

Flex2SQL™

Mertech's ISAM to SQL Database

Connectivity (ISDBC) Drivers For DataFlex®

Contents

- Overview
- Flex2SQL Product Fact Sheet
- Interaction with the DataFlex Runtime
- Handling Key DataFlex Features
- Other Features
- Evaluate Our Product
- Contact Information

Overview

The Internet revolution has underscored the importance of making data available reliably and at a high speed to an ever-growing user base. Corporations continue to consolidate their data into single, enterprise-wide databases so that information about customers, products, and the market can be easily extracted and manipulated. These databases also allow customers and their mobile workforce to update and retrieve information at the same time, from any location. The new, data-intensive demands of today's businesses require a database server that is robust, scalable, gives excellent response time in extracting and manipulating data, has great disaster recovery features and above all provides excellent security features and guarantees 24X7 availability.

Relational or Structured Query Language (SQL) based databases from companies like Oracle, Microsoft, IBM and open-source databases such as MySQL and PostgreSQL meet those criteria. These products cover about 90% of the database market and every new major application supports one or more of these databases. As you can see, relational databases are at the heart of any reasonable-size IT infrastructure.

Mertech's ISAM to SQL Database Connectivity (ISDBC) drivers allow DataFlex developers to offer SQL database servers with their applications without changing source code or compromising on performance. The seamless integration of non-DataFlex databases with a DataFlex application has been made possible largely by a change in the DataFlex runtime architecture. This change was introduced with DataFlex 3.1c and Visual DataFlex 4, which allow multiple databases to work with DataFlex applications without rewriting or recompiling any existing data.

Mertech's Flex2SQL product bundle is comprised of 1) a GUI migration tool that takes your existing DataFlex tables and indexes and migrates them to your target backend, creating the table structures, index structures, and RECNUM values for all the files in the filelist and 2) the database driver DLL or shared-object that works with the DataFlex runtime and handles all client/server connection and database operations.

Flex2SQL Product Fact Sheet

Supported DataFlex Features:


- Complete support for all DataFlex code and database structures
- Support for overlap/partial overlap fields
- Complete support for case-sensitive indexes
- Complete support for mixed ascending and descending index segments
- Preservation of RECNUM values during data migration
- Complete support for transaction with rollback
- Exact replication of DataFlex table structure
- Support for DataFlex file relationships
- Alias file handling
- Support for DataFlex tools like Database Builder and VDF Studio

Product Specific Features:

- Comprehensive database migration tool that eases the job of migrating and managing tables and data
- Access to views and synonyms as DataFlex files
- Record locking instead of file locking
- Seamless integration of non-DataFlex tables into DataFlex applications. All necessary files like .FD and .INT files are generated automatically and the file name is entered into the filelist automatically
- Access to existing data stored in non-DataFlex database systems
- SQL Script generation for creating tables and indexes independently
- A sophisticated file creation and restructuring tool that allows you to commit changes directly to the backend or to generate DataFlex restructure code for distribution to remote sites
- A full index maintenance tool that allows addition and deletion of indexes and verification of indexes on the backend
- Extensive macro commands to help optimize data fetches without changing your DataFlex code
- Unified stored procedure interface that works across all drivers in the same fashion
- Comprehensive SQL command set for embedding SQL statements
- High-speed data migration using native tools like SQL*Loader, Bulk Copy, and DB2Load
- Use of native database transaction logging and rollback capabilities for disaster recovery
- A complete migration report that shows you at a glance the creation of tables, indexes, and the number of records migrated
- Storage analysis tool that allows you to evaluate storage availability before migration

Interaction with the DataFlex Runtime

The DataFlex (DF) runtime engine in Visual DataFlex (VDF) and DataFlex 3.1c and above allows external database drivers to be loaded as DLLs and work as a part of the runtime. The DataFlex database is also handled as a driver and is the default driver.



These changes in the runtime enable a driver to work with existing DataFlex programs without changing a single line of code. All that is needed is a login to the database (if

the database requires it). The login can be handled as a separate dialog which "chain waits" to the main program. The login remains effective for all other programs and "chain waits."

To give you an idea on how these pieces work together, let us look at a simple DataFlex program that opens a file.

