



IBM Software Group

# IBM® WebSphere® Extended Deployment V6.1

## *WebSphere Virtual Enterprise*

*Formerly Operations Optimization*

## *Edition management overview*



@business on demand.

© 2007 IBM Corporation  
Updated June 18, 2008

This presentation will cover the Application Edition Management features of WebSphere Extended Deployment V6.

## Agenda

- Overview
- Operational scenarios



This presentation will first give an overview of WebSphere Extended Deployment's Application Edition Management features, then present a few operational scenarios that illustrate how Edition Management can be used.

## Overview

- WebSphere Extended Deployment supports managing multiple editions of an application in a WebSphere cell
  - ▶ Interruption-free rollout of application updates
  - ▶ Ability to “roll back” to a previous application version
  - ▶ “Validation mode” to verify functionality using a subset of users



WebSphere Extended Deployment allows you to install and manage multiple editions of the same enterprise application within a cell. You can use the Edition Control Center to manage the rollout of new editions to some or all of your servers without service interruption for your users, and you also have the ability to roll back to previous editions.

## Application editions


- An *edition* is a distinct instance of a Java™ 2 Enterprise Edition (J2EE) application (similar to *version*)
- Each edition is numbered (or named)
- An edition is a *deployment* version of an application
  - ▶ May be a distinct *build* version
  - ▶ May be the same build version with different deployment bindings (such as resource-ref)
  - ▶ May be both



An *edition* of an application is the term that is used to refer to a specific instance of a particular application. The term edition is used rather than version because two application editions could actually be the same 'build version' of the application, but deployed with different bindings.


Shared libraries are not handled by the edition manager, and two versions of the same application cannot run within the same Java process.

## Operational scenarios: Grouped rollout



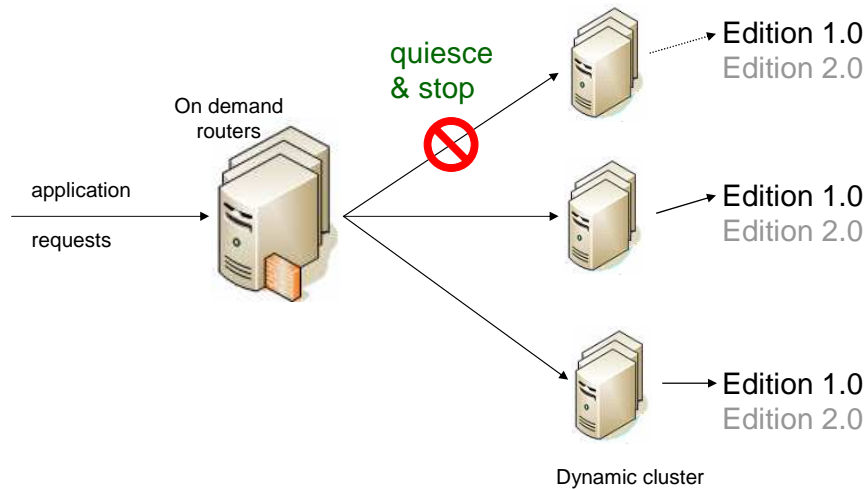
You have several options when rolling out an edition. The first is known as 'group rollout'. In this scenario, the servers in a target cluster are divided into groups and the new application edition is rolled out across your cluster one group at a time. The administrator defines the group size, or specifies the number of nodes to process at a time. This type of roll out is most effective when the cluster is large.

## Operational scenarios: Grouped rollout



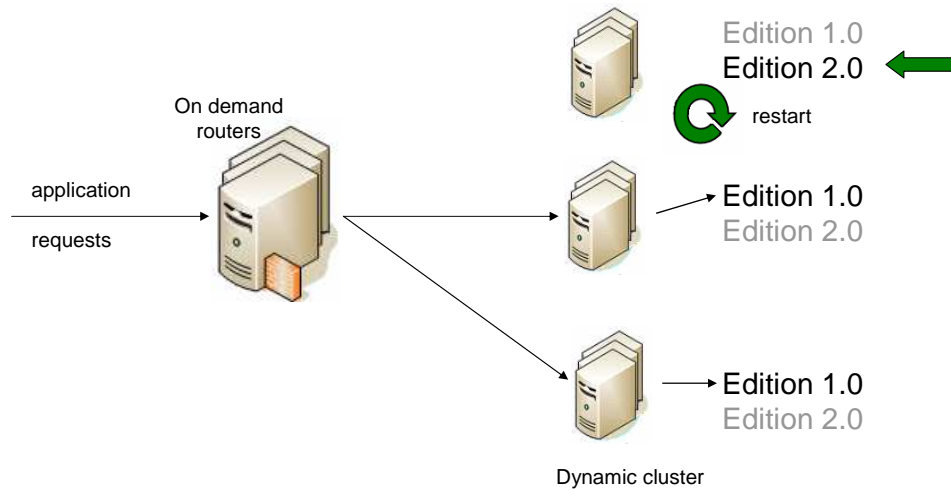
The first step is to install the new application edition. During the installation process the code for the new application edition is distributed to the target nodes in an inactive state. An inactive application edition is installed, but can not be started. Note that even though a new edition is installed, the old edition is still active and serving requests.

