

## Frequently asked questions about WebSphere V5.0

### 1. How do I find out the product version after installing WebSphere?

There are two ways to find out the product version:

1. The file `WAS_install_root\properties\version\BASE.Product` gives information about the version of Application Server that is installed.
2. Run the batch file `WAS_install_root\bin\versionInfo.bat` to find out the version of Application server and the fixes that are installed.

### 2. How can I verify my WebSphere V5.0 installation?

After installing WebSphere successfully, use the Installation Verification Test (IVT) tool to verify your installation. The IVT program starts the server automatically if the server is not running. Once the server initializes, the IVT runs a series of verification tests and reports pass or fail status in a console window. It also logs results to `WAS_install_root\logs\ivt.log`. To run IVT, just execute the `ivt.bat` (in windows system ) or `ivt.sh` (in Unix systems) which resides in `WAS_install_root\bin` directory.

### 3. Can I use XMLConfig and wscp in WebSphere V5.0?

No. XMLConfig and wscp (WebSphere Control Program) are not supported in WebSphere V5.0. XMLConfig and wscp were a part of the WebSphere Application Server V4.0 administration repository support. In V5.0, the repository no longer exists, and the tools that manipulate it are no longer needed. All the configuration and application data are stored in XML files in all editions of WebSphere V5.0.

"wsadmin" is the new scripting interface for V5.0 replacing wscp. wsadmin is based on the Bean Scripting Framework (BSF), which supports a variety of scripting languages to configure and control your WebSphere Application Server V5.0 installation. The WebSphere Application Server V5.0 supports the Jacl scripting language only. The infocenter for V5.0 contains detailed documentation on how to migrate wscp scripts to wsadmin scripts.

#### 4. **How can I create a backup of my WebSphere V5.0 configuration?**

Use the backupConfig utility to backup your configuration. The backupConfig command is a simple utility to backup the configuration of your node to a file. By default, all servers on the node stop before the backup is made so that partially synchronized information is not saved. You can run this command from the bin directory of a WebSphere Application Server installation or a network deployment installation.

Syntax: backupConfig <backup\_file> [options]

where backup\_file represents the file the backup is to be written to.

Note that there is also a restoreConfig utility available to restore the configuration of your node after backing up the configuration using the backupConfig command.

#### 5. **What is a profile in wsadmin?**

A profile is just a script, but it is invoked before the main script or before entering interactive mode. The intention is that it could be used to set up a scripting environment customized for the user or the installation.

The -profile option in wsadmin is used to specify a profile script. The profile script runs before other commands, or scripts. If you are running a wsadmin command (using -c option), the profile executes before it invokes the command. If you are running a script (using -f option), the profile executes before it runs the script. In interactive mode, you can use the profile to perform any standard initialization that you want. You can specify multiple -profile options on the command line, and they invoke in the order that you supply them.

#### 6. **Are the ANT utilities that are shipped with WebSphere V5.0 the open source ANT or is there something additional to that?**

ANT is a powerful scripting tool that can be used to automate build and deployment processes. It is similar to the Make tool, except that instead of using operating system shell-based commands, it uses Java classes to perform its operations. WebSphere Application Server V5.0 includes the Open Source ANT plus some additional WebSphere specific ANT tasks. Some of WebSphere's additional ANT tasks are listed below:

wsInstallApp task enables you to install a new application into a WebSphere Server or Cell.

wsUninstallApp task enables you to uninstall an existing application from a WebSphere Server or Cell.

wsListApps task lists all the applications installed on a WebSphere Server or Cell.

wsStartServer task enables you to start a standalone server instance.

wsadmin task executes the WebSphere command-line administration tool with the specified arguments.

