TIBCO® Data Virtualization

Northbay Example Database Guide

Last Updated: *November 5, 2018*

Two-Second Advantage®

Important Information

Preface

Documentation for this and other TIBCO products is available on the TIBCO Documentation site. This site is updated more frequently than any documentation that might be included with the product. To ensure that you are accessing the latest available help topics, please visit:

- https://docs.tibco.com

Product-Specific Documentation

The following documents form the TIBCO® Data Virtualization(TDV) documentation set:

- *TIBCO TDV and Business Directory Release Notes* Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.
- TDV Installation and Upgrade Guide
- TDV Administration Guide
- TDV Reference Guide
- TDV User Guide
- TDV Security Features Guide
- Business Directory Guide
- TDV Application Programming Interface Guide
- TDV Tutorial Guide
- TDV Extensibility Guide
- TDV Getting Started Guide
- TDV Client Interfaces Guide
- TDV Adapter Guide
- TDV Discovery Guide
- TDV Active Cluster Guide
- TDV Monitor Guide
- TDV Northbay Example

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the TIBCO Product Documentation website mainly in the HTML and PDF formats.

The TIBCO Product Documentation website is updated frequently and is more current than any other documentation included with the product. To access the latest documentation, visit https://docs.tibco.com.

Documentation for TIBCO Data Virtualization is available on https://docs.tibco.com/products/tibco-data-virtualization-server.

How to Contact TIBCO Support

You can contact TIBCO Support in the following ways:

- For an overview of TIBCO Support, visit https://www.tibco.com/services/support.
- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the TIBCO Support portal at https://support.tibco.com.
- For creating a Support case, you must have a valid maintenance or support contract with TIBCO. You also need a user name and password to log in to https://support.tibco.com. If you do not have a user name, you can request one by clicking **Register** on the website.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote on feature requests from within the TIBCO Ideas Portal. For a free registration, go to https://community.tibco.com.

Using the Northbay Example Database

This chapter describes how to install and use the Northbay example database in these sections:

About Northbay

Northbay is an example database with data and schema that simulate a simplified set of tables from a fictitious online auction site.

You can specify the size of the Northbay tables to suit your immediate needs. The number of rows with meaningful data is configurable at installation and may be expanded with another installation. You can expand the number of data rows in the User table (and two other tables) from as small as 1,000 to up to 50,000 rows each. Many billions of rows of data can be quickly generated in the bid table! The Northbay example databases will aid when testing queries and joins on larger tables.

With thousands of standardized rows, Northbay is a useful tool for tests and demonstrations. Though table sizes are specified at installation, the data created is repeatable based upon the user and bid table row count. Northbay example databases created with the same number of user and bid rows will have the identical data. Use of this resource might also aid in diagnosis of on-site difficulties if it can be used to reproduce or demonstrate a problem.

Installing Northbay

The Northbay example database may be installed with defaults to quickly load example data resources with minimal parameter specifications, or several optional parameters may be specified to control the number of rows generated by the installer.

Default Northbay Database

Northbay installed with default values results in user, address, and credit card tables with 1,000 rows each and a bid table with the minimum size of 10,000 rows.

Default Northbay Database	
Table	**Row Count**
address	1000
bid	10000
category	29
creditcard	1000
favoritesearch	100
item	1000
sale	516
user	1000

User-Specified Northbay Database Limits

Northbay data tables may be dynamically expanded at installation by specifying the size of the user and bid tables at installation.

The user table may be expanded up to 50,000 rows and associated tables (address and credit card) are expanded to match the user table row count. Specify the –users parameter to set the number of rows in the User table.

The bid table may be expanded up to 86,400,000,000 rows! Use the –bids parameter to specify the number of rows.

As an aside, the bid table limit is imposed by the number of milliseconds in a day (the theoretical limit on the number of bids a single item could receive in a day) multiplied by a thousand, the number of items in the auction database. If for some strange reason, 86 billion rows are not enough for your testing purposes, then please send support an e-mail and we can accommodate any reasonable requests for larger bid tables.

Loading Northbay into the Default Repository

This procedure describes how to install the Northbay example database with default settings. The command line installation utility will create tables, generate data, and load the Northbay example database onto the default CIS repository database.

