

Sybase® ETL: Optimized Data Loading for Sybase IQ Reporting Environments



PRODUCT DATASHEET

In today's highly competitive marketplace, business intelligence and reporting are the lifeblood of ongoing operational success. Companies create reports for a wide range of purposes in order to better understand and act upon changing business events and conditions. However, query-intensive reports often compete with operational applications for computing resources, and these query workloads can frequently affect the performance of OLTP systems and Enterprise Data Warehouses. In such cases, IT organizations frequently offload critical data to separate report servers to support the reporting needs of the business.

Sybase IQ is the world's leading column-based analytics server, architected for reporting and analytics. Sybase IQ turbo-charges reporting environments by providing extremely high performance for feeding data to reporting tools and dashboards. For purposes of optimized loading of data into Sybase IQ report servers, Sybase also offers Sybase ETL for bulk and incremental data loading. Sybase ETL is a highly efficient tool that provides quick return on investment by delivering fast, reliable data movement from source systems into Sybase IQ by extracting the source data, transforming it to meet specific business rules, and loading it into Sybase IQ with support for all existing Sybase IQ loading interfaces.

Sybase ETL has been specifically designed to extract data from heterogeneous source systems (e.g. Sybase ASE; IBM DB2; Oracle; Microsoft SQL Server; MySQL; Text Files; XML files) and load Sybase IQ as its only target. It utilizes javascript-based transformation components that are easily customized to meet the needs of report server creation, and runs on a scalable grid architecture that aligns well with Sybase IQ's multiplex grid. It supports all the IQ bulk load interfaces and provides bulk mode Upsert and Delete capability that is specifically optimized for IQ to address the incremental load need. These capabilities, along with the fact that it is easy to develop and deploy, and is offered at an affordable price point, make Sybase ETL very attractive to Sybase IQ report server users who need fast time to production, are faced with tough budget constraints, and are looking for a good alternative to developing custom-coded loading scripts.

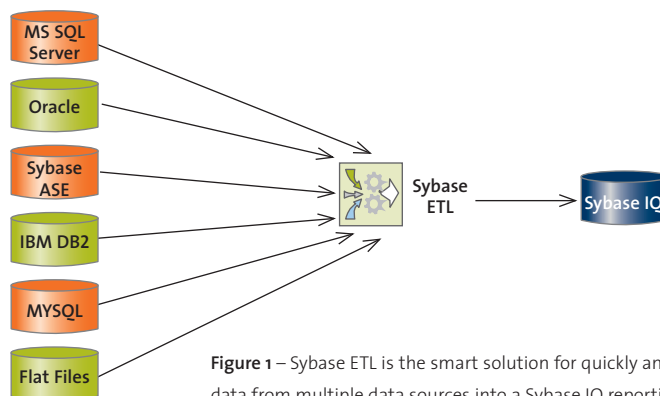


Figure 1 – Sybase ETL is the smart solution for quickly and efficiently moving data from multiple data sources into a Sybase IQ reporting environment.

SYBASE ETL: IDEAL FOR DEPARTMENTAL DEPLOYMENTS

Sybase ETL is ideal for departmental or division level Sybase IQ reporting applications where row-based operational databases are being offloaded, moderate amounts of data schema transformations are required, and batch or incremental loading into Sybase IQ meets the reporting needs. Sybase ETL provides Sybase IQ reporting customers with an excellent opportunity to replace custom-coded loading tools which are often difficult to maintain, and to improve their loading performance into Sybase IQ. As noted previously, Sybase ETL can extract data from all the major source systems that Sybase IQ customers typically use for reporting purposes, and comes with a migration wizard for migrating existing applications into Sybase IQ.

SYBASE ETL: NEW FEATURES AND BENEFITS

For a very reasonable price that fits into departmental budgets, the latest release of Sybase ETL offers a rich and robust set of new features and capabilities that provide outstanding price/performance, low total cost of ownership, and rapid return on investment in conjunction with Sybase IQ reporting environments:

- **Scalability** – because of the Sybase ETL grid architecture, it can support multiple writers in the Sybase IQ multiplex grid. This structure enables the parallel loading of multiple tables which provides an almost linear performance improvement as scaling is needed – administrators just need to plug in a new ETL grid node when needed.
- **Speed** – Sybase ETL provides for bulk and incremental data loading procedures that are optimized for Sybase IQ. Loading data incrementally typically involves updating existing records, inserting new records, and deleting old records. The ability to execute these operations in bulk mode speeds up their performance. Sybase ETL adds this capability to Sybase IQ environments.
- **Usability** – user friendly GUI interface for quick development and deployment of ETL projects. Support for Sybase IQ client side file loading for situations where Sybase ETL and Sybase IQ are running on different servers. There is no need to create mounted file systems or FTP the final data over before loading into Sybase IQ—client side file loading is now accomplished via standard Sybase IQ loading commands.
- **Manageability** – Sybase ETL provides the ability to monitor and control grid nodes from a standard web browser, including viewing grid node status, viewing and downloading system logs, and monitoring running jobs and projects. Administrators can pause, restart, and cancel running jobs, view job execution history, and shutdown grid nodes if needed.
- **Security** – Sybase ETL provides security and compliance support for FDCC (Federal Desktop Core Configuration) and IPV 6. FDCC is a list of security settings recommended by the national institute of standards and technology for general purpose PCs that are connected directly to networks of US government agencies. IPV 6 is a next generation IP protocol that offers larger address spaces, address auto-configuration, and multicasting.
- **Heterogenous support** – Sybase ETL is supported on most of the popular hardware and operating system platforms, including most variants of Unix, Linux and Windows. It is also i8N enabled.

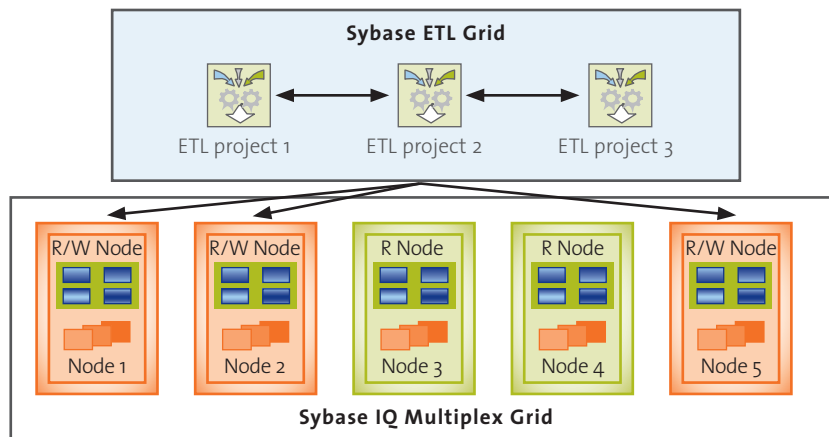


Figure 2 – Parallel loading of multiple tables provides a tremendous performance boost for data loading.