WebSphere Network Deployment 5.1 On Linux for zSeries

WebSphere software

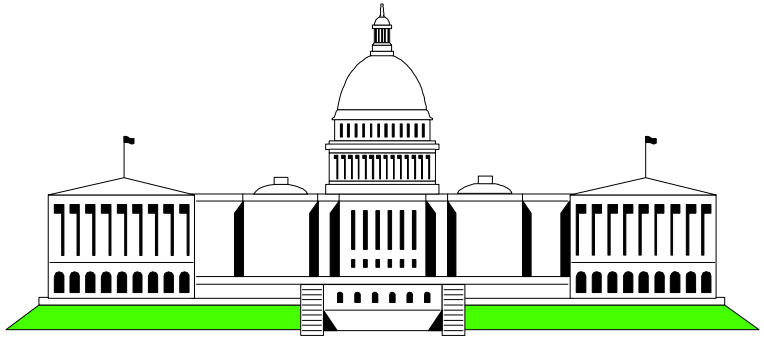

Washington System Center

Acknowledgments

The material in this presentation was taken in part from the WebSphere Tutorials.

Topics

- Facts, features, fit
- Cells, nodes, node agents and clusters
- Installing the product
- Federation
- High Availability and failover

ND Features above base

- Web Services support
  - Industry-leading
  - Advanced Web services features
    - UDDI Registry Acts as a repository that allows storage of business units that describe basic Web services
  - Web Services Gateway Enables Web services invocation by users from outside the firewall with the benefit of robust security protection
- Enhanced Problem resolution capabilities
  - Provides first failure data capture to report and analyze problems as they occur
- High availability and Enhanced workload management
  - Dynamic caching, performance management tools and Edge Server Component
    - Cluster support
    - Failure bypass
    - Intelligent workload distribution across a cluster
    - Centralized security

Edge Server component expanded - p1

WSC
Advanced
Technical Support

- Designed for highly available, high volume environments
  - Includes sophisticated load balancing, caching and centralized security capabilities based on WebSphere Edge Server.
    - Load-balancing component provides a scalable solution for distributing and routing HTTP, servlet, and Enterprise JavaBean™ (EJB) requests.
      - Distribute workload across multiple servers through sophisticated load balancing and clustering capabilities
      - Enables isolation of application servers to avoid single points of failure
    - Custom advisors can be used to load-balance requests based on unique application and platform criteria.
    - Consultants. To extend the load balancing capabilities beyond purely a WebSphere Application Server environment, consultant code can be used to optimize server performance within a Cisco or Nortel infrastructure.

WebSphere software

Edge Server component expanded - p2

WSC
Advanced
Technical Support

- Enhanced Caching (dynamic). edge-of-network caching capability
  - Improves response time by offloading back-end servers and peering links.
    - JSPs etc.
  - Caching proxy provides the following plug-in support:
    - Tivoli Access Manager plug-in
      - A plug-in that allows users to exploit an LDAP-based repository for storing user authentication and authorization information.
        - Allows centralized security
        - WAS ND and Tivoli Access manager more tightly integrated
      - An authentication/authorization plug-in that allows independent software vendors (ISVs) to exploit third party authentication/authorization mechanisms such as RADIUS or SecurID tokens.
- Edge Side Includes (ESI) support.
  - Markup language
  - Dynamic web page assembly

WebSphere software

The big picture

WSC
Advanced
Technical Support

- Where does WebSphere Network Deployment fit in?

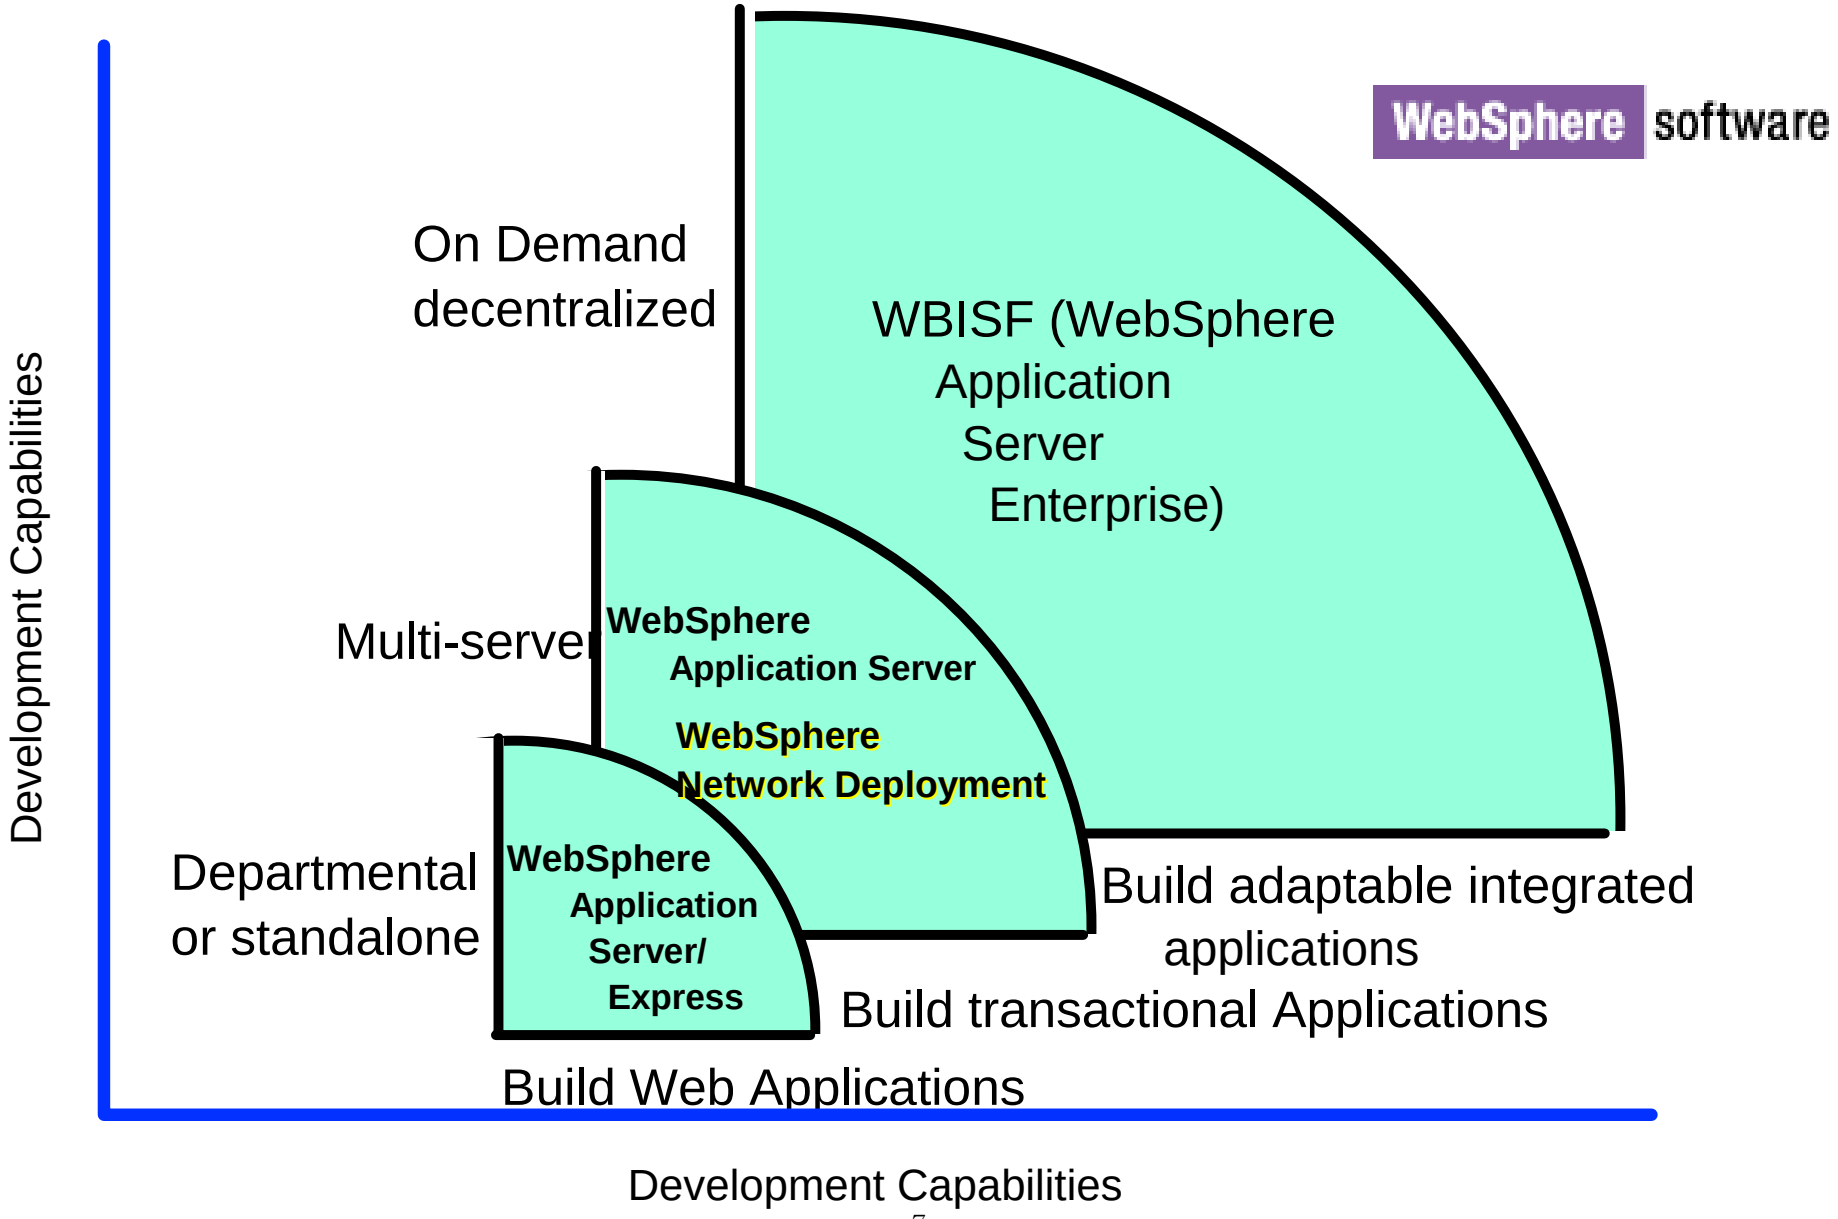

Concepts - Nodes

WSC
Advanced
Technical Support

- A node
  - A logical grouping of the managed servers on a physical machine.
  - Usually corresponds to a physical computer system with a distinct IP host address, but multiple nodes may exist at the same IP address.
    - Node names are usually identical to the host name for the computer, but don't need to be.
  - Names must be unique

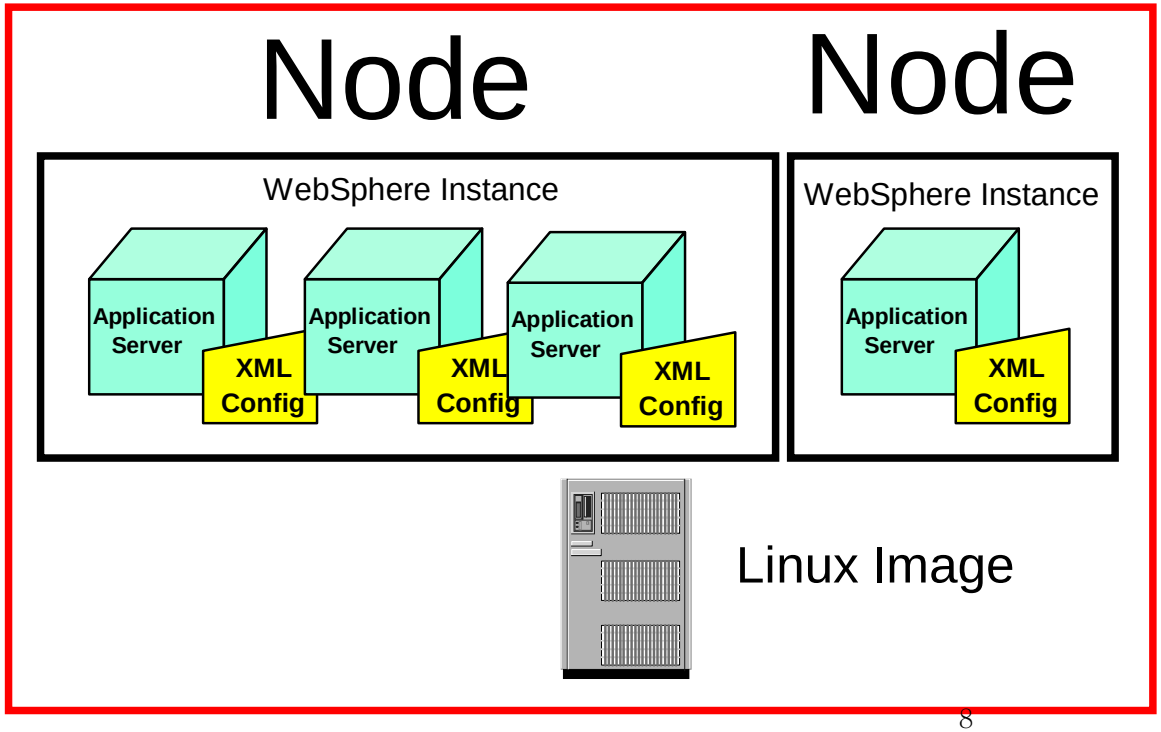

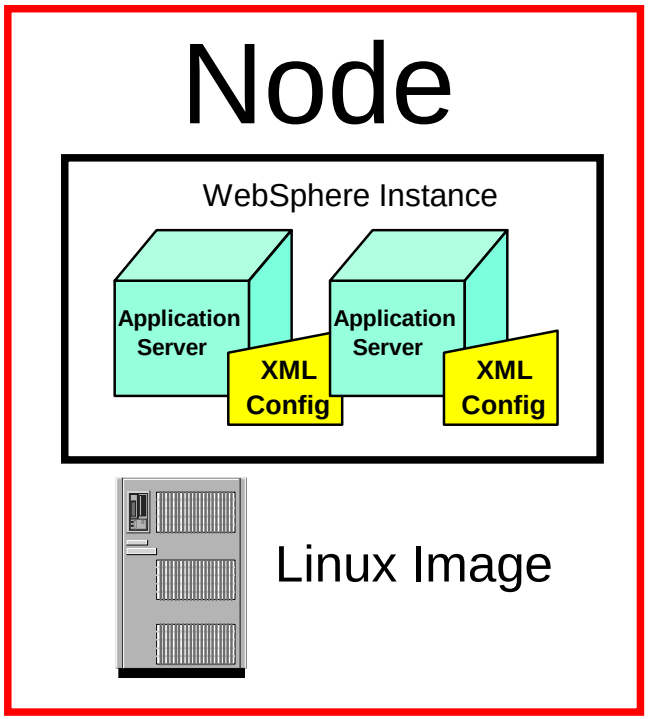

Concepts - Node Agents

WSC
Advanced
Technical Support

- A node agent manages all WebSphere Application Servers on a node.
  - Represents the node in the management cell.
  - Can not exist without WebSphere ND
  - One node agent for each node
  - May be managed from the console session or command line
    - Must be restarted from command line if agent is down

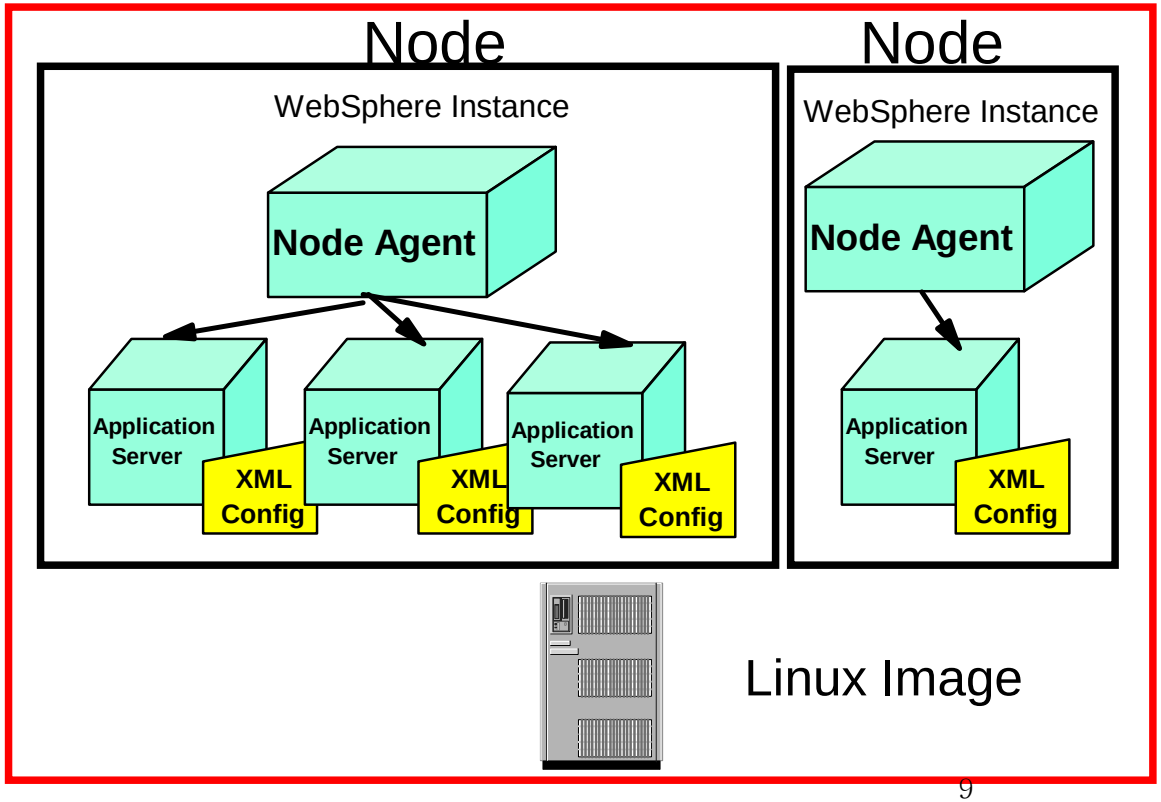

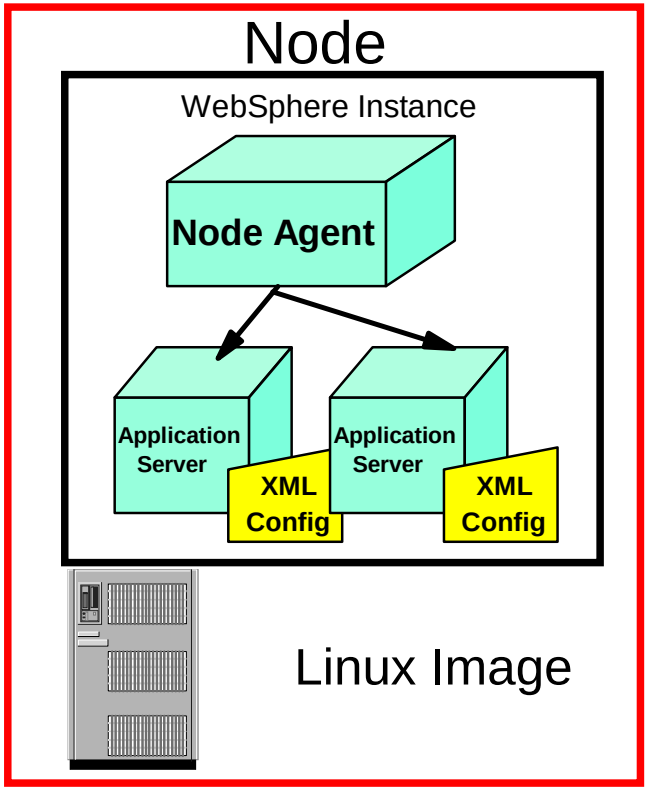

Cells

WebSphere software

- A cell
  - A configuration concept
    - Consists of arbitrary, logical groupings of one or more nodes in a WebSphere Application Server distributed network.
      - A way for administrators to logically associate nodes with one another.
      - Administrators define the nodes that make up a cell according to whatever criteria makes sense in their organizational environments.
      - Administrative configuration data is stored in XML files
      - A cell retains master configuration files for each server in each node in the cell.
      - Each node and server also have their own local configuration files.

