

Ascential **DataStage**

Sybase Enterprise Stage User Guide Version 1.0



Part No. 00D-005ES10
January 2005

This document and the software described or referenced in it, are confidential and proprietary to Ascential Software Corporation ("Ascential"). They are provided under, and are subject to, the terms and conditions of a license agreement between Ascential and the licensee, and may not be transferred, disclosed, or otherwise provided to third parties, unless otherwise permitted by that agreement. No portion of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of Ascential. The specifications and other information contained in this document for some purposes may not be complete, current, or correct, and are subject to change without notice. NO REPRESENTATION OR OTHER AFFIRMATION OF FACT CONTAINED IN THIS DOCUMENT, INCLUDING WITHOUT LIMITATION STATEMENTS REGARDING CAPACITY, PERFORMANCE, OR SUITABILITY FOR USE OF PRODUCTS OR SOFTWARE DESCRIBED HEREIN, SHALL BE DEEMED TO BE A WARRANTY BY ASCENTIAL FOR ANY PURPOSE OR GIVE RISE TO ANY LIABILITY OF ASCENTIAL WHATSOEVER. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL ASCENTIAL BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE. If you are acquiring this software on behalf of the U.S. government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as defined in Clause 252.227-7013 (c) (1) of DFARs.

This product or the use thereof may be covered by or is licensed under one or more of the following issued patents: US6604110, US5727158, US5909681, US5995980, US6272449, US6289474, US6311265, US6330008, US6347310, US6415286; Australian Patent No. 704678; Canadian Patent No. 2205660; European Patent No. 799450; Japanese Patent No. 11500247.

© 2005 Ascential Software Corporation. All rights reserved. DataStage®, EasyLogic®, EasyPath®, Enterprise Data Quality Management®, Iterations®, Matchware®, Mercator®, MetaBroker®, Application Integration, Simplified®, Ascential™, Ascential AuditStage™, Ascential DataStage™, Ascential ProfileStage™, Ascential QualityStage™, Ascential Enterprise Integration Suite™, Ascential Real-time Integration Services™, Ascential MetaStage™, and Ascential RTI™ are trademarks of Ascential Software Corporation or its affiliates and may be registered in the United States or other jurisdictions. Other marks are the property of the owners of those marks.

The software delivered to Licensee may contain third-party software code. See *Legal Notices (LegalNotices.pdf)* for more information.

Contents

1. Organization of this Manual.....	4
2. Documentation Conventions.....	4
3. Preface.....	4
4. Introduction.....	5
5. Accessing Sybase Databases.....	7
6. Must Do's.....	8
6.1. <i>Updating a Sybase Database.....</i>	8
6.2. <i>Loading a Sybase Database.....</i>	9
6.3. <i>Reading a Sybase Database.....</i>	9
6.4. <i>Performing a Direct Lookup on a Sybase Database Table.....</i>	10
6.5. <i>Performing an In-Memory Lookup on a Sybase Database Table.....</i>	10
7. Stage Page.....	10
7.1. <i>Advanced Tab.....</i>	10
7.2. <i>NLS Map.....</i>	11
8. Input Page.....	12
8.1. <i>Input Link Properties Tab.....</i>	12
8.1.1. <i>Target category.....</i>	14
8.1.2. <i>Connection Category.....</i>	15
8.1.3. <i>Options Category.....</i>	16
9. Outputs Page.....	17
9.1. <i>Output Link Properties Tab.....</i>	17
9.1.1. <i>Source Category.....</i>	18
9.1.2. <i>Connection Category.....</i>	19
9.1.3. <i>Options Category.....</i>	19

1. Organization of this Manual

This manual contains information about the installation and use of the Sybase Enterprise Stage provided by Ascential Software.

2. Documentation Conventions

This manual uses the following conventions:

Convention	Usage
Bold	In syntax, bold indicates commands, function names, buttons, actions and options.
Bold Italics []	For information that has to be provided in the syntax Brackets enclose optional items. Do not type the brackets unless indicated.
{ }	Braces enclose non-optional items from which at least one item must be selected. Do not type the braces.
...	Three periods indicate that more of the same type of item can optionally follow.
Courier	Courier indicates examples of source code and system output.
Courier Bold	In examples, courier bold indicates characters that the user types or keys the user presses (for example, < Return >).