Open MYFILE

The runtime engine calls the DataFlex API function call, `dfFileOpen()`, which searches the filelist entries for a file named MYFILE. It should be mentioned that if MYFILE is on a server, the entry in the filelist will be MYFILE.INT or of the type `ORA_DRV:\\servername\schema*tableName`. The "INT" extension means that there is an intermediate file. There is more information on this in the next section. The DataFlex API sees the .INT extension (or the embedded servername) and opens the .INT file. From this .INT file, the API looks for the driver name entry. If the driver was previously loaded, it searches for the driver from a driver table. If it was not previously loaded, then it loads the driver and adds it to the driver table. Once the driver has been loaded, the API calls the driver FileOpen code and passes the control to the driver.

The driver searches the system catalogs to find out if the table name (MYFILE) is a valid table name. If it is, the driver performs the operations that are necessary to open the file. Once the driver is finished it passes control back to the API, which returns control back to the runtime engine.

In a similar fashion, the driver handles all major functions like find, save, edit, delete, transactions, locking, etc. The API calls each of these functions for each specific driver.

The Intermediate File

Flex2SQL generates an intermediate (INT) file for each DataFlex file that is converted to the target database. The intermediate file contains information used by the DataFlex API and the database driver.

When a command in DataFlex is given to open the file with a root name containing an .INT extension, it is an indication to the API that the file is a non-DataFlex file. Instead of opening the DataFlex file, the API opens the intermediate file.

The API searches the .INT file for the DRIVER_NAME token. This tells the API which driver to load to handle the file being opened. The API then loads the appropriate driver and passes control to the driver to handle "opening" of the file.

Handling Key DataFlex Features

RECNUM or RowID Table Style

You can choose to use RECNUM or RowID table style during the migration process. Choosing RECNUM emulates existing DataFlex record processing where a unique numeric identifier is used to identify each record. You can elect to preserve existing DataFlex RECNUM values or let the SQL server assign RECNUM values starting with one.

As long as a DataFlex program accesses a table with the Merteck driver, if the RECNUM field exists, it is automatically incremented. If another program wants to access this table, the Flex2SQL migration utility provides an option to create a trigger on the table to handle the auto-increment feature.

Alternately, you can choose RowID table style to take advantage of new RowID VDF commands. This is the preferred method for new applications.

Locking

Merteck's drivers support record-level locking. Since DataFlex locking is based on files, this is an added benefit for using the client/server solution in medium to large multi-user environments. If a DataFlex program contains reread, the driver re-finds the record in the record buffer and locks the current record until the transaction is completed. Any record not explicitly locked can be updated by other users. This provides a great boost in concurrency over the native DataFlex database.

Transactions

Most relational dbms are configured to automatically start a transaction when save and update operations take place. In DataFlex terms, reread/lock/start_transaction commands all start transactions and unlock/end_transaction commit transactions. Transactions can be aborted or rolled back only by the abort_transaction command. Since most DataFlex programs use reread/unlock combination, there is no automatic provision to rollback transactions in case of an error. However, transaction rollback can be implemented either by using the DataFlex transaction commands or by using commands implemented by Merteck.

Overlaps

All Merteck drivers support overlap fields, even partial overlaps. An overlap field can be thought of as a virtual field which is comprised of one or more actual fields.

Merteck's drivers handle overlaps through the intermediate file. Flex2SQL generates all the necessary descriptions of the overlap fields so the programmer does not have to worry about how the overlap is generated. No change in the code is required to work with overlaps.

Merteck provides triggers that can be enabled to ensure that overlaps are taken care of if data is updated using non-DataFlex tools.

Relationships

An important part of good database design is reducing data redundancy and improving efficiency. To achieve this goal, you define each table so it only stores one kind of data and then define relationships between tables. Flex2SQL provides a utility to synchronize relationships between DataFlex tables and the SQL Server backend. Enforcing relationships on the SQL Server helps ensure data consistency.

Relationships are handled in the same way as overlaps, that is, through intermediate files, and work the same way as a DataFlex relationship does.

Restructuring

You can restructure and manage your database in a variety of ways. One way is to use Flex2SQL's restructuring dialogs and your own restructuring code. Another option is to make SQL Server the main repository for your database structure and use available SQL modeling tools to manage your database. Then, you can run Flex2SQL's Generate .INT File from Table/View/Synonym to bring your changes back into DataFlex.

Mertech also supports restructuring and management of tables using VDF Studio or Database Builder. You can use DataFlex as the master repository for your data and the Mertech drivers and DataFlex will ensure that all updates are reflected in the SQL backend, INT files, and any FD files.

