This presentation will give an overview of WebSphere Extended Deployment version 6.1.

This module references

- WebSphere Extended Deployment Data Grid, which is now called WebSphere eXtreme Scale; and
- 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.

Agenda

- Introduction to WebSphere Extended Deployment
- Compute grid
- Operations optimization
- Data grid

This presentation will introduce WebSphere Extended Deployment and how it is packaged in version 6.1, then it will introduce the major features of WebSphere Extended Deployment.

Section

Introduction

This section will introduce the WebSphere Extended Deployment product.

Overview

- A set of extensions for your existing middleware systems focused around quality of service, performance, resiliency, and manageability
  - Integrates with standard z/OS features
- Three main focus areas
  - Compute grid for batch and compute intensive workloads
  - Operations optimization with support for application virtualization
  - Data grid for high performance transaction processing
- Features can be “turned on” incrementally.

WebSphere Extended Deployment is an add-on product for WebSphere Application Server that provides advanced features for your environment. WebSphere Extended Deployment adds features in the following areas:

Compute Grid functionality enables you to run and manage long-running and compute-intensive J2EE applications in a WebSphere Extended Deployment environment designed for those types of work.

Dynamic operations features enable a virtualized, dynamic environment for goal-oriented workload management, to take better advantage of hardware resources in a distributed environment. It provides several advanced features for enhanced administration, including the ability to manage and seamlessly roll out multiple editions of an application. Operations visualization tools help you better understand what is happening in your environment, health monitoring agents can automatically take action when certain software conditions are detected, and more.

The partitioning facility is a programming framework and runtime environment for implementing high-volume transactional applications that scale linearly with additional hardware. This is done by dividing a large environment into smaller, uniquely addressable ‘partitions’ to increase efficiency and reduce contention. ObjectGrid is a new high-performance, transactional, and customizable cache framework for Java™ objects.

WebSphere Extended Deployment integrates well into z/OS and compliments many of the standard features such as request routing, work flow management and high availability.