3. Preface

This manual is intended for any user who wants to get information about either the whole system configuration or any of its components.

4. Introduction

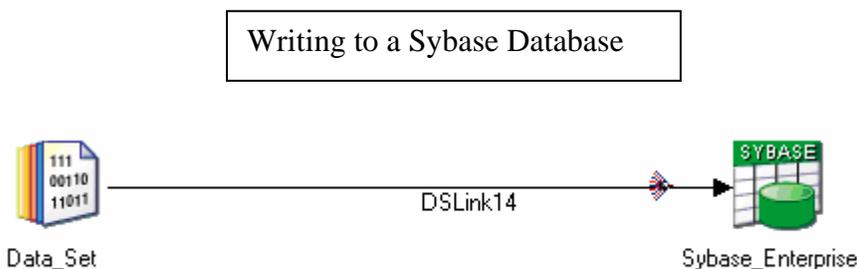
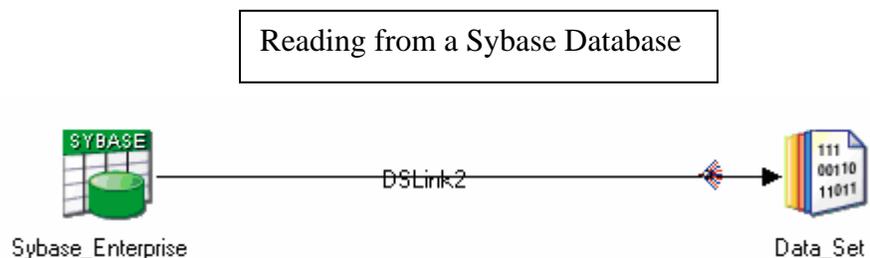
The Sybase Enterprise Stage is a database stage. It allows you to read data from and write data to a Sybase database. It can also be used in conjunction with a Lookup stage to access a lookup table hosted by a Sybase database.

The Sybase Enterprise Stage can have a single input link and a single reject link, or a single output link or output reference link.

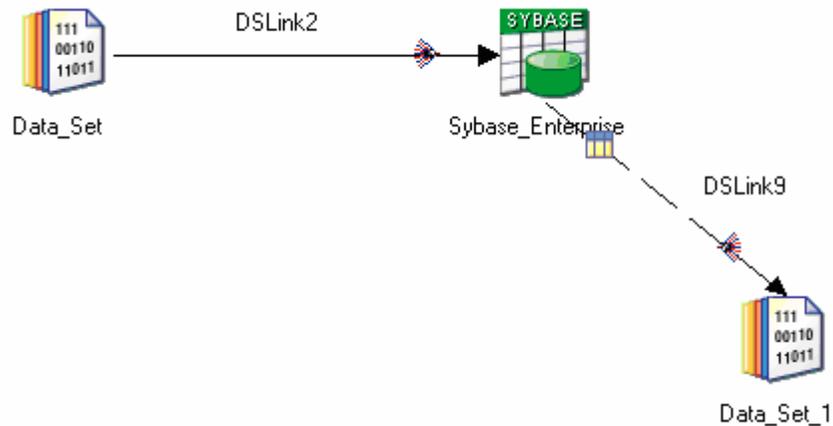
The stage performs one of the following operations:

- Updates a Sybase table using INSERT and/or UPDATE as appropriate.
- Loads a Sybase table (using Sybase “LOAD TABLE”).
- Reads a Sybase table.
- Deletes rows from a Sybase table.
- Performs a lookup directly on a Sybase table.
- Loads a Sybase table into memory and then perform a lookup on it.

