



IBM Software Group

IBM® WebSphere® Application Server V6

Topology Options and Terminology



@business on demand.

© 2005 IBM Corporation
Updated April 26, 2005

This presentation will introduce some topology options and define some of the terminology used when talking about WebSphere Application Server version 6.

Goals

- Introduce different WebSphere topology options
 - ▶ Stand-alone Node:
 - Express or Network Deployment Stand-alone Node
 - ▶ Network Deployment Cell topology
 - Managed and Unmanaged Web Server
- Introduce V6 terminology
- Covered in other presentations:
 - ▶ Creating these topologies will be discussed in the “Installation and Migration” presentation

The goals of this presentation include introducing different WebSphere topology options, such as a stand-alone node, or Network Deployment cell, and introducing version six terminologies.

Agenda

- Stand-alone Node Topology (Express or Network Deployment Stand-alone node)
 - ▶ Application Server topology
 - ▶ Web Server topology

- Network Deployment Cell
 - ▶ Cell Topology
 - ▶ Web Server topology

Symbols used in the presentation



▪ Instance of an Operating System
▪ One Logical OS copy (Windows™) or Logical Partition (like in AIX® or z/OS®)



One Install of WebSphere product = One set of WebSphere install Binaries

Server

The agenda for this presentation includes starting with the stand-alone node topology, and then exploring the Network Deployment cell topology.

Section

Basic Terminology

This section will address basic terminology.


Terminology: Operating System (OS)

- Distributed
 - ▶ Microsoft™ Windows™
 - ▶ UNIX™ – AIX™, Linux™/Intel™, Linux/PPC™, Linux on zSeries™, Solaris™, HP-UX™
- iSeries™: i5/OS™, OS/400™
- zSeries™: z/OS™ and OS/390™

The terms “distributed” or “distributed platforms” includes supported hardware other than the iSeries and zSeries hardware, and encompasses Windows, UNIX, and Linux operating systems.

Terminology: Machine / Logical Partition (LPAR)

- Runs an instance of the Operating System (OS)
- Case 1: Physical machine can run only 1 OS instance
 - ▶ Example: Windows, Linux/Intel
- Case 2: Physical machine can run multiple OS instances, each in its own LPAR
 - ▶ Example: z/OS, AIX, zLinux, Linux/PPC
- WebSphere Application Server is installed on an **instance** of the OS



z/OS, AIX, etc.

Throughout this and the subsequent presentations, reference will sometimes be made to a machine, or to a system or to an LPAR. Logically speaking, these terms are often used in an interchangeable fashion. What is really meant is that entity in which an operating system instance is running. On the distributed platforms, for example Windows or Linux on Intel, this is analogous to a physical machine. However, for the z/OS platform this refers to an individual LPAR running within z/Series hardware, which normally has multiple LPARs configured and running.

Terminology: Application Server

▪ Application Server

- ▶ Runs one or more J2EE applications
- ▶ Provides services required to run applications:
 - Web container, EJB container, JMS services, Security, Transactions, etc.
 - Operational management of the applications
- ▶ Runs as Java™ (JVM) process
 - Single JVM on non-z/OS systems
 - Multiple JVMs on z/OS system



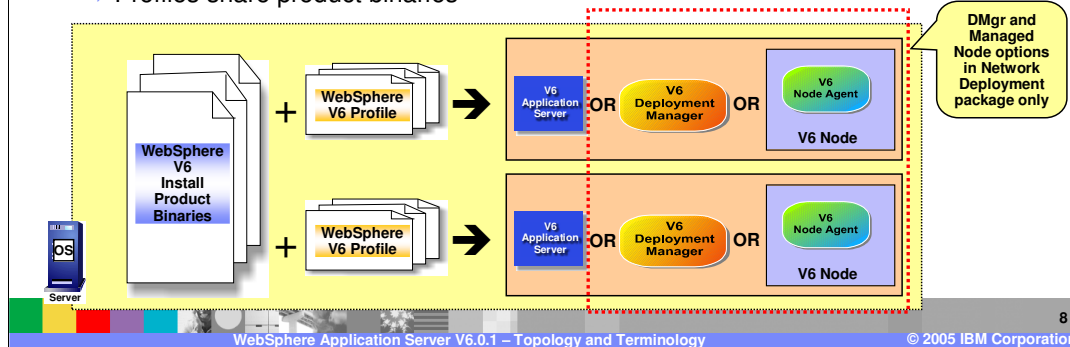
“Server”, means a process which provides some function or functions. The functions can be things like providing a Web or Enterprise Java™ Bean (EJB) Container for hosting Java™2 Enterprise Edition (J2EE) artifacts, providing services such as Java™ Messaging Service (JMS) server or Naming, and providing configuration and operational management support.