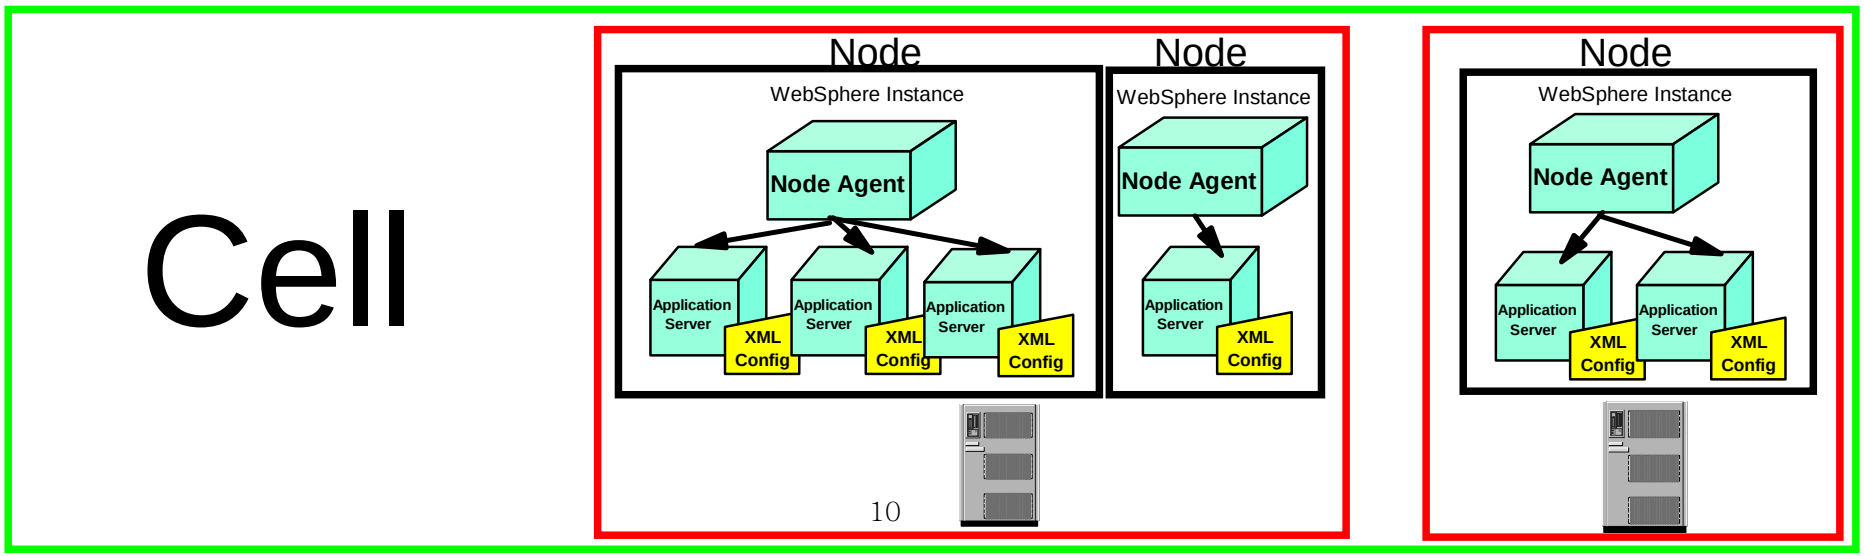

Deployment Manager

- Administrative agents that provide a centralized management view for a cell and all nodes in that cell.
  - Hosts the administrative console.
    - Provides a single, central point of administrative control for all elements of the entire WebSphere Application Server distributed cell.
    - Provides management of clusters and the management of workload balancing of application servers across one or several nodes.

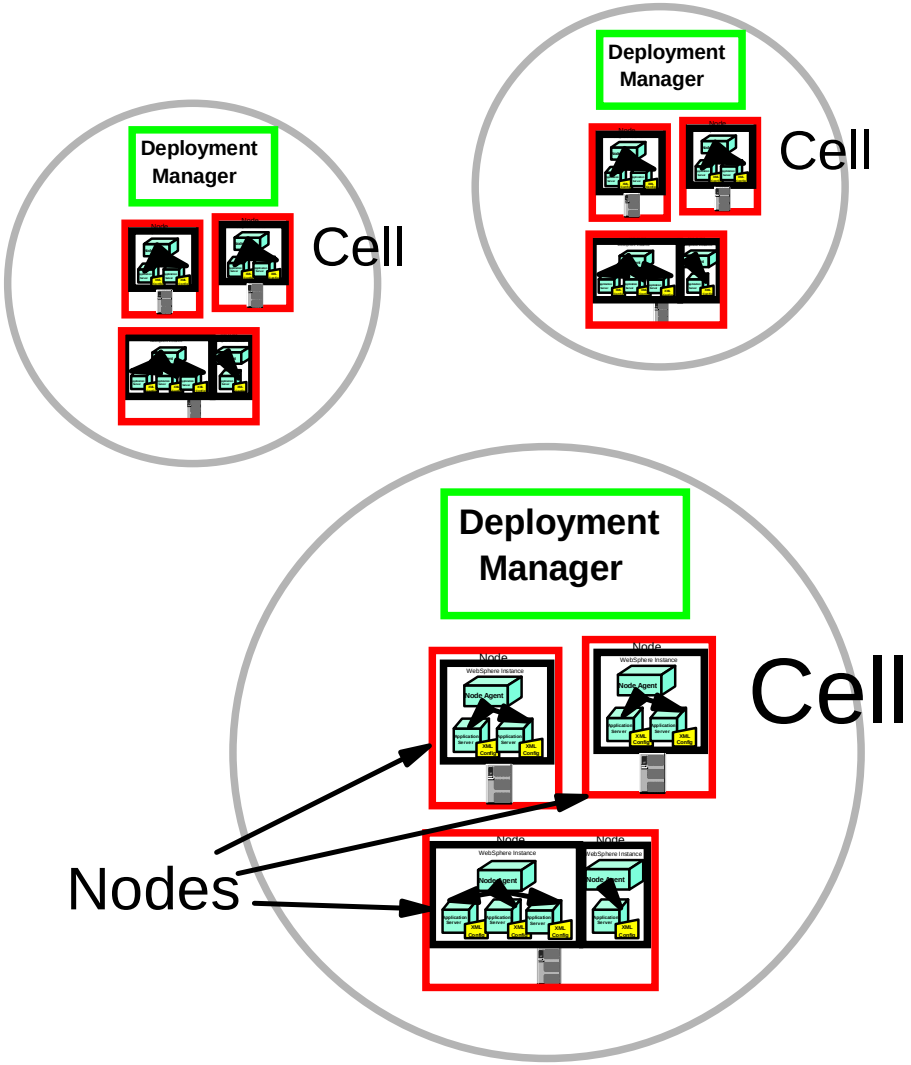

WebSphere ND V5 Model - bringing it all together

WSC Advanced Technical Support

- Cell (Domain in previous releases)
  - Network of nodes with logical view for administration
- Deployment Manager
  - Manages multiple nodes
- Node Agent
  - Resides on single node
  - Manages Servers on node
  - Comes with base, but not used unless ND is used
- Managed Process or Server
  - Each server having own JVM
    - Application Servers
    - JMS Server
    - Generic Servers

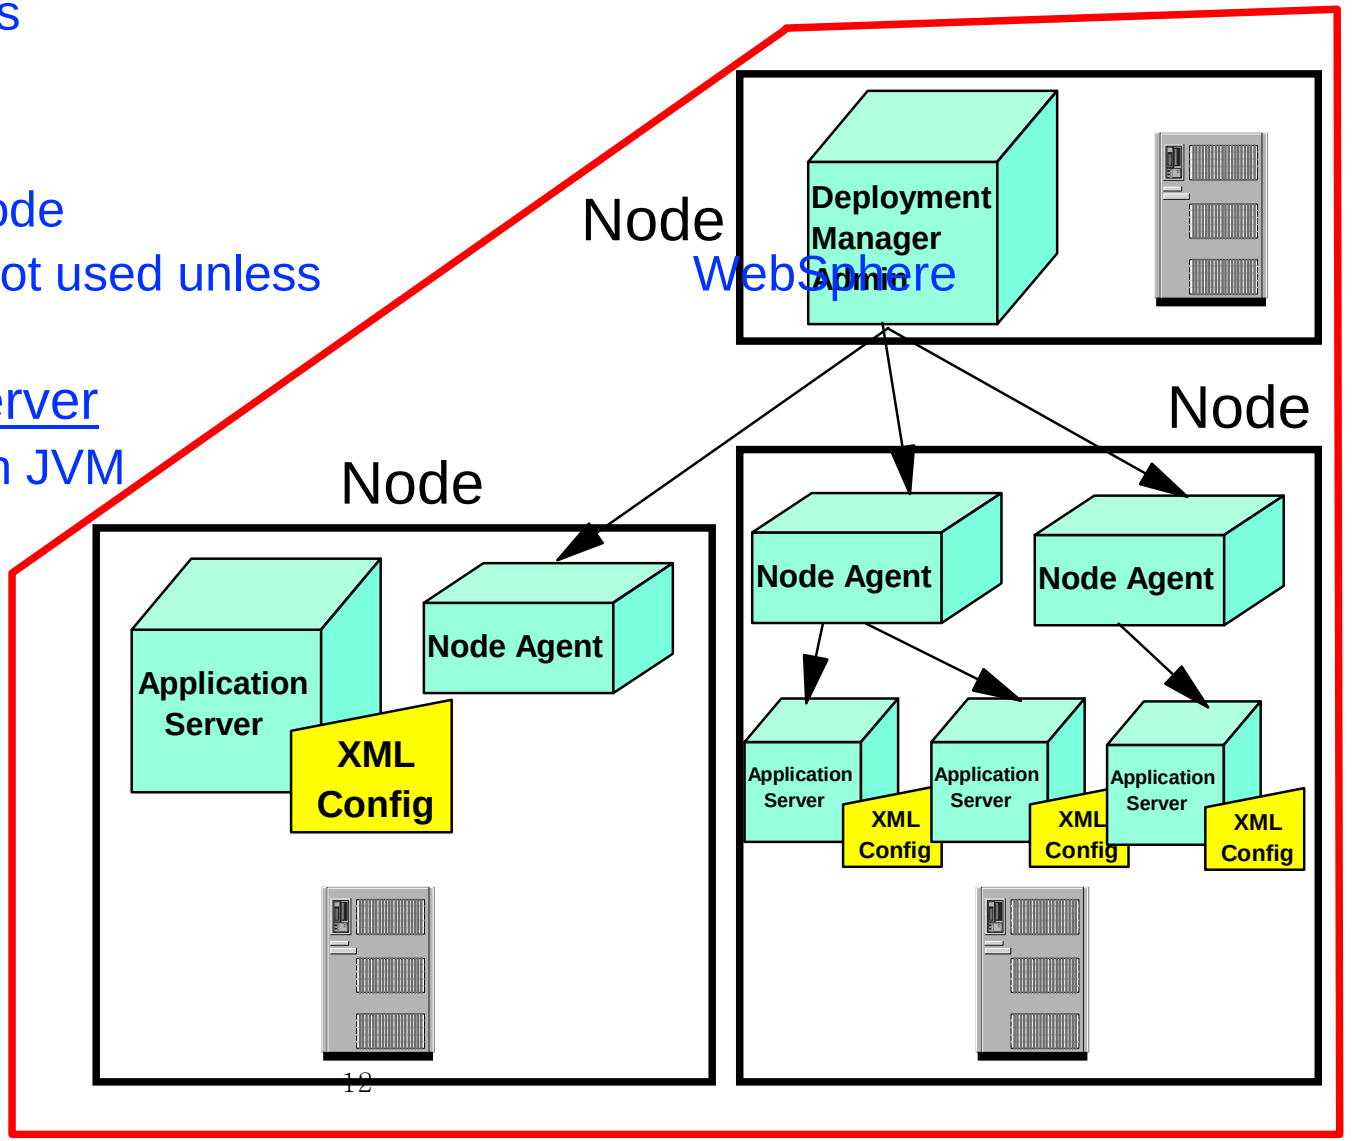

Different models

- WebSphere Network Deployment may be standalone on an image or may coexist with a single or multiple base WebSphere Application Servers
  - The coexistence will involve using alternate ports to avoid conflicts with existing WebSphere servers already installed on the image

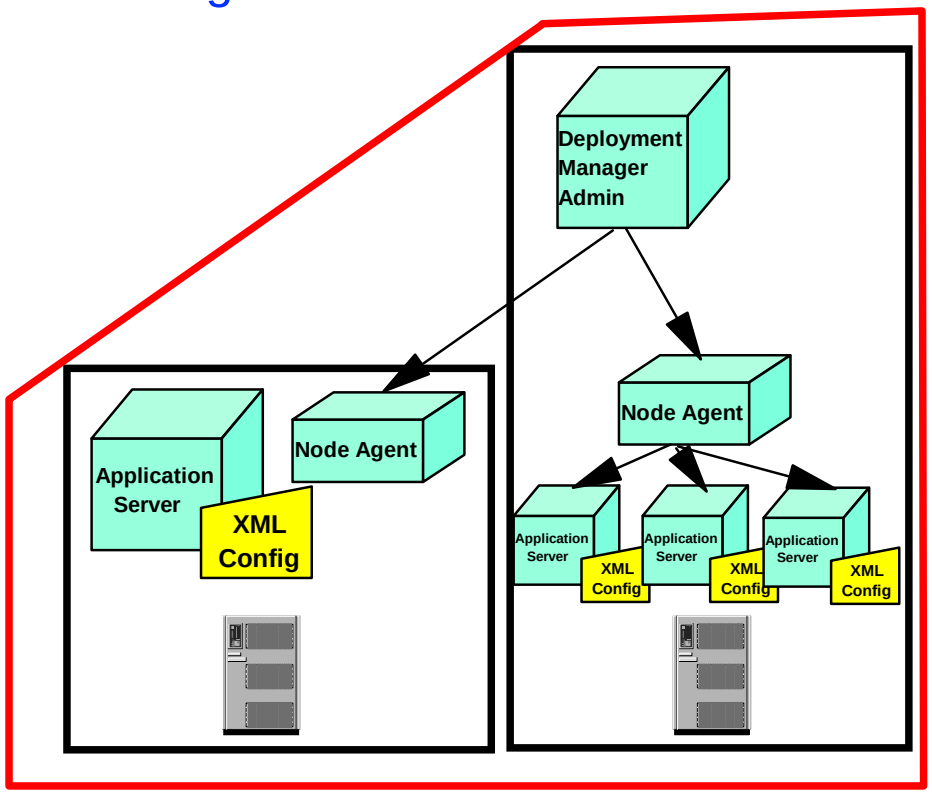

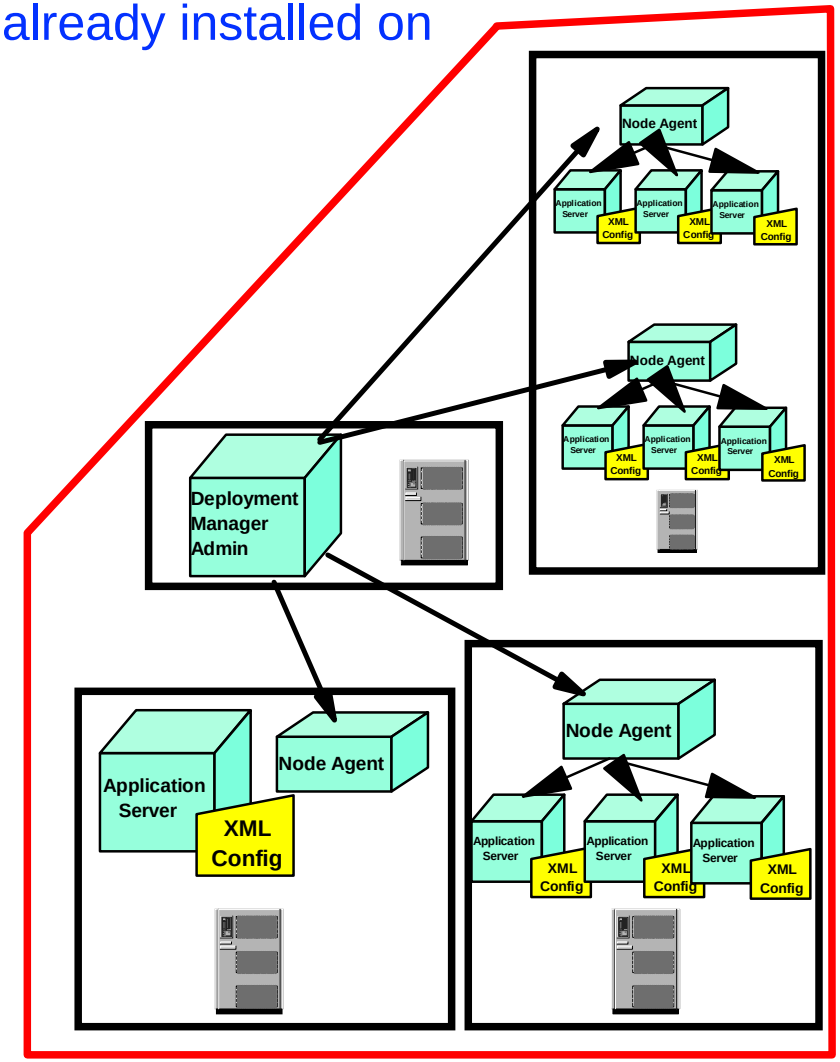

The goal is high availability

High availability (HA)

- Definition - An application or service that is provided in a near-continuous fashion.
  - A solution must include the concepts of
    - **Process availability**
    - **Data availability**

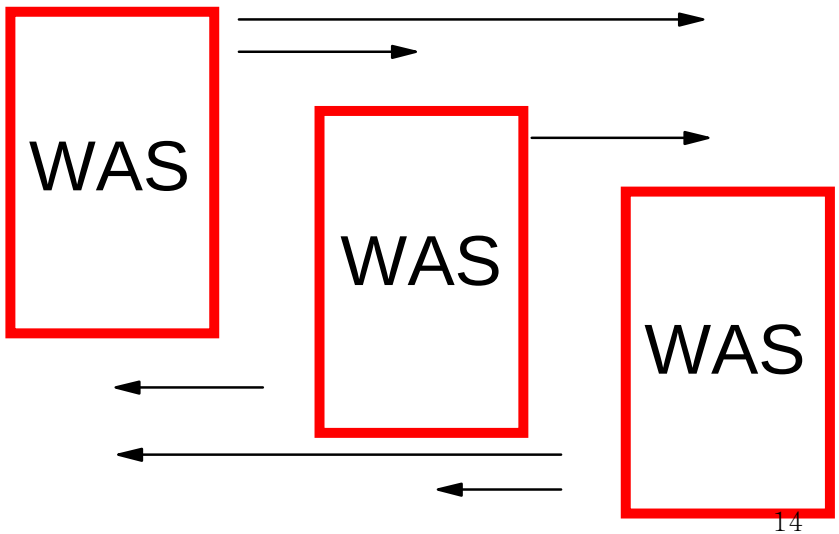

WebSphere software

Process Availability

WSC
Advanced
Technical Support

- A process is considered available when there is a server available to service a request. WebSphere implements process availability as follows:
  - Clusters. The ability to service any given request is supported by a number of similar processes.
  - Failover. When a process becomes unavailable, requests are directed away from the unavailable process and toward the other, available, processes in the cluster.
  - When there are no available processes to service a request, a programming model exists that allows clients to perform recovery, if possible, and to resubmit requests, if necessary. This path may lead to loss of data and to the system being down for some time, depending on the nature and seriousness of the failure.

WebSphere software

Data Availability

- Occurs when data is preserved across process failures and is available for processes that continue to be available. Data availability requires:
  - Failover to redundant copies of data, when the primary source of data for a process is no longer available
  - Procedures for recovery of data, when the catastrophic errors occur in system (e.g. fire)

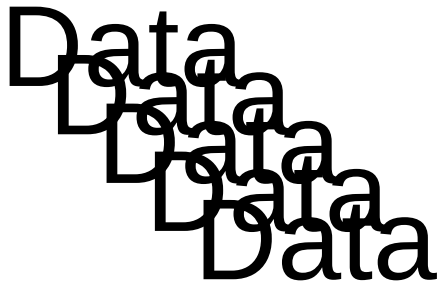

Clusters - what are they in detail?

WSC
Advanced
Technical Support

- A cluster is a collection of servers working together as a single system to ensure that mission-critical applications and resources remain available to clients.
  - Can be visualized as consisting of several logical components:
    - Application servers (processes or “clones”) that service requests
    - Administrative servers (processes) that that administer the application servers
    - Data repositories (databases or native operating system files) that store information used by the Application and Administrative servers process failures and is available for processes that continue to be available.

WebSphere software

Clusters types

- Vertical
  - Refers to the practice of defining multiple clones of an application server on the same physical machine
  - May be on a single application server, which is implemented by a single JVM process and cannot always fully utilize the CPU power of a large machine and drive CPU load up to 100%.
  - Or may include more than one instance of a WebSphere Application Server on the same image to provide a straightforward mechanism to create multiple JVM processes, that together can fully utilize all the processing power available as well as providing process level failover.

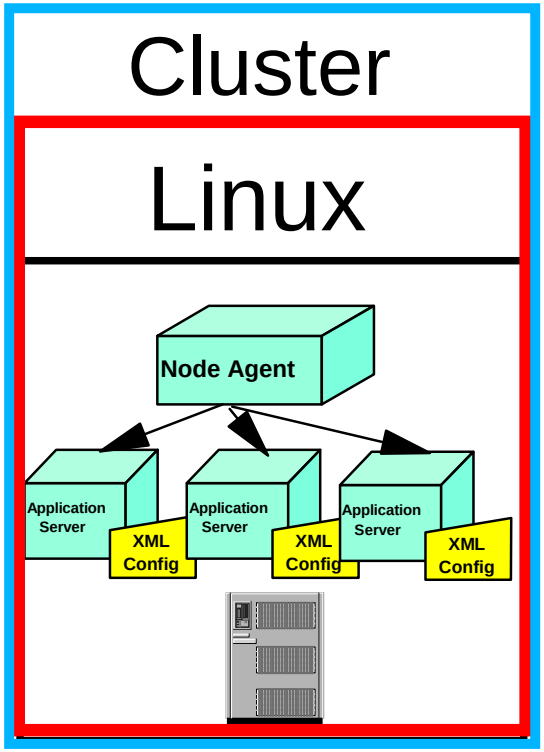

OR

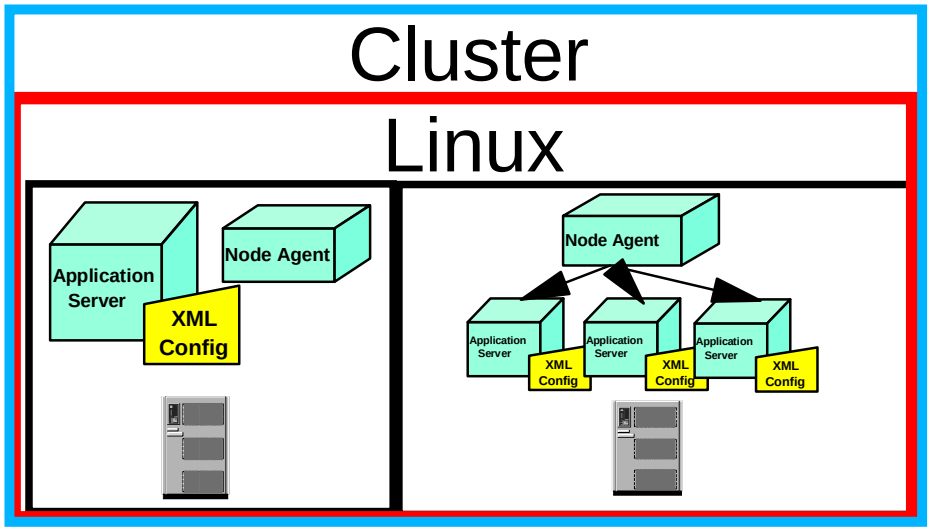

Clusters types continued

- Horizontal
  - Refers to the more traditional practice of defining clones of an application server on multiple physical machines, thereby allowing a single WebSphere application to span several machines while presenting a single system image.
  - Horizontal cloning can provide both increased throughput and failover redundancy.