**7. How do I run my ANT build file in WebSphere V5.0?**

To run your ANT build files in WebSphere use the ws\_ant command line utility. ws\_ant also sets up the correct classpath so that the WebSphere ANT tasks can be used successfully along with ANT's built -in tasks. Run ws\_ant from the bin directory under your WebSphere install directory. For example, to execute the target named install in your build file, use the following command:-

```
ws_ant -buildfile C:\build.xml install.
```

**8. Is Tivoli Policy Director shipped with WebSphere Application Server V5.0?**

No. Tivoli Policy Director is not shipped with WebSphere V5.0. However, you can purchase it separately and integrate it with WebSphere V5.0. Tivoli Policy Director is a robust and secure policy management tool for e-business and distributed applications. It uniquely addresses the challenges of e-business security: escalating costs, growing complexity, and the inability to implement security policies across platforms.

**9. Do we ship the DataDirect JDBC drivers with WebSphere 5.0? If so, why?**

Yes, we ship the DataDirect drivers with WebSphere Application Server 5.0. Although the various Relational Database Drivers provide suitable native drivers, we still ship the DataDirect drivers to give users another option.

**10. Which version of MQ is supported by WebSphere V5.0 as an external JMS Provider? Is MQ shipped with any edition of WebSphere V5.0?**

WebSphere V5.0 supports WebSphere MQ V5.3 . The full blown WebSphere MQ 5.3 is not shipped in WebSphere Application Server 5.0 and WebSphere Network Deployment. WebSphere MQ is shipped with WebSphere 5.0 Enterprise Edition. WebSphere Application Server 5.0 has an embedded JMS provider which is a lighter footprint of WebSphere MQ 5.3.

**11. Do we support Windows 2000 Professional workstation in WebSphere V5.0?**

Windows 2000 Professional WorkStation is supported for development purposes. However, for production purposes, Windows 2000 server and Windows NT server are supported.

**12. Which platforms do not use ISMP (InstallShield for MultiPlatforms) installation wizard?**

The platforms that do not use ISMP are zSeries and iSeries.

**13. Do I have to install something additional to WebSphere Application Server 5.0 to test my JMS Applications?**

No, WebSphere V5.0 has an Internal JMS Provider that can be installed with WebSphere Application Server. This component provides embedded Point-to-Point and Publish/Subscribe support.

**14. Is Cloudscape part of WebSphere Application server install?**

Yes, TheWebSphere Application Server (base) product installation image includes IBM Cloudscape. The sample applications that are shipped with WebSphere uses Cloudcape as the datastore.

**15. Are the browsers Netscape 6.0 and Netscape 7.0 supported in WebSphere V5.0?**

Yes, WebSphere V5.0 supports Netscape 6.0 and Netscape 7.0 as long as it is using the Mozilla engine.

**16. What jdbc drivers are supported for Oracle 9i?**

The thin and oci jdbc drivers provided by Oracle are supported.

**17. Can Oracle be used for HTTP Session persistence instead of DB2?**

Yes

**18. Does WebSphere provide the implementation class for the Database helper classes?**

Yes, WebSphere provides implementation classes for the Database helper classes for all supported databases.

**19. Can I use Soniq MQ as an external JMS Provider for my JMS Applications running in WebSphere V5.0?**

Yes - you can define a JMS Provider for Sonic MQ or any other JMS product. Then you can define connection factories for that provider. The key is that you must use that product's native administrative tool to create the actual Queues and Topics and bind them into JNDI. The support for external, third party JMS providers in 5.0 is not much different from WAS 4.0.

**20. What is Application Server Toolkit? Is it a separate package from WebSphere Application Server 5.0 (base)?**

Application Server Toolkit contains a number of Eclipse based tooling such as the Debugger and Eclipse based Log Analyzer. This toolkit ships with WebSphere V5.0. However it is packaged in a separate CD.

**21. Does the internal WebSphere JMS Provider provide Publish/Subscribe messaging support?**

Yes.

In J2EE 1.3, the Java Messaging Service becomes an integral part of the J2EE platform. In v5.0, WebSphere satisfies this requirement with the Internal WebSphere JMS Provider that is installed with WebSphere Application server. This component provides embedded Point-to-Point and Publish/Subscribe support by introducing a new kind of WebSphere Application Server server, the JMS Server.