These presentations use the term “process” loosely, as it is not always analogous to what is normally known as a process within an operating system. On the distributed platform it is, with a server being an operating system process that runs in a single JVM instance. However, on the z/OS platform, a server is composed of multiple operating system processes, each running in their own JVM. There is one Control Process and potentially many Servant Processes making up a z/OS WebSphere server.

Servers are defined though a set of XML configuration files which are maintained on the file system and read in by the server when starting up.

Terminology: WebSphere Profiles

- WebSphere V6 files are split into 2 categories
 - ▶ Product Files - shared application binaries for WebSphere
 - ▶ User Files - set of user customizations
 - Includes WebSphere configuration, installed applications, resource adapters, properties, log files, transaction log files, etc.
 - Defines a profile
- Each Profile defines a WebSphere runtime environment
 - ▶ Stand-alone Application Server, Deployment Manager (DMgr) and Managed Node
 - z/OS has one 'profile' under each of these configurations named 'default'
 - ▶ Profiles share product binaries



This is a replacement to the WebSphere V5 wsInstance function on the distributed platforms.

To understand Profiles, start by understanding the files that make up WebSphere Application Server. There are two categories of files: Product files and User files. The Product files include the application binaries needed to run the application server. The User files contain information used by the application server. For example, this is where variables are defined, resources are configured, log files are written, and so on. A Profile is a collection of these files, creating a WebSphere runtime environment. When combined with the shared binaries, a profile becomes a complete WebSphere Application Server installation.

This sharing of application binaries, and the separation of configuration files, is an efficient use of disk space. Also, with one set of binaries, updates to the code can be applied in one location per physical machine, even when multiple profiles are configured. This isn't really a new concept on z/OS and you will see that configuration hasn't really changed for z/OS with the introduction of profiles.

This is discussed in more details in a separate WebSphere Profiles presentation.

Section

Stand-alone Application Server Topology

This section covers the topology of the stand-alone WebSphere Application Server installation.

Terminology: Stand-alone Application Server

- Single node that runs one Application Server
- Independent of other nodes and other Application Servers
 - ▶ Cannot be managed by Deployment Manager



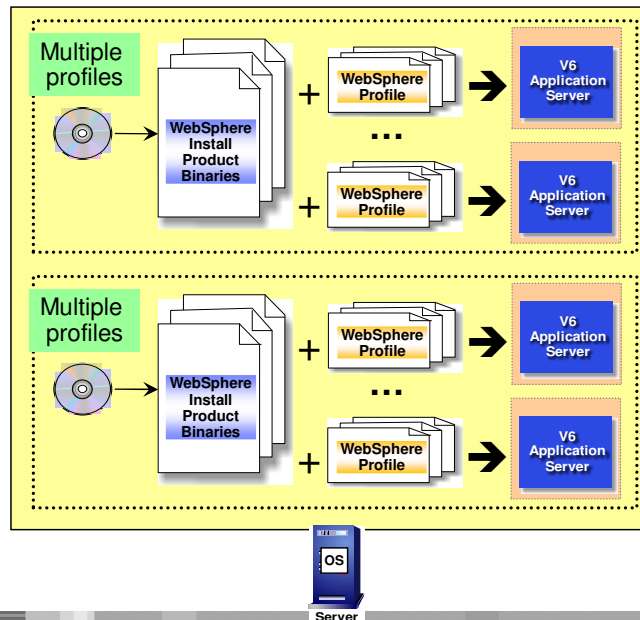
Sometimes, Stand-alone Application Server is also referred to as Stand-alone Node or Single Node



A stand-alone application server installation is a single node that runs one Application Server process independent of a Network Deployment cell.

