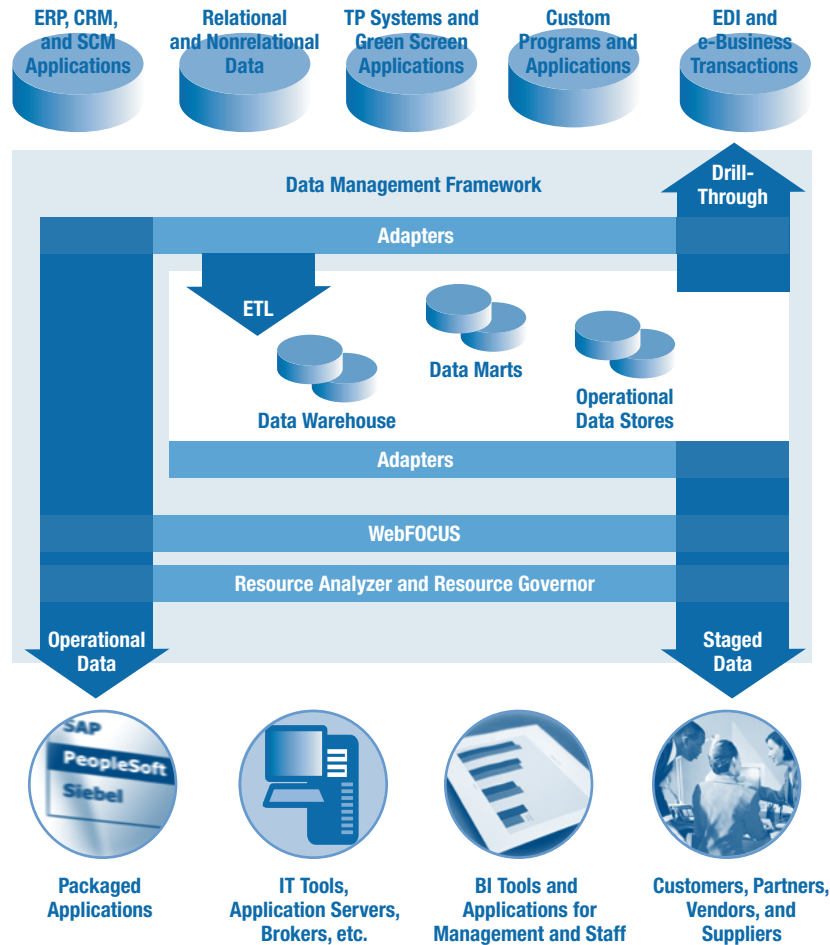


Information Builders is one of the largest independent software companies in the world. We work with over 12,000 customers to ensure their success by putting interactive and actionable information in the hands of everyone in the extended enterprise: employees, partners, suppliers, and customers.

Information Builders makes it easy to address a wide range of unique information needs throughout your enterprise by enabling you to design a powerful and flexible data architecture that fully leverages all your data assets.

Data Management Solutions



Rapidly Build a Powerful and Flexible Data Architecture

Companies today are faced with the challenge of managing more data, more applications, and more information consumers than ever before. In order to effectively leverage their enterprise information assets, organizations must implement a flexible data architecture that provides full support for the different access needs and latency requirements created by various IT initiatives such as enterprise business intelligence, real-time reporting and analysis, business activity monitoring (BAM), and new e-commerce and enterprise application integration projects.

Information Builders offers the broadest range of data management options available today – empowering you to provide direct access to data in operational systems, stage data in warehouses or marts, and allow dynamic drill-through to detail data –

so you can design your architecture to address the unique information needs of all your users. From a comprehensive adapter suite and powerful ETL tools, to intelligent resource analysis and governing solutions, Information Builders provides everything you need to develop an effective, yet economical strategy for accessing, storing, managing, and delivering all your enterprise information assets.