There are separate command line utilities for installing Northbay on either UNIX or Windows based installations of the CIS Server.

1. Search for and navigate to the 'Northbay.' within the <CIS_install_dir>
2. Type and enter the following command:

   For UNIX installations:

```
./northbay.sh –parameter1 <value1> -p2 <v2> -pN <vN>
```

   For Windows installations:

```
./northbay.bat –parameter1 <value1> -p2 <v2> -pN <vN>
```

   The parameter options for both utilities are the same. Parameters and values need not be specified when the default values apply to your case. For example, if your host is "localhost" you don't need to use the parameter: "-host localhost".

   Example: /northbay.bat –user ITDeptAdmin –password zyx098# -port 9410 -users 50000 –bids 1000000000

   Here are the most commonly used optional parameters:

Parameter Syntax	Default
-host <CIS host name >	localhost
-port <base port of CIS Server>	9400
-user <user name>	admin
-password < password>	admin
-domain <user domain>	composite
-users <row count of user, address, and creditcard tables>	min: 1000; max: 50000
-bids <bid table size>	10000; min:1000; max:86400000000

Northbay installer being executed on a Windows platform within a command prompt.

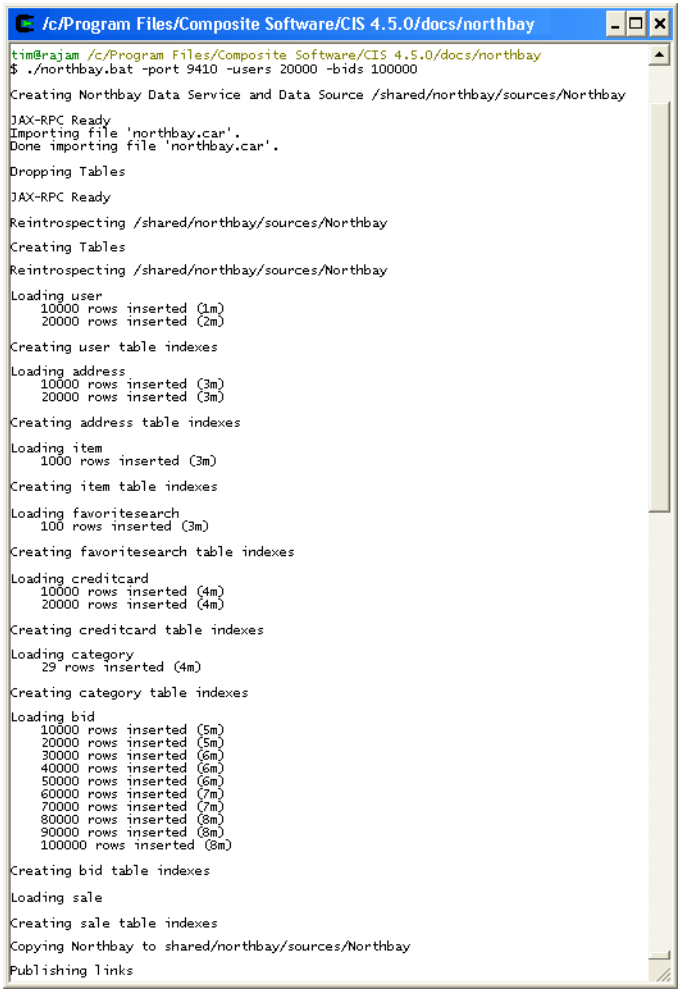

Note: While the example database may be as large as 50,000 users and 86+ billion bids, the generation and insertion of this data might take a very long time.

Loading Northbay into Other Data Sources

The –dsPath parameter enables direct specification to load Northbay into a supported target relational data source. Most relational data sources that support execution of DDL for automated table creation, insertion, and update (requires the ability to drop tables to change the user and bid table sizes) of data may use the command line installation utility to load the Northbay example data source.

The target data source must be introspected by the CIS Server prior to execution of the Northbay installation utility.

Define the target relational data source in CIS Studio or use a previously introspected resource. Note or copy the resource name (path) for use as the value of the –dsPath parameter.