Embedded SQL

Significant performance improvement can be seen by rewriting some batch processes in SQL. This can be either as embedded SQL statements, or as a stored procedure or function for execution directly on the server.

A simple example is a business process that computes the outstanding balance for all customers. In a record-oriented scenario, all the records have to be brought over to the client in a loop that computes a running total. This can instead be turned into a short snippet of SQL code that executes fully on the server, with minimal network traffic and interaction.

Mertech has developed a comprehensive SQL command set for embedding SQL statements that allow you to optimize your programs for faster data access. In addition, Mertech has developed a class layer that allows you to use the ESQL classes available in the Data Access sql.pkg. Mertech's class layer implements a unified ESQL interface across all of our drivers, allowing for applications to migrate to database servers DAW does not support, such as MySQL and PostgreSQL.

Other Features

Dynamic Primary Keys

This powerful feature addresses a shortcoming in VDF when it comes to using server-side generated Universally Unique Identifiers(UUID), sometime called Globally Unique Identifiers (GUID), as primary keys in database operations. This new feature gives developers the ability to create UUID columns with server-side UUID generation and use them seamlessly with VDF applications.

NULL Handling

DataFlex applications have no concept of an SQL NULL value. A NULL value is always mapped to an empty value in a DataFlex application. For example, if you have a DF_ASCII field that is assigned to a VARCHAR table column, and the column is NULL, DataFlex reads the DF_ASCII field as a blank ASCII string. Flex2SQL allows you to choose whether you want a non-index field to be NULLable when creating a new table or restructuring an existing table. If NOT NULL is selected, you can choose an appropriate default value for the field.

Azure

Windows Azure is Microsoft's application platform that gives you the option to deploy in the cloud. Mertech drivers provide transparent support to run your Microsoft Azure SQL Database on Azure. No code changes are required on the Flex2SQL side. Migration to Azure is not currently supported.

Failover Support

Microsoft introduced database mirroring to increase database availability. When mirroring is configured and a failure is detected in a principal server, a standby or mirror server can take over with no loss of committed data. To add support for SQL Server Mirroring, Mertech added the SET_MIRROR_SERVER command to the SQLFlex command set. To use the SQL Server Mirroring feature, developers add the SET_MIRROR_SERVER command to their login module. This identifies the mirrored server to the driver, which in turns negotiates communication with the server. If the principal server fails, the driver automatically connects to the mirror server and continues operation. No additional code is required in DataFlex to check for a lost connection.

Auto-Reconnect

Mertech's driver code that handles the reconnect after a mirroring failover is the same code used to handle the auto-reconnect after other disconnects. Previously, if a server timed out a client connection due to lack of activity, a DataFlex program encountered unexpected errors. With Mertech's auto-reconnect feature, the driver automatically reconnects to the server and continues from the same point where the disconnect occurred.

Supported Platforms

Mertech products support operating systems in which DataFlex or a targeted backend are available. The following table shows the availability of Mertech drivers on various platforms.

Operating System	Oracle	MS SQL Server	MySQL	PostgreSQL
	V9i and above	V2005 and above (including localdb)	V5.0 and above	V8.4 and above
Linux	X		X	X
SCO Open V	X		X	
Windows	X	X	X	X

Evaluate Our Product

You can download and evaluate our product free of charge for 30 days at www.mertechdata.com. After 30 days, all you need is a registration code to continue. All of our product documentation, white papers, and case studies are posted on our web site as well.

Contact Information

If you would like to know more about Mertech's products, please visit our website at www.mertechdata.com or contact us at the addresses below. Pre-sales consultations are free!

Corporate Head Office

Mertech Data Systems, Inc.
18503 Pines Boulevard, Suite 312
Pembroke Pines, FL 33029
USA
Tel: +1 (954) 585 9016
Fax:+1 (866) 228 1213

California Office

Mertech Data Systems, Inc.
114 East Shaw Avenue, Suite 209
Fresno, CA 93710
USA
Tel: +1 (954) 585 9016

Brazilian Office

Mertech Data Systems Brasil Ltda.
Av. Yojiro Takaoka, 4.384 – Cj. 2227, Alphaville
Santana de Parnaiba
SP 06541-038
Brazil
Tel: +55 (11) 4152 8433
Fax:+55 (11) 4152 8433
For more information e-mail:
mertechbrasil@mertechdata.com

Technical and Sales Support

e-mail: techsupport@mertechdata.com
e-mail: sales@mertechdata.com