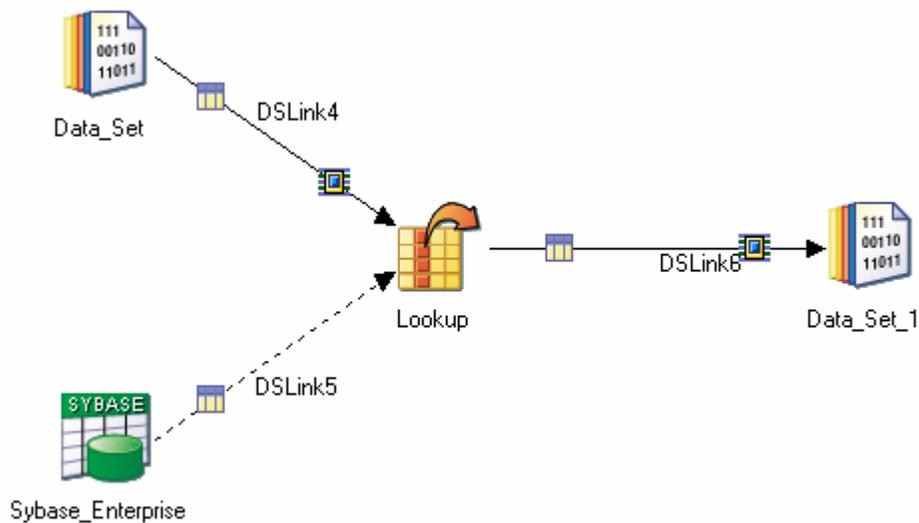
When using a Sybase stage as a source for lookup data, there are special considerations about column naming. If you have columns of the same name in both the source (Table) and lookup data sets, only the source data set column will go to the output data. If you want this column to be replaced by the column from the lookup data source, you need to drop the source data column before you perform the lookup.



Update/Insert/Delete a Sybase Database Table



Using the Sybase Enterprise Stage to access a Lookup Table



When you edit a Sybase Enterprise Stage, the Sybase Enterprise Stage editor appears. This is based on the generic stage editor.

The stage editor has up to three pages, depending on whether you are reading or writing a database:

- **Stage Page.** This is always present and is used to specify general information about the stage.
- **Inputs Page.** This is present when you are writing to a Sybase database. This is where you specify details about the data being written.
- **Outputs Page.** This is present when you are reading from a Sybase database, or performing a lookup on a Sybase database. This is where you specify details about the data being read.

5. Accessing Sybase Databases

Sybase open client software has to be installed on the server side for the functioning of Sybase Enterprise Stage. The configuration details are as below.

- Create the user defined environment variable SYBASE and set this to the \$SYBASE path that specifies the Sybase home directory (e.g., export SYBASE=/disk3/Sybase).
- Create the user defined environment variable SYBASE_OCS and set this to the Sybase open client software installation directory (e.g., export SYBASE_OCS=OCS-12_5).
- Interfaces file: - Add the details about the database server (database name, host machine name or IP address and port number) to the interfaces file located in \$SYBASE directory.
- Add SYBASE/bin to your PATH and SYBASE/lib to your LIBPATH, LD_LIBRARY_PATH, or SHLIB_PATH.
- Have login privileges to Sybase using a valid Sybase user name and corresponding password, server name and database. These must be recognized by Sybase before you attempt to access it.

Note: \$SYBASE/\$SYBASE_OCS/bin must appear first in your PATH. This is to ensure that \$SYBASE/\$SYBASE_OCS/bin/isql being executed always, when user executes “isql” command.

In case of accessing Sybase Databases with NLS, the necessary steps are as follows:

- Create a database using collation of that language you are going to test. (For ex: create database <<database path>> COLLATION 932JPN for Japanese (Shift_JIS) database.
- Install the Data Stage server in that particular language. (For ex: Japanese (shift_jis)). Upgrading the existing Data Stage server will not work as you will not get any option of selecting a support of other languages. You need to uninstall the existing server and install with the language you want.
- The language you want to test should be the default setting in your OS (Desktop), i.e., the machine on which you are going to test through the Data Stage client. You need to select the language through the Control Panel >> Regional settings and the keyboard input should also be set to that language.
- Client Setting: You have an NLS tab in the Enterprise Stage. You need to select an option of the language you want to test. (For example: If your OS is having Japanese as default then in the Data Stage Client the project default will be Shift_JIS which you don't need to select for every job you run).

6. Must Do's

Data Stage has many defaults, which means that it can be very easy to include Sybase Enterprise Stages in a job. This section specifies the minimum steps to take to get a Sybase Enterprise Stage functioning. DataStage provides a versatile user interface, and there are many shortcuts to achieving a particular end, this section describes the basic method, you will learn where the shortcuts are when you get familiar with the product.

The steps required depend on what you are using a Sybase Enterprise Stage for.