## Operational scenarios: Grouped rollout



To begin the rollout process, the on demand router stops routing new requests to a group of servers. The application edition manager will wait a specified amount of time for the servers to finish processing the requests they are currently working on. When all existing requests have completed, edition 1.0 of the application is then stopped on each server in the group.


## Operational scenarios: Grouped rollout



When the edition 1.0 of the application has stopped, it is deactivated and edition 2.0 of the application becomes active and is started. This step is known as "activation".




## Operational scenarios: Grouped rollout



As soon as the edition 2.0 application becomes available, the on demand router resumes routing traffic to that group of servers. New work requests on those servers are now processed by edition 2.0 of the application. The process is repeated across the cluster until all servers are running the latest edition of the application.

Keep in mind that when using grouped rollout, both versions of the application will be simultaneously serving traffic for a period of time.


## Operational scenarios: Atomic rollout



The second method, called 'atomic rollout' uses a 'divide-and-switch' technique to rollout a new edition across *half* of the cluster at a time. This ensures that only one version is actively serving traffic at any given time.


As with a grouped rollout, the first step is to install the new application edition. At this point the 2.0 edition of the application is in the inactive state in which it is installed, but can not be started. Note that even though the new edition is installed, the old edition is still active and serving requests.

## Operational scenarios: Atomic rollout




The on demand router will first quiesce traffic to *half* of the servers in the cluster. When existing requests complete, edition 1.0 of the application is stopped .

## Operational scenarios: Atomic rollout



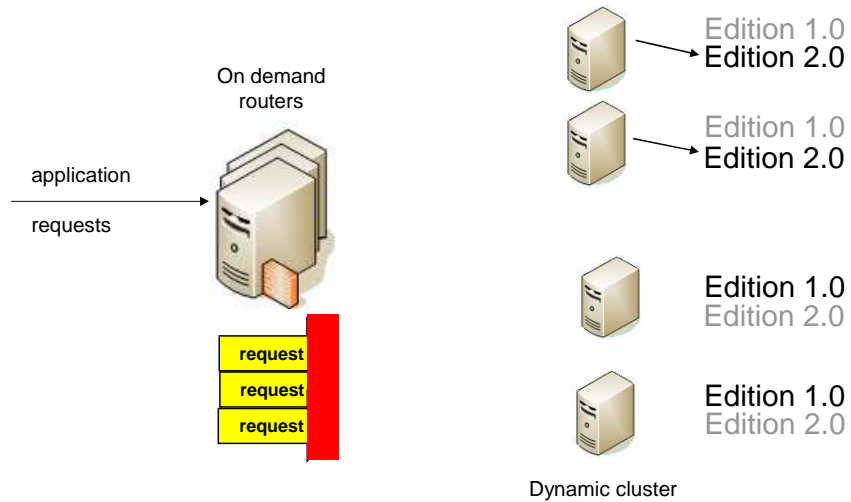
Edition 1.0 of the application is then deactivated and edition 2.0 of the application becomes active and is started.

## Operational scenarios: Atomic rollout



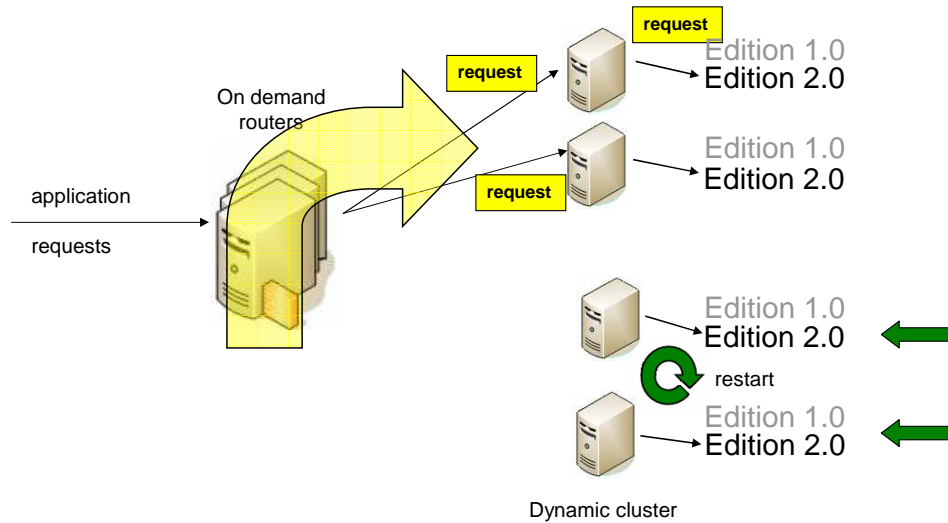
The next step is to replace edition 1.0 on the remaining servers. The first set of servers are held offline until this step completes.

