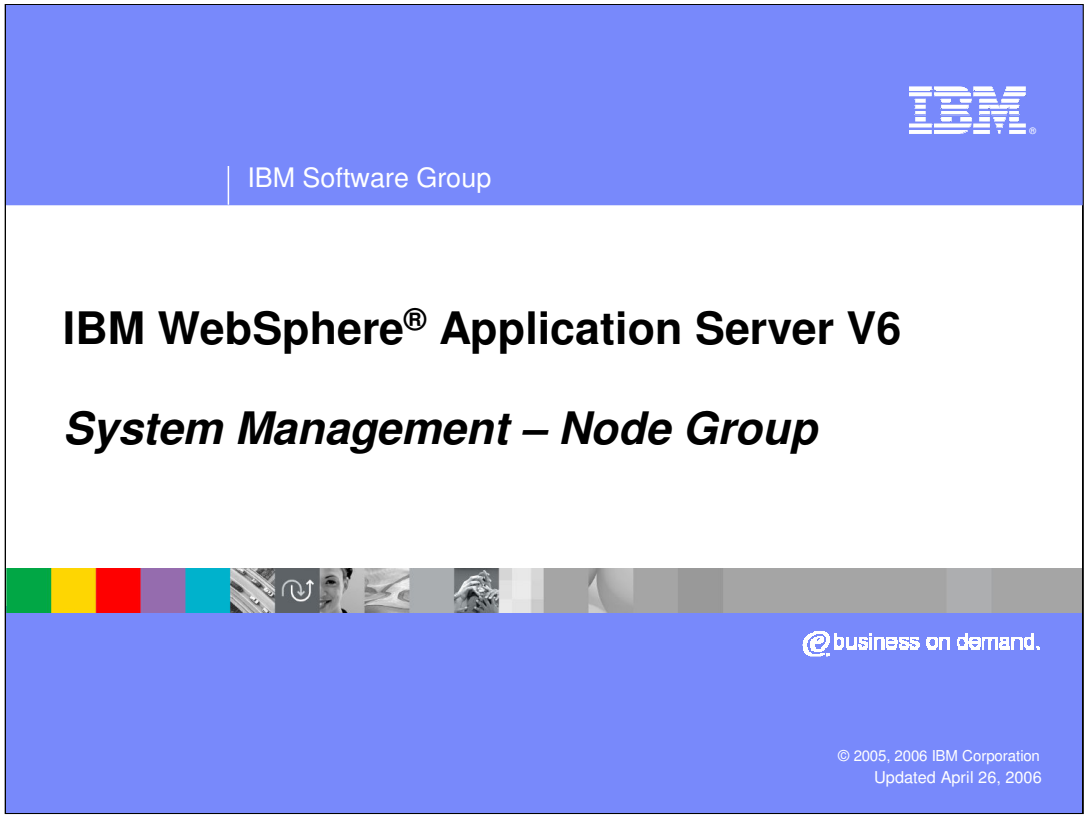

This presentation will focus on the concept of Node Groups as introduced in WebSphere Application Server V6

Agenda

- Node Group Concept
- Node Groups in Distributed or z/OS® only cell
- Resource Checking within a Node Group

The agenda for this presentation will focus on understanding the concept of Node Groups in WebSphere Application Server V6. Within the context of a Node Group you will also learn about the resource checking and validation.

Node group: Rationale

- Enables grouping of nodes with same capabilities
- Allows validation of node capability before performing certain functions
- Default configuration with single node group is sufficient

Node Groups allow nodes of similar capabilities to be logically grouped together within a cell. This grouping of nodes allows for the validation of capabilities of a defined Node Group before performing specific functions. For example, a specific function like adding a node would be checked to see if the node that you are adding has same capabilities as the nodes in the Node Group. There always exists a Node Group definition in V6, and in most cases, the default should be sufficient.

Node group usage in V6

- Define cluster boundary
  - Cluster must be fully contained within a Node Group
  - Node Groups encompass nodes of similar capability
- Optional Resource validation for an application within the boundary of a deployment target's (server or cluster) Node Group

Node Groups define a boundary for server cluster formation. Nodes organized into a Node Group should be enough alike in terms of installed software, available resources, and configuration to enable servers on those nodes to host the same applications as part of a server cluster. If validation is enabled, any Java™ 2 Enterprise Edition (J2EE) resource existing outside the scope of the Node Group for a given application will be reported. For example, if an application is deployed to a server on Node1, but assigned a JDBC data source defined to Node2 which is outside the Node Group, a message will be written to the systemout log.