1. Using the command line, navigate to:
```
.../docs/northbay
```

2. Run the following command:

   For UNIX installations:
```
./docs/northbay/northbay.sh
```

   For Windows installations:
```
./docs/northbay/northbay.bat
```

   The parameter options for both utilities are the same. Omit those parameters that will utilize the default value. Use the -users and the –bids parameters to specify the size of the associated tables.

   Here are the most commonly used optional parameters:

Parameter Syntax	Default
	localhost
	9400
	admin
	admin
	composite
	/shared

Parameter Syntax	Default
	<none>
	<none>
	min: 1000; max: 50000
	10000; min:1000; max:86400000000

Changing the Size of the Northbay Database

Change the size of the Northbay database by reinstallation of Northbay with new parameter values for -users and -bids. When the installer re-installs Northbay, the tables in the –dsPath are dropped, where the -dsPath default points to the CIS repository.

The Northbay installer recreates the tables and populates them with generated data according to the new -users and -bids parameter values.

Northbay Database Data Consistency

The Northbay database installer generates reliably consistent tables and data based upon the sizes of the user and bid tables. When Northbay is deployed in different locations with the same number of user and bid rows, the data generated is identical with only minor data type differences based on the target RDBMS.

When increasing only the number of users or increasing just the size of the bid table, most of the rest of the pre-existing data set will remain the same while those tables directly associated with user or bid table sizes will be amplified.

For example, if the bid table size should be increased while maintaining data consistency in the other tables, then the Northbay installer should be run with the larger -bids parameter value, while specifying the previous initial user table size. Changing only one setting ensures that the same data for the other unassociated tables are restored to the original.

Because this is accomplished by dropping existing tables and restoring the originally generated data, any data rows added after initial installation are not maintained during this kind of table amplification.

Restoring the Northbay Database

Restore the Northbay database to an initial state by simply reinstalling the database with the same -users and -bids parameter values. The Northbay installation utility will drop all Northbay tables and then recreate those tables and the data according to the parameter values.

Northbay Table Overviews

The Northbay database contains eight tables which are described in here:

user Table

This table lists all Northbay users with some account and personal information. The table row count is variable and may have a minimum of 1,000 users and up to a maximum of 50,000 users. Specify the size of the table at installation with the –users parameter option.

The column, non_english_name, contains a mixture of non-English characters that may or may not be rendered for proper display, depending on your local computing environment.

To ensure that international characters are displayed properly in Studio:

1. Create a new folder called “fallback” in the directory:
`<CIS_install_dir>jre\lib\fonts`

2. If you are working on a Windows based system, copy the Arial Unicode MS (TrueType) file: ARIALUNI.TTF from the Windows directory C:/Windows/fonts or C:/WINNT/fonts and paste it into the new fallback directory just created.

3. If Studio is already running, restart it so that the new Arial Unicode fonts may be used.

Column Name	Description	Data Type	Nullable
user_id	Unique user identification number	INT	N
first_name	User's first name	VARCHAR	N
middle_initial	User's middle initial	VARCHAR	N
last_name	User's last name	VARCHAR	N
non_english_name	User's names in another language	VARCHAR	Y
email	User's email	VARCHAR	N
opt_in	Whether or not the user wants to receive mail	BIT	N
phone	User's phone number	VARCHAR	N
birthday	User's birthday in range of 1940 through 1988	DATE	N
gender	Specifies user's gender. M is male, F is female.	VARCHAR	N
user_name	User's Northbay site user name log in	VARCHAR	N
password	User's password to log into Northbay site	VARCHAR	N
create_date	User's account creation date in range between Jan. 1, 2000 and Jul. 15, 2007	TIMESTAMP	N
primary_address_id	User's primary address ID if user has more than one address. NULL if user has no address on file	INT	Y

address Table

This table holds the users' address information. Some users have multiple addresses, others have none given on file. The row count of the address table is equal to that of the user table.