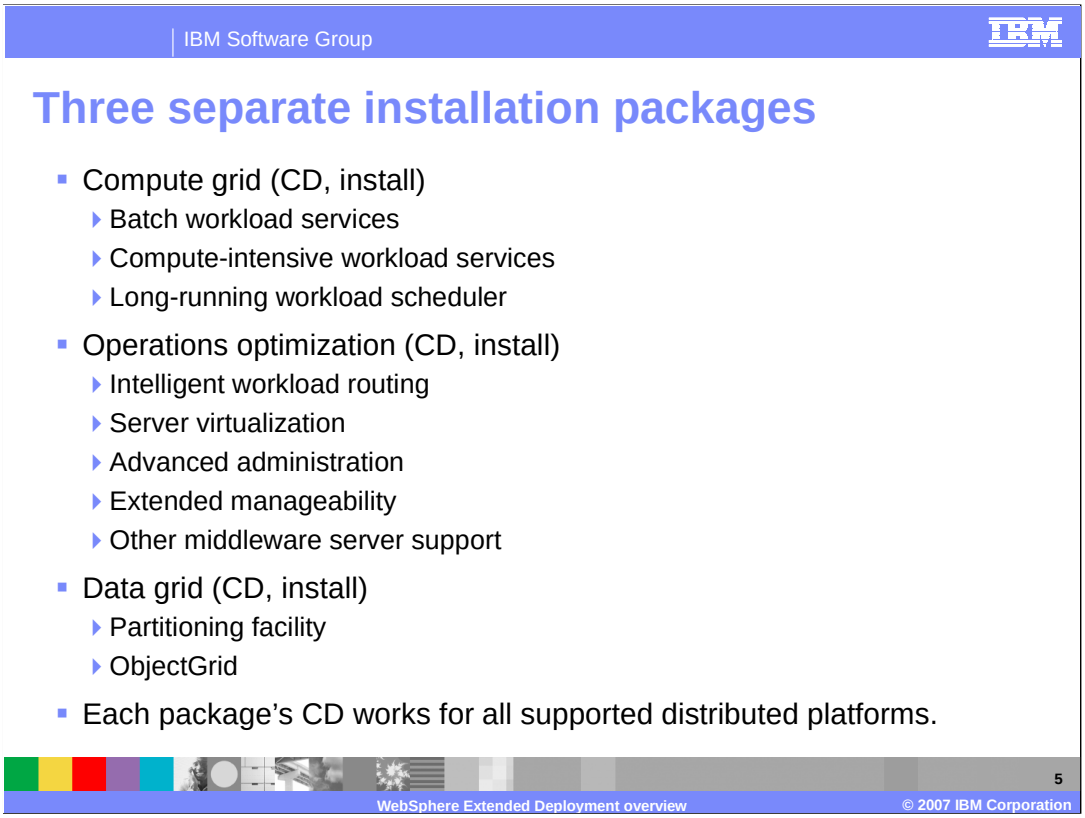

WebSphere Extended Deployment version 6.1 is composed of three separate packages:

Compute Grid extends WebSphere application servers to allow long running work to run within the WebSphere J2EE environment.

Operations Optimization contains many performance, scalability, and administrative enhancements, and extends the WebSphere administrative domain to include support for middleware servers on nodes not running WebSphere Extended Deployment.

Data Grid enables high-volume data-intensive applications.

Each of these packages is provided on a separate CD and must be installed separately.

New concepts and features

- Platform consistent
- Integrated into z/OS platform
- WebSphere business grid
  - Goal-based on-demand autonomics
  - J2EE batch programming
- Enhanced dynamic operations environment
  - Enhanced multi-platform workload management and additional classification rules
  - Consistent workload management model
  - Improved multi-system workload distribution
  - Workload balancing for WebSphere messaging

WebSphere Extended Deployment is built from a single shared common code base which assures platform portability and consistency between z/OS and other distributed platforms. This also ensures applications and skills are portable across platforms. At the same time WebSphere Extended Deployment for z/OS has been integrated into the z/OS platform. WebSphere Extended Deployment for z/OS leverages platform Quality of Service features such as Sysplex, WLM, RRS, RACF (SAF), SMF, and RMF. WebSphere Extended Deployment for z/OS optimizes local access to z/OS data resources. In addition, WebSphere Extended Deployment brings several interesting new features to WebSphere for z/OS. J2EE batch is part of WebSphere business grid and parallels MVS batch. The dynamic operations enhances Sysplex distributor and z/OS WLM in interesting ways.

Overview

- WebSphere Extended Deployment V6. 1 is an ‘add-on’ product for WebSphere Application Server V6.1
  - Installation is a simple delta to an existing environment
  - Concepts and features can be “turned on” incrementally.
  - Integrates seamlessly into a network deployment environment
  - Integrates with standard z/OS features
  - UNIX

WebSphere Extended Deployment is an add-on product for WebSphere Application Server version 6.1 that provides advanced features for your environment. The installation is a simple delta to your existing environment and does not require “migration” or restructuring of your current installation. WebSphere Extended Deployment has many new functions such as J2EE batch and dynamic horizontal scaling. These features can be selectively incorporated into your system on a schedule suitable to your organization. WebSphere Extended Deployment extends the administration console and scripting, simplifies application management, facilitates the creation of routing rules, and enhances the performance of WebSphere. WebSphere Extended Deployment integrates well into z/OS and compliments many of the standard features such as request routing, work flow management and high availability. Each of these topics will be discussed briefly in this presentation, and covered in more detail in the following presentations.

