



New Features Summary

Sybase IQ InfoPrimer 15.3

Microsoft Windows, UNIX, and Linux

DOCUMENT ID: DC01641-01-1530-02

LAST REVISED: July 2011

Copyright © 2011 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

Replication Server and Sybase IQ InfoPrimer	
Integration	1
Sybase IQ InfoPrimer - An Option to Sybase IQ	1
New Project Types - Extract and Load Project and SQL Transformation Project	2
Database Connection Profiles for Extract and Load Project and SQL Transformation Project	4
Multiple-Object Transfer Support for Easy Migration of Repositories	4
Performance and Scalability Enhancements	5
Repository Performance Enhancements	5
Enhanced Support for Datasources	5
JavaScript Transformation Enhancements	6
Navigator Enhancements	6
Sub-Capacity Licensing	6
Extended Platform Support	7
Installation Location	7

Contents

Replication Server and Sybase IQ InfoPrimer Integration

You can now use Replication Server® 15.6 ESD #1 with Sybase® IQ InfoPrimer 15.3 to replicate and transform data from Adaptive Server Enterprise (ASE) to Sybase IQ.

The integrated Replication Server and Sybase IQ InfoPrimer solution performs a nonatomic bulk materialization of data from an ASE primary database to a replicate IQ database.

Sybase IQ InfoPrimer creates staging tables on the replicate IQ database and performs the data-extract step of the materialization process on each primary database table. It also transforms and loads the data into end-user tables by executing generated stored procedures. These transformation stored procedures are also executed during replication.

See the *Replication Server 15.6 New Features Guide* for a full-feature description, *Sybase IQ InfoPrimer 15.3 Users Guide* for the user interface changes in Sybase IQ InfoPrimer, and *Replication Server 15.6 Release Bulletin* for known issues in the data flow.

Sybase IQ InfoPrimer - An Option to Sybase IQ

Sybase IQ InfoPrimer, previously known as Sybase ETL, is now packaged and delivered as a separately licensed option to Sybase IQ 15.3, and must be installed using its own installer.

Sybase IQ InfoPrimer provides mechanisms for:

- Extracting data from a datasource, transforming and cleaning the data, and loading it into a Sybase IQ data target using a comprehensive set of transformation functions.
- Extracting and loading data from data sources into Sybase IQ and then performing a set of transformations on the data directly in the Sybase IQ data targets.

Sybase IQ InfoPrimer supports extracting data from multiple heterogeneous data sources namely, Adaptive Server® Enterprise, IBM DB2, Oracle, Microsoft SQL Server, Sybase IQ, SQL Anywhere®, MySQL, text files, and XML files.

Sybase IQ InfoPrimer offers these new capabilities:

- **Extended high speed extract and load capabilities**
 - Supports large objects (LOB) data from enhanced datasources namely, Sybase ASE, Oracle, Microsoft SQL Server, and files.
 - Enables high-speed bulk extraction from Sybase Adaptive Server, Microsoft SQL Server, and Oracle datasources.
 - Enhances support for generating **LOAD TABLE** commands for existing bulk- load files if they are:

New Project Types - Extract and Load Project and SQL Transformation Project

- Sybase IQ compatible binary format
- Text files with delimited records
- Files with fixed length values and variable length records
- Files in Apache Web log formats (Common Log Format and Extended Log Format)
- LOBs
- Text files generated by the extract phase of execution.
- Continues to support XML, SQL Anywhere, MySQL, Sybase IQ, and IBM DB2 as datasources.
- Supports parallel loading and transformation of multiple tables.
- Provides better support for large tables, wide tables, and large number of tables.
- Supports incremental data changes in existing tables (**insert**, **update**, **delete**).
- **Ease of use**
 - Enhances repository performance.
 - Enables fast and easy migration of repositories.
 - Supports full initialization from source to target with and without applied transformation.
 - Supports incremental loads to existing tables with and without transformation.
 - Enhances life cycle management with database connection profiles.
- **Seamless integration with Sybase IQ**
 - Supports table-based, schema-based, and column-based transformation.
 - Supports generation of Sybase IQ **LOAD TABLE** input from a variety of data sources.
 - Enhances support for **Insert Location** capability of the Sybase IQ host.
 - Enhances JavaScript-based transformation support.
 - Supports multiple numbers of transformations within a single load.
 - Supports multiple numbers of load or transformation jobs concurrently.
- **Enhanced licensing capabilities**
 - Supports SySAM sub-capacity licensing.

See *Sybase IQ InfoPrimer Users Guide*.