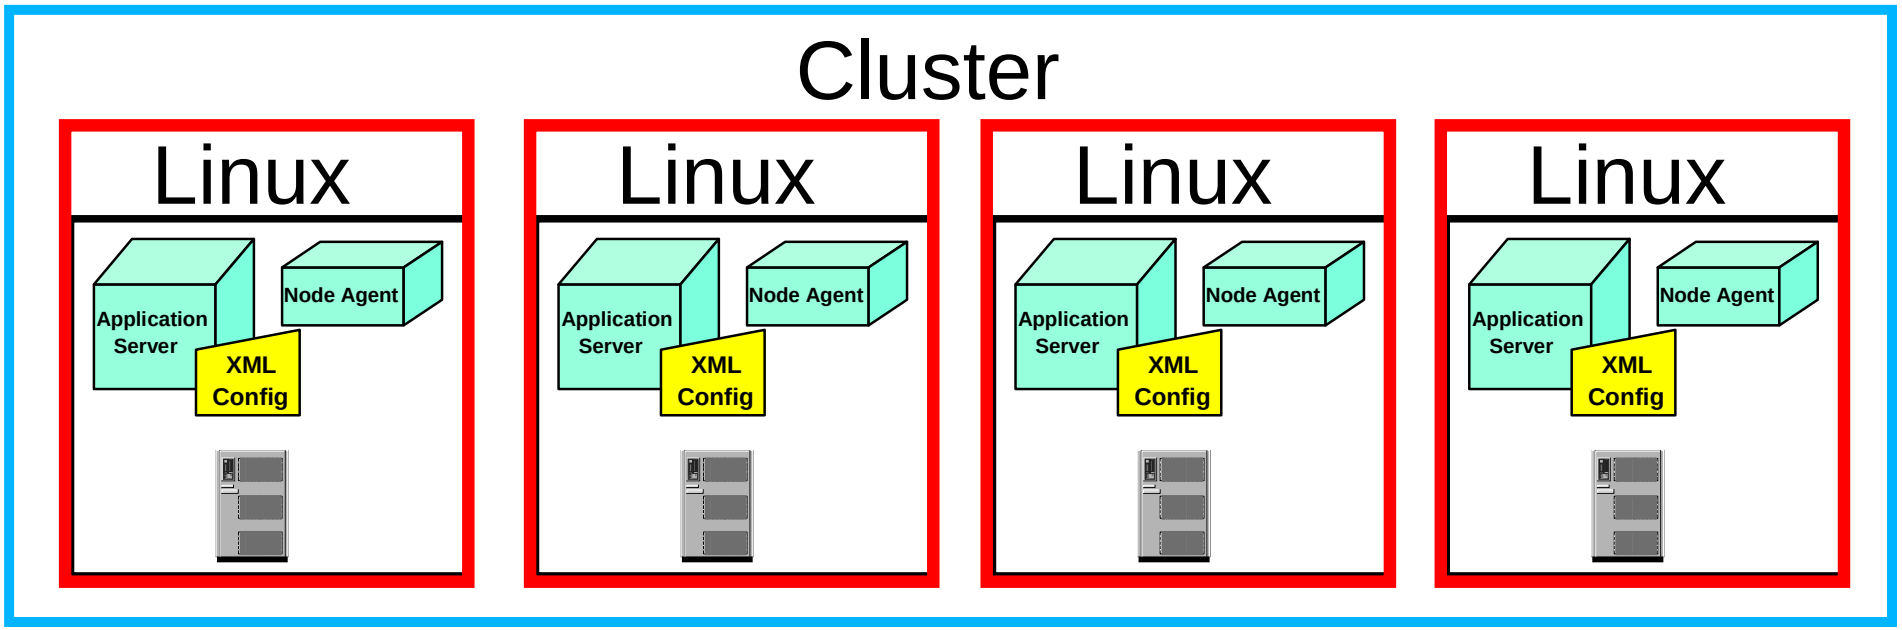

Clusters - the big picture

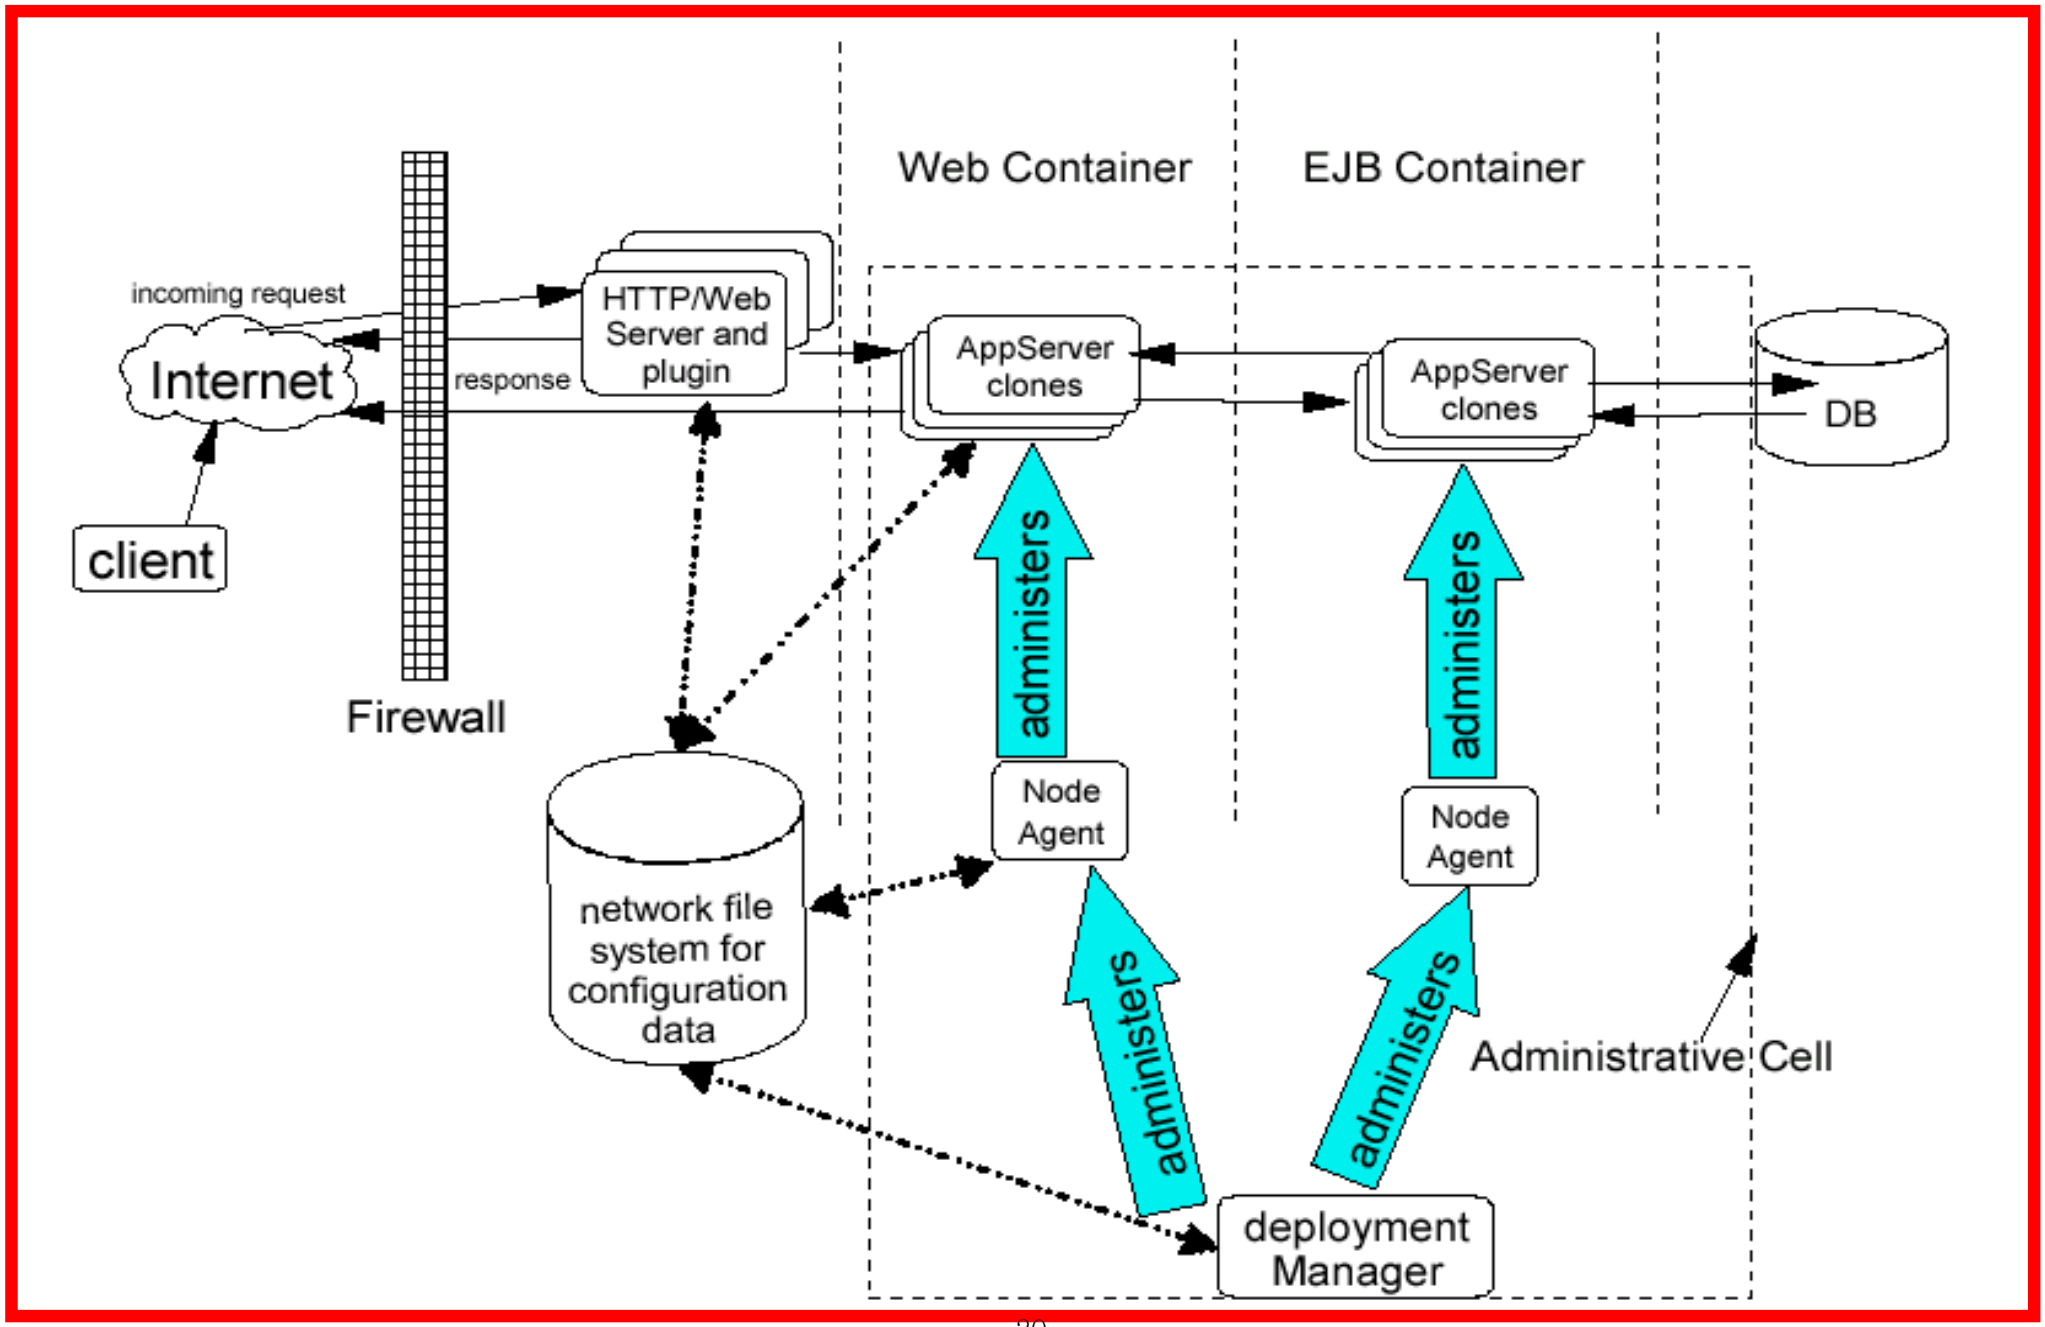

Installation

- WebSphere Network Deployment installs in the same manner as the base WebSphere Application Server.

