle ODBC driver

1. Download and install Microsoft Visual Studio 2017. If other versions are present, you will need to uninstall those first.
2. In a browser, navigate to the Oracle Instant Client download page: *https://www.oracle.com/database/technologies/instant-client/downloads.html*.
3. For Microsoft Windows (x64), click Instant Client.
4. To download the Instant Client Basic package, expand the version you need, and click the zip file.
5. Extract the .zip file.
6. From the same location, download the ODBC Package .zip file.
7. To the Instant Client Basic directory downloaded in step 3, extract the files in the ODBC directory.
8. Edit your environment variables Path, and add the location of the downloaded directory.
9. In the instant client directory, right-click odbc_install.exe, and run as administrator.
10. Once the driver is installed, navigate to the network directory in the instant client directory.
11. From your Oracle DB server, create tnsnames.ora and listener.ora files, and copy the contents of those files.

Installing the Snowflake ODBC driver

1. In a browser, navigate to the Snowflake ODBC download page: *https://www.snowflake.com/en/developers/downloads/odbc/*.
2. Download the latest Windows driver.
3. Double-click the installer file.
4. Click Next twice.
5. Click Install.
6. Click Finish.

Creating a DSN for PostgreSQL, Oracle, and Snowflake

1. Press Start, and type `ODBC Data Sources (64-bit)`.
2. Select either the User DSN or System DSN tab, and click Add.
3. Select the PostgreSQL ODBC driver, and click Finish.
4. To test the connection, fill out the following fields, and click Test:
   - Data source name
   - Database
   - Server
   - Port
   - UserName
   - Password
5. If the connection is successful, click Save.
6. Click Add.
7. Select the Oracle ODBC driver, and click Finish.
8. To test the connection, fill out the following fields, and click Test:
   - Data Source Name
   - TNS Service Name
   - User ID
9. Enter the password for the connection, and click OK.
10. If the connection is successful, click OK.
11. Click Add.
12. Select the Snowflake ODBC driver, and click Finish.
13. To test the connection, enter the following information, and click Test:
    - Data Source
    - User
    - Password
    - Server
    - Database
    - Schema
14. If the connection is successful, click OK.

Connecting to the PostgreSQL database in TDS

1. Next to Connect → New Connection, click the drop-down menu.
2. From the Group drop-down menu, select PostgreSQL.
3. Click Connect.

Connecting to the Oracle database in TDS

1. Next to Connect → New Connection, click the drop-down menu.
2. From the Group drop-down menu, select Oracle.
3. In the Database name drop-down menu, select the container database.
4. Enter the user and password information.
5. Next to Schema, click the three dots.
6. When prompted that selecting a Schema requires making a connection, click Yes.
7. From the drop-down menu, select DS2.
8. Click Connect.

Connecting to Snowflake in TDS

1. Next to Connect → New Connection, click the drop-down menu.
2. Enter the host, user, and password information.
3. Select the Warehouse and Database, and click Connect.

Running the use case scenarios

Once we connected all of our databases in TDS, we ran the following use case scenarios.

Scenario 1: Pre-migration research

For this scenario, we used TDS to create a visual representation of the database with links to object references and dependencies for PostgreSQL and Snowflake.

1. Next to the database connection → PostgreSQL/Snowflake, click the drop-down menu.
2. In the Status bar, click Diagram.
3. To highlight all the tables, simultaneously hold the shift key, click them, and drag them into the diagram pane.

Scenario 2: Finding data integrity issues

For this scenario, we used TDS to compare the schema differences between a production and development database. We started with a fresh 5GB DVD Store 2.1 database on the production and development instances. On the development database instance, we made a few schema changes: We changed the data type in two columns in the customers table, added a new column to the customer history table, made a column on the orders table nullable, and created a new table. After making these changes, we used TDS to compare the schemas.

1. Next to Compare → Schema Compare for PostgreSQL, click the drop-down menu.
2. On the Wizard welcome screen, click Next.
3. Select the production and development databases, and click Next.
4. Click Next.
5. Click Compare.

Scenario 3: Detecting data differences

For this scenario, we used TDS to compare the table data between a production and development database. We started with a fresh 5GB DVD Store 2.1 database on the production and development instance. We then made changes to the development instance DB by adding 10 customers to the customers table, adding five orders for one of the 10 created customers in the orders table, and adding two orderlines for each new order in the orderlines table. Finally, we used the data compare tool in TDS to find the differences in the data between the two database instances.

1. Next to Compare → Data Compare, click the drop-down menu.
2. Enter a name for the comparison, select the source database, select the target server, select the target database, and click Next.
3. Confirm that the checkboxes next to each table are checked, and click Next.
4. Click Compare.

Scenario 4: Freeing up resources

For this scenario, we used the Automation tool in TDS to automate Scenario 2. We made the same changes from Scenario 2 to the development database schema before following the steps below.

1. Click Automate.
2. Under Toolbox, click Execute a Toad Schema Compare for PostgreSQL project file, and publish output file (report).
3. To create a new comparison, in the Schema compare field, click the far right icon.
4. Enter a name, and click Save.
5. Select the source database, the target connection and database, and click Confirm.
6. In the Summary report field, click the three dots.
7. Enter a file name, and click Open.
8. In the detail report field, click the three dots.
9. Enter a file name, and click Open.
10. Click Run.
11. Enter a file name for the Automation file, and click Save.

Scenario 5: Avoiding data conversion complexities

For this scenario, we used the Import Export tool in TDS to copy the tables and data from an Oracle DVD Store 2.1 database to a PostgreSQL database. Once we copied the tables, we compiled and ran a SQL script from the DVD Store 2.1 postgresql load directory. We used this script to alter the table structures and create foreign keys, indexes, and stored procedures.

1. Click Import Export → Import Export Data.
2. Click Datasource table (Existing or new table), open the drop-down under Put data in, select the DB, and click Next.
3. Select the categories table, and click Next.
4. Select A single new table, enter a table name, select the schema, and click Next.
5. Click Add table.
6. Repeat steps 3 - 5 for the remaining 7 tables.
7. Click Next.
8. Click Finish.
9. Once the table transfer is complete, right-click the DB schema → New SQL Editor.
10. Copy the SQL queries from the PostgreSQL load directory into the SQL editor, and press F5 to run.

Read the report at https://facts.pt/oH7nrLG ▶

This project was commissioned by Quest Software.