6.1. Updating a Sybase Database

- In the **Input Link Properties Tab**, under the Target category specify the update method as follows:
 - Specify a Write Method of Upsert.
 - Specify the Table you are writing.
 - Specify the Auto-generate Upsert SQL = Yes/No, this allows you to specify whether to use a statement automatically generated by DataStage or specify your own.

Choose the Upsert Mode, this allows you to specify whether to insert and update, or update only, and whether to use a statement automatically generated by DataStage or specify your own.

- Choose the Upsert Mode, this allows you to specify whether to insert/update/delete only or insert then update/update then insert/delete then insert
- If you have chosen Auto-generate Upsert SQL = No, specify the Insert/ Update/ Delete SQL statement to use according to the Upsert Mode selection.
- If you have chosen Auto-generate Upsert SQL = Yes, DataStage provides the auto-generated statement as a basis, which you can edit as required.

Under the **Connection category**, you can either specify a connection details for server, user and password. If you do not want the values to be hard coded, you can insert it from the Job Parameter Table, or by creating the environment variable either at the job level or at the project level.

- By default, DataStage connects to Sybase default database, but you can specify a different database if required.

Under the **Options category**:

- If you want to send rejected rows down a rejects link, set Output Rejects to True (it is false by default).
- Ensure column meta data has been specified for the write.

6.2. Loading a Sybase Database

This is the default write method.

- In the **Input Link Properties Tab**, under the Target category:
 - Specify a Write Method of Write.
 - Specify the Table you are writing.
 - Specify the Write Mode (by default DataStage appends to existing tables, you can also choose to create a new table, replace an existing table, or keep existing table details but replace all the rows).

Under the Connection category, you can either specify a connection details for server, user and password. If you do not want the values to be hard coded, you can insert it from the Job Parameter Table, or by creating the environment variable either at the job level or at the project level.

- By default, DataStage connects to Sybase default database, but you can specify a different database if required.

6.3. Reading a Sybase Database

- In the **Output Link Properties Tab**:
 - Choose a Read Method. This is Table by default, but you can also choose to read using auto-generated SQL or user-generated SQL. The read operates sequentially on a single node.
 - Specify the table to be read.
 - If using a Read Method of user-generated SQL, specify the SELECT SQL statement to use. DataStage provides the auto-generated statement, which can be edited.

Under the **Connection category**, you can either specify a connection details for server, user and password or if you do not want the values to be hard coded, you can insert it from the Job Parameter Table, or by creating the environment variable either at the job level or at the project level.

By default, DataStage connects to Sybase default database, but you can specify a different database if required.

- Ensure column meta data has been specified for the read.

6.4. Performing a Direct Lookup on a Sybase Database Table

- Connect the Sybase Enterprise Stage to a Lookup stage using a reference link.

In the **Output Link Properties Tab**:

- Set the Lookup Type to Sparse.
- Choose a Read Method. This is Table by default (which reads directly from a table), but you can also choose to read using auto-generated SQL or user-generated SQL.
- Specify the table to be read for the lookup.
- If using a Read Method of user-generated SQL, specify the SELECT SQL statement to use. DataStage provides the auto-generated statement as a basis, which you can edit as required.

Under the **Connection category**, you can either specify a connection details for server, user and password or if you do not want the values to be hard coded, you can insert it from the Job Parameter Table, or by creating the environment variable either at the job level or at the project level.

- By default, DataStage connects to Sybase default database, but you can specify a different database if required.
- Ensure column meta data has been specified for the write.

6.5. Performing an In-Memory Lookup on a Sybase Database Table

This is the default method. It has the same requirements as a direct lookup, except:

- In the **Output Link Properties Tab**:
 - Set the Lookup Type to Normal.

7. Stage Page

The **General tab** allows you to specify an optional description of the stage. The **Advanced tab** allows you to specify how the stage executes. The **NLS Map** tab appears if you have NLS enabled on your system; it allows you to specify a character set map for the stage.