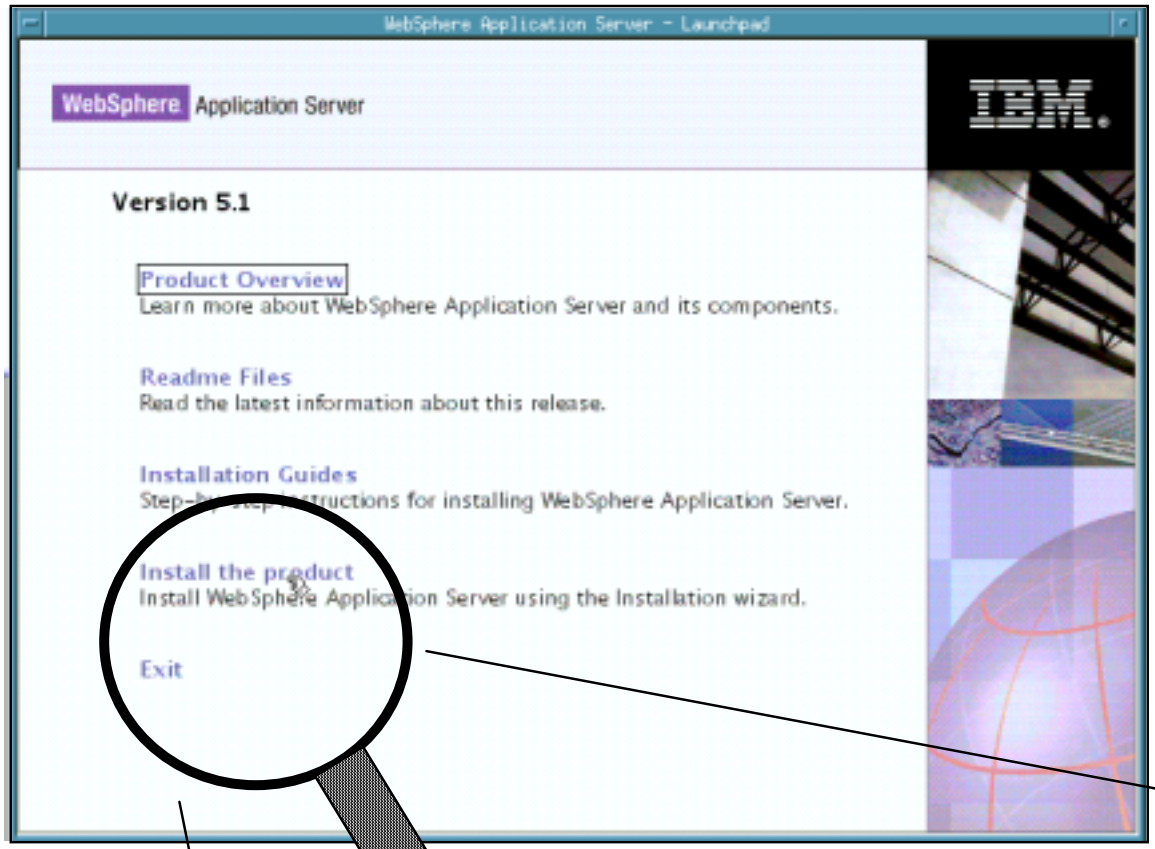

Xterm

>launchpad.sh

Getting started

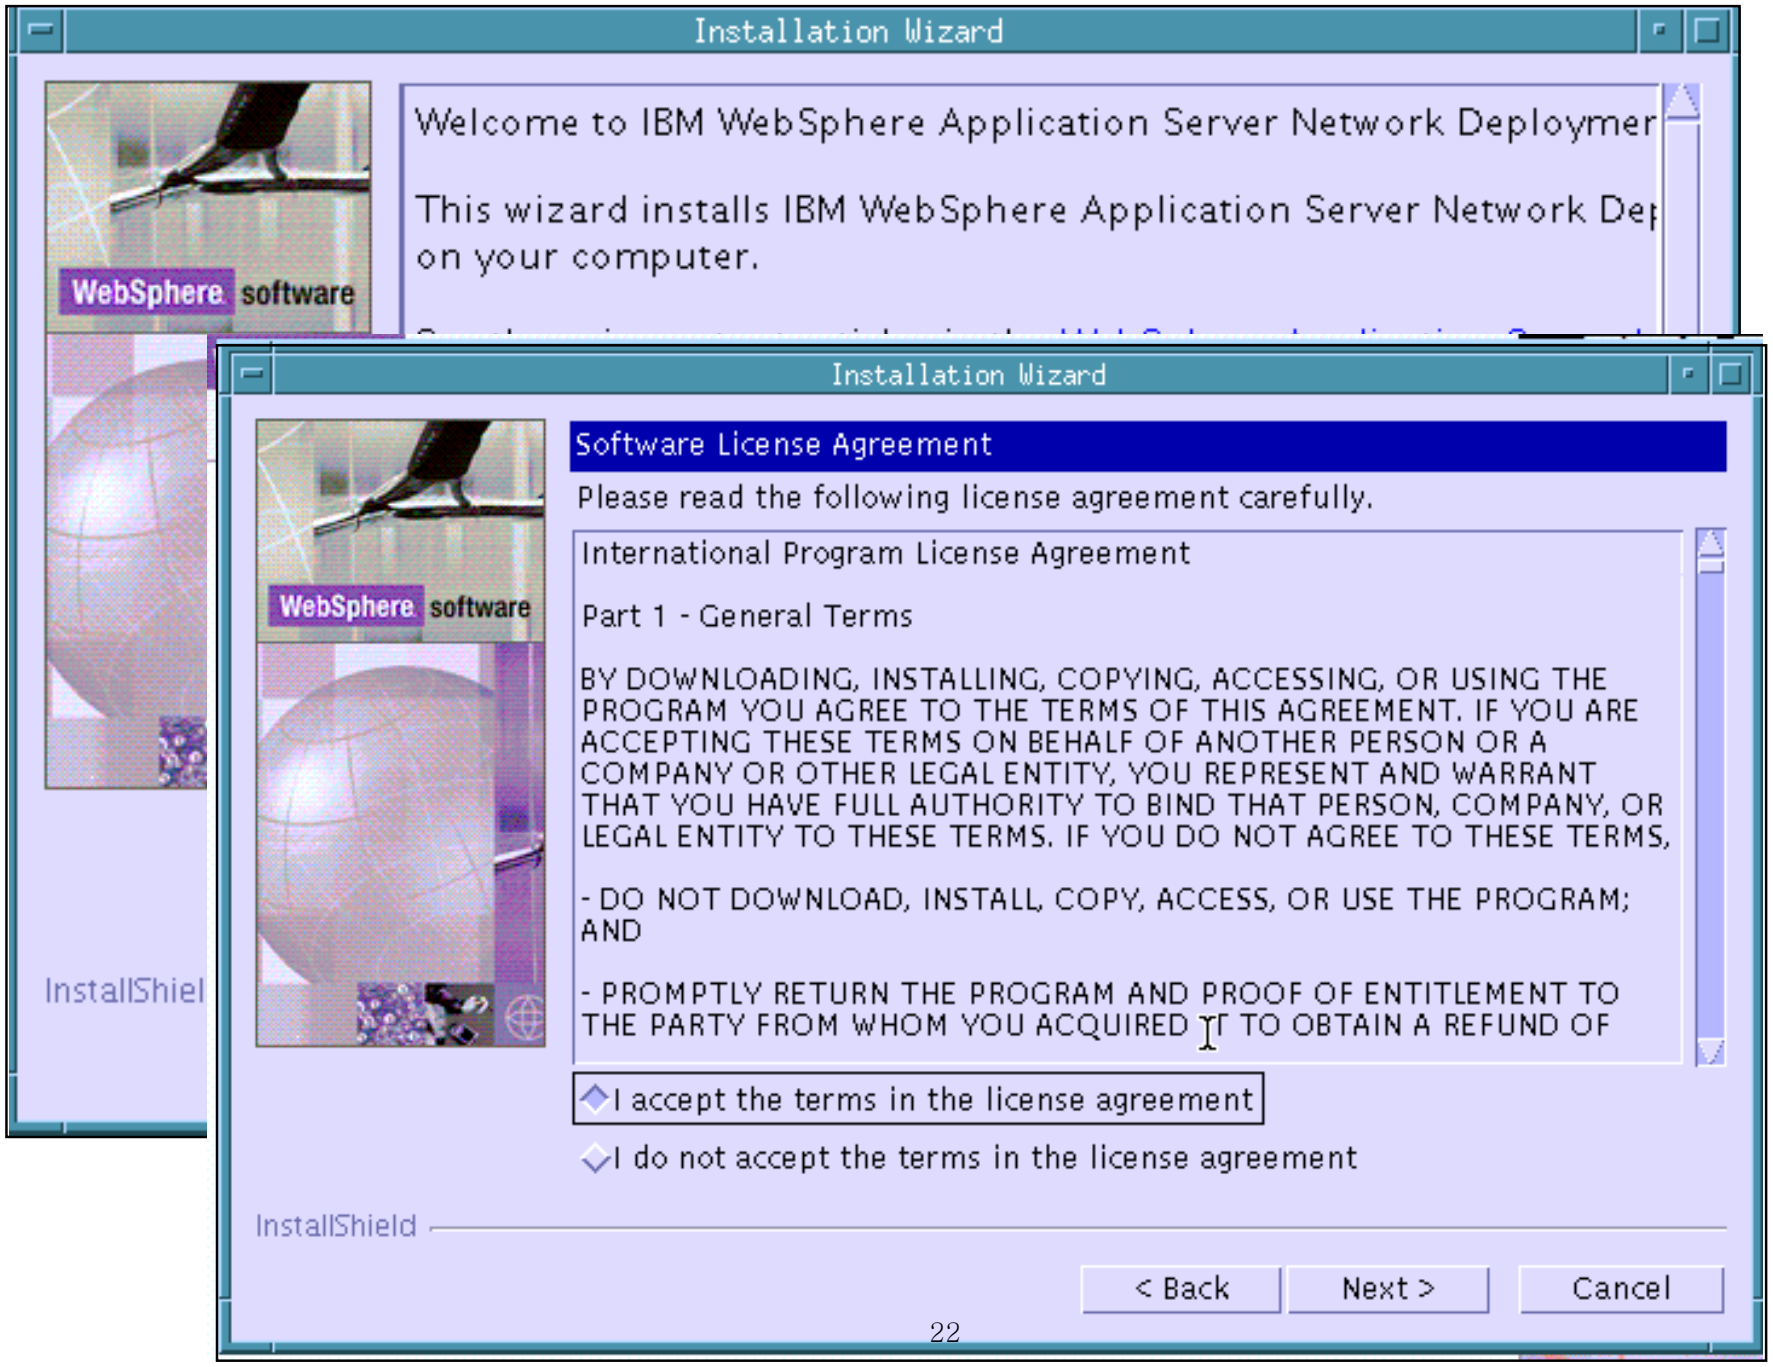

Choose what to install

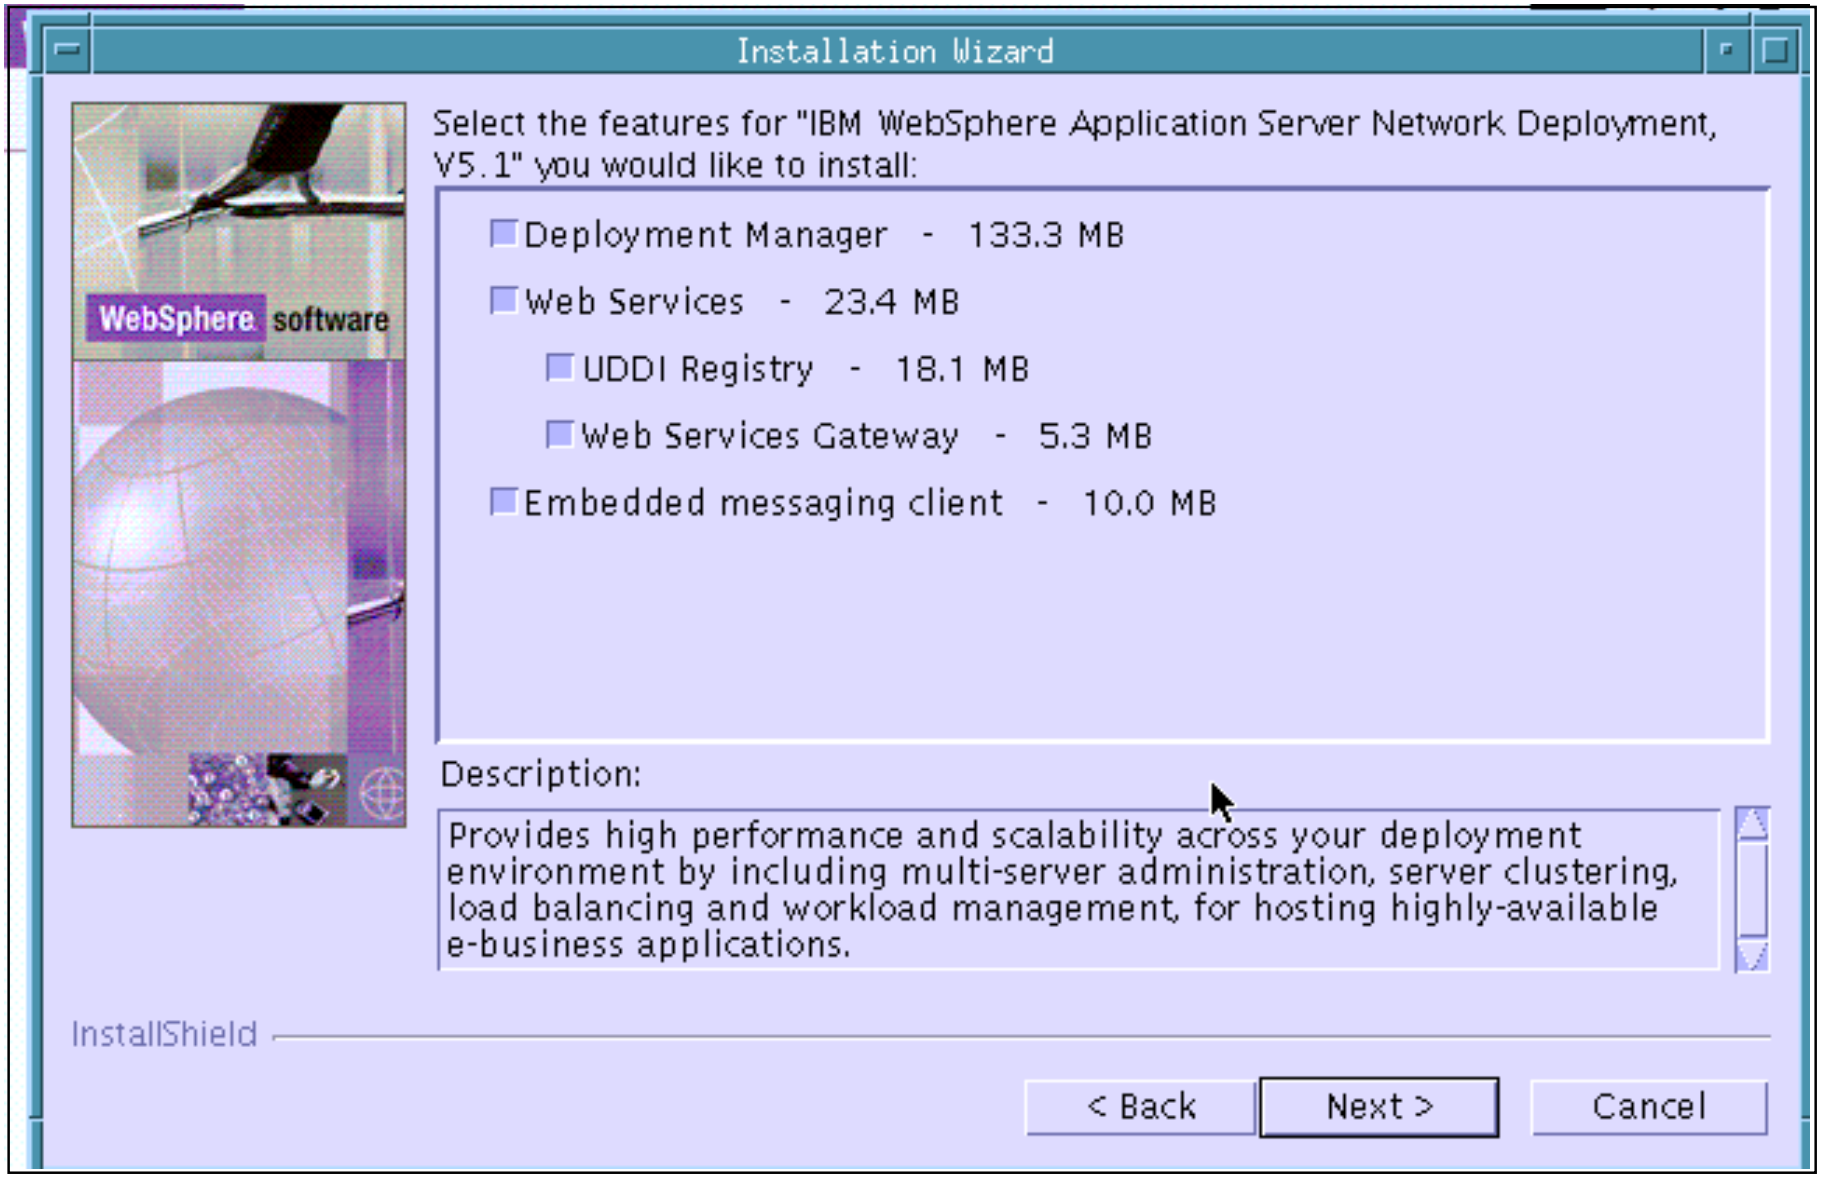

Installation

- May run
  - Standalone
  - With other WAS Application Servers
    - Coexistence required