Stand-alone Application Server: Topology

- Can have combination of multiple installs or multiple profiles or both
- Multiple profiles benefits
 - ▶ Uniform fixes applied to all profiles
 - ▶ Requires less disk space and less installs
- Multiple install benefits
 - ▶ Can manage fixes separately
 - ▶ Requires more disk space



By default, when you install the stand-alone server on distributed platforms, it creates the first Instance of the Application server. On z/OS, after you have done the SMP/E install of the product binaries, you will use the ISPF Customization Dialogs to install WebSphere instances and will have the choice to install a Network Deployment Cell, a Stand-Alone Server or an 'empty managed node'.

You are limited to one Application Server within the a stand-alone install. If you need multiple servers within the same install, use the Server profiles mechanism on the distributed platform or configure additional Stand-Alone Application Servers using the ISPF Customization Dialog on z/OS.

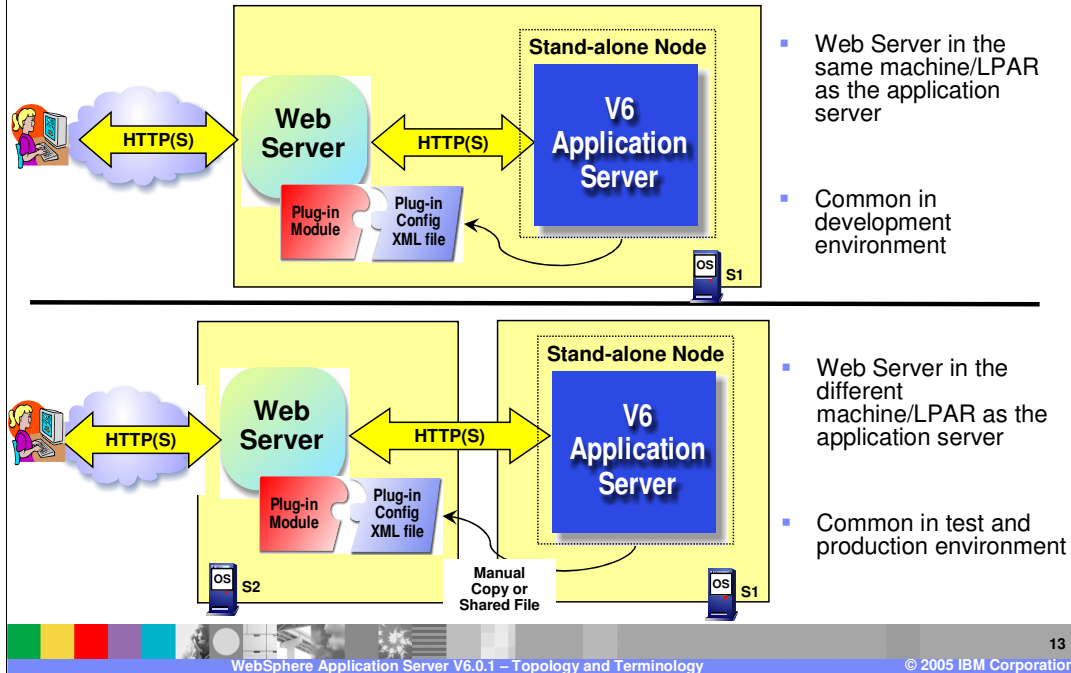
When running multiple Servers on the same machine, assign unique ports to avoid port conflicts.

Section

Web Server Topology for Stand-alone Application Server

This section will discuss the Web Server topology in a Stand-alone Application Server environment.

Web Server Topology Options: Single Server



Web Server can be on the same local machine as the Application Server or can be on a remote machine – remote to the Application Server. The local topology is more common in development environments and small shops.

Web Server Definition in Stand-alone Server

- Web Servers are now defined in WebSphere Stand-alone Application Server topology
- This allows generation of custom plug-in configuration files
- Only one Web Server can be defined

New in WebSphere Application Server V6 is the ability to define a Web Server node with the Application Server topology. This allows creating custom plug-in configuration files for the defined Web Servers. The Stand-alone Application Server can have only one Web Server definition, whereas the Network Deployment cell can have multiple Web Server definitions (explained later).