Point-to-Point messaging employs a Queue as the JMS destination. Each message in the queue will be delivered only to a single consumer. Publish/Subscribe service employs a Topic as the JMS destination. Subscribers register their interest in a particular topic. Publishers create messages that are distributed to all of the subscribers. Multiple subscribers may subscribe to this topic, and each subscribing client will receive the message. Each WebSphere Application Server Embedded Messaging JMS Server is responsible for managing a Queue Manager and a Broker. The Broker is the provider of the Publish/Subscribe service.

**22. What is the difference between a durable subscription and a non-durable subscription? Can a message driven bean be a durable subscriber?**

A durable subscription to a topic means that a JMS subscriber receives all messages, even if the subscriber is inactive. If a message is sent to a topic that has an inactive durable subscriber, the message will be persisted and delivered when the durable subscriber is once again active. A non-durable subscription to a topic means the subscriber only receives messages that are published while the subscriber is active. Any messages delivered while the subscriber is inactive will be lost.

Since message-driven bean containers are JMS consumers, the container can register itself as a durable or non-durable subscriber to messages published to a topic. Durability allows persistent messages to be sent to a topic even though the application server hosting the message-driven bean consumers has crashed.

**23. Is Message Driven Beans support new in WebSphere 5.0?**

EJB 2.0 specifications introduced a new type of bean called the Message Driven Bean. WebSphere V5.0, being fully compliant to EJB 2.0 specification, therefore supports Message Driven Beans. Message Driven Beans are new in WebSphere V5.0. However, WebSphere Enterprise Edition V4 provided extended messaging support, which had Message Beans. Message Beans and Message Driven Beans are similar in that they both implement the MessageListener interface, but Message Beans are in fact Stateless Session Beans. WebSphere V5.0 provides a tool to migrate version 4.x Message Beans to Message Driven Beans. This migration utility is called mb2mdb. It is a batch file in Windows systems and a shell file in Unix systems.

**24. Which CSD (Corrective Service Diskette) is required when using WebSphere MQ 5.3 as an external JMS Provider with WebSphere v5.0?**

WebSphere MQ5.3 CSD01 is the level of WebSphere MQ that is supported with WebSphere v5.0

**25. What is WSIL?**

WSIL stands for WebServices Inspection Language. The gist of the WebServices Inspection specification can be summarized as follows:  
"WSIL defines how a service requestor can discover an XML Web Service description on a Web server, enabling such requestors to easily browse Web servers for XML Web Services." Thus, using WSIL, a service provider can make services available for discovery and use by service requestors. This is very similar to UDDI's mission statement and should come as no surprise since WSIL is meant to be the stepping-stone to UDDI. WSIL defines an XML based language by which services can be advertised to interested requestors.

WSIL 1.0 is supported in WebSphere V5.0.

**26. Can we instrument an application so that it provides its own metrics to the Tivoli Performance Viewer?**

The Tivoli Performance Viewer only monitors the provided PMI metrics. You can, however, create your own monitoring application that would be capable of gathering information through the PMI API as well as any API that you declare for your own application. You can then use the data collected to create any new metrics category that you want.

**27. Is JCE (Java Cryptology Extensions) supported with WebSphere Application Server V5.0?**

Yes, JCE is supported by WebSphere Application Server V5.0. JCE allows for the use of cryptography devices with in a Java application. The Java Cryptography Extension(JCE) is a set of packages that provide a framework and implementations for encryption, key generation and key agreement, and Message Authentication Code (MAC) algorithms.

**28. Can the WebSphere Test Environment for WebSphere Studio Application Developer be WebSphere V5.0 Network Deployment?**

No - The WebSphere Test Environment for WebSphere Studio Application Developer is WebSphere Application Server (base) V5.0. It is not WebSphere V5.0 Network Deployment. However, it is possible to remotely install an application from Application Developer to WebSphere Network Deployment.

**29. What is a RAR file?**

RAR is a Resource Archive file. It contains all the information necessary for installing, configuring and running a JCA Resource Adapter. A resource adapter is a system-level software driver that is used by a Java application to connect to an enterprise information system (EIS). A resource adapter plugs into an application server and provides connectivity between the EIS, the application server, and the enterprise application.

