




IBM Software Group

# IBM® WebSphere® Extended Deployment V6.1

## *Overview of new features*



@business on demand.

© 2007 IBM Corporation  
Updated June 12, 2008

This presentation provides an overview of the new features in version 6.1 of WebSphere Extended Deployment.

This module references

WebSphere Extended Deployment Operations Optimization, which is now called WebSphere Virtual Enterprise, and

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

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

## Agenda

- New packaging model
- Enhancements
  - ▶ Operations optimization
  - ▶ Compute Grid
  - ▶ Data grid



This presentation will first cover changes to packaging for WebSphere Extended Deployment Version 6.1, followed by an overview of changes to the major functional areas.

## New packaging model

- Version 6.1 has new name and packaging
- Operations optimization
  - ▶ Workload routing (HTTP, SOAP, JMS)
  - ▶ Workload “traffic shaping” (policy-based)
  - ▶ Virtualization
  - ▶ Operations dashboard
  - ▶ Application placement
  - ▶ Health management
  - ▶ Application edition management
- Compute grid
  - ▶ Batch workload services
  - ▶ Compute-intensive workload services
  - ▶ Long-running workload scheduler
- Data grid
  - ▶ Partitioning facility
  - ▶ ObjectGrid



The WebSphere Extended Deployment Version 6.1 product is divided into three functionally independent packages, providing more choices to address your business challenges. In addition to the full product containing all features, you have the option of ordering the following separately chargeable components:

**Operations optimization** provides application server virtualization, resource management, and a host of advanced operational facilities, such as performance visualization, health monitoring, and application version management. This combination of capabilities is sometimes referred to collectively as dynamic operations.

**Compute grid** offers flexible support for running and controlling batch-like and compute intensive jobs.

**Data grid** provides high end caching and transaction partitioning capabilities through the ObjectGrid, and the partitioning facility.

This division of functions enhances WebSphere Extended Deployment’s flexibility for different usage scenarios, allowing more options for how you deploy in your production environment. While these packages can be installed individually, they complement each other when used in combination.

## Section

# ***Operations optimization server management***



This section describes server management changes to WebSphere Extended Deployment dynamic operations. These changes are incorporated in the operations optimization package.

## Other middleware server support

- Extended Deployment V6.0 supports non-WebSphere environments already
  - ▶ Mostly for advanced routing, and traffic control and monitoring
- V6.1 expands the non-WebSphere support to include most critical Extended Deployment functions
  - ▶ Application placement through dynamic clusters
  - ▶ Health management
  - ▶ Application version handling (in some instances)
- V6.1 unifies the management model into a new system called middleware servers and applications
  - ▶ Basic server management (start, stop, status) from WebSphere administration console
- Support for other runtimes
  - ▶ BEA, tomcat, JBoss, Geronimo, WebSphere Community Edition, .NET, ...



WebSphere Extended Deployment 6.0 provided some support for non-WebSphere servers through generic server clusters. This support was limited to HTTP request routing and charting for a subset of available performance metrics.

Version 6.1 lets you manage all of the servers in your environment through the WebSphere Extended Deployment administrative domain, including servers that are not created within WebSphere Extended Deployment. You can easily create representations of servers that were created outside of the WebSphere administrative domain and the Web applications deployed on those servers.

The version 6.0 mixed server environment monitoring agent is replaced by a middleware agent, which provides enhanced management and monitoring capabilities for other middleware servers. This includes the ability to start and stop the servers, monitor server status, and monitor an expanded set of performance and health metrics.

This enables version 6.1 to provide the full range of dynamic operations capabilities previously available only within the WebSphere Extended Deployment environment.

## Complete management for PHP servers

- Production management
  - ▶ Servers
    - Create, remove, start, stop, status, configuration management
  - ▶ Applications
    - Deployment, replication and management
    - Application version management
  - ▶ Automated on-demand router configuration
- Quality of service enhancements
  - ▶ Dynamic operations support
    - Dynamic clusters
    - Health monitoring



Beyond the additional support provided for other middleware servers, WebSphere Extended Deployment 6.1 provides further enhancements for external PHP servers running as an extension to Apache HTTP server. This includes the complete set of server and application life cycle management capabilities previously available only for WebSphere Application Server managed servers.

In WebSphere Extended Deployment version 6.1 you can create PHP servers and dynamic clusters, you can deploy PHP applications, and fully manage these applications from WebSphere Extended Deployment.

