



SAS® Enterprise Data Integration Server

A complete solution designed to meet the full spectrum of enterprise data integration needs

What does SAS Enterprise Data Integration Server do?

SAS Enterprise Data Integration Server, featuring DataFlux technology, is a powerful, configurable and comprehensive solution that can:

- Access virtually all data sources.
- Extract, cleanse, transform, conform, aggregate, load and manage data.
- Support data warehousing, migration, synchronization and federation initiatives.
- Support both batch-oriented and real-time Master Data Management solutions.
- Create real-time data integration services in support of service-oriented architectures.

Why is SAS Enterprise Data Integration Server important?

SAS Enterprise Data Integration Server enables organizations to manage data integration projects on an enterprise scale in a timely, cost-effective manner and meet the high data quality expectations of information consumers.

For whom is SAS Enterprise Data Integration Server designed?

It is designed for organizations in all industry sectors that are implementing one or more data integration projects, dealing with changing business landscapes and business-driven IT initiatives, trying to meet regulatory requirements, or implementing data governance.



**THE
POWER
TO KNOW®**

SAS Enterprise Data Integration Server provides a comprehensive solution to handle the challenges of distributed and rapidly increasing data volumes, inconsistently defined data across disparate IT systems, and the high expectations of data consumers across the enterprise who depend on data to be correct, complete and available when they need it. It does this in a timely, cost-effective manner, ensures credibility and consistency, and gives organizations the ability to efficiently manage data integration projects on an enterprise scale.

Key benefits

- **Leverage/reuse work by others.** A common repository enables the centralized storage, management and reuse of work based on user authorizations, reducing both development and maintenance time.
- **Work collaboratively in groups or teams.** Leveraging the common repository, multi-user development capabilities are supported. The advanced GUI environment and centralized metadata repository capture and manage work in a way that is easy to see and understand.
- **Improve productivity.** Easy to learn and work in, the GUI environment provides a standard interface for building and documenting work. Collaboration is encouraged and manual coding, though available when needed, is minimized. New team members can get up to speed quickly on work done by others, which is important when documentation is inadequate or missing.
- **Always access the data you need.** From older legacy systems to the latest ERP applications, data from virtually any hardware platform or operating system can be accessed and processed. New source systems can easily be added and security managed centrally. This saves time, shortens learning curves and gives decision makers the information they need.
- **Manage security and administration at all levels.** Reusable templates make it quick and easy to provide role-based authorizations and administrative privileges at the user, department or enterprise level.
- **Meet time constraints even within decreasing windows of availability.** SAS processes data fast! Organizations can take advantage of the grid-enabled load-balanced, multithreaded parallel processing architecture that can quickly process, transform and move data between different platforms and systems.
- **Deliver consistent, trusted and verifiable information.** Consistently getting correct data when and where it is needed provides increased confidence in the accuracy and timeliness of operational and business information. Data quality auditing tools monitor the quality of data in processes and source systems. Users can see where data originated and how it was transformed. Optional enrichment components can add value and ensure the best possible data.
- **Eliminate overlapping, redundant tools and systems with one solution.** SAS offers the only comprehensive enterprise data integration solution that is built from the ground up to meet the full spectrum of data integration needs. It eliminates the piecemeal approach of linking and managing technologies from different vendors, and provides lower overall cost, reduced risk and faster results.

Product overview

SAS offers the only comprehensive enterprise data integration environment that is built from the ground up to meet the full spectrum of enterprise data integration needs. Instead of linking and managing technologies from different vendors, SAS Enterprise Data Integration Server provides a collaborative design environment promoting object reuse and sharing, administrative controls, wizard-driven design process workflow, and ease of use and maintenance. This flexible, reliable solution can access data from any system in any form, transform and cleanse data even in real time, and handle data migration, synchronization and federation projects all through a versatile services environment that is easy to deploy and maintain.

