

IBM WebSphere DataStage for z/OS, Version 7.5.1

Highlights

- ***Offers a complete integration platform that runs under UNIX® System Services on the IBM® z/OS® platform***
- ***Brings the speed of IBM WebSphere® DataStage® parallel processing to z/OS***
- ***Accesses IBM z/OS UNIX files, QSAM data sets, VSAM files, IBM DB2® software and Teradata***
- ***Includes a powerful graphical user interface to design extract, transform and load (ETL) processes that run natively on the mainframe***
- ***Extends easily to support enterprise data quality***
- ***Ports easily to and from the IBM AIX® operating system***

Enterprise data integration

Mainframe computers continue to play a key role in many enterprises. They process huge volumes of data and represent the most reliable platform to run corporate applications for day-to-day business functions.

What is needed is a data warehouse. But integrating information from multiple sources with different formats onto a mainframe environment can be difficult. Hand coding programs to integrate the data is expensive and requires maintenance for many resources. Hand-coded routines may not be fast enough to meet the processing requirements. If the warehouse needs to be ported to a distributed environment, all of the integration work must be re-created. The key to success in this situation is to have a data integration platform that can handle data in its native form, leverage a powerful parallel processing framework and provide a graphical, codeless design tool.

IBM WebSphere Information Integration

The IBM WebSphere® Information Integration platform integrates and transforms any data and content to deliver information you can trust for critical business initiatives. It provides breakthrough productivity, flexibility and performance, so your company and your customers and partners have the right information for running and growing their businesses. It helps companies understand, cleanse and enhance information, while governing its quality to ultimately provide authoritative information. Integrated across the extended enterprise and delivered when users need it, this consistent, timely and complete information can enrich business processes, enable key contextual insights and inspire confident business decision making.

An integration solution for z/OS

IBM WebSphere DataStage is the extract, transform and load (ETL) component of the WebSphere Information

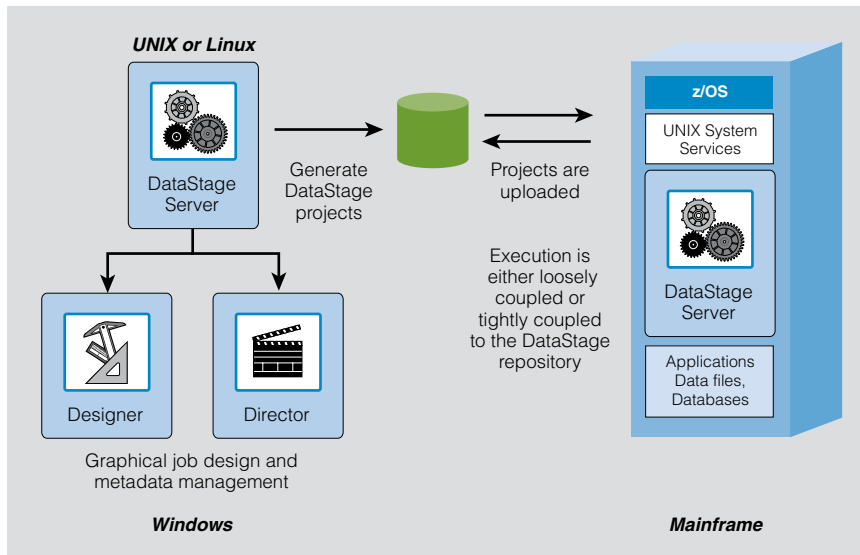


Figure 1: WebSphere DataStage for z/OS architecture

Integration platform. WebSphere DataStage for z/OS is specifically designed to run on UNIX System Services running under z/OS. This offering makes it easy to integrate and to enrich mainframe data to create a centralized data store for Business Intelligence. In addition, it helps integrate and stage mainframe data so that data can be moved to a distributed platform and integrated with other enterprise data.