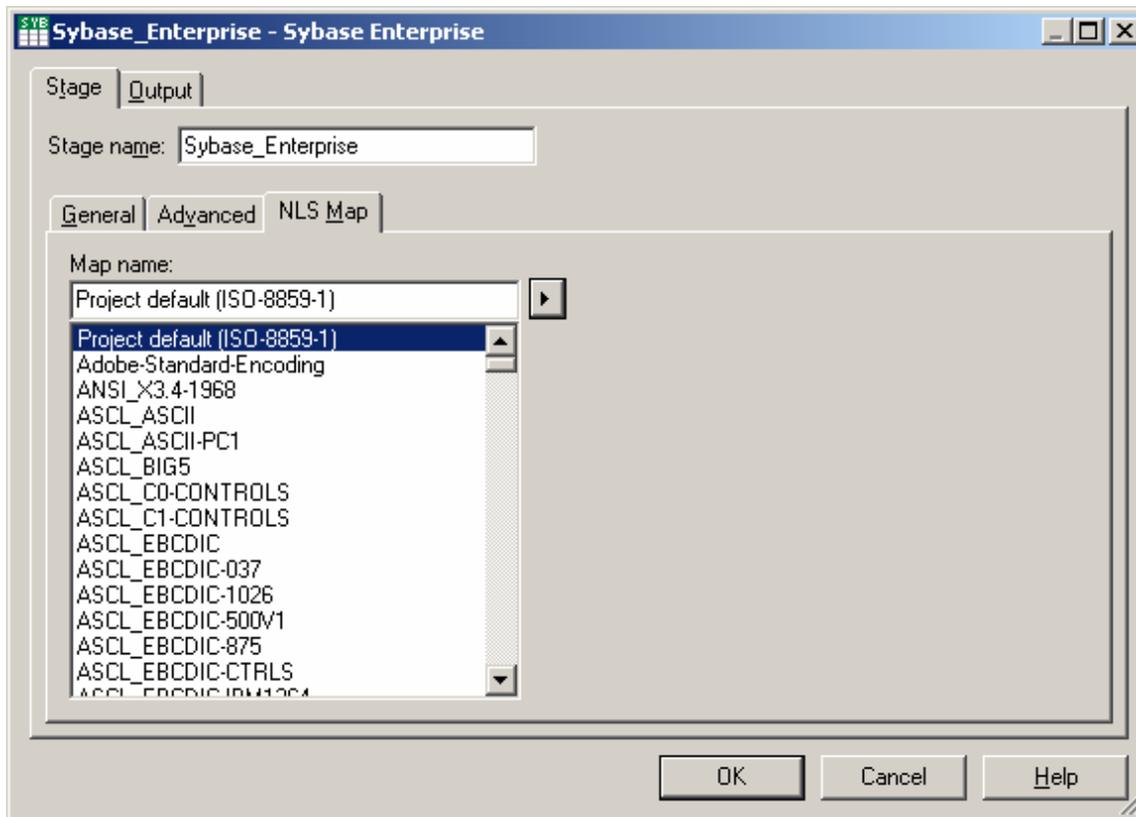
7.1. Advanced Tab

This tab allows you to specify the following:

- **Execution Mode.** The stage can execute in parallel mode or sequential mode. In parallel mode the data is processed by the available nodes as specified in the Configuration file, and by any node constraints specified on the Advanced tab. The conductor node processes the data in sequential mode. Sybase Enterprise stage runs in **sequential mode**.
- **Combinability mode.** This is **Auto** by default, which allows DataStage to combine the Stages that underlie parallel stages so that they run in the same process if it is sensible for this type of stage.
- **Preserve partitioning.** You can select **Set** or **Clear**. If you select **Set**, read operations will request that the next stage preserves the partitioning as is (it is ignored for write operations). Note that this field is only visible if the stage has output links.
- **Node pool and resource constraints.** Select this option to constrain parallel execution to the node pool or pools and/or resource pool or pools specified in the grid. The grid allows you to make choices from drop down lists populated from the Configuration file.
- **Node map constraint.** Select this option to constrain parallel execution to the nodes in a defined node map. You can define a node map by typing node numbers into the text box or by clicking the browse button to open the Available Nodes dialog box and selecting nodes from there. You are effectively defining a new node pool for this stage (in addition to any node pools defined in the Configuration file).

7.2. NLS Map

The **NLS Map** (National Language Support) allows you to define a character set map for the Sybase Enterprise stage. This overrides the default character set map set for the project or the job. You can specify that the map be supplied as a job parameter if required.