Section

Network Deployment Cell Terminology

Next few pages will discuss the terminology of a Network Deployment Cell environment.

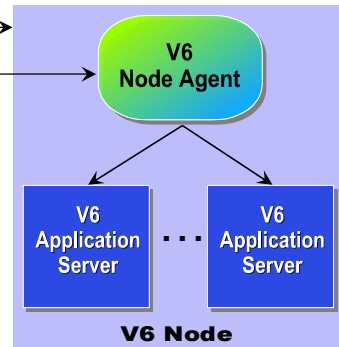
Terminology: Managed Node and Node Agent

Managed Node

- ▶ Grouping of Application Servers for configuration and operational management on a **single** machine/LPAR
- ▶ Contains one Node Agent
- ▶ Can have multiple nodes on a single machine

Node Agent

- ▶ JVM process that manages the servers on a Node
- ▶ Node Agent is created automatically when you add (federate) a Stand-alone Application server node to a Cell




In WebSphere Version 5 there is the concept of a node. A node is a grouping of servers for the purpose of configuration and operational management. A node's boundaries are limited to within a machine or LPAR, and cannot span beyond those boundaries. Also, it is permissible to have more than one node within the same machine or LPAR. Note that this is different from the concept of a node as it was in WebSphere Version 4, on both the distributed and z/OS platforms.

The stand-alone Application Server is a single server used for deploying J2EE applications. It contains a Web Container, and EJB Container and provides Naming services. It can also optionally provide JMS services. Normally it will also contain the Administrative Console application which allows it to be configured and managed using the browser based administrative console.

Terminology: Deployment Manager (DMgr)

- Manages one or more nodes in a distributed topology
- Communicates with the Node Agents for configuration and operational management for that node



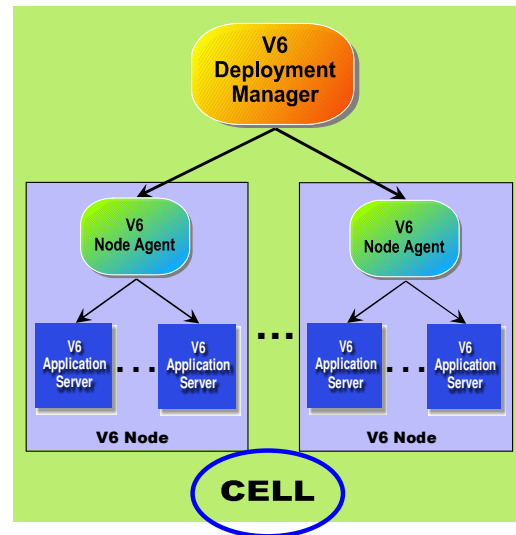
The Deployment Manager contains the master repository.

Each Node has the XML Configuration files and the J2EE Application files to start all the WebSphere processes for that Node.

The XML Configuration and J2EE application at the Node level is a subset of the files at the Deployment Manager, containing only the files that are required for that node

Terminology: Cell

- Network of multiple Nodes and one Deployment Manager (DMgr) over one or more machines
- Provides a single logical administration domain
- DMgr and the Nodes can be on the same or different Machine/LPAR



The next level up the hierarchy from a node is the cell. A cell is a network of nodes which is used to provide a single point of administration for the nodes. On the z/OS platform, a cell's boundaries must be within a single Sysplex. On the Distributed platforms, a cell may be made up from nodes on multiple heterogeneous systems.

Terminology: Unmanaged Node

- New in V6 is the notion of Unmanaged Node that you can define in a Cell or a Stand-alone Application Server
- Difference between Unmanaged Node and Managed Node
 - ▶ A Managed node has a Node Agent
 - ▶ An Unmanaged Node does not have a Node Agent
- Deployment manager cannot manage an Unmanaged Node
- Unmanaged Nodes are useful for creating Web Server definitions
 - ▶ Enables the creation of custom Web Server plug-in file for that Node

An Unmanaged Node is defined with the Network Deployment Cell topology as a remote node that WebSphere does not manage. However, the Unmanaged Node provides the capability of defining other servers like Web Servers within the topology, as unmanaged servers. Providing this capability allows the system administrator to do specific functions like create custom plug-in configuration file for the Web Server on the Unmanaged Node.