Advanced development and maintenance

WebSphere DataStage for z/OS features a powerful architecture—shown

in Figure 1—that gives developers maximum speed, flexibility and effectiveness in building, deploying and managing their data integration infrastructure. The productivity-enhancing features in WebSphere DataStage reduce learning curves, simplify administration and optimize the use of development resources, resulting in a decreased development and maintenance cycle for data integration applications. The bottom line is that WebSphere DataStage enables companies to spend less time developing their integration and more time reaping the benefits of using the data.

Complete development environment

WebSphere DataStage employs a “work as you think” design metaphor. Developers use a top-down data flow model of application programming and execution, which allows them to create a visual, sequential data flow.

A robust graphical palette helps developers diagram the flow of data through their environment via simple GUI-driven drag-and-drop design components. WebSphere DataStage features one of the industry’s most extensive data integration development environments with a library of prebuilt functions and routines. This complete set of data transformation capabilities makes it easy to map data from source to target and enrich it along the way.

A scalable platform to satisfy the most demanding environments

WebSphere DataStage enables companies to solve large-scale business problems through high-performance processing of massive data volumes. By leveraging the parallel processing capabilities of multiprocessor hardware platforms, WebSphere DataStage for z/OS can scale to satisfy the demands of ever-growing data volumes and ever-shrinking batch windows.

WebSphere DataStage for z/OS helps to minimize the time-processing requirements. By fully leveraging the parallel processing capabilities, it can linearly increase speed of data throughput for integrating massive amounts of data. It also boosts developer productivity by eliminating the need to code new applications to run in parallel—a costly process that often requires the expertise of specialists. Developers use sequential logic, and the deployment configuration automatically adds the desired degree of parallelism.

Comprehensive connectivity to source and target systems

WebSphere DataStage for z/OS supports numerous heterogeneous data sources and targets in a single job, including:

- *IBM z/OS UNIX files*
- *QSAM data sets*
- *VSAM (ESDS, KSDS and RRDS)*
- *IBM DB2 software*
- *Teradata*

Flexible deployments

The development environment for WebSphere DataStage for z/OS runs on a Microsoft® Windows® client connected to a UNIX or Linux® DataStage server. Job compilation and execution is processed using either a connected, tightly coupled model or a loosely coupled mechanism.

The tightly coupled mode allows jobs to be designed, compiled and run under the control of the WebSphere DataStage clients. The WebSphere DataStage repository provides login and monitoring information. Remote shell (rsh) and FTP are used automatically to connect to the mainframe. The loosely coupled mode does not require remote shell server to be enabled on the mainframe. Job scripts can be automatically sent to the mainframe using FTP or can be sent manually. The jobs are then compiled on the mainframe and executed using command-line interfaces or a mainframe scheduler. Job monitoring and logging information is not returned to the WebSphere DataStage repository in this mode.

Portability to and from AIX

Since WebSphere DataStage for z/OS utilizes UNIX System Services to run on z/OS, all job scripts and integration work can be easily ported to the IBM AIX platform. This helps to ensure that the data warehouse can be moved to AIX without re-creating the integration routines developed for z/OS, if this is required at a future date. It also means a data warehouse that has outgrown the AIX platform can easily migrate the data integration routines to z/OS.

For more information

To learn more about WebSphere DataStage for z/OS or other IBM WebSphere Information Integration offerings for the mainframe or other enterprise platforms, contact your IBM marketing representative or IBM Business Partner, or visit ibm.com/software/data/integration



© Copyright IBM Corporation 2006

IBM Software Group
Route 100
Somers, NY 10589
U.S.A.

Printed in the United States of America
September 2006
All Rights Reserved

IBM, the IBM logo, AIX, DataStage, DB2, WebSphere and z/OS are trademarks of International Business Machines Corporation in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. Offerings are subject to change, extension or withdrawal without notice.

All statements regarding IBM future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.

The information contained in this document is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this document, it is provided "as is" without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this document or any other documents. Nothing contained in this document is intended to, nor shall have the effect of, creating any warranties or representations from IBM Software.

TAKE BACK CONTROL WITH **Information Management**