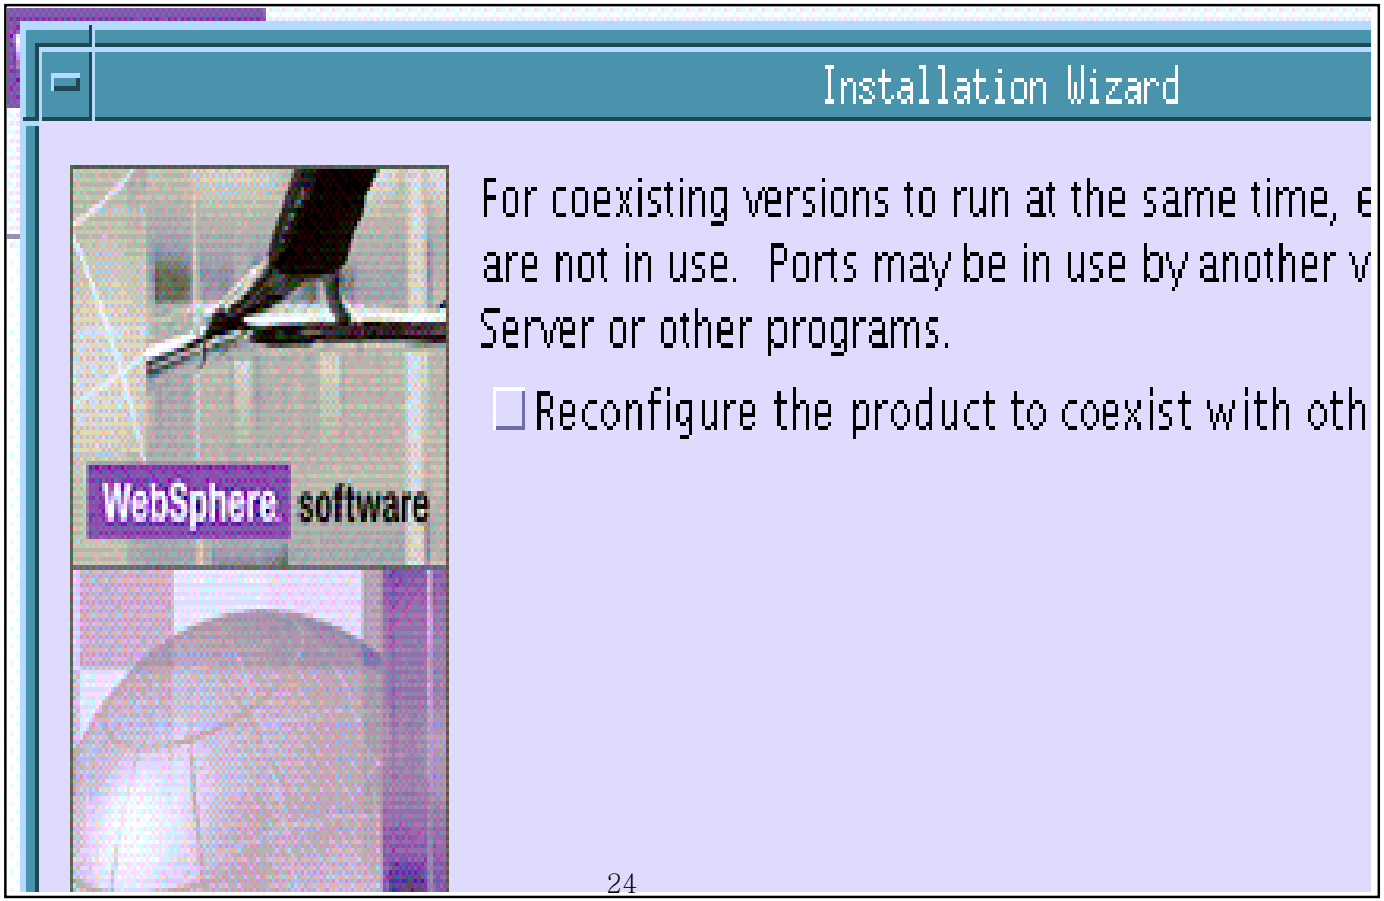

Coexistence

- Must use alternate ports
  - Default ports in use by one or more base servers

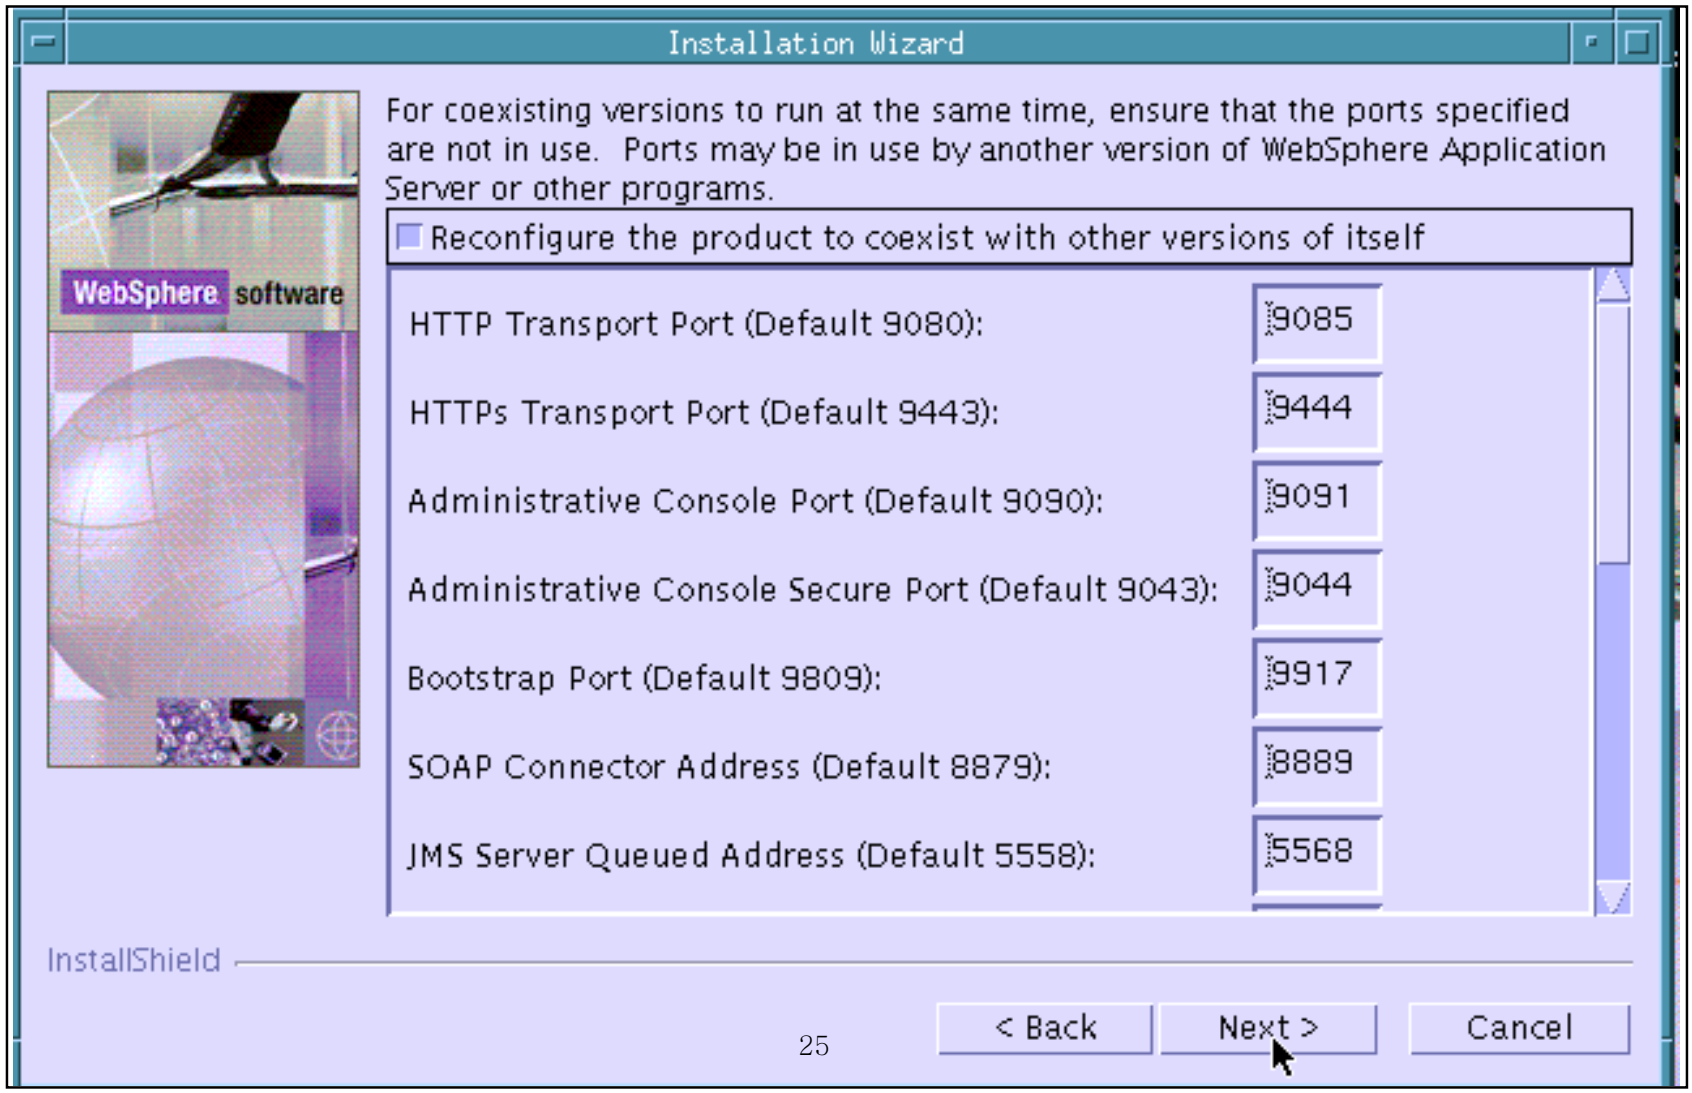

Select Source Location

- May use an existing directory as the directory root or any directory of choice

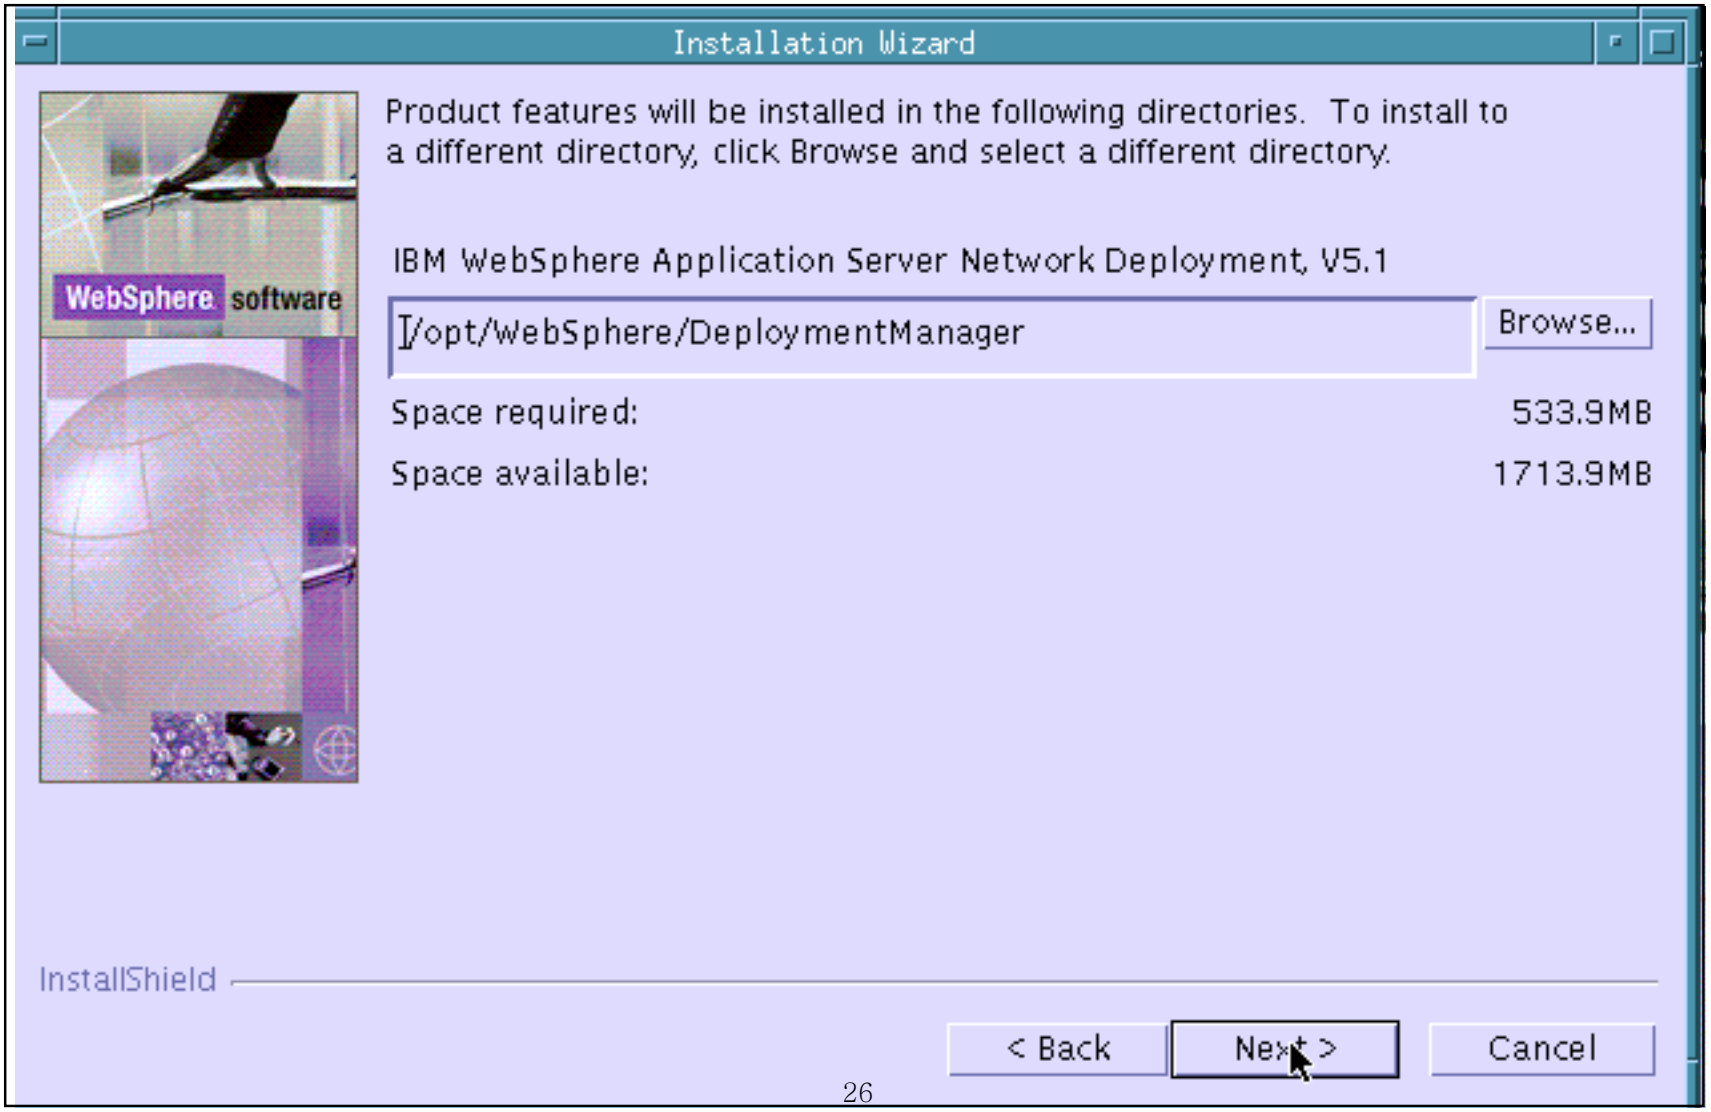

Select Identity Names

- May use suggested defaults or any names of choice

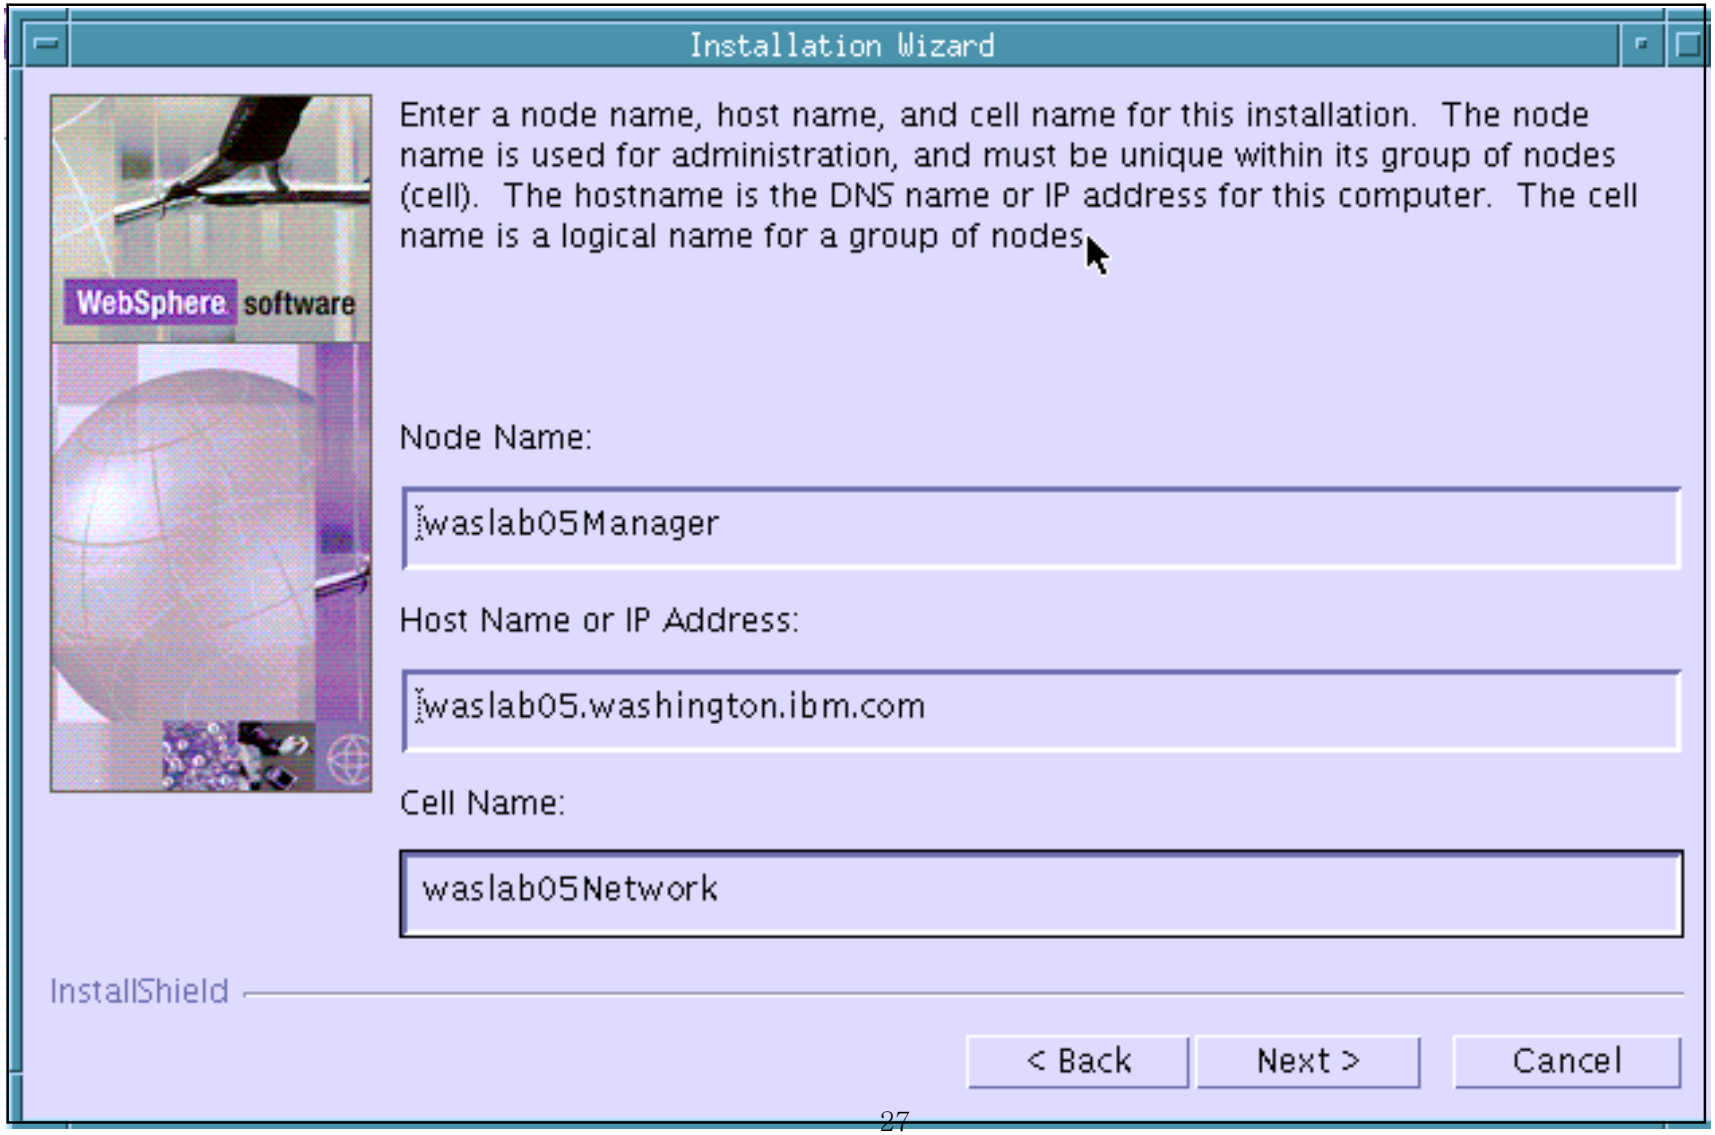

Start the server

- Deployment Manager Server
  - May be started and verified from firststeps screen
  - May also be started from
    - /opt/WebSphere/DeploymentManager/bin startManager.sh

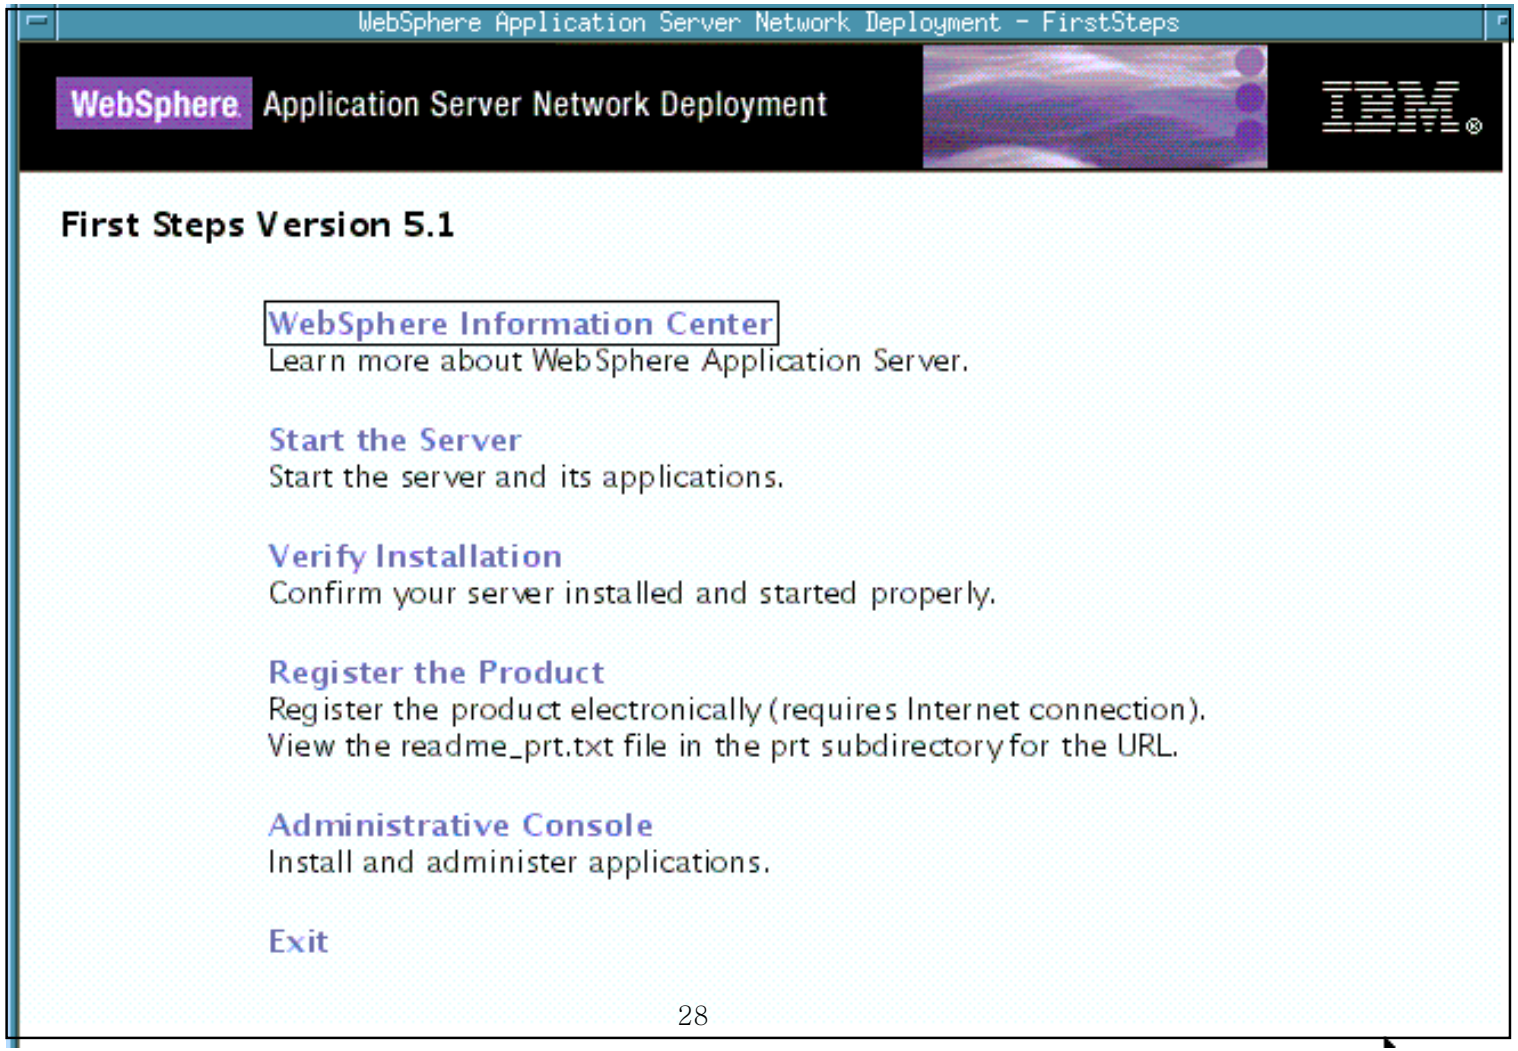

Federating Nodes

WSC
Advanced
Technical Support

- The console session is brought up using a browser session just as with a base application server. However, the port may vary depending on if other servers exist on the same image.

ip-address/909x/admin

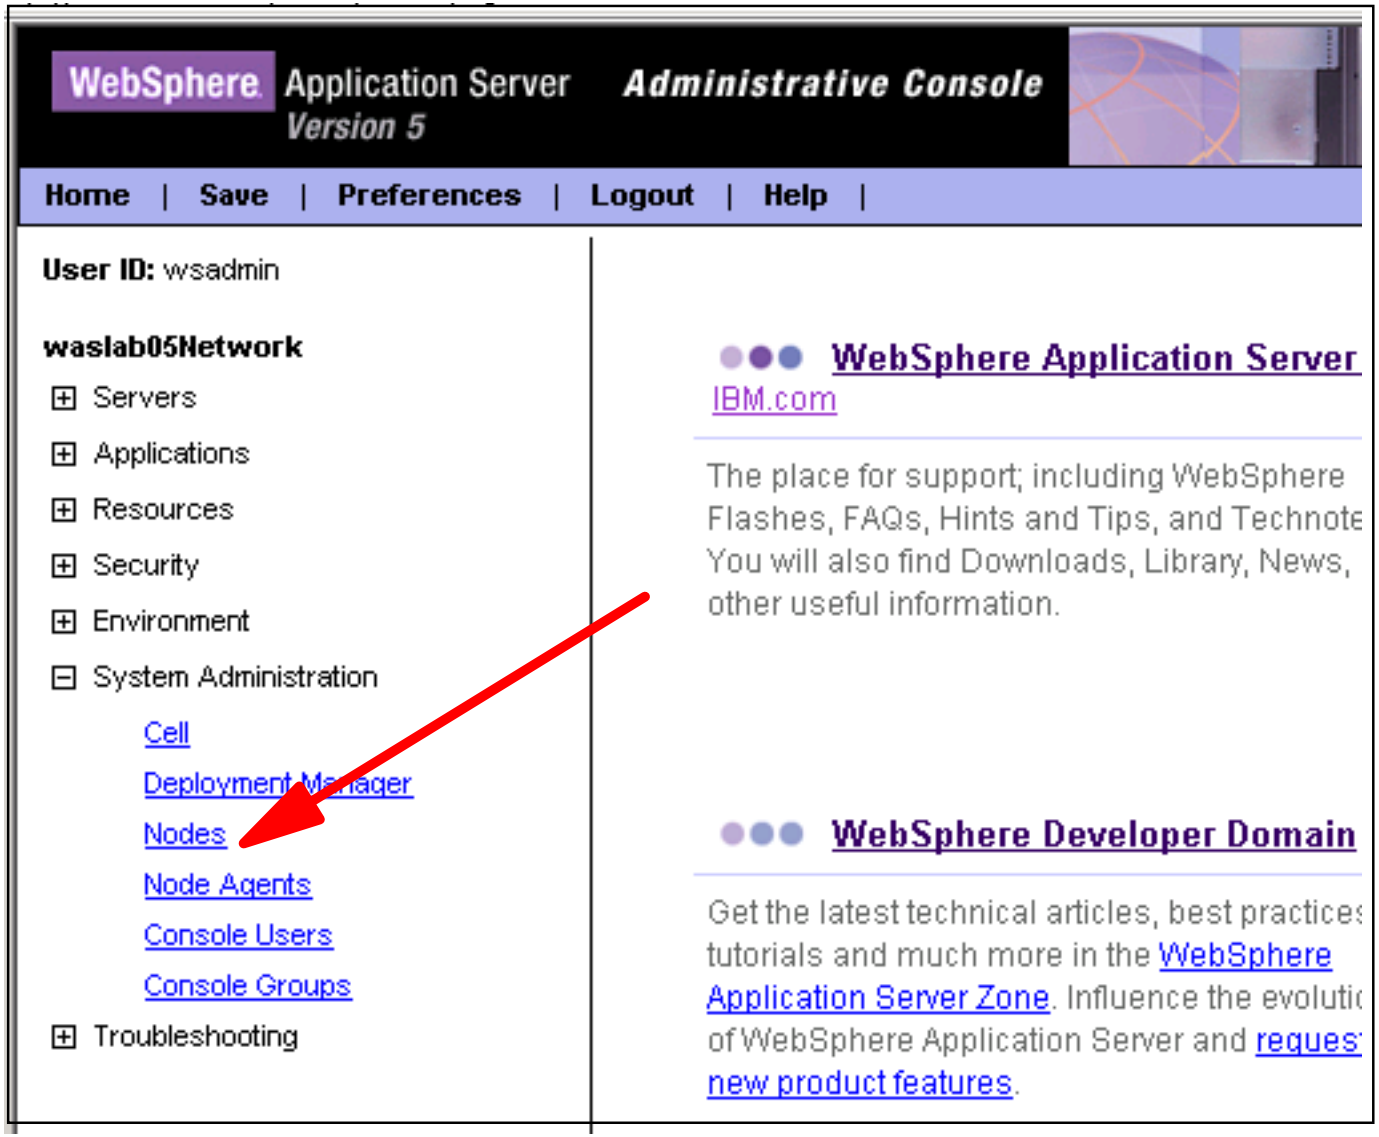

Add Nodes continued

WSC
Advanced
Technical Support

- One node will already be defined as specified during the install process
- Add a node for each base WebSphere on various images