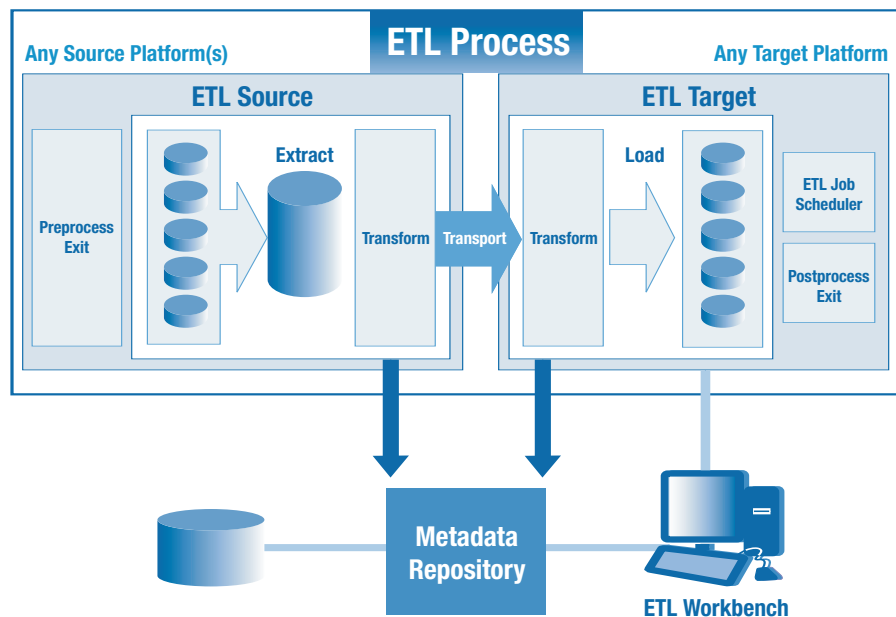
Direct Data Access

Codeless plug-and-play adapters provide fast, efficient, and reusable access to more critical information assets than any other vendor. Our adapters form a universal adapter framework that requires just a single skill set for anything-to-anything connectivity – any application, any platform, any data structure, any message type. Once acquired, adapters and adapter skills can be used and reused across a variety of IT initiatives such as EAI, application development, data warehousing, business intelligence, and more.

Robust Extract, Transform, and Load (ETL) Technology

DataMigrator is a powerful and comprehensive set of fully automated tools designed to dramatically simplify the creation, maintenance, and expansion of data warehouses, data marts, micro marts, and operational data stores. With its intuitive, easy-to-use interface, DataMigrator enables fast, flexible end-to-end ETL process creation involving heterogeneous data structures across disparate computing platforms and allows users to efficiently:

- Aggregate, join, merge, and apply selection criteria to information from any combination of back-office systems
- Simplify data movement from back-office systems to e-business platforms, using automatically generated and managed FTP scripts or most native transport protocols
- Transform data from raw forms into structures that are more suited to your application



- Simplify loading of data into a target database through the automatic invocation of bulk loaders or row-at-a-time inserts
- Execute, schedule, review, manage, audit, and create dependencies among ETL requests

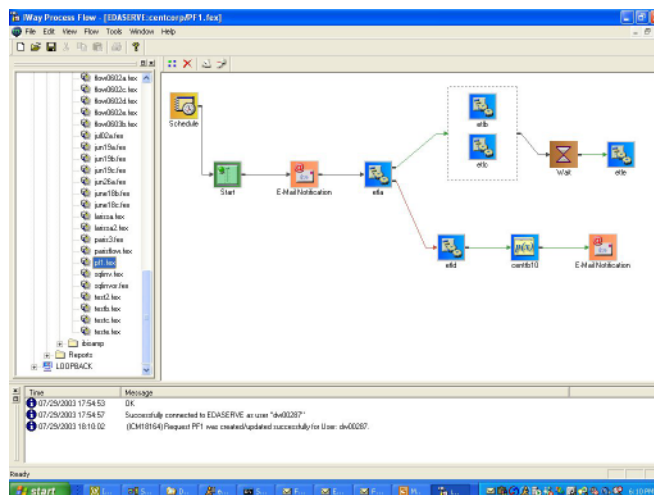
Unparalleled Data Access and Migration

DataMigrator provides greater reach and range than any other ETL offering on the market today. It employs our award-winning adapters to provide seamless access to and migration of virtually any enterprise data. DataMigrator can also utilize ODBC, JDBC, XML, and IBM WebSphere MQ connectivity for additional database access. And because the ETL Manager integration engine provides complete data structure transparency, all sources accessed by the servers on your network appear as a set of relational tables – so cross-platform, multiple-database extractions and joins are just a mouse-click away.