A JCA resource adapter is any resource adapter conforming to the J2EE Connector Architecture Specification. WebSphere Application Server supports any resource adapter that implements version 1.0 of this specification. Although not part of WebSphere Application Server, IBM supplies resource adapters for enterprise systems such as CICS, HOD, IMS, SAP, and Crossworlds as separate products.

**30. Can RAR files be directly installed into WebSphere Application Server (base) v5.0?**

Yes, Resource Adapter Archive files can be directly installed into WebSphere Application Server(base) v5.0.

**31. When looking up an EJB, can I use global JNDI name instead of EJB references?**

Yes , You can use global JNDI names to lookup the EJBs as seen below:

```
Object myEJB=initialContext.lookup("ejb/AccountEJB");
```

However it is a recommended programming practice to lookup the EJBs using EJB References. With EJB references, the Application Component providers calls the EJB whatever they want to, and write the code accordingly.

```
Object myEJB = initialContext.lookup("java:comp/env/ejb/myEJB");
```

The EJB references are linked to the actual JNDI names using WebSphere bindings which can be specified at deployment time.

**32. In EJB 2.0, should all the bean classes be abstract?**

Yes.

A container-managed entity bean class is defined as an abstract class. The abstract bean does not define how its state is maintained. The class does not have any fields. The "virtual", container-managed properties are accessed only through accessor-methods which are defined as abstract methods. The concrete bean class is created by the deployment tools.

**33. Where does the actual java class for the EJB home method reside?**

The purpose of EJB home method is to develop business logic that is not specific to a bean instance but related to the EJB Component. The home methods can be declared in the remote or local home interface. They are implemented in the bean class with an "ejbHome" prefix.

**34. If an EJB is build with local interfaces, can its name be viewed in the JNDI name space?**

Yes. An EJB's name can be viewed in the JNDI name space irrespective of whether it has local interfaces or remote interfaces or both.

**35. Can EJB 2.0 Entity beans have relationships with other entity beans in different JVMs?**

No, Container Managed Relationship (CMR) requires both entity beans to reside in a single JVM as one of the entity beans has to have a local interface to allow for navigability over the relationship. Basically, the target EJB in a unidirectional relationship must provide local interfaces. In bi-directional relationships, both EJBs must have local interfaces.



**36. Can a standalone Java client lookup a connection factory in the built in JMS server?**

Yes, a stand alone Java client can lookup a JMS Connection factory in the the built-in JMS Server. The client code can do a java:comp/env lookup on the connection factory. The reference is a resource environment reference and the reference can be bound to the global JNDI name for the connection factory using the Application Assembly Tool.

**37. In WebSphere V5.0, Is a Message-driven bean bound to the topic name that it should process messages from?**

Message-driven beans are bound to the JNDI name of a listener port configured to monitor messages available at a particular destination which could be a queue or a topic. A listener port is used to simplify administration of the association between a connection factory, destination, and deployed message-driven bean.

**38. Is it possible to have a Message Driven Bean associated to two Listener ports? Also, is it possible to have a Listener port associated to two destinations?**

No, it is not possible to associate a Message Driven Bean to multiple destinations. It is also not possible to associate a Listener Port to multiple destinations. There is a 1:1 relationship between Destinations and Listener Ports. There is a 1: N relationship between Listener Ports and MDBs.

**39. Is it possible to send test messages to the Embedded Messaging Provider's queue using the Admin Console?**

In v5.0, it is not possible to send test messages to the Embedded Messaging queue using the Admin Console. However an Mbean interface will be available at Global Availability which can be accessed through the command line scripting interface wsadmin.

**40. How do I find out the number of messages in the queue if I am using the Embedded JMS Provider?**

An Mbean will be provided with WebSphere v5.0 for Embedded JMS Provider operations such as Clear, Depth, get, put and browse. This Mbean will be exposed through wsadmin scripting tool for WebSphere Application server v5.0.