New Project Types - Extract and Load Project and SQL Transformation Project

In addition to the Extract, Transform, and Load (ETL) Project type that continues to provide the same functionality as in the previous releases, two new project types are added in Sybase IQ InfoPrimer.

While the existing ETL project type provides the capability for extracting data from a data source, transforming and cleaning the data, and then loading it into a Sybase IQ data target, the new project types, namely Extract and Load (EL) and SQL Transformation projects, allow

New Project Types - Extract and Load Project and SQL Transformation Project

you to extract data from source into IQ stage tables, transform it, and load it into IQ destination tables. This helps to leverage the data processing power of Sybase IQ.

- **Extract and Load project** – is used for designing high speed bulk extraction and bulk loading processes, with a single instance of Sybase IQ as the target. The EL project lets you extract data from any supported datasource and load it into a Sybase IQ database or to a Sybase IQ load file.

Unlike the ETL project, an EL project facilitates extracting multiple tables in Sybase IQ from a single database source or from multiple files. An EL project allows you to migrate data from Oracle, Adaptive Server Enterprise, Microsoft SQL Server, flat files, and other supported datasources to a Sybase IQ database, if minor or no schema changes are needed.

The data file types supported as a source by an EL project are:

- Sybase IQ compatible binary format
- Text files with delimited record values (including Apache Web log formats)
- Files with fixed length values and variable length records
- Large objects (LOBs)

Note: Files with fixed-length values and fixed-length records are not supported. For this particular file type, use ETL project instead.

Extracting LOB data is supported only for Adaptive Server, Oracle, Microsoft SQL Server, and LOB files (binary files or character files).

Note: The EL project does not support XML as a source. To load XML sources, use the XML via SQL Data Providers component of the ETL project.

See Sybase IQ InfoPrimer Users Guide > Extract and Load Projects.

- **SQL Transformation project** – performs SQL transformations where the source and target are both on the same instance of Sybase IQ. It generates transformations to be performed on data that exists in a Sybase IQ database. As compared to JavaScript transformations, SQL transformations executed in Sybase IQ can enhance performance and scalability.

A SQL Transformation project defines the Sybase IQ connection at the project level. The entire transformation project is translated into a single SQL script, so all transformation tasks use the same connection. A SQL Transformation project controls execution flow, based on the connections between its components. The components required to handle execution flow are:

- Start – represents the beginning of a control flow.
- Error – marks the end of a failed control flow.
- Success – marks the successful end of a control flow.
- Decision – routes execution flow based on SQL rules.
- Task Group – defines ordered groups of transformation tasks.

See Sybase IQ InfoPrimer Users Guide > SQL Transformation Projects.

- **Extract Transform and Load (ETL) project** – extracts data from multiple heterogeneous data sources and loads it into one or more data targets using a comprehensive set of transformation functions. All functionalities of the ETL projects

remain the same as in Sybase ETL version 4.9. See *Sybase IQ InfoPrimer Users Guide > Extract, Transform, and Load Projects*.

Database Connection Profiles for Extract and Load Project and SQL Transformation Project

Sybase IQ InfoPrimer includes a Connection Profile Manager to help you create a connection profile and store it in the repository for use in EL and SQL Transformation projects.

A connection profile contains the information needed to connect to a database runtime instance. It manages the configured connections defined for the installed database instance running on the server. When you create a connection profile, you specify standard data source configuration parameters, such as a connection URL, user name, and password. These parameter values are stored as secure metadata for the created connection profile, which can be used for connectivity checks. You can use either default connection profiles, or you can create and configure a new connection profile.

If you define a connection profile using the Connection Profile Manager from the Tools menu of Sybase IQ InfoPrimer Development interface, you can use it when creating an EL or a SQL Transformation project and are not required to provide the database connection details again. EL and SQL Transformation projects can also subscribe to profiles thus ensuring they always use the current connection settings stored in the repository when executed.

See *Sybase IQ InfoPrimer Users Guide > Administration Tasks > Managing Connection Profiles*.

Multiple-Object Transfer Support for Easy Migration of Repositories

Sybase IQ InfoPrimer allows you to select and transfer multiple objects from a source repository to a target repository.

In earlier versions of ETL, you could transfer objects only if two or more connections were open, by selecting the **Transfer** option in the context menu of projects and jobs in the Navigator. However, in Sybase IQ InfoPrimer, you can also select the **Transfer** option from the context menu at the repository level or from the File menu, regardless of whether any connections are open. You can open the connections from within the Transfer window. Connections opened from within the Transfer window close automatically when the window is closed.