Section

Network Deployment Cell Topology Options

The next section will discuss the topology of a Network Deployment Cell environment.

Network Deployment Install Options/Distributed Platforms

- Network Deployment package allows the creation of the following WebSphere Profiles:
 - ▶ Deployment Manager
 - ▶ Stand-alone Application Server
 - ▶ Custom Profile (sometimes called Managed Node)
 - This is similar to creating a stand-alone Application Server and federating (adding) the node to a DMgr
- No need to install a separate Deployment Manager package
 - ▶ Create a profile of type “Deployment Manager”



During the installation of Network Deployment, you can create any one of the valid Profiles – Stand-alone Application Server, Deployment Manager or a Custom Profile. The Stand-alone Application Server profile is no different than what has been discussed earlier in this presentation. Except for licensing considerations, it is the same as the Express Application Server. The Custom Profile allows you to create a Managed Node (already federated to a Deployment Manager), where you can then define Application Servers or Web Server nodes.

One big difference from version 5 install, is that in version 6, you have only one install image from which you can create all the runtimes. In version 5, there are two distinct install image one for the Deployment Manager and another for the Application Server.

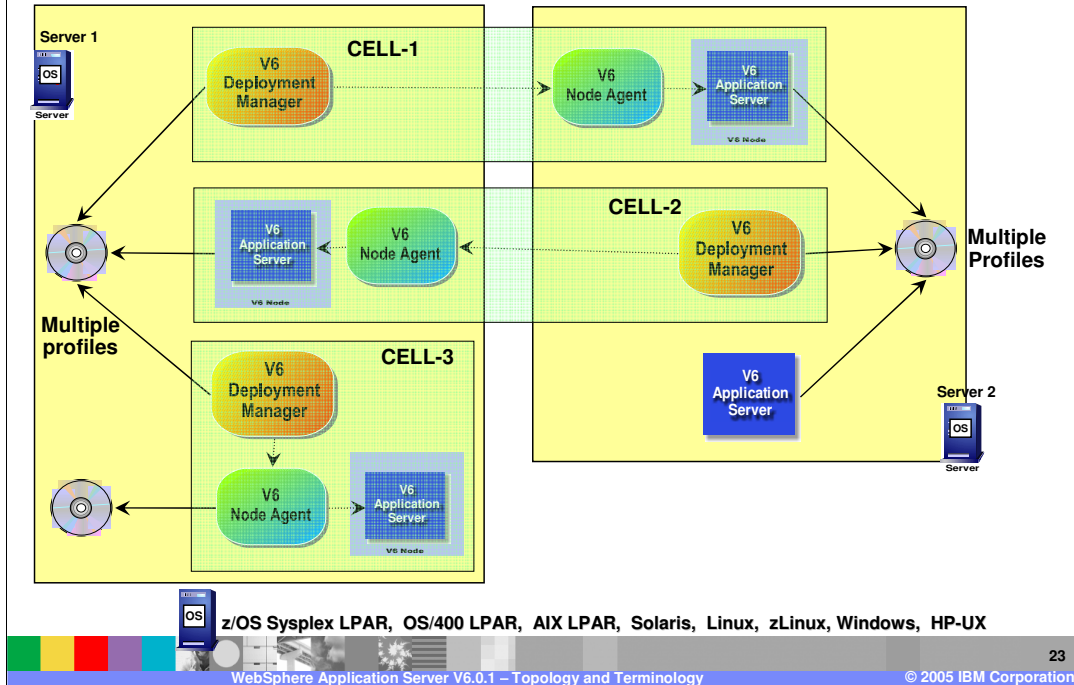
z/OS Configuration Options

- z/OS is similar to the Distributed Platform Network Deployment package
- Once the code is SMP/E installed, you have the options to configure:
 - ▶ Deployment Manager cell
 - ▶ Stand-alone Application Server
 - ▶ Empty Managed Node (known as a Custom Profile on Distributed Platforms)
 - This is similar to creating a stand-alone Application Server and federating (adding) the node to a DMgr

z/OS is a little different from the Distributed Platforms in how it is installed because of the use of SMP/E to install the product binaries. This naturally separates the installation of the product binaries from the configuration of the actual Application Server and you'll see that nothing has really changed from V5 other than the addition of a new 'Empty Managed Node' option. So, once the code has been SMP/E installed, you will run the ISPF Customization Dialog to configure one of the options shown on the foil. The way that z/OS distinguishes the various configurations under the one install image is thru the 'home directory'. Each time you configure a different option, you will specify either a separate HFS with its own mount point or a separate directory within a common HFS. The Distributed Platform uses the concept of Profiles to distinguish the various configurations. On the following foils, when you see 'Profiles', think of these different configuration options.

