



IBM Software Group

IBM® WebSphere® Extended Deployment V6.1 for z/OS®

Compute grid overview



@business on demand.

© 2007 IBM Corporation
Updated June 19, 2008

This presentation will give an overview of WebSphere Extended Deployment Compute Grid Version 6.1.

This module references WebSphere Extended Deployment Operations Optimization, which is now called WebSphere Virtual Enterprise.

Though the module uses the previous names, the technical material covered is still accurate.

Compute grid overview

- WebSphere Application Server has traditionally focused on transactional applications
 - ▶ Applications typically are designed to handle large volumes of relatively small tasks
 - ▶ Not all applications fit this type of design
- Long-running work often requires different capabilities
 - ▶ Work Manager (asynchronous beans) and message-driven beans support some of these needs



Applications designed to run on WebSphere Application Server tend to be transaction-oriented applications, made up of several relatively short-lived tasks. However, not all types of work fit this model. WebSphere Application Server has some features that can be beneficial to long-running applications, such as the work manager and message-driven beans for running asynchronous tasks in the EJB container. WebSphere Extended Deployment expands on these features by providing a full set of capabilities for running and managing long-running work.

Compute grid capabilities

- Compute Grid provides further support for these needs
 - ▶ Asynchronous submission and execution
 - ▶ Greater separation of submitter and execution environment
 - Supports work being submitted from outside the WebSphere environment
 - ▶ Work is persisted to highly available and non-volatile data stores
 - ▶ Administrative controls to monitor and manage units of work
 - ▶ Capability to prioritize and schedule work

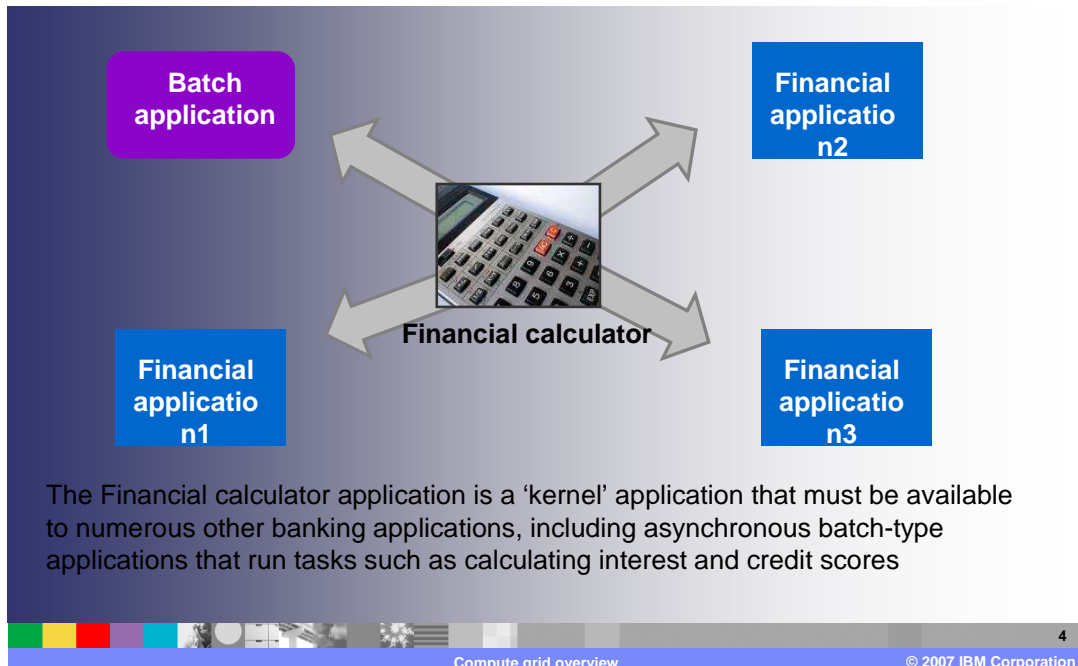
Compute Grid enhances WebSphere Application Server to support long-running applications

- ▶ Provides capabilities to deploy long-running applications along with transactional applications
- Compute Grid balances the work across different nodes within a cell
- Fully integrated with traditional batch processing