It is also possible to use the utility runmqsc to find the Depth (the number of messages in the queue). To use runmqsc , first run the executable file dspmq.exe from <WebSphereEmbeddedMessaging\_install\_directory>/bin to find out the Queue Manager Name. Then enter runmqsc <QueueManagerName>.

At the runmqsc prompt, enter display qlocal (<Queue Name>). This will provide information such as CURDEPTH which shows the number of messages in the queue.

**41. Is it possible to co-exist embedded JMS Provider and the WebSphere MQ JMS provider?**

No, it is not possible to co-exist embedded JMS provider and MQ JMS provider on the same node since the internal as well as MQ JMS providers use the same JMS transport. However you can have them on different nodes.

**42. Can the Embedded JMS Servers participate in a cluster?**

No, Features such as Queue Manager Clustering are not supported for the Embedded JMS Provider. Upgrade to WebSphere MQ if you need Queue Manager Clustering.

**43. Can more than one Message Driven Bean consume messages from a queue?**

More than one Message Driven Beans can listen on a queue. However once a message is consumed by one MDB, then no more MDBs get it. The MDBs compete for the message in such a case.

**44. Can Tivoli Performance Viewer monitor the Embedded JMS server?**

No.

**45. What is the HandleDelegate SPI?**

The `javax.ejb.spi.HandleDelegate` service provider interface defines methods that enable portable implementations of `Handle` and `HomeHandle` that are instantiated in a different vendor's container to serialize and deserialize `EJBObject` and `EJBHome` references. The `HandleDelegate` interface is not used by enterprise beans or J2EE application components directly. EJB, web and application client containers provide implementations of the `HandleDelegate` interface. The `HandleDelegate` object is accessible in the client J2EE component's JNDI namespace at the reserved name "java:comp/HandleDelegate". Portable implementations of `Handle` and `HomeHandle` must look up the `HandleDelegate` object of the container in which they are instantiated using JNDI at the name "java:comp/HandleDelegate" and use the `HandleDelegate` object to serialize and deserialize `EJBObject` and `EJBHome` references.

**46. Where are EJB QL queries entered?**

EJB QL queries are entered in the EJB Deployment descriptor (`ejb-jar.xml`) file. WebSphere Studio Application Developer or Application Assembly tool can be used for this task.

**47. Are filters servlets?**

Filters are not servlets; they do not actually create a response. They are preprocessors of the request before it reaches a servlet, and/or postprocessors of the response leaving a servlet. In a sense, filters are a mature version of the old "servlet chaining" concept.

**48. Can I have WebSphere Application Server 5.0 (base) servers and WebSphere Enterprise 5.0 servers in the same cell?**

No, you can only have homogenous environments as part of a cell.

**49. Does the addNode command carry over applications from the Base Application Server to the cell?**

The main purpose of the addNode tool is to create a node agent for the new node and incorporate the new node into the cell. addNode also renames the servers so that there are no duplicate named servers within the cell, and creates a single JMSServer process for the node. By default, no applications are carried over from the Base Application server up into the cell. So after addNode, you get a node, node agent, JMSServer, any application servers and their configuration files, but no applications. You have to install the application(s) using the console or wsadmin against the Deployment Manager so that they are properly installed and replicated to the node. This is the DEFAULT behavior.

However, there is a "-includeapps" option for addNode. This option tells addNode to make a "best attempt" to carry over the applications on the servers on the node. If the applications don't exist in the cell already, on some other node, then using the -includeapps option will upload them into the cell

**50. After adding a node to the cell using the addNode utility, where is the Admin Console web application installed?**

By default, after addNode, the Admin Console web application is only installed on the Deployment Manager. It is possible to go back and reinstall the console on individual application servers but such consoles can only administer the individual application server, not the cell as a whole. It is also technically possible to install the console on the node agent but this is complicated as the Node Agent does not have a Web Container.

**51. Can I have more than one nodes installed on one host machine?**

Yes, a host may have more than one nodes installed. However a node must reside wholly on a single host.