## Operational scenarios: Atomic rollout



While waiting for those requests to finish, there may be a brief build-up of requests in the on demand router's dispatch queue, since no servers are receiving traffic.

## Operational scenarios: Atomic rollout



15

Edition management overview


© 2007 IBM Corporation

After the application is fully offline, the first half of the cluster serving edition 2.0 of the application is brought back online and can begin to handle work requests.

Once the existing requests have completed, requests, including those that built up in the queue, will start being routed to the first set of servers, which is now running the new edition of the application.

While this is happening, the second half of the cluster transitions from the previous edition to edition 2.0 and the new edition is started.

## Operational scenarios: Atomic rollout



Once they have been restarted, the entire cluster will be running the new application edition. Note that since your cluster will effectively be running at half capacity during this process, it is wise to perform an edition upgrade during a low-traffic time period.

In both rollout scenarios, the old edition is still installed on the servers, though inactive. This makes it possible to easily roll-back to a previous application edition if necessary.



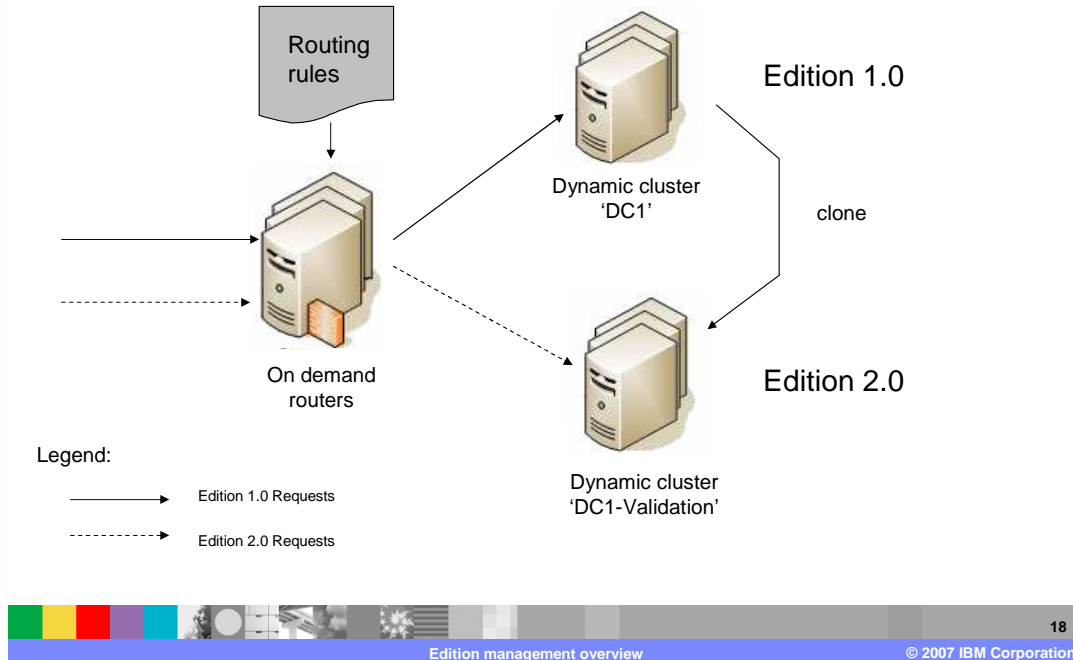
## Operational scenarios: Validation mode

- Creates a 'clone' of the deployment target (Dynamic Cluster) in production environment and deploys new version
- Use routing policies to control edition visibility
  - ▶ Restrict new edition to only certain users, for example
- Use rollout function to move edition from validation mode to production
  - ▶ Edition deactivated on 'clone' environment
  - ▶ Edition rolled out on original deployment targets