See the separate WebSphere Profiles presentation to understand how the concept of Profiles fits into the z/OS Platform.

Topology Flexibility - Example



You have total control over how you want to create the topology of your Network Deployment Cell or Stand-alone Application Server.

You can configure multiple Cells sharing the same hardware or different profiles or different installs. You can decide if the Deployment Manager and Nodes share the same machine or are on different machines, and so on.

In this example, there are multiple profiles on both the Servers 1 and 2. There are multiple installs on Server 1, and Multiple Nodes and multiple Deployment Managers on the same Server. The Stand-alone Application Server profile coexist with the other Network Deployment cells.

This diagram also has examples of multiple profiles and multiple installs, multiple Nodes on the same machine, multiple Deployment Managers on the same machine, and so on.

The only requirement is that when running multiple WebSphere processes on the same operating system, the port numbers must be unique.

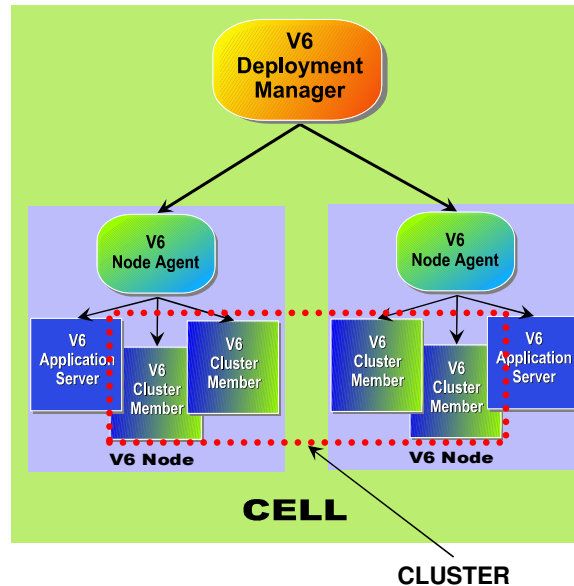
Section

Network Deployment Cell Cluster Topology

This next section will contains an examples of a Network Deployment Cluster topology.

Terminology: Cluster and Cluster Member

- A Cluster is a grouping of application servers, called cluster members, running the same set of J2EE Applications
- Cluster provides scalability and failover capabilities
- A Cell can have zero or more clusters
- Can span machine or LPAR boundaries
 - Vertical - Same machine/LPAR
 - Horizontal – Different machine/LPAR
- Can span different OS types – like AIX and Windows
 - Exception – cluster members cannot span distributed and z/OS



In this example, Cluster and Cluster member definitions are the same as in version 5. Horizontal Clusters provide a fail over of a machine that is not provided by vertical clustering.

Section


Network Deployment Cell

Mixed Version Nodes

The next slide will discuss the Network Deployment cell with mixed version nodes.

Mixed V6 and V5.x Version Nodes in a Cell

- Support of mixed V6 and V5.x nodes in a single cell
 - DMgr must be at V6 level
- V5.x nodes can be on different platforms
- Some or all nodes, including Deployment Manager node, may share physical machine/LPAR
- Configuration files for V5.x code follows the V5.x definition
- Supports mixed V5 and V6 cluster members
- Some limitations apply to a mixed Node
 - Please review these limitations, listed in the Migration presentation



The Migration presentation will cover the steps performed to go from a full V5 cell, to a mixed cell to a full V6 cell.

Some of the restrictions of mixed version cell are:

- Cannot add a new V5 node
- Cannot add a new V5 Application server
- Applications using JDK 1.4 cannot run on V5.0.2 or below

However, if you anticipate and add new V5 nodes and WebSphere Application Servers before migrating to V6, and do not start them, you can use them later.