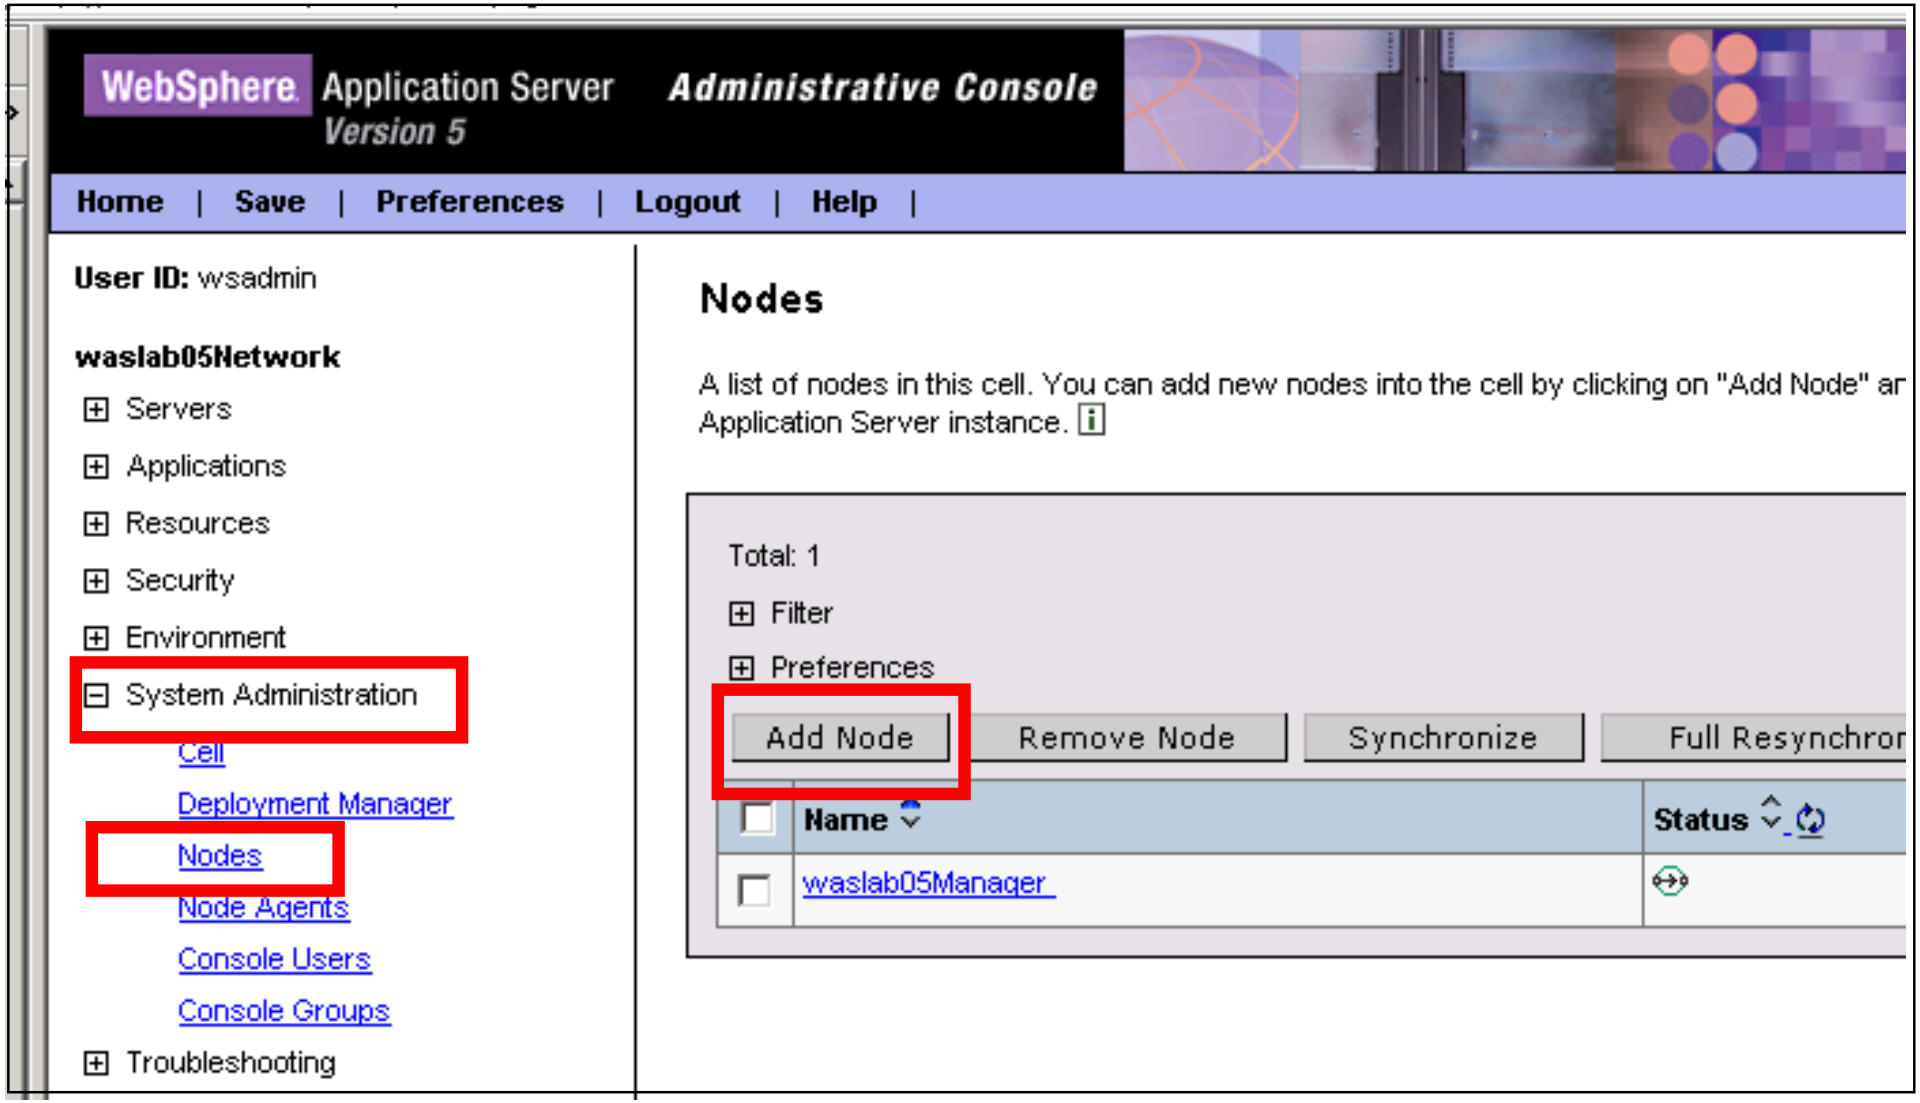

Add Nodes continued

- Provide the hostname
- Include applications is desired

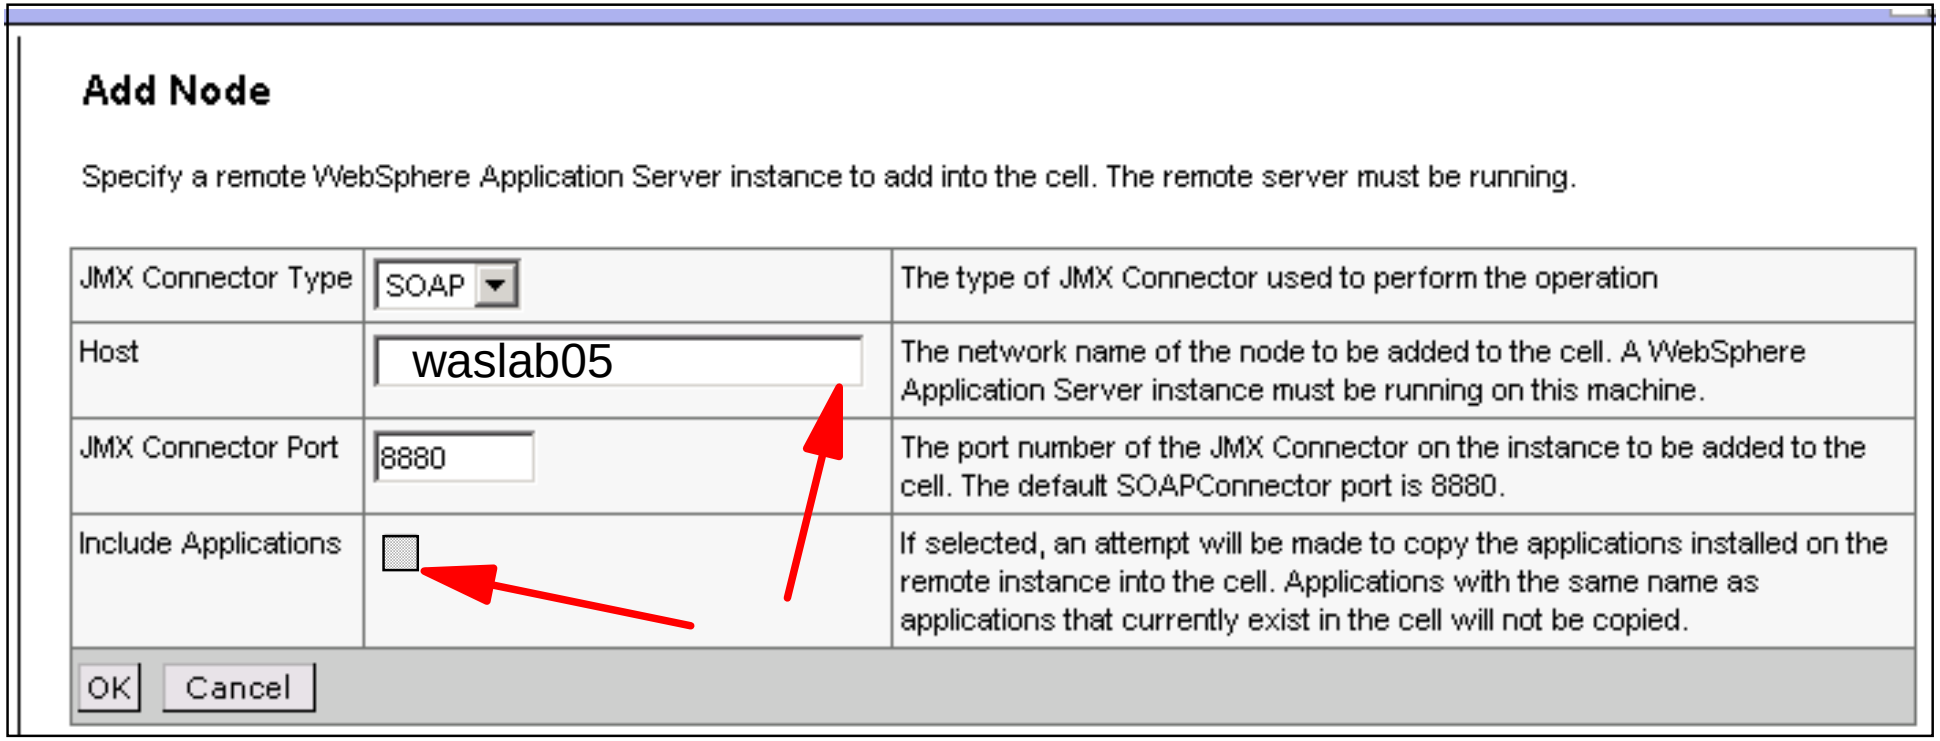

Add Nodes continued

WSC
Advanced
Technical Support

- Messages
- New nodes added

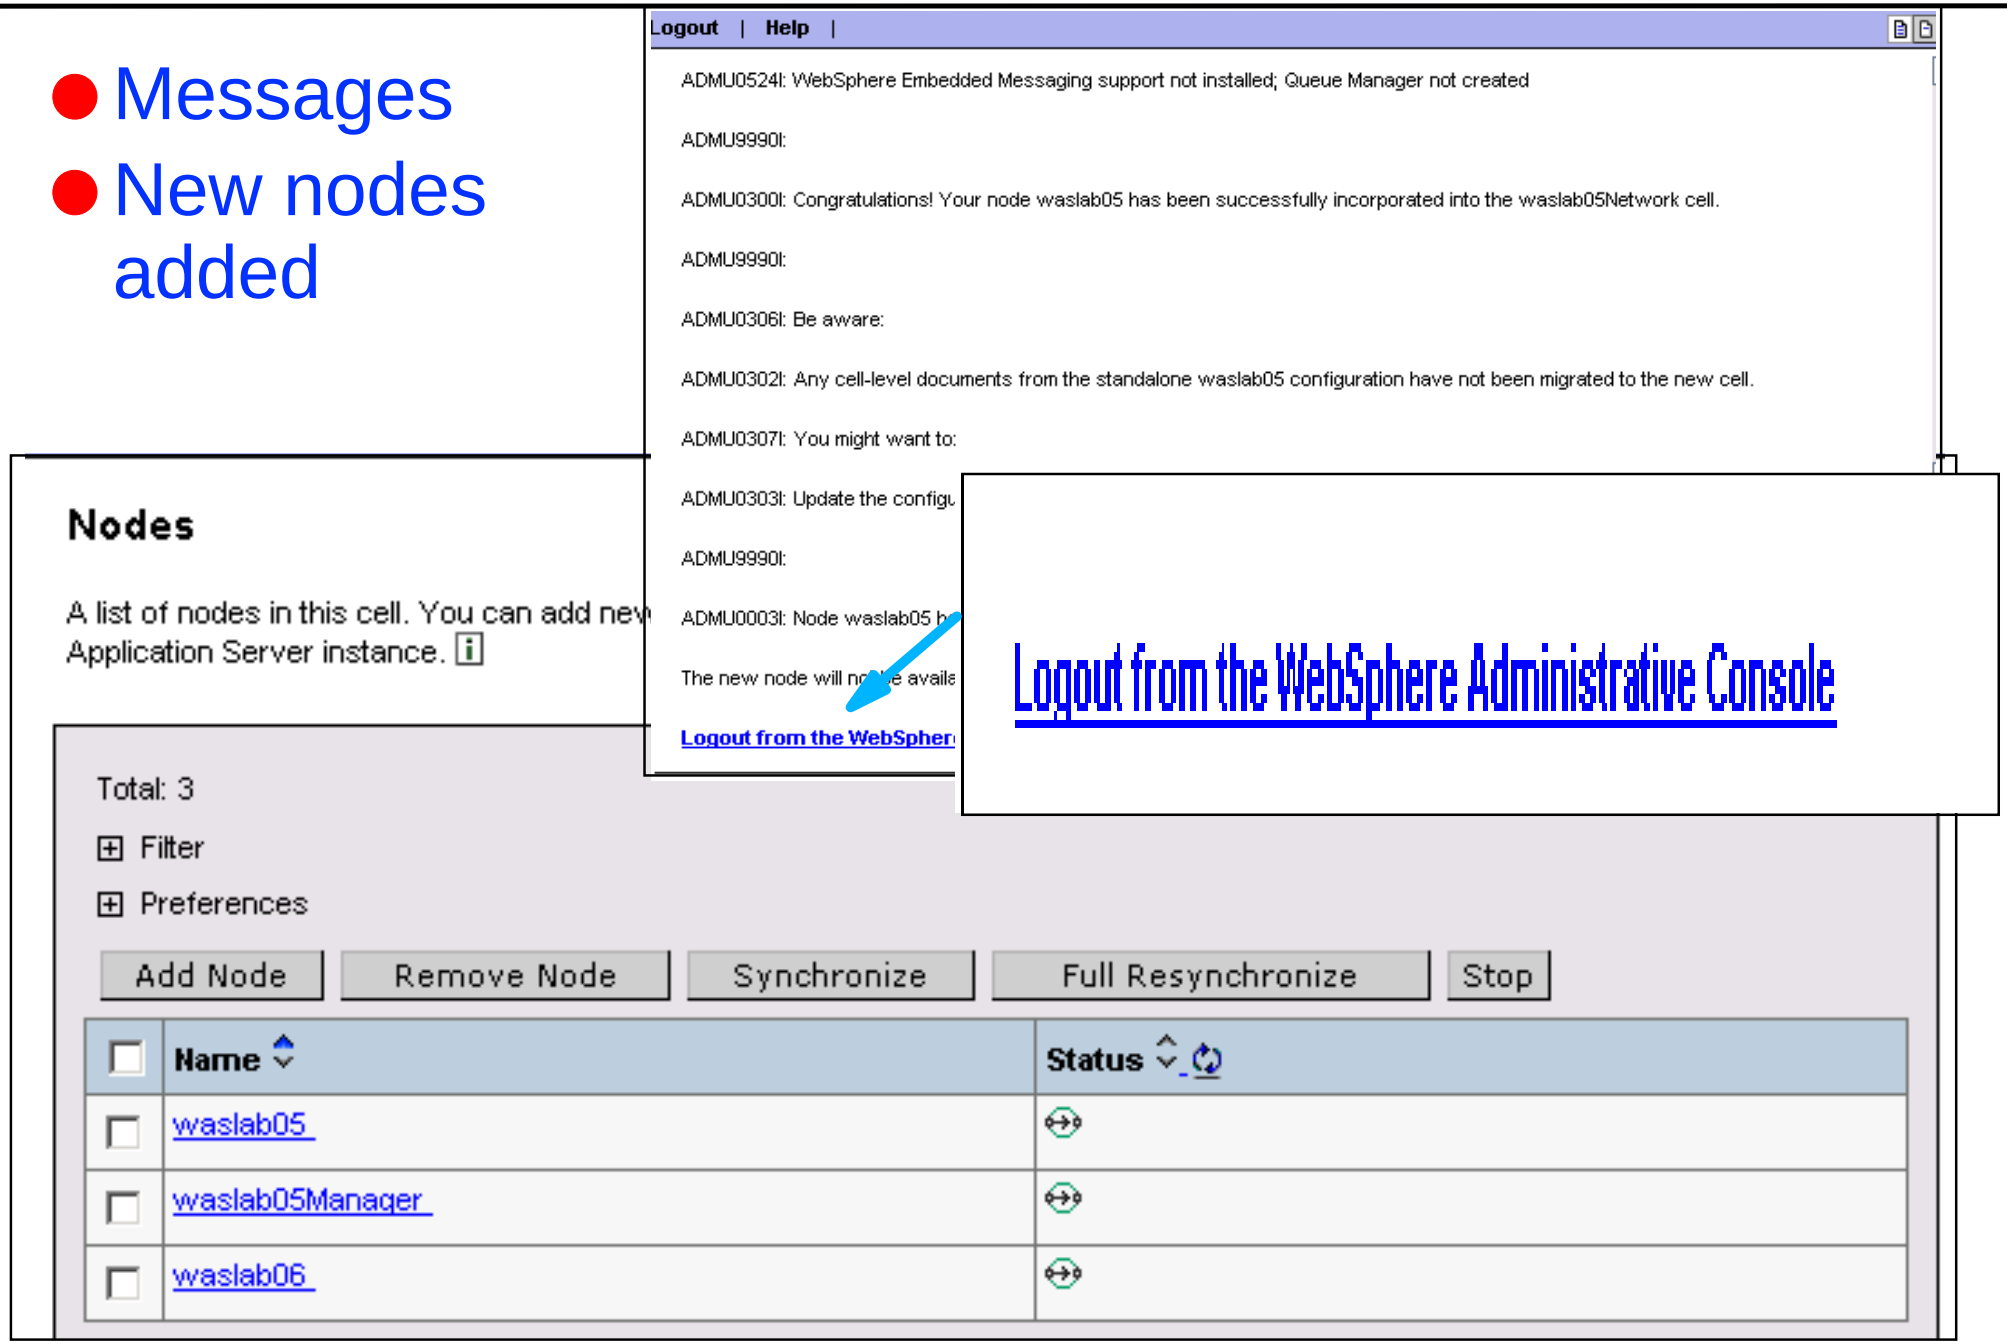

Add Nodes continued - servers

WSC
Advanced
Technical Support

- View of servers after nodes added
- Note that servers were stopped before nodes were added
- Server can then be restarted.