## Changing the virtualization model

- Extended Deployment 6.0.2 virtualization model based on pre-defined resource pools
  - ▶ Create a **node group** containing a set of machines
  - ▶ Define a **dynamic cluster** that lives on that group
- Extended Deployment 6.1 uses an expression-based resource pool
  - ▶ Node groups not necessary
  - ▶ Dynamic cluster node membership based on an expression
    - Example: *Node\_Name like "xdblade\*\*"*
  - ▶ Membership rule builder part of define dynamic cluster task

A dynamic cluster is similar to the familiar concept of a 'cluster' from WebSphere Application Server, but can be resized dynamically within a node group. As demand for applications running on a dynamic cluster increases or decreases, instances of that dynamic cluster can be started or stopped on nodes within the bounding node group to accommodate the changes in load. Each node in the bounding node group has a configured instance of the dynamic cluster that is ready to be started dynamically when needed.

In versions 5.1 and 6.0, dynamic clusters could only be created on nodes within a statically defined node group. Version 6.1 eliminates this requirement, allowing you to specify a *membership policy* during dynamic cluster creation. This policy specifies node properties to evaluate against all of the nodes in the cell. The membership expression can specify node name, host name, node group membership, or node custom property values. Complex expressions can be created using Boolean operators AND, OR, and NOT.

## Isolation policies for dynamic clusters

- Dynamic clusters can be defined with isolation
- Isolation policies restrict which dynamic clusters can run simultaneously on a given node
  - ▶ No isolation
    - Dynamic cluster's instances can run on the same node as instances of a different dynamic cluster
  - ▶ Strict isolation
    - Only a single dynamic cluster's instances can run on any given node
  - ▶ Group isolation
    - An explicitly grouped list of dynamic clusters can run on any given node

8

Overview of new features

© 2007 IBM Corporation

Dynamic cluster isolation allows you to isolate applications from other applications that are deployed in the cell. This new dynamic cluster parameter specifies the isolation requirements that the application placement controller applies when determining where an instance of the dynamic cluster should be placed. There are three options when configuring the dynamic cluster isolation requirements:

**“No isolation requirements”** is the default isolation policy. This policy tells the application placement controller that instances of the dynamic cluster can be co-located with any other running process when placed on a node. This is the same behavior seen in earlier versions of WebSphere Extended Deployment.

**“Strict isolation”** tells the application placement controller that when an instance of the dynamic cluster is placed on a node, it must not be co-located with a running instance of any other dynamic clusters. In other words, it can only be co-located with other vertically stacked instances of itself.

**“Associate with isolation group”** is a convenient way to specify a collection of one or more dynamic clusters whose running instances can be co-located with each other.

Isolation policies take precedence over placement decisions made for handling demand.



## Server maintenance mode

- Server maintenance mode provides the capability to isolate a running server (of any type) from production traffic
  - ▶ Allows for problem determination or other maintenance without disruption to production traffic
- Options for maintenance mode:
  - ▶ Stop server
  - ▶ Leave server running and keep affinity
  - ▶ Leave server running and break affinity
- May start a new dynamic cluster member first to assure the minimum policy met

9

Overview of new features

© 2007 IBM Corporation

WebSphere Extended Deployment version 6.0 allowed you to instruct the on-demand router to stop routing requests to all servers on a given node by placing the entire node into maintenance mode. In Version 6.1, you can now set maintenance mode on individual servers.

When a server is in maintenance mode, HTTP traffic from the on-demand router to the server will be stopped. If affinity is kept, existing requests will still be routed; if affinity is broken, all traffic will be halted to the server and the requests will fail over if session replication is configured. All *new* IIOP and JMS traffic will be halted to the server, though existing transactions will be allowed to complete.

If the server is a member of a dynamic cluster, a new cluster member can be started before the server is placed into maintenance mode in order to assure the minimum policy on the dynamic cluster is met.

Use maintenance mode to perform diagnostics, maintenance, or tuning on a node or server without fear of disrupting incoming traffic to that node or server in a production environment.

## Section

# ***Operations optimization health management***



This section describes changes to health management in WebSphere Extended Deployment operations optimization.

## Custom health conditions

- Can create expressions defining what “unhealthy” means in your environment
- Custom expressions can be built using operands which represent
  - ▶ Metrics from the on-demand router
  - ▶ Base PMI metrics
  - ▶ MBean operations and attributes
  - ▶ URL return codes
- Complex expressions using a mix of operands is supported
- Other middleware server types can leverage the on-demand router metrics, and URI return code operands.

WebSphere Extended Deployment Version 6.0 provided a limited set of statically defined health condition that could be monitored. For situations where none these predefined health conditions fit your needs, version 6.1 allows you to define a custom condition. A custom condition is an expression that is tested against metrics in your environment.