**52. Can I install Deployment Manager on the same host machine as WebSphere Application Server 5.0 is installed, and then do addNode to add the node to the cell?**

Yes - the Deployment Manager can be installed on the same machine as your WebSphere base server is. You can then use the command - addNode localhost <Deployment Manager Port> to add the node to the cell.

**53. Why do resources in WebSphere Network Deployment have scope?**

Resources in Network Deployment have scope. The configuration definition for the resources is stored in a resources.xml file that can be scoped at the cell, node or server level. The reason is administrative convenience. Particularly in a clustered environment, you probably want to define your resources at the cell or at least the node level. This means that you define and configure them in one place and the File Synchronization ensures that the definition is available to all servers in the cell (or node if defined at that scope). Resources defined at the server scope are only available to that one process. Defining resources scoped at the server level would not be appropriate for clusters, because it would mean that any change to the resource configuration would have to be made several times.

**54. Can a WebSphere V5.0 cell be compared to a WebSphere V4.0 domain?**

Yes - All servers that share a common repository are said to be in the same cell.

**55. Does the Node Agent have to be running for an Application Server in the node to be started?**

A Network Deployment application server cannot start unless the node agent is running. This is because the ORB in a Network Deployment application server has to register with the Location Service Daemon which only runs in the node agent. So the node agent always has to be running before an application server will run in Network Deployment.

**56. Does the Deployment Manager and Node Agent only exist in Network Deployment?**

The Deployment Manager is only installed by the Network Deployment package. The configuration for a node agent is only created in the Application Server nodes when you run the addNode utility. The addNode utility only works if there is a running Deployment Manager available to communicate with. So, in that sense, the node agent only comes into existence when Network Deployment has been installed and the Deployment Manager is running.

**57. Do I first need to install WebSphere Application Server (base) on a host machine before installing Deployment Manager?**

No - the Deployment Manager can be installed separately on a host machine. After installing the Deployment Manager, you can add nodes to the cell using the addNode utility.

**58. In WebSphere v5.0 Network Deployment, after an application update(redeploy) is performed, is the entire deployed archive propagated to the node?**

No, Intelligent binary distribution is done whereby only the changes to the application are propagated to the node.

**59. How is Deployment Manager fail-over handled in Network Deployment?**

Deployment Manager fail-over is not addressed in Network Deployment. Deployment Manager does not handle any client requests. Since Nodes have copies of all configuration information, the effect of a Deployment Manager failure is minimal. The only real problem with the Deployment Manager failing would be the inability to broadcast configuration changes to the Node Agents.

**60. When we create a cluster based on an existing Application Server, are the applications propagated to the cluster members?**

If you use the template application server as the first cluster member, then the application configuration is propagated to all the cluster members. However, if you use the application server as a template only and do not make it the first cluster member, then the applications are not propagated to the cluster members.

**61. Does the Deployment Manager need to be up and running in order to use Work Load Management? Does the Deployment Manager need to be running to start the individual cluster members?**

In order to change the load weighting for cluster members, the Deployment Manager needs to be running. But as far as things running steady-state, it is not required. Even though it is convenient to use the cluster start operation defined on the Mbean running in the Deployment Manager, the cluster members can be started individually even if the Deployment Manager is down. If a cluster member crashes, the clients will fail-over due to request timeouts when directed to that server. So things will continue to work. But they work better when the deployment manager is running and available to update the routing table and drive it out through the node agents, application servers, and eventually to the clients.

**62. How many Mbean Servers are there in a JVM?**

One - Each JMX enabled JVM contains a Mbean Server (Managed Bean Server) that registers all the Mbeans in the system. External programs access Mbeans registered on the Mbean Server via Connectors/Adapters.

**63. Are there any tools to administer WebSphere 5.0 using JMX?**

The WebSphere web based Admin Console and the scripting interface wsadmin administers WebSphere using JMX.

**64. What are the options for session persistence in WebSphere v5.0?**

The options for session persistence are Database Persistence and memory-to-memory session replication. When using Database Persistence, session information is persisted in a back-end database.