8. Input Page

The **Input page** allows you to specify details about how the Sybase Enterprise Stage writes data to a Sybase database. The Sybase Enterprise Stage can have only one input link writing to one table. The **General tab** allows you to specify an optional description of the input link. The **Properties tab** allows you to specify details of exactly what the link does. The **Partitioning tab** allows you to specify how incoming data is partitioned before being written to the database. The Columns tab specifies the column definitions of incoming data. The **Advanced tab** allows you to change the default buffering settings for the input link. Details about Sybase Enterprise Stage properties, partitioning, and formatting are given in the following sections.

8.1. Input Link Properties Tab

The **Properties tab** allows you to specify properties for the input link. These dictate how incoming data is written and where. Some of the properties are mandatory, although many have default settings. Properties without default settings appear in the warning color (red by default) and turn black when you supply a value for them. The following table gives a quick reference list of the properties and their attributes. A more detailed description of each property follows.

Category/Property	Values	Default	Mandatory?	Repeats?	Dependent of
Target/Table	String	N/A	Y	N	N/A
Target/Write Method	Upsert, Write	Write	Y	N	N/A
Target/Write Mode	Append, Create, Replace, Truncate	Append	Y	N	N/A
Target/Auto-generated Upsert SQL	Yes/No	Yes	Y	N	N/A
Target/Insert SQL	String	N/A	N	N	N/A
Target/Insert Array Size	Number	1	N	N	Insert SQL
Target/Update SQL	String	N/A	N	N	N/A
Target/Delete SQL	String	N/A	N	N	N/A
Target/Write Method	Upsert, Write	Write	Y	N	N/A
Target/Upsert Mode	Delete only, Delete then Insert, Insert only, Insert then Update, Update only, Update then Insert.	Insert then update	Y	N	N/A
Connection/User	String	N/A	Y	N	N/A
Connection/Password	String	N/A	Y	N	N/A
Connection/Server	String	N/A	Y	N	N/A
Connection/Database	String	N/A	Y	N	N/A
Options/Truncate Column Name	True/ False	False	Y	N	N/A
Options/Default String Length	Number	1	N	N	N/A
Options/Open Command	String	N	N	N	N/A
Options/Close Command	String	N	N	N	N/A
Options/Length to Truncate	Number	N	N	N	N/A
Options/Drop Unmatched Field	True/False	False	N	N	N/A
Options/Row Commit Interval	Number	1	N	N	Insert Array Size
Options/ Output Reject Records	True/False	False	N	N	N/A

8.1.1. Target category

Table: Specify the name of the table to write to. You can specify a job parameter if required.

Write Method: Choose from Upsert or Write (the default). Upsert allows you to provide the insert, update and delete SQL statements. Write sets up a connection to Sybase and inserts records into a table, taking a single input data set. The Write Mode property determines how the records of a data set are inserted into the table.

Write Mode: This only appears for the Write, Write Method. Select from the following:

- Append. This is the default. New records are appended to an existing table.
- Create. Create a new table. If the table already exists an error occurs and the job terminates. You must specify this mode if the table does not exist.
- Replace. The existing table is first dropped and a new table is created in its place.
- Truncate. The existing table attributes (including schema) are retained, but any existing records are discarded. New records are then appended to the table.

Auto-generated Upsert SQL: You can specify one of the two values **Yes** or **No**. Set to **Yes** to view auto-generated SQL statements and **No** to specify your own.

Upsert mode: This only appears for the Upsert write method. Allows you to specify how the insert, update and delete statements are to be derived.

1. If **Auto-generated Upsert SQL** = YES, Choose from:

- **Delete Only:** Select this to view an auto-generated Delete statement, DataStage generates a delete statement for you, based on the values supplied for table name and column. The statement can be viewed by selecting the **Delete SQL** properties.
- **Delete then Insert:** DataStage generates delete and insert statements for you, based on the values supplied for table name and column. The statements can be viewed by selecting the **Delete SQL** or **Insert SQL** properties.
- **Insert Only:** Select this to view an auto-generated **Insert** statement, DataStage generates an insert statement for you, based on the values supplied for table name and column. The statement can be viewed by selecting the **Insert SQL** properties.
- **Insert then Update:** DataStage generates insert and update statements for you, based on the values supplied for table name and column. The statements can be viewed by selecting the **Insert SQL** or **Update SQL** properties.
- **Update Only:** Select this to view an auto-generated **Update** statement, DataStage generates an update statement for you, based on the values supplied for table name and column. The statement can be viewed by selecting the **Update SQL** properties.