Column Name	Description	Data Type	Nullable
address_id	Unique address identification number – used to distinguish address	INT	N
user_id	user_id of user who possesses address	INT	N
street_address	Address street address	VARCHAR	N
city	City where address is located	VARCHAR	N
state	State where city is located	VARCHAR	N
zip_code	Zip code of address	VARCHAR	N
country	Country where address is located – all addresses are located in USA	VARCHAR	N
address_id	Unique address identification number – used to distinguish address	INT	N
user_id	user_id of user who possesses address	INT	N

creditcard Table

This table holds users' credit card information. None of the credit card numbers are valid. The length of the credit card table is equal to that of the user table.

Column Name	Description	Data Type	Nullable
user_id	User ID of user who possesses credit card	INT	N
credit_card_number	Credit card number	VARCHAR	N
credit_card_expiration_date	Credit card expiration date	DATE	N

item Table

Table of objects that have been offered for sale on Northbay. Table is initialized at 1000 rows.

Column Name	Description	Data Type	Nullable
item_id	Unique item identifier in other tables	INT	N
seller_id	Seller identifier equal to user_id	INT	N
category_id	Category_id of category item belongs in	INT	N
title	Name and description of product	VARCHAR	N
starting_bid	Starting price for this item	DECIMAL	N
entry_date	When this item was entered. Lies between the seller creation date and Jul. 15th, 2007	TIMESTAMP	N
expiration_date	When this item expires. All items are expired or will expire in the year 2020	TIMESTAMP	N
shipping_amount	Shipping cost of item	DECIMAL	N
picture	Picture of item	BLOB	Y
ext_description	Extended description of product	CLOB	Y

category Table

Category reference for the item table's category_id column. Table is initialized at 29 rows.

Column Name	Description	Data Type	Nullable
category_id	Unique, category identifier	INT	N
parent_id	Parent category organizes	INT	Y
category	Category type	VARCHAR	N

favoritesearch Table

This table is a list of 100 users and their favorite searches.

Column Name	Description	Data Type	Nullable
user_id	User who possesses favorite search	INT	N
favorite_search	User's favorite query	VARCHAR	N

bid Table

This table documents the time and amount of every bid execute on an item. The minimum length is 10,000 rows and the maximum is 86,400,000,000.

Column Name	Description	Data Type	Nullable
bid_id	Unique, bid identifier	INT	N
Item_id	Item_id of item bid is on	INT	N
bidder_id	User_id of bidder	INT	N
amount	Size of bid	DECIMAL	N
time	Bid time – The value is within the range between the item entry date and July 15, 2007	TIMESTAMP	N

sale Table

This table documents all the auctions that have ended. Its length is dependent on the number of bids because for each sale item, there needs to exist an initial bid on that item.

The sale table size is 516 rows with the default bid table size. The table is completely derived from the bid table and it is limited by the maximum number of items minus some number of items that are not yet sold.

Column Name	Description	Data Type	Nullable
sale_id	Unique, to identify the sale	INT	N
item_id	item_id of sold item	INT	N
item_description	Description of item	VARCHAR	N
buyer_id	user_id of bidder that had last bid	INT	N
buyer_name	Buyer's whole name	VARCHAR	N
buyer_non_english_name	Buyer's non_english_name	VARCHAR	Y
street_address	Street address of user that won bid	VARCHAR	N
city	City where street is located	VARCHAR	N
state	State where city is located	VARCHAR	N
zip_code	Zip_code where address is located	VARCHAR	N
country	Country where address is located. Limited to USA	VARCHAR	N
credit_card_number	User's credit card number	VARCHAR	N
credit_card_expiration_date	User's credit card's expiration date	DATE	N
purchase_date	Sale time and date	TIMESTAMP	N
item_amount	Final or winning bid for the item	DECIMAL	N
shipping_amount	Shipping cost	DECIMAL	N
seller_id	User Id of seller	INT	N
seller_name	Seller's full name	VARCHAR	N
buyer_feedback	Feedback rating the seller gave to the buyer on their performance in this transaction. Value 1-5	TINYINT	N

Column Name	Description	Data Type	Nullable
seller_feedback	Feedback rating the buyer gave to the seller on his performance in this transaction. Value 1-5	FLOAT	N
ship_date	Reported date of seller shipment of item	TIMESTAMP	N
pay_date	Date buyer paid the seller	TIMESTAMP	N