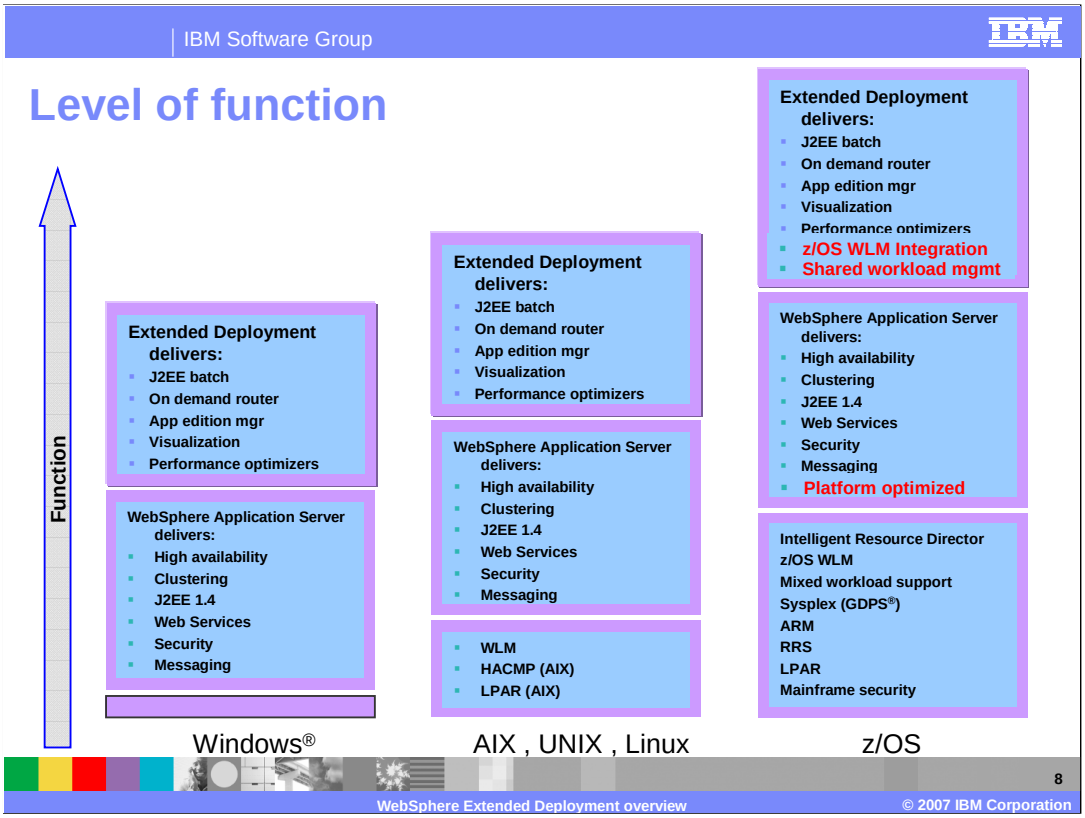

WebSphere Extended Deployment enhances all the WebSphere platforms. These platforms each have their basic capabilities and offer different levels of quality of service.

WebSphere Application Server brings J2EE programming.

WebSphere Extended Deployment enhances WebSphere base.

WebSphere Extended Deployment for z/OS brings forward the basic z/OS features.

Section

Compute grid

This section will introduce WebSphere Extended Deployment Compute Grid.

Compute grid overview

- WebSphere Application Server
  - Traditionally focused on transactional applications
  - Applications typically designed to handle large volumes of relatively small tasks
  - Not all applications fit this type of design

WebSphere Application Server has evolved some features that can be beneficial to long-running applications, such as the WebSphere work manager and message-driven beans for running asynchronous tasks in the EJB container. Still, applications designed to run on WebSphere Application Server tend to be transaction-oriented applications, made up of several relatively short-lived tasks.

Compute grid

- Compute-intensive applications focus on large amounts of processor-bound work
- Batch applications focus on doing large amounts of work based on a specific task, for example record processing
- Non-WebSphere applications

Long-running work can generally be divided into three types, all of which are supported by WebSphere Extended Deployment.

Compute-intensive applications are long-running tasks that focus on performing a large amount of work that tends to be processor-bound, such as large-scale simulation work. In this case WebSphere Extended Deployment provides a runtime environment and a thread of execution, and all other logic is implemented by the application.

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.