- **Update then Insert:** DataStage generates update and insert statements for you, based on the values supplied for table name and column. The statements can be viewed by selecting the **Update SQL** or **Insert SQL** properties.

2. If **Auto-generated Upsert SQL** = No, Choose from:

- **Delete Only:** Select this to enter your own delete statement. Then select the **Delete SQL** property and edit the statement proforma.
- **Delete then Insert:** DataStage generates delete and insert statements for you, based on the values supplied for table name and column. The statements can be viewed by selecting the **Delete SQL** or **Insert SQL** properties.
- **Insert Only:** Select this to enter your own insert statement. Then select the **Insert SQL** property and edit the statement proforma.
- **Insert then Update:** DataStage generates insert and update statements for you, based on the values supplied for table name and column. The statements can be viewed by selecting the **Insert SQL** or **Update SQL** properties.
- **Update Only:** Select this to enter your own Update statement. Then select the **Update SQL** property and edit the statement proforma.
- **Update then Insert:** DataStage generates update and insert statements for you, based on the values supplied for table name and column. The statements can be viewed by selecting the **Update SQL** or **Insert SQL** properties.

Insert SQL: Only appears for the Upsert write method. This property allows you to view an auto-generated Insert statement, or to specify your own (depending on the setting of the Update Mode property). It has a dependent property:

Insert Array Size Specify the size of the insert host array. The default size is 1 record. If you want each insert statement to be executed individually, specify 1 for this property.

Delete SQL: Only appears for the Upsert write method. This property allows you to view an auto-generated Delete statement, or to specify your own (depending on the setting of the Update Mode property).

Update SQL: Only appears for the Upsert write method. This property allows you to view an auto-generated Update statement, or to specify your own (depending on the setting of the Upsert Mode property)

8.1.2. Connection Category

User: Enter the User name. This option is mandatory.

Password: Enter the Password. This option is mandatory.

Server: Specify the Server name to be used for all Sybase database connections. This option is mandatory.

Database: Enter the database. This is optional and if not specified connects to the default database.

Note: If you do not want the connection strings to be hard coded, you can insert it from the Job Parameter Table. This can be done in two ways.

- By entering values in the job parameter table, and
- By creating the environment variable either at the job level or at project level.

8.1.3. Options Category

Truncate Column Names: This only appears for the Write Method of Write. You can set one of two values, either True or False. Set **True** to truncate field names to the size allowed by the SYBASE. Set **False** to disable the truncation of field names.

Default String Length: This is an optional property and only appears for the Write Method of Write. It is set to the size of 1 byte by default. Sets the default string length of variable-length strings written to a table.

The maximum length you can set is 2000 bytes. Note that the stage always allocates the specified number of bytes for a variable-length string. In this case, setting a value of 2000 allocates 2000 bytes for every string. Therefore, you should set the expected maximum length of your largest string and no larger. In case of NLS maximum = 2000 / number of bytes per NLS character.

For example: For Unicode Database, Japanese character occupies 3 bytes.

Open Command: Use it to specify a command, in single quotes, to be parsed and executed by the Sybase database on all processing nodes before the table is opened. You can specify a job parameter if required.

Close Command: Use it to specify any command, in single quotes, to be parsed and executed by the Sybase database on all processing nodes after the stage finishes processing the table. You can specify a job parameter if required.

Length to Truncate: Enter the length to truncate columns names to.

Drop unmatched field: You can set one of two values, either **True** or **False**. Set to **True** to silently drop all input columns that do not correspond to columns in an existing table. Otherwise the stage reports an error and terminates the job. The default value is False.

Row Commit Interval: Specify the number of records to be committed before the start of a new transaction. This option is considered only if Insert array size is set to 1 otherwise Row Commit Interval is set to Insert Array Size.

Output Reject Records: This only appears for the Upsert write method. It is **False** by default, set to **True** to send rejected records to the reject link.

8.2. Partitioning Tab

Sybase IQ does not support partitioning and hence you can ignore this tab, and let the stage take the default values.

9. Outputs Page

The **Outputs page** allows you to specify details about how the Sybase Enterprise Stage reads data from a Sybase database. The Sybase Enterprise Stage can have only one output link. Alternatively it can have a reference output link, which is used by the Lookup stage when referring to a Sybase lookup table. It can also have a reject link where rejected records are routed (used in conjunction with an input link). The Output Name drop-down list allows you to choose whether you are looking at details of the main output link or the reject link.