Interactive data integration development environment

A graphical user interface (GUI) provides developers with an easy-to-use, point-and-click environment that uses an intuitive set of configurable windows for managing multiple data integration development processes.

Enterprise connectivity

Connectivity is the foundation of enterprise data integration. SAS Enterprise Data Integration Server provides connectivity to virtually all types of data sources and types, operating systems and hardware environments using both native access and open standards. It also supports the reading and writing of data from message queues and the sending and receiving of data to and from Web services.

Data quality

Providing clean, credible and correct information is one of IT's foremost charges. From removing duplicate records to standardizing master data, SAS provides an enterprise approach that lets you create and share data rules and processes between projects and across the entire data integration solution. With SAS Enterprise Data Integration Server, you get a seamless environment for profiling and rules creation, data integration processing and results monitoring.

Extraction, transformation and loading (ETL)

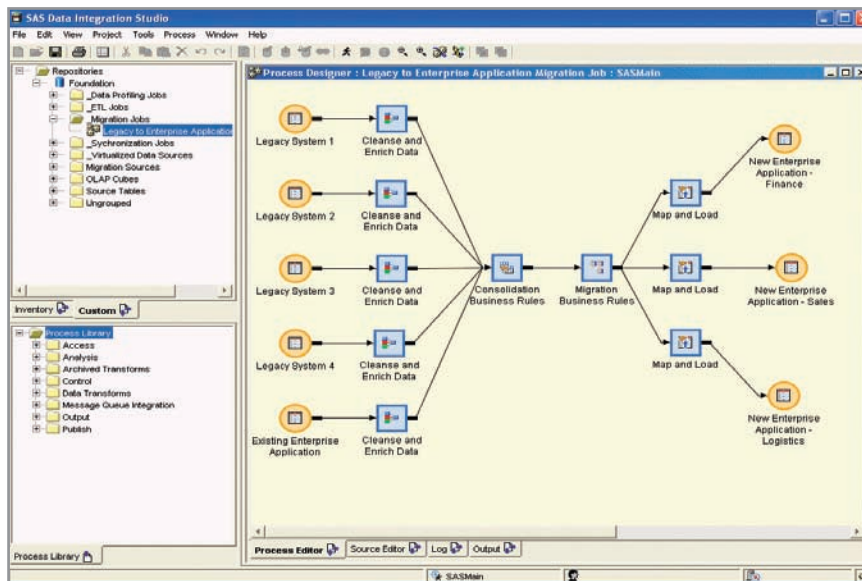
Loading data warehouses and data marts within their allotted time windows, quickly building analytical marts for special projects, and creating extract files for reporting and analysis applications are tasks IT organizations face each day. SAS Enterprise Data Integration Server includes an intuitive point-and-click design that allows developers to easily build logical process workflows, read and reuse the metadata from the input and output data stores, and create business rules and transformations in metadata, thus enabling the rapid generation of reusable data integration services.

Metadata management

SAS provides a shared metadata environment that is both independent (for data integration) and part of SAS' comprehensive platform. Technical, business, process and administrative metadata is stored and managed in a way that leverages and facilitates reuse of existing table definitions, business rules and more. Shared metadata provides a consistent definition across data sources to speed integration projects, simplify design and reduce maintenance costs.

Migration and synchronization

Moving data from system to system is a constant activity in most organizations. Mergers and acquisitions result in multiple, overlapping systems containing information that often needs to be synchronized and ultimately migrated. Moving legacy data during upgrades and conversions is an ongoing process, as is the movement of data into and out of ERP systems. SAS Enterprise Data Integration Server provides the capability to migrate, synchronize and replicate data across different operational systems and data sources. The point-and-click process design editor makes it easy to docu-



SAS Enterprise Data Integration Server provides the capability to cleanse, enrich, migrate, synchronize and replicate data across different operational systems and data sources. The point-and-click process design desktop makes it easy to document processes in workflows that can be deployed in a variety of ways and executed in real time or batch.

ment migration and synchronization processes in workflows that can be reused and modified for other projects. Powerful data transformations are available for altering, reformatting and consolidating information during these processes. Also, it is easy to reuse business rules ensuring that bad data is never spread from system to system. In this way, information delivered across all applications, systems, environments and geographies is up-to-date, consistent and accurate.