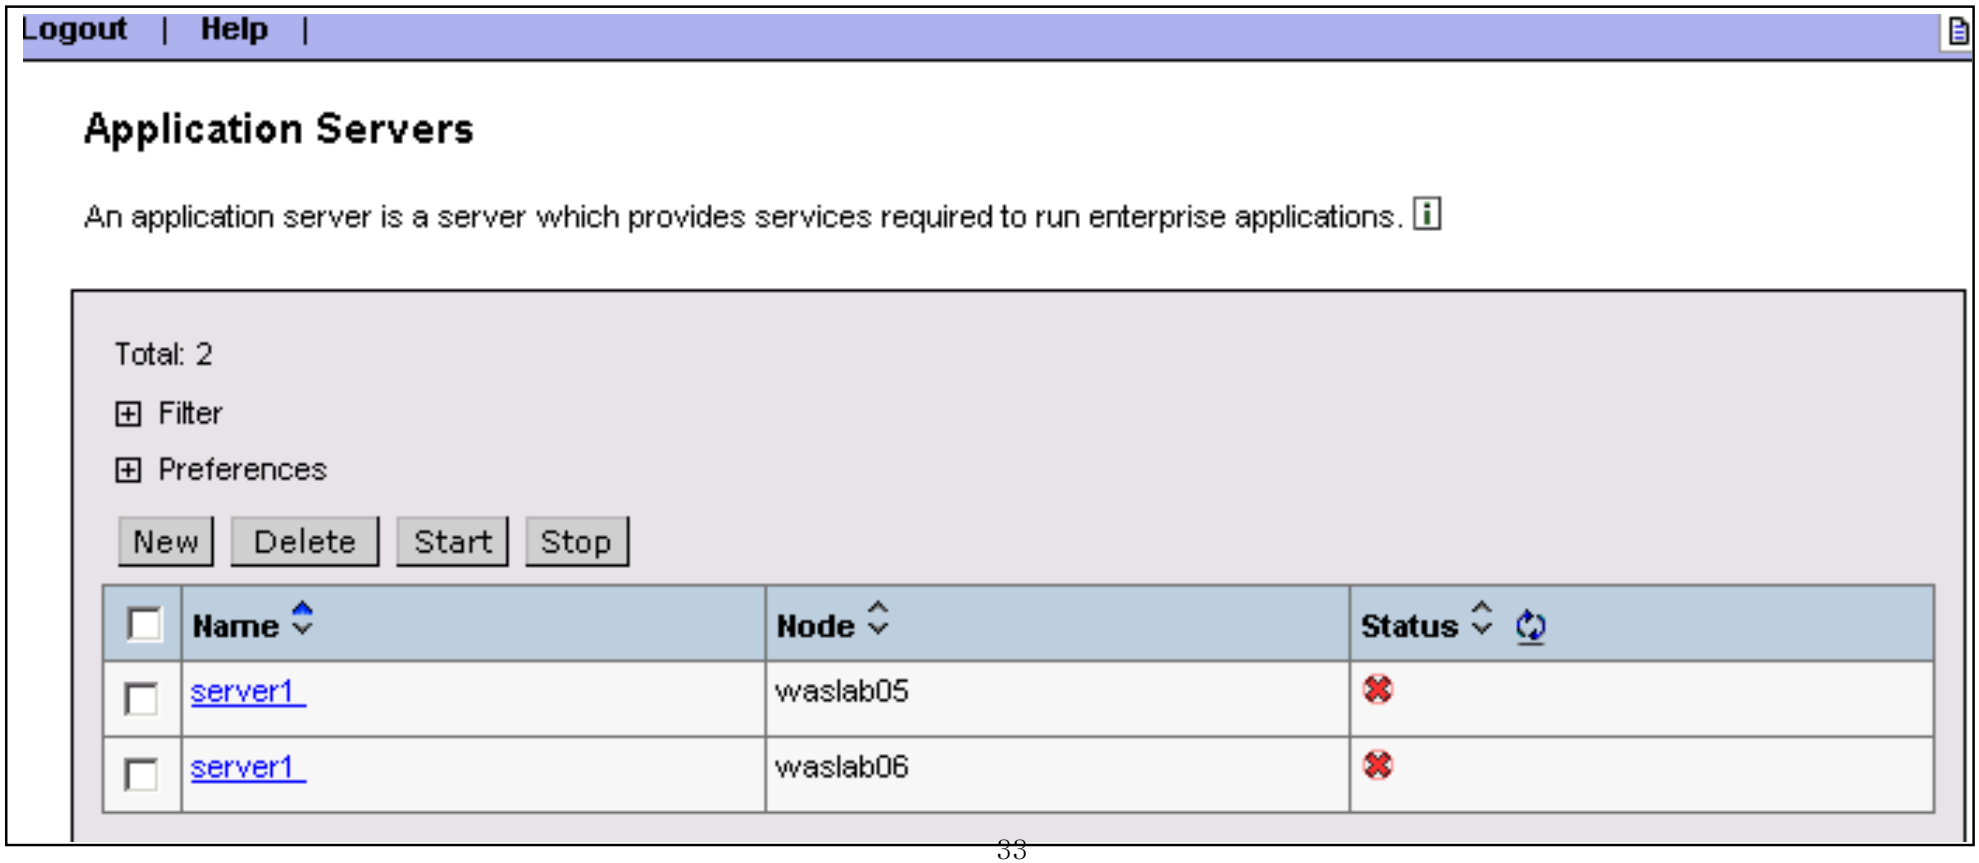

Add Nodes continued - Ports

WSC
Advanced
Technical Support

- If more that one node exists on a particular IP
  - Must modify ports for nodes using defaults
  - JMS and JMX ports must also be modified if embedded Messaging used

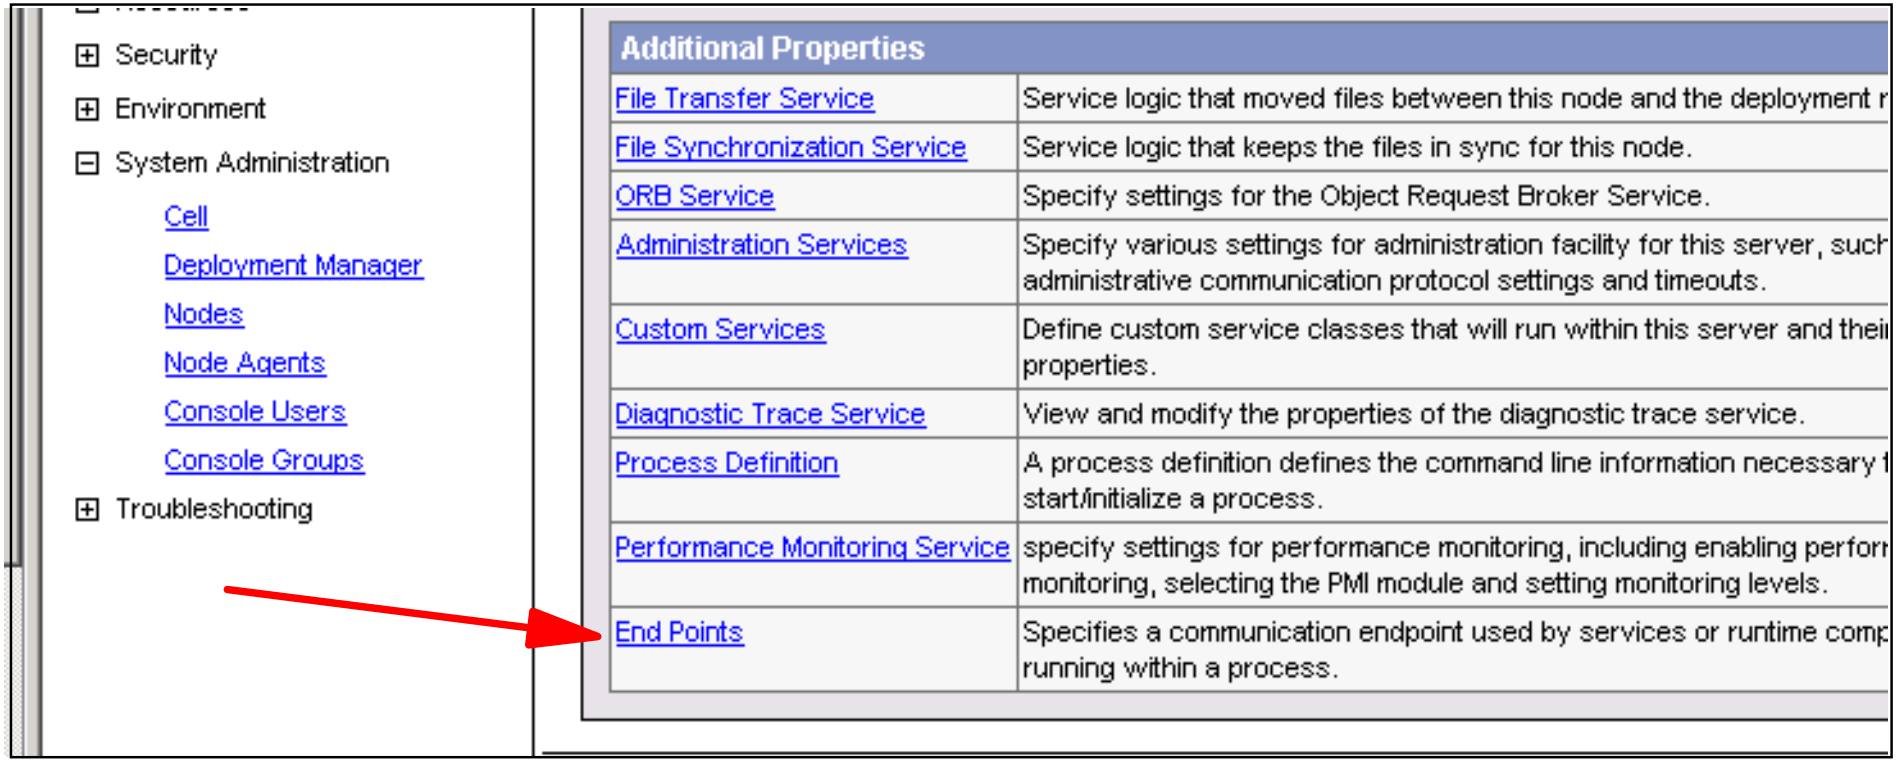

Add Nodes continued - End Points

WSC
Advanced
Technical Support

- Modify each entry to use unique ports

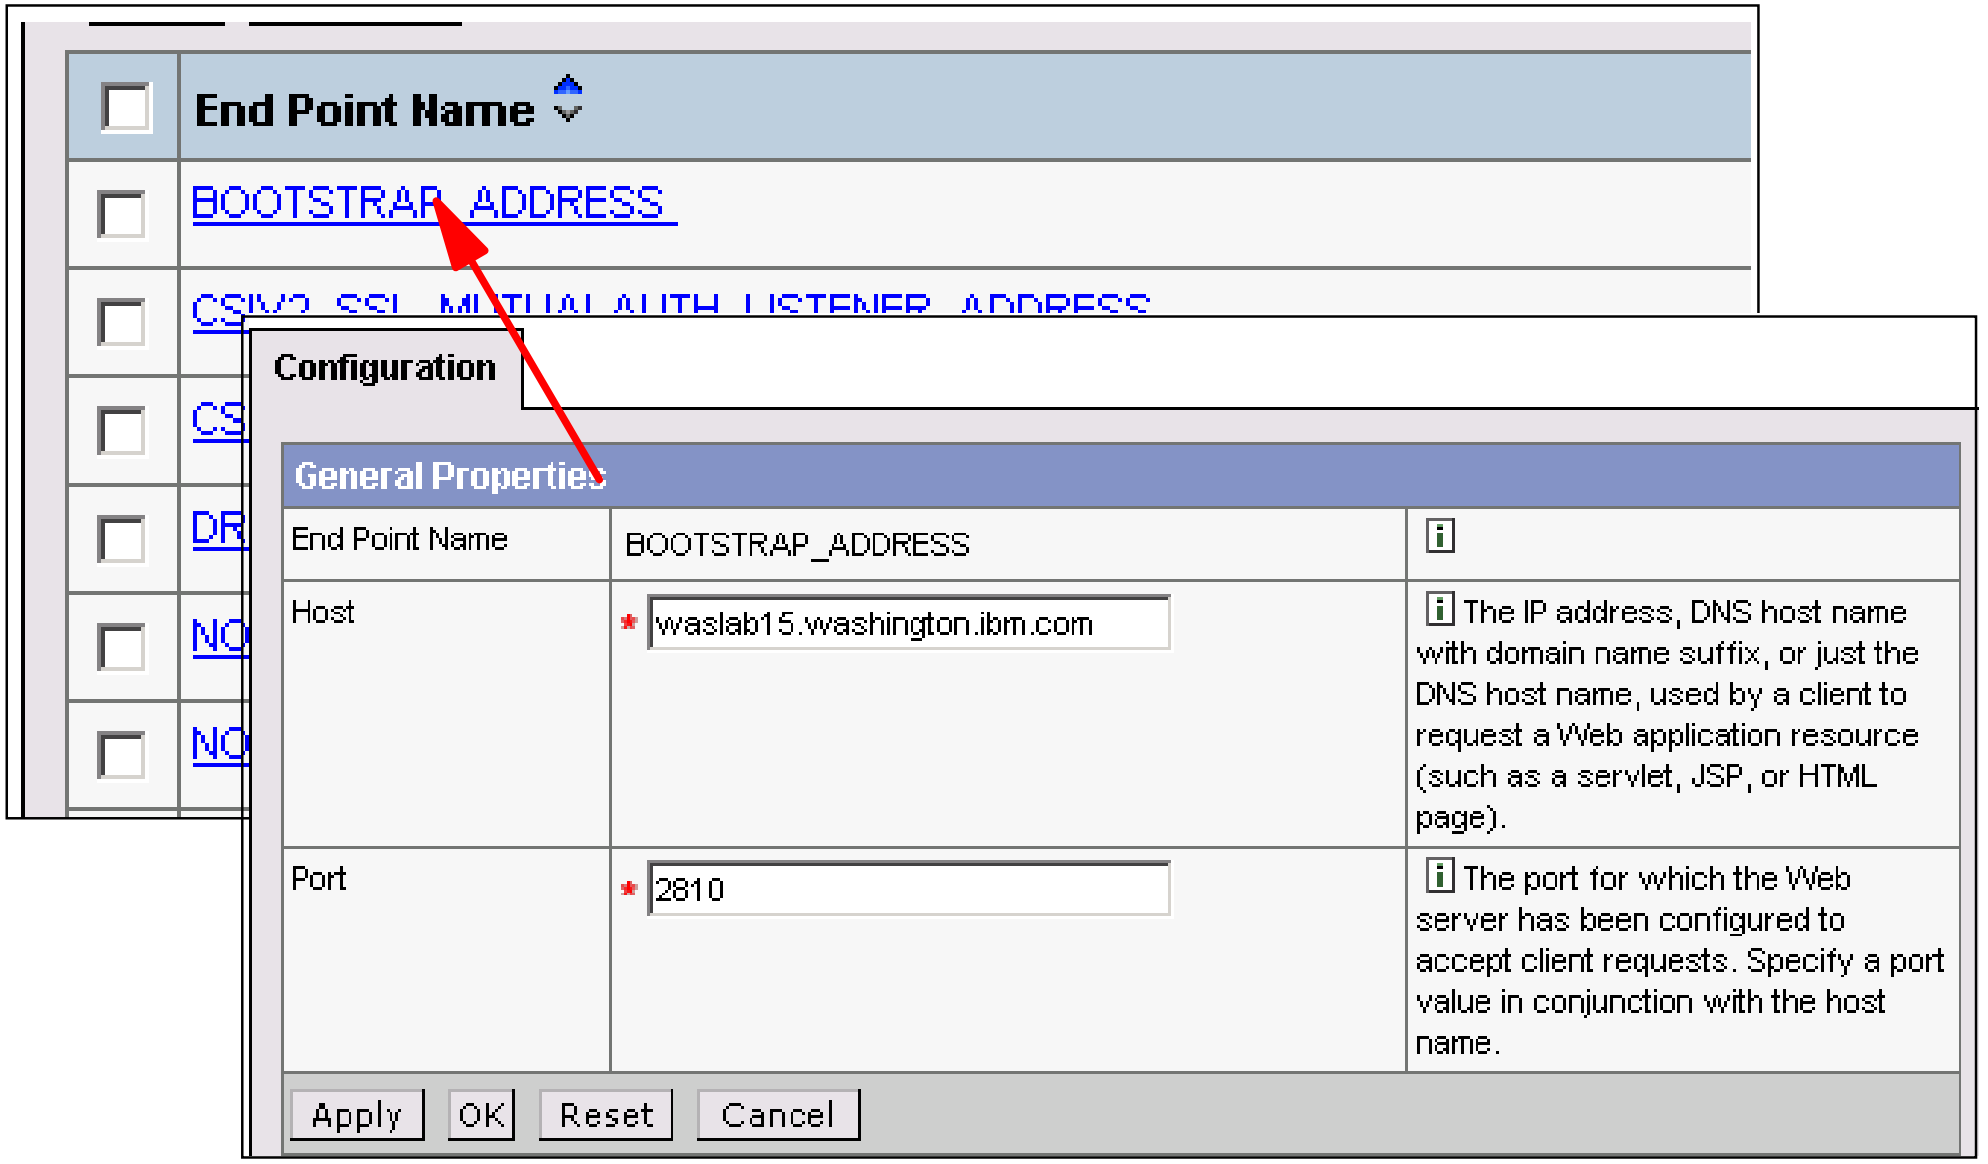

Add clusters for high availability

WSC
Advanced
Technical Support

- Select Servers>Clusters
- Select New

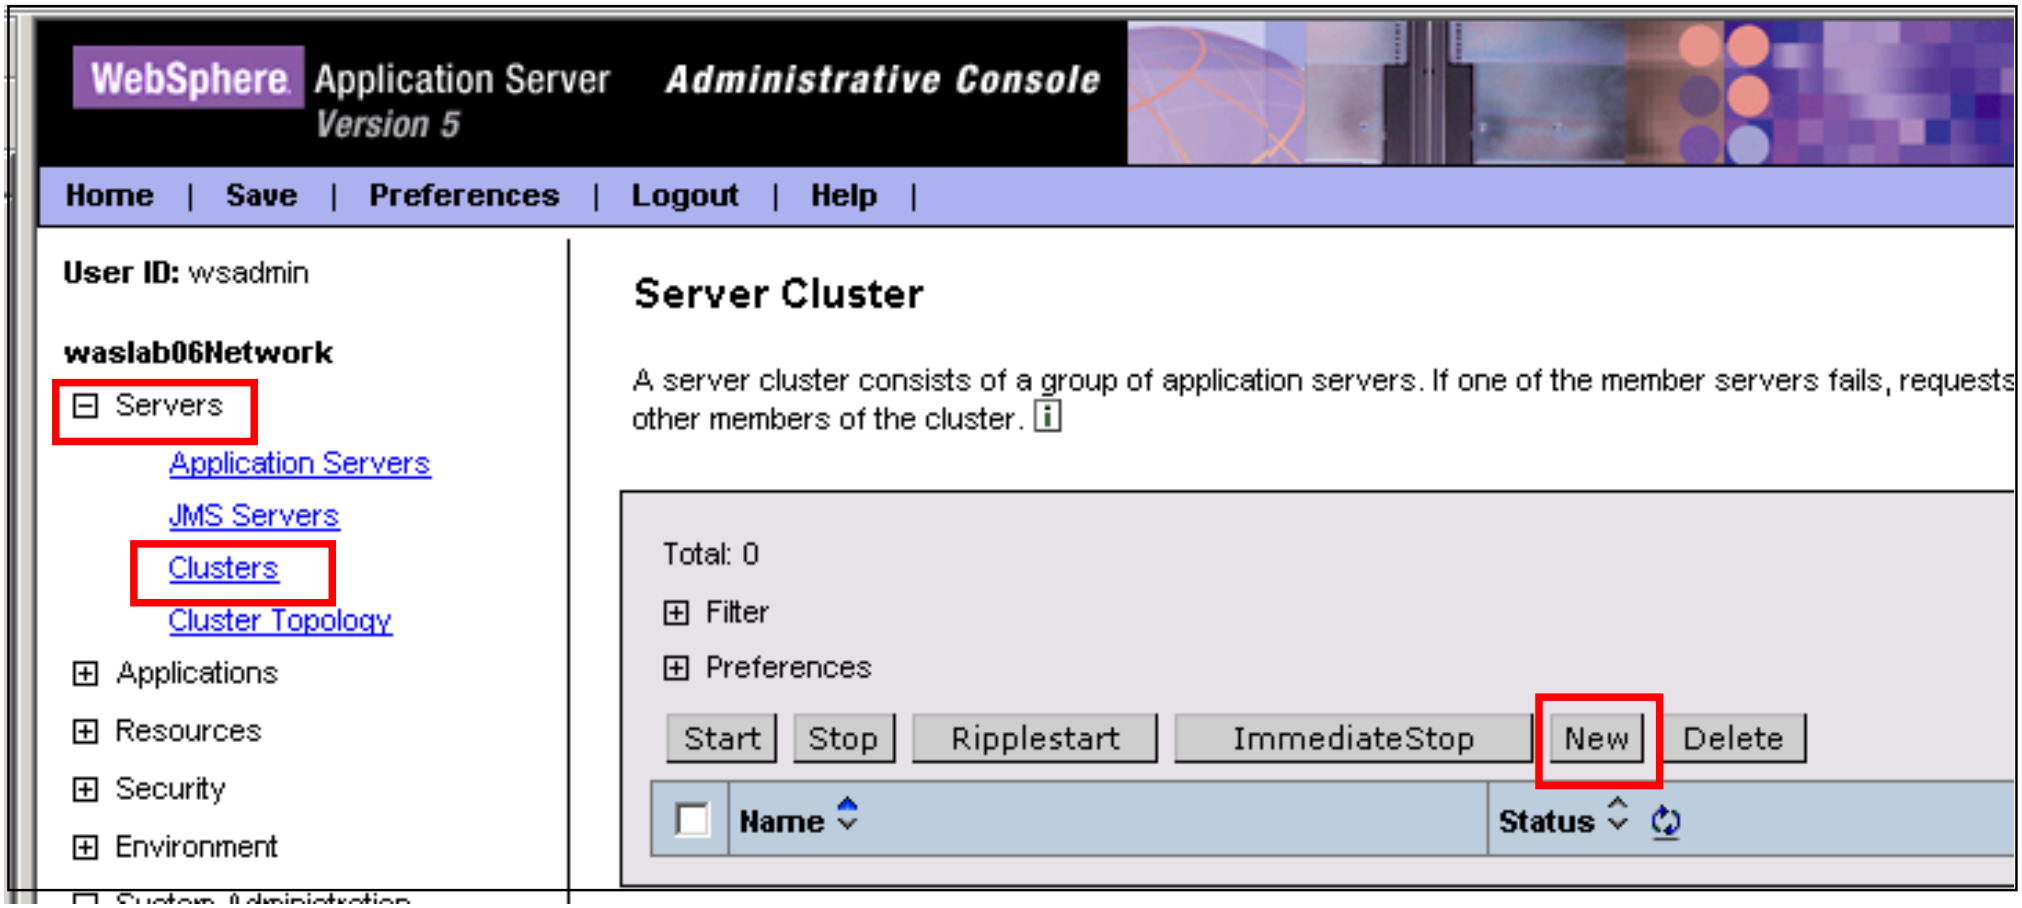

Define the cluster structure

WSC
Advanced
Technical Support

- Choose a cluster name

Step 1 : Enter Basic Cluster Information

Cluster name:	* Cluster1	The name of this cluster.
Prefer local:	☑ Prefer local enabled	Enable or disable Node scoped routing optimization.
Internal replication domain:	☐ Create Replication Domain for this cluster	If this option is selected, a Replication Domain will be created and the name will be set as the Cluster name
Existing server:	○ Do not include an existing server in this cluster ◉ Select an existing server to add to this cluster Choose a server from this list: waslab06Network/waslab061/server1 Weight: 2 ☐ Create Replication Entry in this Server	Choosing existing Server as a Cluster Member. A list of Servers which are not already a part of existing Clusters is provided. You can specify the weight for this Cluster Member. You can also choose if a Replication Entry needs to be created in this Server for internal replication.