Note: To transfer repository objects from a source repository to a target repository, all repositories involved must be accessible from the same Sybase IQ InfoPrimer Development instance.

Sybase IQ InfoPrimer transfers only project and job definitions. It does not support the transfer of additional data like parameter sets, execution properties, and database connection profiles.

See *Sybase IQ InfoPrimer Users Guide > Administration Tasks > Transferring Objects Between Repositories*.

Performance and Scalability Enhancements

Review the performance and scalability enhancements made in Sybase IQ InfoPrimer 15.3.

Repository Performance Enhancements

The repository performance has been enhanced to support large metadata configurations for projects, which allows support for large numbers of tables and columns.

Sybase IQ InfoPrimer minimizes queries on the repository database each time it requires metadata information. When accessing the repository database, Sybase IQ InfoPrimer can get as much metadata information as possible in each SQL query, thus eliminating the bottleneck caused by the execution of excessive queries.

The improved repository service performance reduces the time you spend on repository operations such as opening, reading, writing, or saving a project.

Enhanced Support for Datasources

Sybase IQ InfoPrimer offers enhanced support for high-speed bulk extraction, for some types of datasources. It also supports large object (LOB) data from these datasources.

- Sybase Adaptive Server Enterprise
- Oracle
- Microsoft SQL Server
- Files – support for:
 - Sybase IQ compatible binary format
 - Loading a portion of the data. For example, skip the first n number of rows and load only m number of rows.
 - Delimiter-separated records
 - Fixed-length records

Navigator Enhancements

- Web log formats

You can also apply a SQL statement against the load source for filtering the data to extract.

See *Sybase IQ InfoPrimer Users Guide*.

JavaScript Transformation Enhancements

In Sybase IQ InfoPrimer, the transformation components, namely Data Splitter JavaScript and Data Calculator JavaScript, have been rearchitected to improve throughput and performance.

None of these changes impact the user interface.

See *Sybase IQ InfoPrimer Users Guide > Extract Transform and Load Projects > ETL Components > Transformation Components*.

Navigator Enhancements

The Navigator supports two new features.

- You can now sort projects and jobs in the Navigator either in alphabetical order or based on their creation or modification dates. Click the **Select navigator display preference** icon in the top-right corner and select the appropriate option.
- You can now define how you want projects to appear in the Navigator. A flat presentation shows all project types under a common project folder. A hierarchical presentation (default), shows each project under its respective project type folder. Click the **Select navigator display preference** icon in the top-right corner and select the appropriate option.

See *Sybase IQ InfoPrimer 15.3 Users Guide > Getting Started > Using the Sybase IQ InfoPrimer Development Interface > Navigator*.

Sub-Capacity Licensing

Sybase now offers sub-capacity licensing, which lets you license Sybase products on a subset of the CPUs available on a physical machine.

You can license Sybase IQ InfoPrimer 15.3 for fewer than the total number of processors on a multiprocessor machine. Subject to the terms of your software license agreement and product documentation, you may be able to license a subset of the processors available on a physical machine using the SySAM sub-capacity license functionality.

See *Sybase IQ InfoPrimer Installation Guide > Planning Your Installation > Obtaining a License > Sub-capacity Licensing*.

Extended Platform Support

Sybase IQ InfoPrimer is now available on new platforms and operating system configurations.

- Microsoft Windows 7
- SuSE Linux Enterprise Server
 - SuSE Linux Enterprise Server 11 – 32bit , 64-bit
 - SuSE Linux Enterprise Server 11 on POWER – 64-bit
- Red Hat Enterprise Linux 5.3 on POWER – 64-bit
- Red Hat Enterprise Linux 5.3 on x86 – 64 bit

For a complete list of supported platforms and other system requirements, see *Sybase IQ InfoPrimer Installation Guide > Planning Your Installation > System Requirements*.

Installation Location

Sybase IQ InfoPrimer must be installed in its own subdirectory under the root directory.

The root directory can either be a directory containing other Sybase products or the directory specified in the SYBASE environment variable. If the SYBASE environment variable has not been defined, install Sybase IQ InfoPrimer in the recommended installation location for Sybase IQ 15.3 or set the SYBASE environment variable before installation.

The default subdirectory name for Sybase IQ InfoPrimer Development is IQIPD-15_3 and for Sybase IQ InfoPrimer Server is IQIPS-15_3.

See *Sybase IQ InfoPrimer Installation Guide > Planning Your Installation > Installation Directory*.

Installation Location