## Health actions

- Action plan
  - ▶ Set of actions to occur when a policy is breached
- New health actions:
  - ▶ Put server into maintenance mode
  - ▶ Put server into maintenance mode and break HTTP and SIP request affinity to the server
  - ▶ Take server out of maintenance mode
- Custom health actions
  - ▶ Customer scripts
  - ▶ Executable files
  - ▶ Java™ programs



When a health condition is breached in WebSphere Extended Deployment 6.0, you are given the option of creating a Java heap or thread dump, and restarting the server. Version 6.1 lets you place the server into maintenance mode, so requests are not routed to the server during the problem determination process. In addition, you may have specific actions that need to occur when an unhealthy situation is detected. WebSphere Extended Deployment 6.1 allows you to define custom actions in addition to the standard set provided with the product. A custom action defines a program file to run when the health condition breaches. You can specify both custom and standard actions in a list of actions that are run sequentially when the unhealthy situation is detected.

## Health monitoring for other middleware

- Monitoring for other middleware servers
  - ▶ Age based
  - ▶ Work based
  - ▶ Excessive response time
  - ▶ Storm drain
  - ▶ Custom health conditions



Through the middleware agent for operations optimization, WebSphere Extended Deployment 6.1 provides a subset of health monitoring capabilities for other middleware servers that have been added to the WebSphere Extended Deployment administrative domain.

## Section

# ***Operations optimization runtime operations***



This section describes changes to the WebSphere Extended Deployment runtime operations and visualization features.

## Runtime visualization views

- Removed
  - Tree-based runtime topology
  - Runtime map
- Replaced with a series of summary and detail view
- Custom charts still available



15

Overview of new features

© 2007 IBM Corporation

The WebSphere Extended Deployment version 6.0 runtime topology and runtime map views have been removed in version 6.1. They have been replaced by a series of summary and detail views. Custom reports and charts are still available to monitor statistics such as availability, response time, traffic, and throughput.

## Advanced visualization overview

- The administrative console is enhanced with a set of detailed views
  - ▶ Operations tab provides insight into
    - the stability of the resource
    - how work for the resource is actively being managed
    - outstanding tasks that need operators to act upon
    - where the resource is currently placed (running).
  - ▶ Reports tab allows for in-depth charts to be viewed to understand the performance of the environment
  - ▶ Operations and Reporting tabs off the detail view of
    - Servers
    - Clusters
    - Applications
    - Service policies
  - ▶ Can save chart groups

WebSphere Extended Deployment 6.1 provides a new set of views to help you understand the dynamic goals-directed environment. The detail view of servers, clusters, applications and service policies in the administrative console now includes tabs for operations and reporting. The Operations tab provides insight into the stability of the resource, how work for the resource is being actively managed, outstanding tasks that need operators to act upon them, and where the resource is currently running. The Reports tab lets you see in-depth, custom charts to understand the performance of the environment. You can save these report definitions as chart groups so that you can easily access the same data set the next time you log in.



## Monitoring Extended Deployment itself

- System summary view provides visual alerts of Extended Deployment component stability
  - ▶ Core runtime components can be tracked (location and stability)
  - ▶ State and stability of on-demand routers
  - ▶ Core group stability
  - ▶ Node state and stability
- Reporting summary view
  - ▶ Operator-configured sets of in-depth charts
  - ▶ Can be viewed at any time
  - ▶ Real-time snapshot of environment's performance



WebSphere Extended Deployment 6.1 introduces an Extended Deployment summary view to let you monitor and receive visual alerts when the system stability becomes questionable or unstable. This new view in the administrative console provides a dynamic view of the status and performance of your cell. You can see the runtime information for the on-demand routers, nodes, core groups, and core components in your environment. Core components include several autonomic controllers and managers.

A reporting summary view lets you configure sets of in-depth charts into groups that can be viewed at any time for a real-time snapshot of the environment's performance.

## Visualization data logging

- Extended deployment provides the ability to capture historical performance metrics to the file system
- User can control frequency and amount of data captured
- Designed to be used in production
- Simple comma separated variable text file format
- Can be imported into excel, data warehouse or reporting tool
- Can be used for:
  - ▶ Capacity planning
  - ▶ To compute charge-back for system usage



The visualization data service logs historic data in text files for reuse with other charting programs. Historic data is comma-delimited, with time stamps in standard Java format. By using the visualization data service, you can log historical data, calculate charge back values, or perform capacity planning.

## Section

# ***Operations optimization centralized installation management***



The centralized installation manager is a new feature in WebSphere Extended Deployment operations optimization Version 6.1.

## New installation approach

- Single installation to the deployment manager
  - ▶ Includes all code for all of the scenarios described for all platforms
- “Push” of installation from deployment manager to
  - ▶ Select a set of nodes and push Extended Deployment to those endpoints
  - ▶ Install appropriate endpoint code based on type of endpoint
- Agent-less endpoint for installation
- Centralization of patch management



The centralized installation manager can simplify the tasks of deploying product packages to your WebSphere Application Server Network Deployment Version 6.1 cell or other server environments. This feature allows you to push WebSphere maintenance packages, including interim fixes, fix packs, and refresh packs, from the deployment manager to any node in the cell.