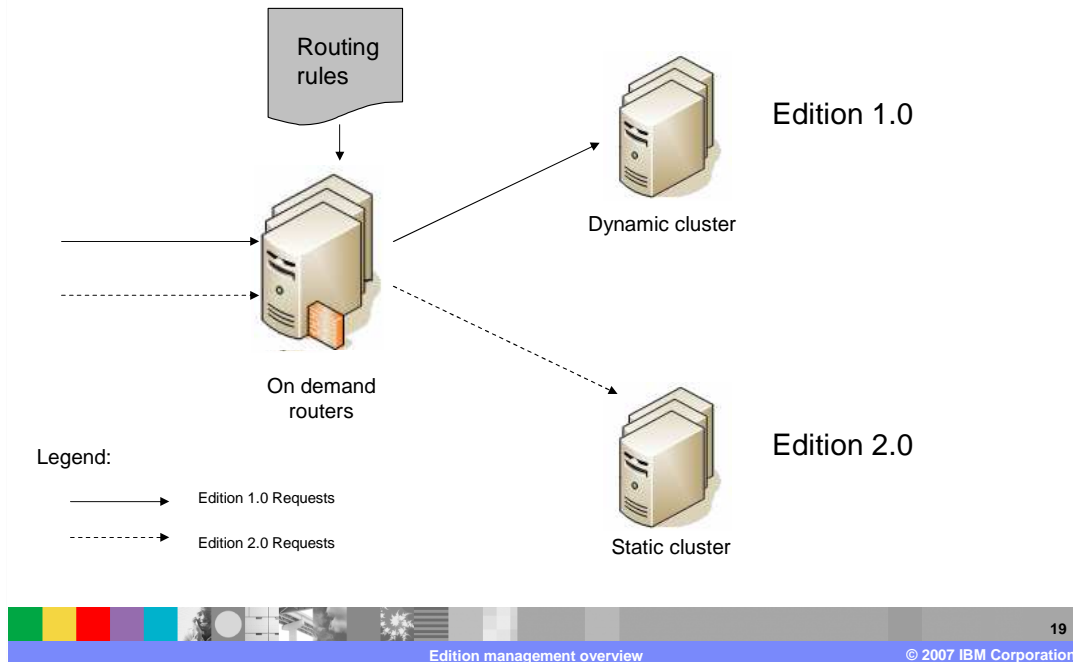
Another option is to 'validate' a new edition before rolling it out to your cluster. Validating an edition creates a clone of your dynamic cluster in your production environment, and deploys the new edition to that 'cloned' cluster, which is made available only to a limited subset of traffic. You use routing policies to limit the validation edition to a certain group of users, a certain range of IP addresses, or many other limiting criteria. This gives you the ability to verify functionality of the new edition in your production environment without exposing it to all of your users. You can later use the rollout feature to roll out the application into regular production mode and stop the validation cluster.

## Operational scenarios: Validation mode



This diagram illustrates the scenario that was just described. The on demand router controls the subset of requests that get routed to the validation cluster based on routing policies that you have defined.

## Operational scenarios: Concurrent activation



It is also possible to create a similar situation without using validation mode to 'clone' a cluster. You can choose to rollout the new edition to a cluster that you use for testing purposes, and configure routing policies in the on demand router that route only certain traffic to the new application edition. An edition can be rolled out to either static or dynamic clusters.

## Rollout considerations

- **Scope of Interruption-free operation**
  - ▶ Dynamic cluster, static cluster, single server
  - ▶ Only to the tier immediately behind the on demand router
  - ▶ HTTP-based requests only
- **Compatibility**
  - ▶ Backward compatible changes only
- **Session Affinity**
  - ▶ Option to wait or expire sessions during quiesce
- **Restart granularity**
  - ▶ Application or server



There are several factors to keep in mind when rolling out a new edition of an application. Since the management of incoming requests is handled by the on demand router, interruption-free rollout is only possible for HTTP and HTTPS requests. In a multi-tiered environment, interruption-free rollout is only possible for the application tier directly behind the on demand router. It is also important that you keep application compatibility in mind. A new edition that expects an incompatible database schema or different data types in the HTTP session will not roll out seamlessly.

## Summary

- WebSphere Extended Deployment enables management of multiple editions of an enterprise application
- New editions can be rolled out without service interruptions
- Validation mode
  - ▶ Verification of edition functionality within the production environment
- Routing policies
  - ▶ Rules used by the on demand router to determine which edition should handle the request



In summary, WebSphere Extended Deployment enables you to manage multiple versions of an enterprise application within a cell. New editions can be rolled out seamlessly, without application downtime, and validation mode lets you verify that a new edition works properly in your production environment before fully rolling it out. The on demand router uses routing policies that you have defined to decide how to route individual requests to any concurrently deployed editions of an application.

## 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 Edition Management 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:

IBM                    WebSphere

J2EE, Java, 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.