Section

Web Server Topology in a Network Deployment Cell


The next section will cover Web server topology within a cell.

Web Server in a Cell

- One or more Web Servers can be defined in a cell topology as a Web Server node
- Applications can be associated with one or more defined Web Servers
- Allows generation of custom plug-in configuration files for each Web Server
 - ▶ Allows Web Servers to target specific applications running in a Cell
- Web Servers are defined on a Managed Node or Unmanaged Node

As discussed before, WebSphere Application Server V6 has the ability to define a Web Server node with the Application Server topology. In Network Deployment, a cell can have multiple Web Server definitions which can be defined on a Managed or Unmanaged Node.

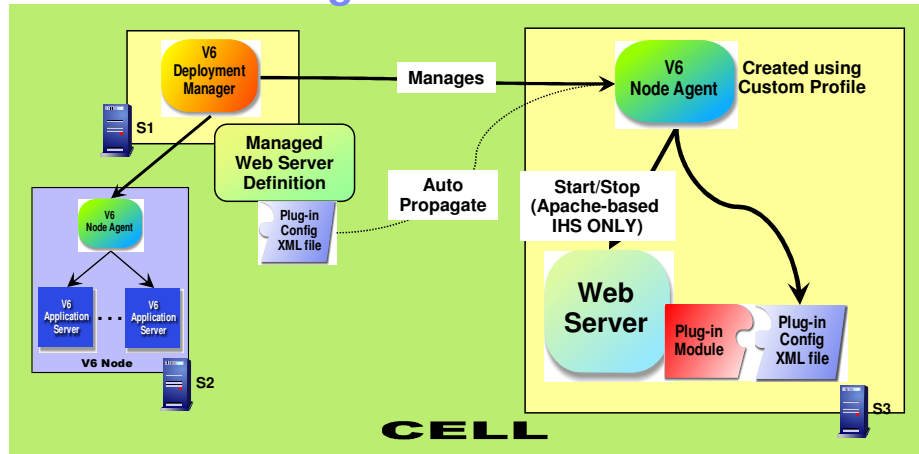
Network Deployment Cell: Web Server Topology



- Web Server on S1 sends request to applications running on Application Server A & D
- Web Server on S2 sends request to applications running on Application Server B & C
- Web Server uses plug-in configuration file (plugin-cfg.xml) to determine which requests to forward and where - Automatically generated by the Deployment Manager

This illustrates an example with two Web Servers targeting different applications within the Cell. By defining the two Web Servers in the Network Deployment Cell topology and mapping the applications to those specific Web Servers, you can create custom plug-in configuration files that provide this functionality.

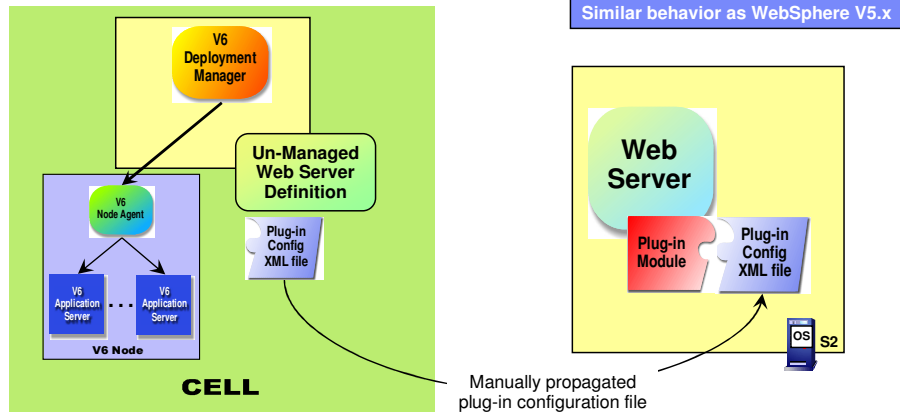
Web Server: Managed Node in a ND Cell



- Web Server defined on a Managed node and therefore are managed by WebSphere DMGr
- Requires WebSphere Managed Node to be created on the Web Server machine
- Node Agent receives commands from DMGr to administer the Web Server
- Automatic push of the plug-in configuration file and in case of IBM Http Server, provides the ability to start, stop an Apache-based IHS
- Common for web servers installed behind a firewall where a WebSphere Node can be installed

This slide shows an example of a web server in a managed node. A Node Agent is present on the server where the Web server is installed. The Web server is managed by the WebSphere Deployment Manager through the node agent. The node agent provides the ability to start and stop the Web Server (for IBM Apache-based Http Server only), as well as to automatically push the plug-in configuration file to the Web Server. The z/OS IHS is Domino-Go based. This scenario is most common for behind a firewall where a WebSphere Node can be installed without as many security concerns.