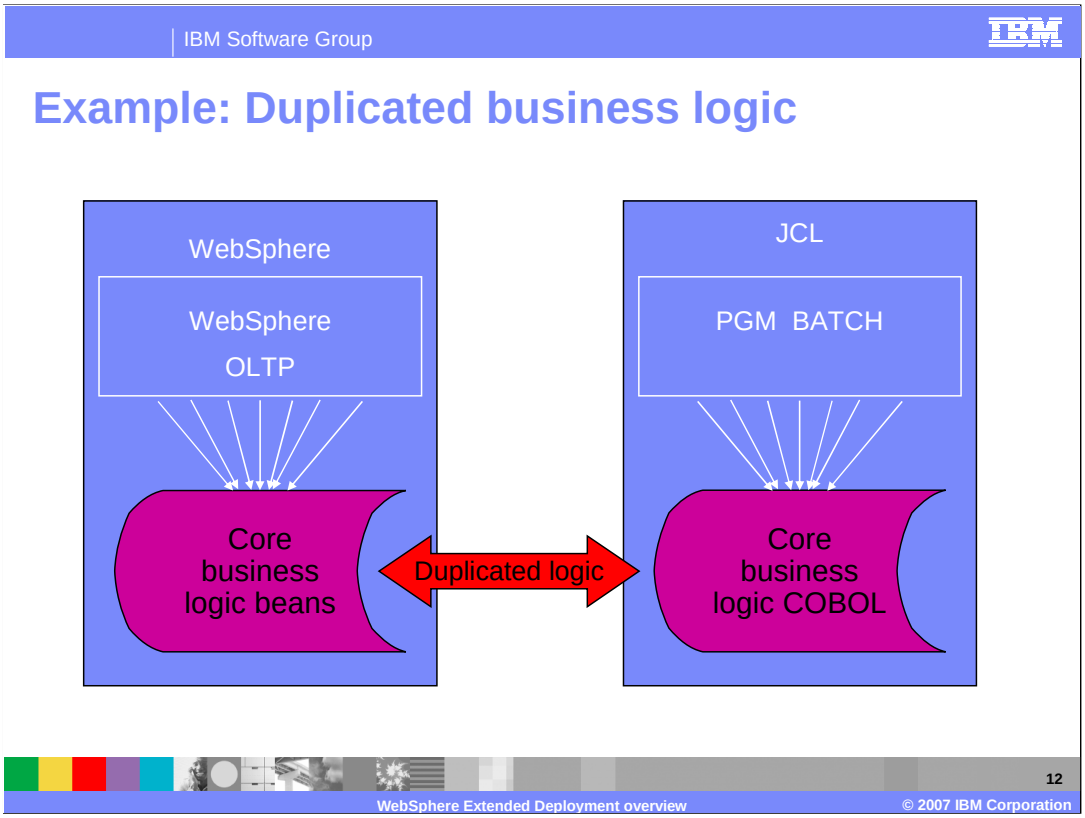

There is strong incentive to reuse core business logic wherever possible. The logic could be for calculation of sales commissions, account interests, stocking or pricing policies, stock market financials or any number of other functions. These functions are often used in both On Line Transaction Processing (OLTP) and batch processing. Often the business logic starts off coded in COBOL and duplicated into Java for the OLTP processing. Real problems can occur just ensuring the functional results are up to date and consistent. J2EE batch offers a solution to this problem. The same core business logic beans can be used for J2EE batch and OLTP. One solution would be to move the batch processing entirely over to J2EE batch. This has the draw back that the investment in the COBOL before and after the batch processing is lost. For example, the data that is used to create the batch input stream and interpret and save the batch output might already be coded in COBOL. In this case, the supporting COBOL can be pretty much left alone and with all the batch processing performed by J2EE batch. In this scenario the COBOL program would generate the batch input, pass it off to a J2EE batch and then catch the output and write the results. This could be attempted with WebSphere without Extended Deployment by substituting OLTP calls for the batch piece, but all the batch characteristics such as check pointing would be lost.

Section

Operations optimization

This section will introduce WebSphere Extended Deployment Operations Optimization.

Operations optimization: Dynamic operations

- Dynamic workload management
  - Virtualized
  - Goal-oriented
  - Policy-based
- More consistent quality of service for critical applications
- More efficient utilization of hardware resources
- On-demand router
- Dynamic application placement

WebSphere Extended Deployment provides a virtualized environment that allows hardware resources to be shared and allocated dynamically based on performance goals that you have defined for your WebSphere applications. This capability can help you provide a more consistent quality of service for your critical applications in times of excessive load, and more efficiently utilize the hardware resources that you already have.