It also allows you to install WebSphere Extended Deployment to any node. If the target node is a WebSphere Network Deployment node it will extend the WebSphere installation. Otherwise it will install the WebSphere Extended Deployment middleware agent.

The centralized installation manager installs selected software packages from its repository, which is located on the deployment manager, to specified nodes.

## Section

# **Compute grid**

This section describes changes to the compute grid component of WebSphere Extended Deployment. In version 5.1 and 6.0 this component was known as business grid.

## Compute grid

- Simplified install and configuration
  - ▶ Components installed with WebSphere Application Server as system application
  - ▶ Enabled through the administrative client
- New dedicated job management console
- POJO programming model
- Expand the applications beyond WebSphere
  - ▶ Java “main” applications that run in an independent JVM
  - ▶ Native applications (C, C++, FORTRAN, and so on) running as their own process
  - ▶ CIWork outside WebSphere containers
- Support for non-WebSphere endpoints
  - ▶ Lightweight grid processing nodes with no WebSphere server
- “Standardize” the scheduler
  - ▶ Support for grid standards such as JSDL
  - ▶ Integration with other IBM schedulers (Tivoli® Workload Scheduler, LoadLeveler®)

Configuration of compute grid has been greatly simplified. The grid execution environment is silently installed along with the first grid application on a deployment target. The long running scheduler is now completely configured through the administrative console, and includes a dedicated Web-based job management console that is completely separate from the WebSphere Extended Deployment administrative console.

The “batch” programming model has also been simplified. Transactional batch programs can be now implemented as simple Plain Old Java Objects , or POJOs, rather than requiring Enterprise Java Bean programming. The batch application must still be packaged into Enterprise Application Archive (.ear) files for deployment.

WebSphere Extended Deployment workload management technology can support a diverse mix of application types while ensuring service levels are met for priority requests. A compute grid can now be composed of both WebSphere Application Servers and native execution endpoints. Native code can run on nodes without WebSphere installed through the WebSphere Extended Deployment middleware agent for compute grid.

WebSphere Extended Deployment Version 6.1 provides an integration capability with external workload schedulers, such as Tivoli Workload Scheduler. A compute grid integration layer enables Tivoli Workload Scheduler and similar products to dispatch and monitor compute grid activities. This integration layer may also be used by other workload scheduler products.

## Section

# *Data grid*



This section describes changes to the WebSphere Extended Deployment ObjectGrid.

## ObjectGrid new features

- Significantly increases the number of Java Virtual Machines (JVM) and size of data sets supported
- Dynamic runtime configuration (add/remove servers from a replication group)
- Snapshot Sessions
- Indirect partitioning (fixed set of partitions with partitions mapped to a replication group using a catalog)
- Synchronized replication queue feature - Enabled server to client event/invalidation callbacks.
- Entity Manager style APIs for ObjectGrid
- More efficient HTTP Session Filter (pooling of objects and commits)
- Query support



Version 6.1 of the ObjectGrid significantly increases the number of Java Virtual Machines and size of data sets supported, supporting potentially thousands of servers. Very large data sets are supported by partitioning the data set across multiple servers.

Also new for ObjectGrid version 6.1 is support for plug-n-play type configuration. The dynamic configuration capability of the ObjectGrid makes it easy to create a dynamic deployment topology, and allows automatic balancing of load across ObjectGrid servers. Additional servers can be added to an ObjectGrid without restarting the ObjectGrid. This allows for very simple or small deployments and for very large terabyte sized deployments where thousands of servers are needed. The static deployment topology introduced in earlier versions, which uses a declarative approach to defining the topology is still available.

WebSphere Extended Deployment ObjectGrid version 6.1 introduces a query capability for objects and entities stored in the ObjectGrid. These queries allow objects to be retrieved using non-key attributes and can perform simple aggregations such as sum, and average against all the data that matches a query. The new Stream Query API allows you to perform continuous queries over streaming data.

ObjectGrid version 6.1 also contains many other performance and programming API enhancements.



## Summary

- New packaging
- Performance enhancements
- New manageability features
- ObjectGrid enhancements
- Simplified and extended compute grid



In summary, WebSphere Extended Deployment version 6.1 is now packaged as three separately installable components: Operations Optimization, Data Grid, and Compute Grid. Each package contains numerous fixes and enhancements to improve performance, increase programming power, and simplify administration.

## 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 XD61\\_Whats\\_New.ppt](mailto:iea@us.ibm.com?subject= Feedback about XD61_Whats_New.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:

IBM      LoadLeveler      Tivoli      WebSphere

Java, 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.