The **General tab** allows you to specify an optional description of the output link. The Properties tab allows you to specify details of exactly what the link does. The Columns tab specifies the column definitions of the data. The Advanced tab allows you to change the default buffering settings for the output link.

Details about Sybase Enterprise Stage properties are given in the following sections.

9.1. Output Link Properties Tab

The Properties tab allows you to specify properties for the output link. These dictate how incoming data is read from what table. Some of the properties are mandatory, although many have default settings. Properties without default settings appear in the warning color (red by default) and turn black when you supply a value for them.

The following table gives a quick reference list of the properties and their attributes. A more detailed description of each property follows.

Category/Property	Values	Default	Mandatory	Repeats	Dependent of
Source/Lookup Type	Normal/Sparse	Normal	Y	N	N/A
Source/Read Method	Table/User-defined SQL	Table	Y	N	N/A
Source/Query	String	N/A	N	N	N/A

Source/ Table	String	N/A	N	N	N/A
Source/ Selectlist	String	N/A	N	N	Table
Source/ Where	String	N/A	N	N	Table
Connection/ User	String	N/A	Y	N	N/A
Connection/ Password	String	N/A	Y	N	N/A
Connection/ Server	String	N/A	Y	N	N/A
Connection / Database	String	N/A	N	N	N/A
Options/ Open Command	String	N	N	N	N/A
Options/ Close Command	String	N	N	N	N/A
Options/ Fetch Array Size	Number	1	N	N	N/A

9.1.1.Source Category

Lookup Type: Sybase Enterprise Stage is connected to a Lookup stage via a reference link, Sybase Enterprise Stage will provide data for an in-memory look up (Lookup Type = Normal) or lookup is set to access the database directly (Lookup Type = Sparse). If the Lookup Type is **Normal**, the Lookup stage can have multiple reference links. If the Lookup Type is **Sparse**, the Lookup stage can have one reference link.

Read Method: This property specifies whether you are specifying a table or a query when reading the Sybase database, and how you are generating the query.

- Select the Table method in order to use the Table property to specify the read.
- Select Auto-generated SQL to have DataStage automatically generate an SQL query based on the columns you have defined and the table you specify in the Table property.
- Select User-defined SQL to define your own query. A user-defined or auto-generated SQL will be read sequentially on one node.

Sybase Enterprise Stage will provide data for an in-memory look up (Lookup Type = Normal) or lookup is set to access the database directly (Lookup Type = Sparse). If the Lookup Type is **Normal**, the Lookup stage can have multiple reference links. If the Lookup Type is **Sparse**, the Lookup stage can only have one reference link.

Query: Optionally allow you to specify an SQL query to read a table. The query specifies the table and the processing that you want to perform on the table. This statement can contain joins, views, database links, synonyms, and so on.

Table: Specifies the name of the table. The table must exist and you must have SELECT privileges on the table.

Table has dependent properties:

- **Where**

Stream links only. Specifies a WHERE clause of the SELECT statement to specify the rows of the table to include or exclude from the read operation. If you do not supply a WHERE clause, all rows are read.

- **Select List**

Optionally specifies an SQL select list, enclosed in single quotes that can be used to determine which columns are read. You must specify the columns in *list* in the same order as the columns are defined in the record schema of the input table.

9.1.2. Connection Category

User: Enter the User name. This option is mandatory.

Password: Enter the Password. This option is mandatory.

Server: Specify the Server name to be used for all Sybase database connections. This option is mandatory.

Database: Enter the database. This is optional and if not specified connects to the default database.

Note: If you do not want the connection strings to be hard coded, you can insert it from the Job Parameter Table. This can be done in two ways.

- By entering values in the job parameter table, and
- By creating the environment variable either at the job level or at project level.

9.1.3.Options Category

Close Command: Use it to specify any command to be parsed and executed by the Sybase database on all processing nodes after the stage finishes processing the Sybase table. You can specify a job parameter if required.

Open Command: Use it to specify any command to be parsed and executed by the Sybase database on all processing nodes before the Sybase table is opened. You can specify a job parameter if required.

Fetch Array Size: Specify the number of rows to be retrieved during each fetch operation. Default is 1.