Data federation

Fast access to the most current operational data drives the success of functions with real-time reporting needs, such as call centers and fraud detection activities. In addition, business users need self-service reporting across multiple data sources. SAS Enterprise Data Integration Server provides the ability to query and use data across multiple systems without the physical reconciliation or movement of source data. By avoiding unnecessary data replication and movement, it is possible to quickly and cost-effectively deliver up-to-date data that is consistent and accurate.

Integrated support services

The features and capabilities provided by SAS Enterprise Data Integration Server are available as an array of supporting services. These supporting services, as part of the SAS®9 platform, all leverage the integrated SAS metadata environment and administrative functions. In addition, with SAS Enterprise Data Integration Server's built-in wizards, users can easily create their own data integration services that can be called in real-time, near real-time and batch modes.

Key Features

Interactive data integration development environment

- Wizards for accessing source systems, creating target structures, importing and exporting metadata, and building and executing data access, transformation and load process flows.
- Dedicated GUI to profile data and identify and repair source system issues, while retaining the business rules for later use in ETL processes.
- Multi-user design environment supports collaboration on large, enterprise-wide projects.
- Ability to distribute data integration tasks to nearly any platform and to connect virtually any source or target data store.
- Integrated workflow scheduling, automatic load balancing and grid computing support.

Enterprise connectivity

- Enterprise Applications: Oracle Applications, PeopleSoft, SAP BW, SAP R/3, Siebel and more.
- Database sources: AS400, ODBC, IBM DB2/UDB, Informix, Microsoft Access, Microsoft Excel, Microsoft SQL Server, MySQL, Netezza, DATAlegro, HP NeoView, OLE/DB, Oracle, Sybase, Sybase ASIQ, SAS data sets, SAS Scalable Performance Data Server, SAS Scalable Performance Data Engine, Teradata, and more.
- Mainframe data sources (OS/390 and z/OS): ADABAS, CA-Datcom, CA-IDMS, COBOL, IBM DB2, IMS-DL/I, ISAM files, Oracle, SYSTEM 2000, Teradata, VSAM (KSDS and ESDS), and other file formats.
- File formats: CSV, XLS, Access, WKS, text/flat files, XML, COBOL Copybooks, and FTP and URL-based sources. Reads and writes data representations such as ASCII, Binary, EBCDIC, Hexadecimal and Octal.
- Support for Message-Oriented Middleware, including WebSphere MQ from IBM, MSMQ from Microsoft and Tibco's Rendezvous.
- Support for unstructured and semi-structured data to parse and process files.
- Access to static and streaming data for sending and receiving via Web services.

Data quality

- Data quality embedded into batch, near-time and real-time processes.
- Data quality rules are callable through message queues, Web services and custom exits.
- Data cleansing provided in native languages with localizations for more than 20 worldwide regions.
- Out-of-the-box standardization rules, or you can build customized rules for special situations.
- Metadata built and shared across the entire SAS Enterprise Intelligence Platform provides an accurate trail of actions.
- Add value to existing data by generating and appending postal addresses, geo-coding, demographic data or facts from other sources of information (DataFlux option).
- Data stewards can profile operational data and monitor ongoing data activities with an interactive GUI designed specifically for their needs.
- Simple process for institutionalizing data quality business rules. Apply basic or complex rules to validate data according to the specific business requirements of a particular process, project or organization.
- Data quality monitoring. Continuously examine data in real time and over time to discover when quality falls below acceptable limits. Issues alerts when corrective action is needed.