Advanced administration features

- Advanced administration features ease management of large environments
- Visualization tools
- Edition control center
- Health monitoring tools
- Extended repository service
- High availability deployment manager

WebSphere Extended Deployment includes many features that make life easier for the administrator of a WebSphere environment. The visualization components enhance the Administrative Console to provide live data on the performance and health characteristics of the entire cell. The Administrative Console also features the Edition Control Center, which enables management of multiple versions of an application, including interruption-free rollout. WebSphere Extended Deployment also has health monitoring capability, and can take remediation steps when software health problems are detected. The extended repository service provides the ability to create checkpoints of your WebSphere cell configuration, providing you with administrative “undo” capability. High availability deployment manager environment eliminates the single point of failure for administrative functions.

Support for non-WebSphere environments

- Support for other runtimes
  - BEA, Tomcat, JBoss, WebSphere Community Edition, .NET, …
- Complete Extended Deployment dynamic operations capabilities
  - Health management and monitoring
  - Application placement for resource control

The on-demand router can route traffic to any HTTP endpoint. This means that you can use it as a front end to applications running on other vendor's servers, including Tomcat, or .NET; or other versions of WebSphere such as WebSphere Application Server Community Edition. This provides you with all of the benefits of the on-demand router's request prioritization and flow control features.

Section

Data grid

This section will introduce WebSphere Extended Deployment Data Grid.

Data grid

- Partitioning facility allows you to 'partitioned' an application across multiple servers.
  - A partition is a unique endpoint within a cluster
  - Requests for certain data or certain Enterprise Java Beans (EJBs) are always routed to the same partition
- ObjectGrid
  - A high performance highly scalable cache framework for storing Java objects

Data Grid is composed of two primary components. The partitioning facility is an extension to WebSphere that provides a programming framework and runtime environment that makes it possible for high transaction-volume applications to scale linearly by adding hardware capacity. To accomplish this, an application is partitioned across multiple servers in a cluster. Each partition is a uniquely addressable endpoint within the cluster, to which requests for certain EJBs or certain data are always routed. Partitioning solves some of the traditional challenges of very large clustering, because it can reduce data contention and reduce the overhead of replicating shared data, like caches or state information.

ObjectGrid provides a high-performance, transactional cache framework for storing Java objects. An ObjectGrid can be used as a generic object cache, and optionally persisted to a hardened store. In the same fashion, it can also be used as a local cache for objects stored in a database. ObjectGrid is a highly customizable feature, with interfaces provided for custom data loaders, invalidation and size management schemes, and more. ObjectGrid can be used independently of the rest of WebSphere Extended Deployment.

Summary

- WebSphere Extended Deployment is an ‘add-on’ product for WebSphere Application Server that provides many advanced features
- An environment for running and managing compute-intensive and batch styleJ2EE applications
- Advanced administrative capabilities
  - Administrative visualization provides new ways to look at performance data about your environment
  - Application edition management features enable seamless rollout of multiple editions of an application
  - Health monitoring can trigger alerts and take action when software health problems are detected

In summary, WebSphere Extended Deployment is an add-on product that provides several advanced features for WebSphere Application Server. The Compute Grid provides an environment for running and managing compute intensive and batch style J2EE applications within a WebSphere Extended Deployment cell. The advanced administration features provided by WebSphere Extended Deployment include advanced visualization tools for monitoring the performance of your environment, the Edition Control Center for the management and seamless deployment of multiple editions of an application, and health monitoring features that can detect software health problems and take actions on your behalf, and much more.

Summary

- Data grid
  - The partitioning facility provides a means to create a highly scalable environment for high-volume transaction processing
  - ObjectGrid is a high-performance, transactional, and extensible cache framework
- Dynamic operations
  - Dynamic operations features create a virtualized, goal-based environment for workload management

The Data Grid package contains the partitioning facility, which provides a programming model and runtime environment for implementing highly scalable solutions for transaction processing, and ObjectGrid which is a high-performance, transactional, and extensible object cache for Java applications.

Lastly, the dynamic operations features enable a virtualized, dynamic environment for better utilizing shareable hardware resources based on user-defined performance goals.

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_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:

AIX GDPS IBM WebSphere z/OS

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

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

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

UNIX is a trademark of The Open Group 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 NONINFRINGEMENT. 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.

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.