Exceptional Ease of Use

A comprehensive ETL Workbench with a point-and-click interface dramatically simplifies and accelerates the creation and execution of ETL jobs.

- **Data Flow Designer.** The intuitive flow-control interface allows for visual creation of joins, rapid table and column searches, and dynamic error notification. Additionally, real-time request testing can be easily and quickly conducted with just a click of the mouse. Users can also add COBOL routines, or routines under a C wrapper to the ETL function list, and load multiple tables on a single pass of an extract.
- **Process Flow Designer.** The ability to graphically develop elaborate process flows makes it easy to design complex jobs such as building large data warehouses with intricate load requirements. The execution of ETL requests can be fully controlled and can be based on a schedule defined by the administrator, or can be dependent on the success or failure of previously executed requests. Jobs composed in the Data Flow Designer, or with your own custom programs or scripts, can be launched at any point – or multiple points – in the process flow.



- **Web Console.** Remote review and administration provides efficient management of ETL jobs in a secure environment. Administrators and users can easily review job statistics and drill down into job logs, review the job queue, start and stop servers and server processes, schedule jobs and configure e-mail nodes, manage adapters and metadata, administrate application directories, and create and edit scripts and user exits.

Step-by-Step ETL Operation

Workflows of simple, yet powerful steps fully automate the process of accessing and migrating data:

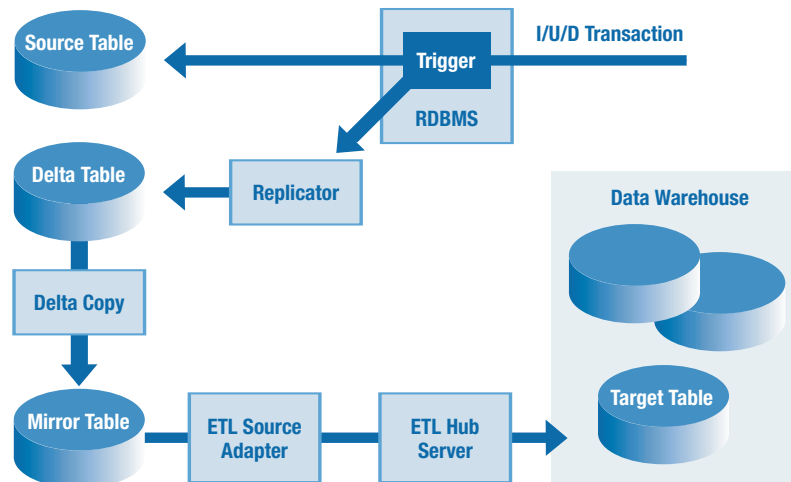
- **Extraction.** Users can efficiently select sources for extraction, join the data, and select columns for extraction. The transformation calculator can also be used for extract transformations or virtual transformations performed on the source platform.
- **Transformation.** DataMigrator is unique in its ability to define virtually any transformation and perform it as efficiently as possible, without custom extraction routines. It can easily perform transformations outside the limits of SQL and can incorporate routines containing IF-THEN-ELSE logic. It also allows for advanced transformation functions such as the ability to create 3GL subroutines, add them to a library of subroutines for reuse, and add comments to document their business use.
- **Loading.** Users can easily choose a target table or create one dynamically, and efficiently map source columns to target columns. DataMigrator simplifies data maintenance tasks and provides maximum flexibility for handling exceptions like duplicate inserts. As a result, users have the ability to reject records and update or delete records with the same keys for complete control over the entire load process.
- **Scheduling.** A platform-independent scheduler lets you specify when the ETL job runs. Additionally, ETL execution can be initiated by an external scheduler, application program, mainframe JCL, or platform-specific batch utility.

No other ETL product gives you this much flexibility: you can extract, transform, transport, and load data from one platform to another, scheduling aggregations to take place as soon as the ETL process is complete. You can even have DataMigrator automatically generate dependent data marts from your data warehouse as soon as the main warehouse has been populated. Convenience and power like that make data warehouse creation a well-synchronized effort, rather than a resource drain.