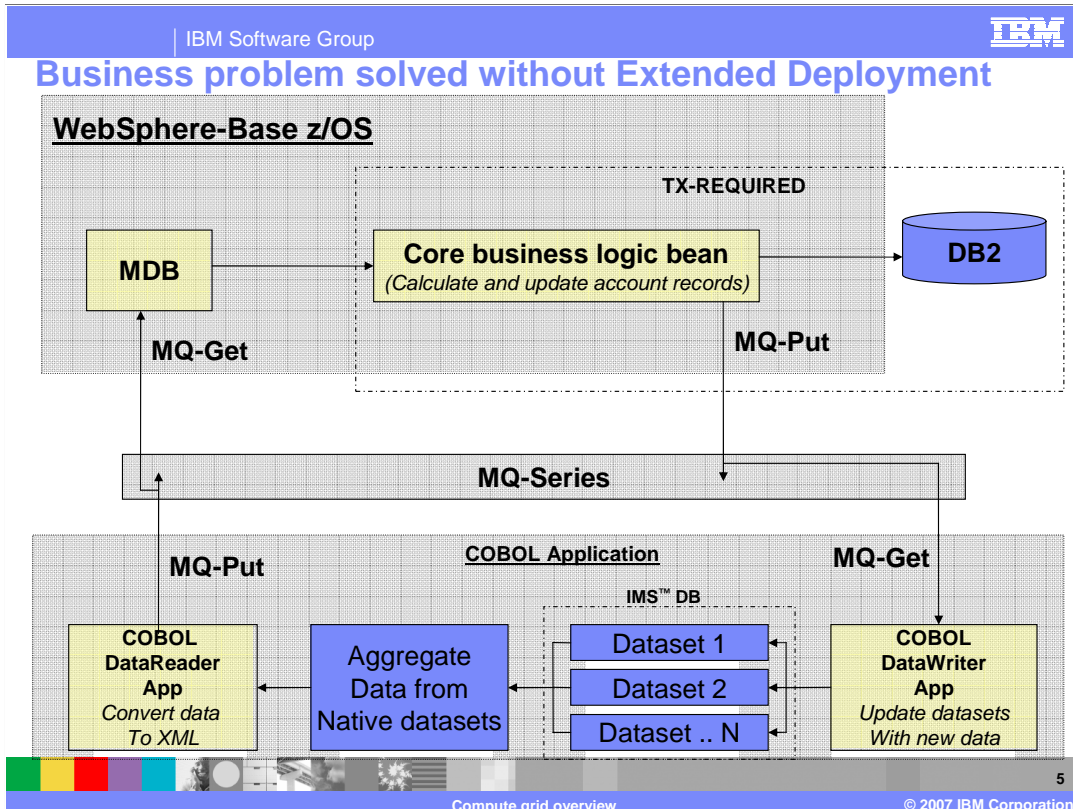


Compute Grid expands the types of work that WebSphere Extended Deployment is designed to handle. Compute Grid adds to WebSphere Extended Deployment the capability to run and manage long-running applications in the same cell as your transactional applications. WebSphere Extended Deployment dynamically manages both workload types and allocates server resources as necessary. This feature includes a special execution environment for long-running work, a separate job submission and scheduling environment, and the administrative tools to monitor and manage long-running work. Compute grid batch processing can be integrated as a job step in traditional batch processing. This feature allows you to place a J2EE™ batch step among conventional batch steps such as COBOL or CICS® processes.

Compute grid example: Financial application



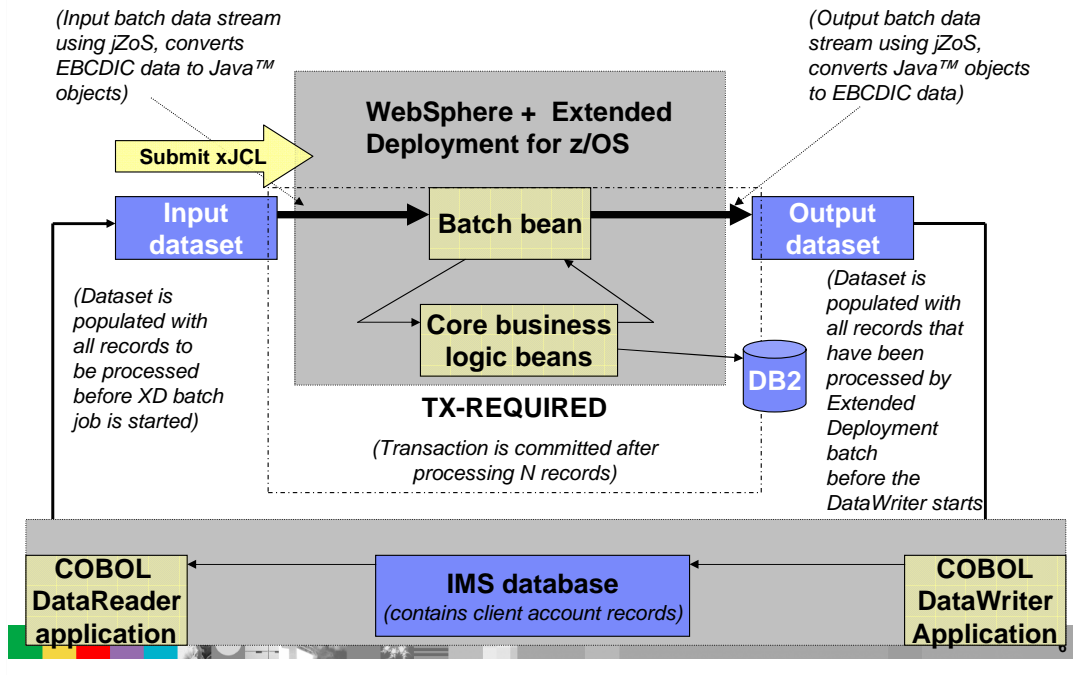
A company often has multiple pieces of business logic within its set of applications. Some of the applications contain duplicate business logic and customers would like to keep a single copy of the business logic. As an example, you may have some core piece of business logic that needs to be completed, such as a credit score rating or financial calculator. To reduce maintenance costs you would like to share the logic across applications and the batch workloads. Without WebSphere Extended Deployment, you have few choices of how to accomplish this.