Extraction, transformation and loading (ETL)

- Comprehensive data transformation library with more than 300 predefined table and column-level transformations.
- Author your own transforms using wizard for metadata driven transformations
- Transformations are reusable and are tracked and registered in metadata.
- Publish/Subscribe services for publishing information to archives, a publishing channel, e-mail or various message queuing middleware.

(Continued on next page)

SAS® Enterprise Data Integration Server

Client environment

Clients include SAS Data Integration Studio, SAS Management Console, and DataFlux components: dfPower Quality, dfPower Customize and dfPower Profile. Additional DataFlux data cleansing modules are available.

- Windows (x86-32): Windows 2000 Professional, Windows XP Professional
- Internet Explorer 5.5+

Server environment

SAS Servers, including Base SAS and SAS Metadata Server, can be installed on one or more hardware systems in a multi-tier configuration.

- AIX (64-bit), Release 5.1+
- HP-UX PA-RISC, Release 11i+
- HP-UX Itanium, Release 11i+
- Linux for Intel (x86-32): Red Hat Linux 8.0, RHAS 2.1, RHEL 3.0 and 4.0, SuSE 8 and 9
- Linux for Itanium (64-bit): Red Hat RHEL 3.0
- Solaris on SPARC, Version 8, 9, 10
- Solaris on x64, Version 10
- Windows (x86-32): Windows NT 4 Server, Windows 2000 Server, Windows Server 2003, Windows XP Professional
- Windows (64-bit on Itanium): Windows Server 2003
- z/OS, Version 1

Required/optional software

- SAS includes a reference implementation of Apache Tomcat.
- SAS client and mid-tier require JRE 1.4.2_09.

Key Features (continued)

- Metadata is captured and documented throughout transformation and data integration processes and is available for immediate reuse.
- Transformations can run on all supported platforms with all data sources.
- Flexible deployment. Transformation and data integration processes can be deployed easily as embedded business logic or as Web services for use by other applications as part of a service-oriented architecture.

Metadata management

- Sophisticated metadata mapping technologies for quickly propagating column definitions from sources to targets and for creating automated intelligent table joins.
- Impact analysis for assessing the scope and impact of making changes to existing objects such as columns, tables and process jobs.
- Data lineage (reverse impact analysis) is critical for both validating processes and building user confidence in data.
- Multi-user collaboration support includes object check-in and check-out and DEV/TEST/PROD promotion management.
- Wizard-driven metadata import and export.
- Metadata-driven deployment flexibility. Using process metadata, process jobs can be deployed for batch execution, as reusable stored processes and as Web services.

Migration and synchronizations

- Metadata-driven access to sources and targets.
- Extensive library of predefined transformations for migration and synchronization that can be extended and shared with other integration processes.
- Embedded, reusable data quality business rules can clean data as it is moved.
- Ability to migrate or synchronize data between database structures, enterprise applications, mainframe legacy files, text, XML, message queues and a host of other sources.
- Optional integrated scheduler allows changes made in one or more systems to be propagated on a scheduled basis out to other systems in a reliable manner.
- Data can be cleansed in real time as it is being moved, replicated or synchronized.
- In addition to SAS' own change data capture (CDC) capabilities, when combined with technology from partner Attunity, enhanced CDC recognizes changes to key fields and replicates/synchronizes changes across multiple heterogeneous databases. CDC capabilities provided by Oracle and DB2 are also supported.

Data federation

- Access to database structures, enterprise applications, mainframe legacy files, text, XML, message queues and a host of other sources.
- Ability to join data across data sources for real-time access and analysis.
- Instant access to a real-time view of the data using the built-in data viewer.
- Query optimization is provided both automatically as part of DBMS requests, and manually within the advanced SQL editor, and can be used for both homogenous and heterogeneous data sources.



THE
POWER
TO KNOW.

SAS Institute Inc. World Headquarters +1 919 677 8000

To contact your local SAS office, please visit: www.sas.com/offices

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies. Copyright © 2007, SAS Institute Inc. All rights reserved. 102451_398534.0507