Next Cancel

Define the cluster structure - Add members

WSC
Advanced
Technical Support

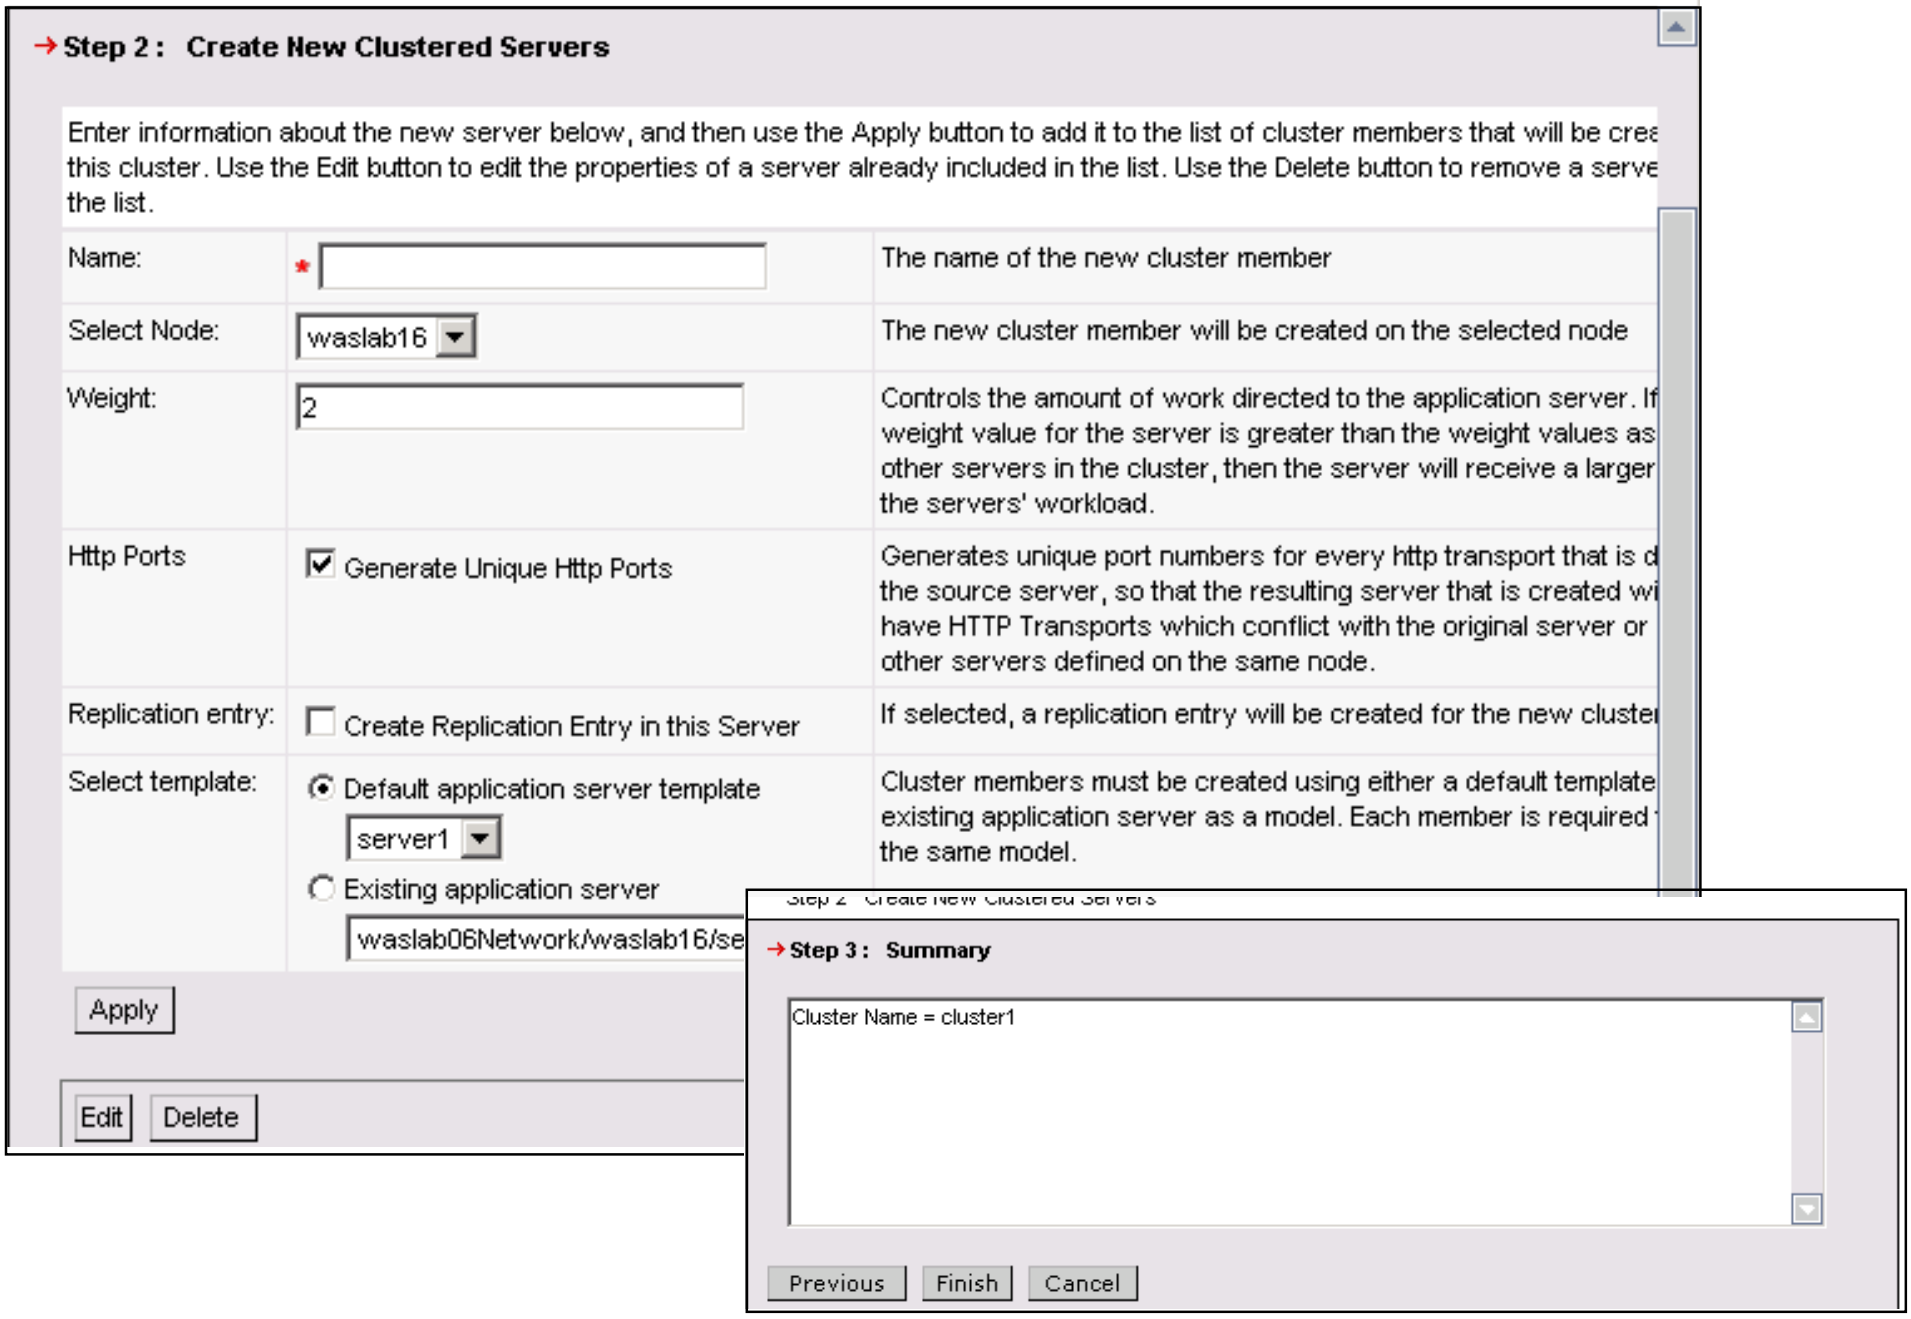

Clusters in cell topology

WSC Advanced Technical Support

- Typical cell topology

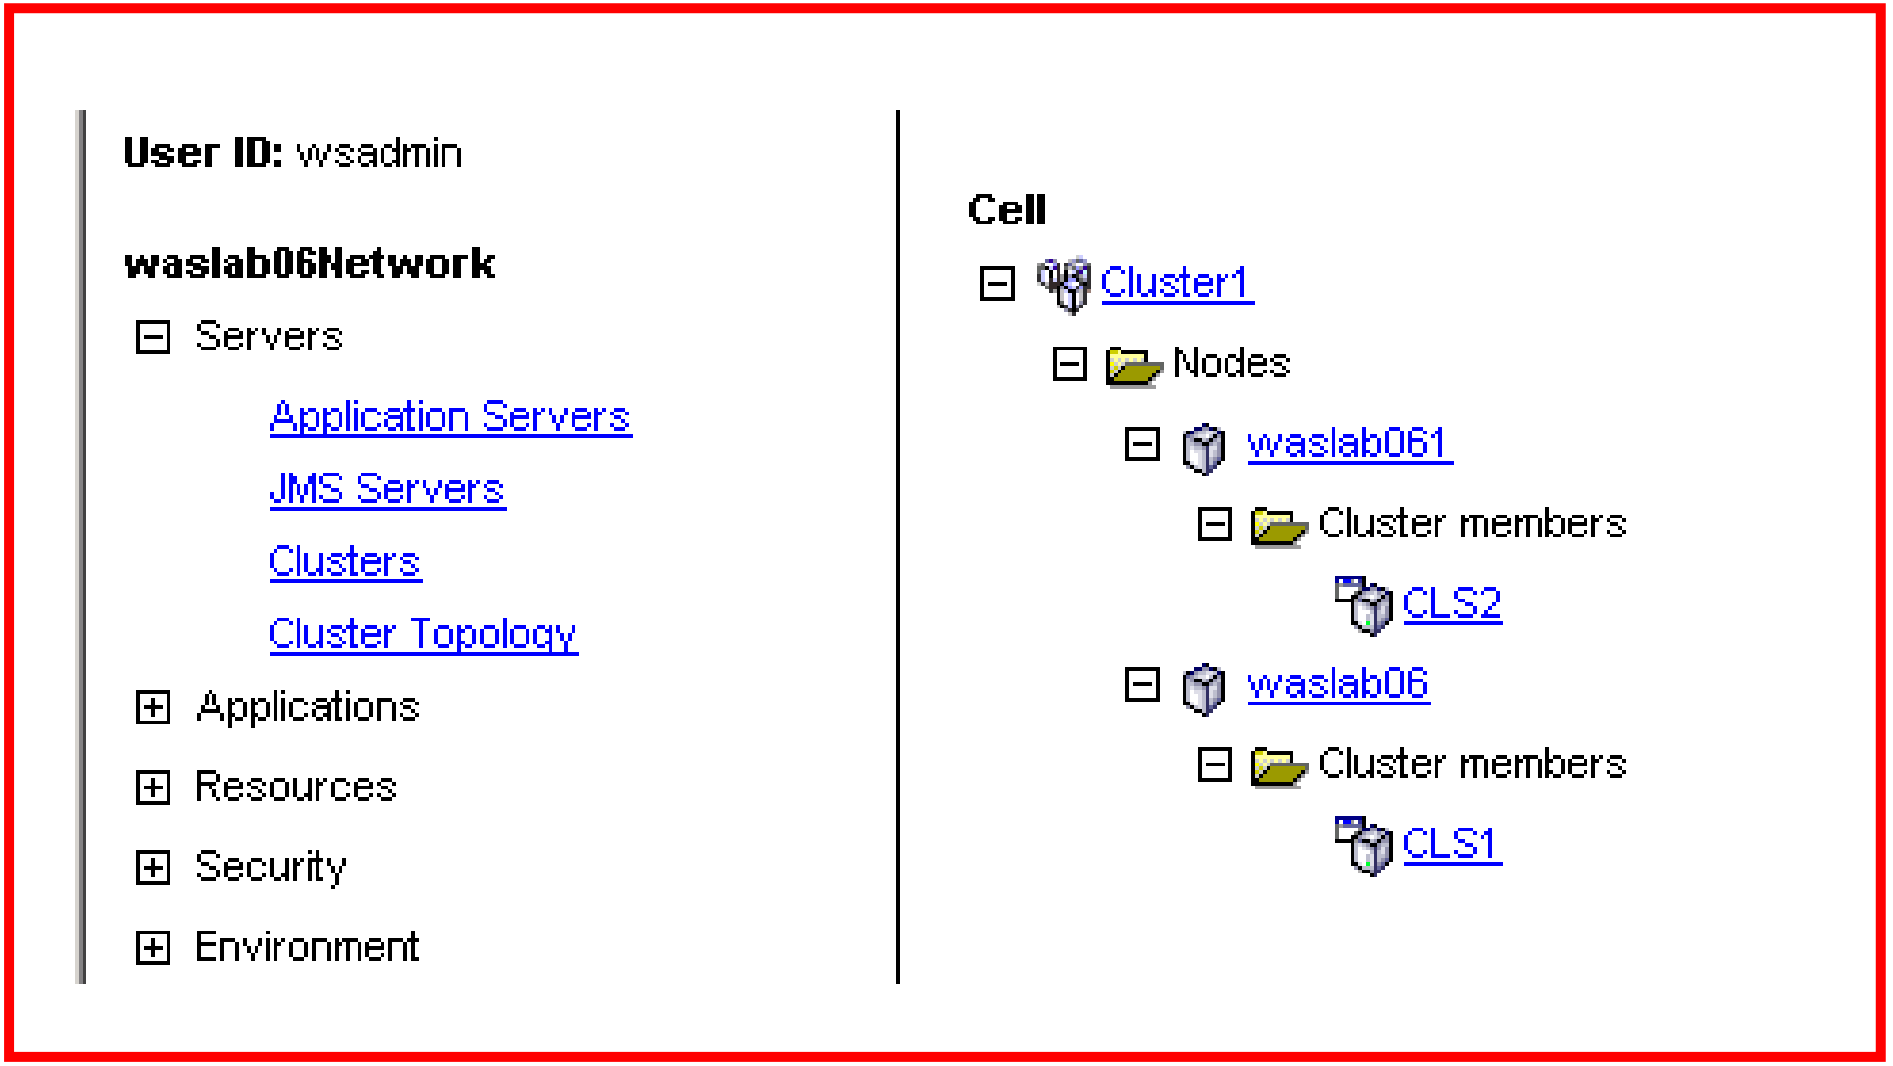

Application to Server Relationships

WSC
Advanced
Technical Support

- Applications are tied to a server, servers, or clusters
- Imported nodes do not duplicate applications
  - Once imported with a node
    - Applications owned by ND node
    - Assigned to run on server node they were imported from
- May be updated to run on other or multiple servers or clusters
    - Map modules to application servers

Specify the application server where you want to install modules contained in your application. Modules can be installed on the same server or dispersed among several servers.

Clusters and Servers:

WebSphere:cell=waslab07Network,cluster=cluster1
WebSphere:cell=waslab07Network,node=waslab07,server=server1
WebSphere:cell=waslab07Network,node=waslab00,server=server1

Apply

	Module	URI	Server
☐	Increment Enterprise Java Bean	Increment.jar,META-INF/ejb-jar.xml	WebSphere:cell=waslab07Network,cluster=cluster1
☐	Default Web Application	DefaultWebApplication.war,WEB-INF/web.xml	WebSphere:cell=waslab07Network,cluster=cluster1

OK Cancel

Ports

WSC
Advanced
Technical Support

- Ports defined in
  - Environments>Virtual Hosts
  - Usually defined to hostname * (clusters)
  - Shown in WebSphere plug-in for HTTP Server
- Ports increment per server on node starting with port 9080
  - Default ports of 80 and 9080 are not defined after node is federated
  - Plug-in must be regenerated after port change

Total: 6

Filter

Preferences

New Delete

Host Name	Port
*	9085
*	81
*	9444
*	9080
*	80
*	9081

Plug-ins

- HTTP Server to WebSphere interface
- Written in XML
  - Environments>Update web server plug-in configuration
    - Generate new plug-in
    - View plug-in
- After plug-in is regenerated
    - When ND is used, the plug-in must be manually copied to the appropriate application server directory

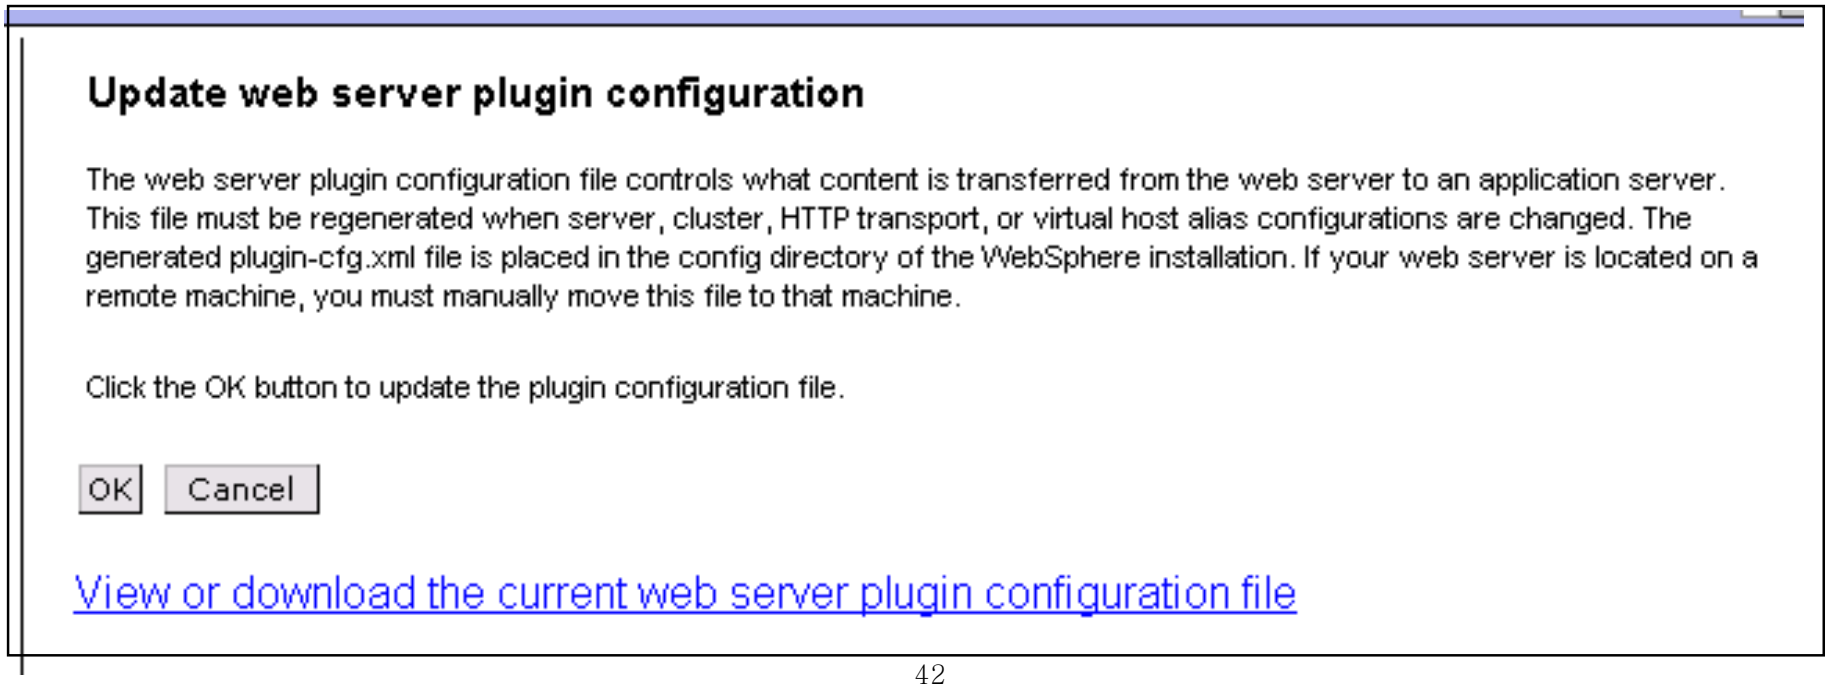

Plug-ins and ports

WSC
Advanced
Technical Support

- Ports and other info such as cluster names are defined in the plug-in
- Ports must be defined
- HTTP Server and plug-in may be run on single or multiple images in the same cluster

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
- <Config ASDisableNagle="false" AcceptAllContent="false" IISDisableNagle="false"
    IgnoreDNSFailures="false" RefreshInterval="60" ResponseChunkSize="64">
    <Log LogLevel="Error"
      Name="/opt/WebSphere/DeploymentManager/logs/http_plugin.log" />
    <Property Name="ESIEnable" Value="true" />
    <Property Name="ESIMaxCacheSize" Value="1024" />
    <Property Name="ESIInvalidationMonitor" Value="false" />
  - <VirtualHostGroup Name="default_host">
      <VirtualHost Name="*:9085" />
      <VirtualHost Name="*:81" />
      <VirtualHost Name="*:9444" />
      <VirtualHost Name="*:9080" />
      <VirtualHost Name="*:80" />
      <VirtualHost Name="*:9081" />
    </VirtualHostGroup>
  - <ServerCluster CloneSeparatorChange="false" LoadBalance="Round Robin"
      Name="cluster1" PostSizeLimit="-1" RemoveSpecialHeaders="true"
      RetryInterval="60">
    - <Server CloneID="vkhlhucd" ConnectTimeout="0" ExtendedHandshake="false"
        LoadBalanceWeight="2" MaxConnections="-1" Name="waslab07_clone1"
        WaitForContinue="false">
```

Summary

- WebSphere Network Deployment extends the abilities of the WebSphere Application Servers to include centralized administration and robust failover and load balancing capabilities.

References

- C3214020.pdf
- WebSphere On-line Documentation
- www.ibm.com/websphere