The WebSphere product also supports session replication to another WebSphere server instance. This support is referred to as memory-to-memory session replication. In this mode, sessions can be replicated to one or more WebSphere instances to address a single-point of failure (SPOF).

**65. Is memory-to-memory session replication for session persistence available in WebSphere Application Server v5.0 (base)?**

No - Memory-to-memory session replication is a Network Deployment option. With WebSphere Application Server v5.0 (base), use database persistence to persist your sessions.

**66. Does memory-to-memory session replication provide all the performance related options (such as time base write)?**

Yes, memory-to-memory replication provides all the performance related options such as time based write or write only when changes are made.

**67. Is the WebSphere Embedded messaging used for memory-to-memory session replication?**

No - If you install WebSphere Application Server (base) and Network Deployment without selecting WebSphere Embedded Messaging, the memory-to-memory session replication would still work. The memory-to-memory session replication uses an Internal Messaging system and it does not have any dependencies on the WebSphere Embedded JMS Provider.

**68. Do I have to restart the server after enabling Security?**

Yes - Server has to be restarted after enabling or disabling Global Security.

**69. What is the difference between groups and roles?**

J2EE security roles, are associated with an application. The distinction between roles and groups, is that groups are associated with the location where the application is installed. Groups are different from roles. The roles stay the same no matter where the application is deployed. Users and groups are defined in the user registry. Roles are defined when the application is assembled using the Application Assembly Tool.

For example, a banking application might have roles like teller, manager, customer, janitor, and so on. These stay the same no matter where the application is deployed. In Bank A there might be users Alice and Bob, and a small group of people who are tellers in this bank in Charlie's department. So Charlie's department is a group. When you deploy the banking application in Bank A you associate Charlie's department with the Teller role. Later, when one of the people in Charlie's department tries to access a resource, WebSphere checks the roles that are required to access the resource, and checks the roles the particular user is in, and if there is an overlap the user is allowed access.

**70. Does INS (Interoperable Naming Service) mandate the Distributed cosNaming in WebSphere v5.0?**

The distributed cosNaming (ie naming split across many servers with single logical view) is purely WebSphere design. It is not mandated by INS. The corbaloc/corbaname URLs and the protocol for bootstrapping into a cosNaming server are part of the OMG specs. The distributed naming provides provides higher availability and more extensibility. The distributed naming removes dependency on a single, central name server.

**71. While looking up an EJB in a cluster, if my provider URL is corbaloc::h01, h02. Is there an implied sequence in the way the lookup is going to take place?**

Yes, there is an implied sequence, from left to right.

**72. How do I test my datasource connection in v5.0?**

The wsadmin scripting interface can be used to test connections to the datasource, The testConnection command is part of the AdminControl object because it is an operational command. This particular type operational command takes a configuration id as an argument, so we invoke the getid command on the AdminConfig object to get it:

Example:

```
set myds [$AdminConfig getid /JDBCProvider:mydriver/DataSource:myds/]  
$AdminControl testConnection $myds
```

**73. Why would one use an unshareable connection?**

Access to a resource marked as Unshareable simply means that there is a 1 to 1 relationship between the connection handle a component is using and the managed connection the handle is associated with. This implies that every call to getConnection returns a connection handle (ManagedConnection) solely for use by the requesting user. Typically, users must choose unshareable if they might do things to the connection that could result in unexpected behavior occurring to another application that is sharing the connection (for example, changing the isolation level).

**74. Is there any migration tool available that automates the migration of wscp scripts to wsadmin?**

No, There is no migration tool available with v5.0 to automate the migration of wscp scripts to wsadmin. The InfoCenter has lots of documentation and examples that detail the changes.

**75. Does WebSphere have to be started to run wsadmin to install an application?**

No, you can install an application using wsadmin without starting the server. The AdminApp commands such as install and installInteractive can be done in a local mode without connecting to the server. To enter wsadmin interactive mode in local mode specify the -conntype NONE option as shown below:

```
wsadmin -conntype NONE
```

Once you enter wsadmin , then execute the \$AdminApp install c:/my.ear to install the application locally.