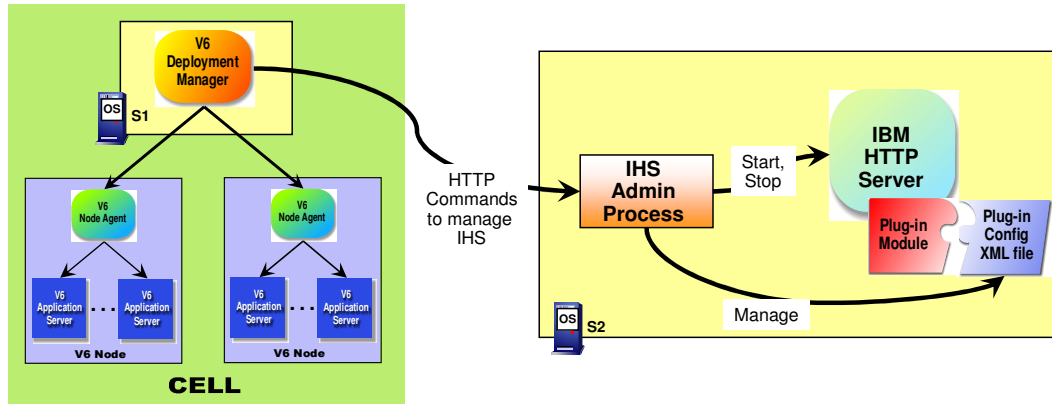
Web Server: Unmanaged Node in a ND Cell



- Web Server not managed by WebSphere - Registered as Unmanaged Node in WebSphere configuration
- Manually propagate the plug-in configuration file from the DMgr machine to the Web Server machine (ftp/copy the file or use Shared file system)
- Still allows WebSphere System Administrator to create custom plug-in files for specific Web Server
- Common for Web servers installed outside firewall or in DMZ, where no WebSphere Node Agent is installed or wanted

In this example, the Web Server is defined in an unmanaged node. This allows a WebSphere System Administrator to create custom plug-in files for that Web Server. This is covered in more detail in the plug-in presentation. When the plug-in is created it must be manually copied to the Web server. The deployment manager has no capability to direct the Web server. The IBM HTTP Server allows for more management in an unmanaged node.

IBM HTTP Server as Unmanaged Node: Special Case



- IBM Http Sever (IHS) V6 can be managed directly by Deployment Manager through the IHS Admin process
 - ▶ Provides ability to start, stop IHS and automatically push the plug-in configuration file to IHS machine

Special code is added to IBM Http Server (IHS) and the Deployment Manager to provide some level of integration. If IHS is defined on an unmanaged Node, you still have the ability to manage IHS through the IHS Admin process. Through the IHS Admin process, the Deployment manager sends commands to start or stop IHS, and to push the plug-in configuration file to the Web Server's location. Again, this only works for the IBM Apache-based IHS. At this point, the starting and stopping of the IHS is not available on the Domino-Go based IHS found on z/OS.

Section

Summary and Reference

The last section is the Summary

Summary

- Multiple Stand-alone Application Server instances can be installed on a single machine
- Network Deployment Cell topology provides flexible topology scheme
 - ▶ Supports multiple platforms in a single cell
 - ▶ Mixed version Node support enables flexible migration
- Web Servers can be centrally managed by Deployment Manager

WebSphere Application Server V6 builds on top of the topology supported in Version 5, and adds some significant enhancements to it in terms of supporting the Web Server nodes on managed or unmanaged nodes. The install has been simplified to easily create the topology that you need.

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
eIogo business	DB2	Series	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.

© Copyright International Business Machines Corporation 2004. 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.