Until WebSphere Extended Deployment Compute Grid was delivered, there was no direct method to provide J2EE-centric batch processing. WebSphere z/OS customers typically utilized a design pattern based on WebSphere MQ that provided some batch-like behavior. However, the major issue with this design pattern was the overhead of the infrastructure relative to the actual business processing for each record. For example:

- The data must be converted from EBCDIC to XML;
- A global transaction is required to place a record onto the queue;
- Global transaction is required to retrieve the record from the queue by a Message-Driven Bean.
- The record is then passed internally from the Message-Driven Bean to the business logic for processing.
- A global transaction is then used to place the resulting data onto the queue.
- A process outside of WebSphere also requires a global transaction to read the data from the queue.
- The data is finally converted from ASCII back to EBCDIC and inserted back into the native database.

Business Problem solved with Extended Deployment

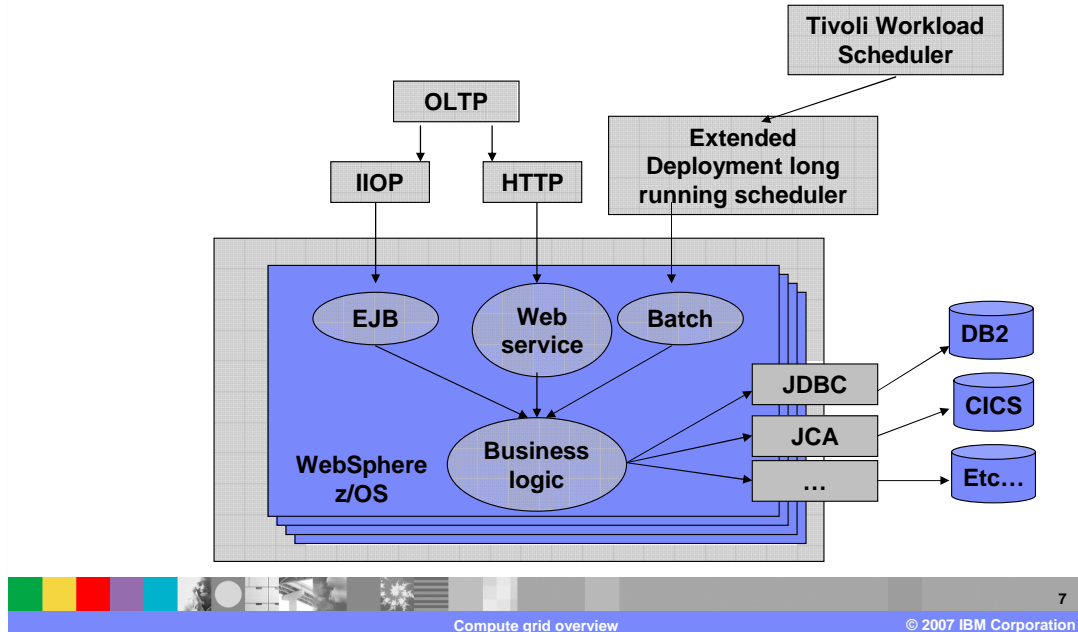


WebSphere Extended Deployment Batch provides a J2EE-centric batch processing environment. This environment allows the business logic to directly access native data and therefore eliminate much of the infrastructure overhead introduced in the WebSphere MQ-based solution. Performance data has shown that this solution can reduce both the MIPS and processing time by approximately 50%. Additionally, most of the MIPS consumed were zAAP-eligible, therefore significantly reducing the overall cost of processing a single record.