DataMigrator also provides full auditability of the ETL process. Each process generates a complete record of results, including a log of significant events that occurred during its run. For example, you can choose to view – in a message window at the bottom of your screen – either summarized messages such as whether the copy request has been submitted and was successful, or detailed messages for purposes such as debugging SQL statements. Job summaries and detailed logging are available from the Web Console.

Change Data Capture for Active Data Warehouses

Through our partnership with Lakeview Technologies, DataMigrator can be integrated with powerful change data capture (CDC) capabilities. CDC transfers only changed records (inserted, updated, or deleted), increasing labor efficiency by requiring only a single ETL job



to be composed to handle both initial load and delta processing, dramatically reducing job processing time and CPU cycles, and improving the timeliness of information in the data warehouse. Changed records can be pulled as frequently as every 5 minutes, so you can enable near-real-time or active data warehouses.

Enterprise-Wide Resource Performance and Usage Analysis

Knowledge really is power. When you know how your enterprise information is being used, you have the power to optimize it, secure it, and modify it to achieve consistent, outstanding performance. Resource Analyzer provides over 80 detailed reports and graphs to help you discover how your e-business applications use your information, including:

- Usage characteristics such as heaviest usage times, most requested content and aggregations, and applications using the most resources
- Performance characteristics such as average or maximum CPU consumption per request, average or maximum response time, number of records returned, and maximum cumulative CPU or response time per request

Resource Analyzer can help your e-business applications operate at peak efficiency for maximum return on investment. By identifying specific bottlenecks and inefficiencies in your environment, you can eliminate resource contention and:

- **Long-running requests.** These requests use excessive resources – a particular problem when dealing with operational systems – and result in long waits for consumers. Target these requests for optimization first.
- **Dormant data.** Unused data in e-commerce databases can make response times slower and batch updates more cumbersome. Eliminate it.
- **Excessive joins.** Joining too many tables is expensive and slow. Stage data from many tables into one for better performance and navigability.

Find Out More

To find out how our solutions can help your company succeed, talk to an Information Builders representative today. Contact your local Information Builders office, visit us at www.informationbuilders.com, or in the U.S. and Canada, call (800) 969-4636.

- **Repeated aggregations.** Aggregations can cause a bottleneck during query execution. Improve performance by precalculating the aggregations and staging them in summary tables.
- **Poor selection criteria.** Indexes increase selection speeds but impede update and insertion performance. Identify those fields that really need to be indexed and eliminate those that don't.
- **Times of heaviest workload.** All e-business systems experience peaks and valleys of demand. Manage peak loads and justify infrastructure expenditures by identifying your heaviest resource consumption times by hour, day, week, or month.
- **Query volumes.** Untuned queries frequently return more data than necessary. Identify these queries so you can target them for optimization.

Preemptive Query Governing

Resource Governor eliminates complex, long-running, or superfluous queries before they reach the database. Resource Governor rapidly and preemptively identifies problem queries by gathering a set of sophisticated metrics about each query. These metrics are created for every type of data source, regardless of underlying database or file type.

Based on usage and performance criteria and business rules that you define, the Resource Governor will cancel the query, issue a warning message, or let the query pass. Query cancellation is preemptive – it occurs before any resources are used to get data from any back-office system. And Resource Governor's intelligent engine allows it to know whether or not a query will break your business rules, even if it has not seen that exact query before.

Using our graphical wizards or unique Business Rules Language (BRL), you can define what a problem query is and solve problems such as:

- Long-running queries that interfere with operational systems
- Queries that do not use correct indexes
- Queries that use improper aggregation levels
- Queries without conditional predicates
- Queries originating from an unauthorized source



Corporate Headquarters Two Penn Plaza, New York, NY 10121-2898 (212) 736-4433 Fax (212) 967-6406
www.informationbuilders.com askinfo@ibi.com

Canadian Headquarters 150 York St., Suite 1000, Toronto, ON M5H 3S5 (416) 364-2760 Fax (416) 364-6552

Copyright © 2003 by Information Builders, Inc. All rights reserved. [21]

All products and product names mentioned in this publication are trademarks or registered trademarks of their respective companies.

DN7504292.0804

For International Inquiries +1(212) 736-4433



Printed in the U.S.A.
on recycled paper