Default node group

- Automatically created based on the Deployment Manager (DMgr) platform
  - Called “DefaultNodeGroup”
  - New nodes are automatically added
- Cannot be deleted
  - DMgr is always part of this Node Group
- For most topologies, the default Node Group will suffice

In V6 there always exists a Node Group and each node must belong to one or more Node Groups. When installing the Deployment Manager, a default Node Group called DefaultNodeGroup is created based on the platform of the Deployment Manager. Any new nodes that are added to the cell will be in this Node Group unless otherwise specified. This default Node Group cannot be deleted and the Deployment Manager is always part of it.

Node group membership rules

- A Node Group can only contain “Managed Nodes”
- A Managed Node must be a member of one or more Node Groups
- All nodes in a z/OS sysplex must be contained in a single Node Group
  - A Node Group cannot span multiple z/OS sysplexes

Nodes can only be members of Node Groups if they meet the following requirements. First, a node must be a Managed Node, meaning there is a Node Agent process active for the node and the Deployment Manager can configure and administer. Second, a Managed Node must belong to one or more Node Groups. Third, all nodes in a z/OS sysplex must be contained in a single Node Group, and a node group cannot span multiple z/OS sysplexes.

Node group administration

- Administrative functions
  - Create new Node Group
  - Delete Node Group
    - It must not have any Nodes in its membership
  - Add nodes to the Node Groups
  - Remove nodes from Node Groups
- New Node Group will have to be administratively created, prior to adding a node

Administering Node Groups is quite simple. There are only a few administrative functions that can be performed. These functions are creating and deleting a Node Group and adding and removing nodes from a Node Group. Any new Node Group will have to be created before adding any nodes to it. For example, if you have a cell with nodes that are defined in the DefaultNodeGroup, and you wish to add a new node but do not want it to belong to the DefaultNodeGroup, you must create a new Node Group prior to adding the new node.

addNode command changes

- addNode has a new option [-nodegroup <name>] to allow adding a node to a specific Node Group
  - If not specified, “DefaultNodeGroup” is assumed
- addNode ensures that the node being added has the same capability as the Node Group

With the addition of the Node Group function, the option of choosing a Node Group has been added to the addNode command. If you do not specify anything while adding a node, the node will be added to the default Node Group. If you try to add a node of a different capability without specifying a compatible Node Group, an error will appear and the node will not be added to the cell.

Use case: Homogenous cells

- All nodes on the same platform type
  - One Node Group, ”DefaultNodeGroup”, for entire cell
- Nodes could be located on different machines

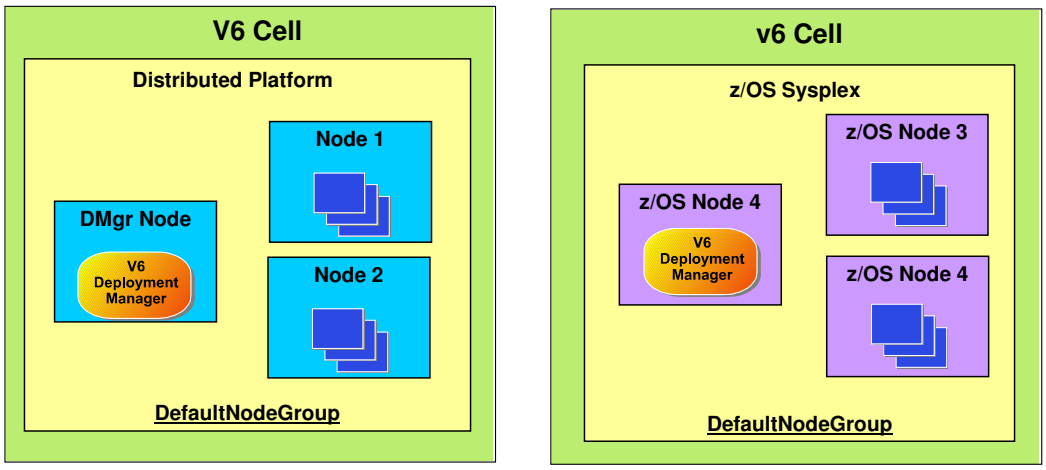

Here is a use case for Node Groups. There are two homogenous environments; one is the distributed and the other is z/OS. For each there is one Node Group defined called DefaultNodeGroup. Again, this was automatically created when installing the Deployment Manager for a given platform. In this scenario there is no extra Node Group configuration that needs to be done.

Summary

- Node Groups provide a way to group together nodes of similar capabilities

In summary, this presentation covered Node Groups and how they allow you logically group together nodes of similar capabilities within the same cell.

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	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

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

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

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel 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.

Linux is a registered trademark of Linus Torvalds.

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

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 and/or changes in the product(s) and/or program(s) 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 (e.g., 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.