The Compute Grid solution provides many advantages :

- Flexibility of adjusting the checkpoint algorithm without having to modify any application code. This in turn affects the transaction scope and resource locking schemes.
- Flexibility of changing data sources by defining new Batch Data Streams and isolating the application changes to specific sections of the code, namely the Batch Data Stream definitions.
- Eliminating the need for intermediaries to message the data. Compute Grid is able to parse the raw data in the Batch Data Stream definition and convert that raw data straight to Java™ objects, as opposed to converting data to XML, or ASCII.
- Ability to recover from the last check-point position in case of system failures.
- Ability to temporarily suspend or cancel jobs in case of resource contention.
- Ability to assign service policies to jobs.
- Simplified system administration of the batch environment since Compute Grid integrates with WebSphere Administrative Console to provide a view of batch jobs in the system.
- Ability to quickly create a conditional flow of job steps using xJCL.
- Ability to use health management features of WebSphere Extended Deployment to proactively manage failures, such as memory leaks, by enabling automatic restart of jobs after recovering from a server failure.

Running mixed workloads- OLTP and batch



The interesting point here is that Extended Deployment allows the business logic to be independent. It can be used either in transactional (OLTP) or batch.

WebSphere Extended Deployment for z/OS enables the processing of these long-running J2EE workloads to occur within a shared infrastructure together with transactional workloads without negatively impacting transactional requests. This enables operating efficiencies to be realized when transactional and long-running workloads are consolidated into a single WebSphere environment. A unified environment provides for common development, administration, and management models, further decreasing costs and increasing operational stability.

Types of long running work

- WebSphere Extended Deployment V6.0 supports two styles of long-running work
- Batch applications focus on doing large amounts of work based on a specific task, for example record processing
 - ▶ Application will provide logic for a single unit of work (process one record)
 - ▶ Container provides support for transactions and mechanisms checkpoint and restart work
- Compute-intensive applications focus on large amounts of processor-bound work
 - ▶ Container provides thread of execution and has limited contact with the work after it is started
 - ▶ Application provides all other logic
- Native applications
 - ▶ Java “main” applications that run in an independent JVM
 - ▶ Native applications (C, C++, FORTRAN,...) running as their own process



Long-running work can generally be divided into two types, both of which are supported by WebSphere Extended Deployment. A batch application is designed to repeatedly perform some specific unit of work, such as processing a database record, over a large set of data. If an application provides the logic for performing the work, WebSphere Extended Deployment can manage the transaction support and life cycle aspects, such as restarting the work at the point it was stopped in the case of an interruption. Compute-intensive applications are also long-running tasks with little user interaction, but focus on performing a large amount of work that tends to be processor-bound, such as large-scale simulation work.

WebSphere Extended Deployment version 6.1 supports running compute intensive work within a WebSphere application server, and outside of WebSphere as native executable. When run within an application server WebSphere Extended Deployment provides a runtime environment and a thread of execution, and all other logic is implemented by the application. Compute intensive work outside of WebSphere can consist of Java programs or native executables.

Compute grid packaging

- WebSphere Extended Deployment Compute Grid V6.1 is an 'add-on' product for middleware platforms
- Extension to WebSphere Network Deployment
- Adds new components to WebSphere



WebSphere Extended Deployment Compute Grid is installed as an extension to WebSphere Network Deployment, and adds components to the base WebSphere system to enable you to run and manage long-running and compute-intensive J2EE applications in a WebSphere Extended Deployment environment. When installed in conjunction with Operations Optimization, the Compute Grid component integrates with that package's dynamic operations features to enable a virtualized, dynamic environment for goal-oriented workload management, and to take better advantage of hardware resources in a distributed environment.

Summary

- Compute grid
 - ▶ An environment for running and managing compute-intensive and batch style J2EE applications



In summary, WebSphere Extended Deployment Compute Grid is an add-on product that provides several advanced features for WebSphere Application Server. The Compute Grid package provides an environment for running and managing compute intensive and batch style J2EE applications within a WebSphere cell.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject= Feedback about XD61z_CG_Overview.ppt



You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

CICS DB2 IBM IMS WebSphere z/OS

EJB, J2EE, Java, JDBC, JVM, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.