

IBM WebSphere Application Server, Version 5



Applications

Note

Before using this information, be sure to read the general information under “Trademarks and service marks” on page ix.

Compilation date: November 21, 2002

© Copyright International Business Machines Corporation 2002. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Trademarks and service marks	ix	Developing Web applications	42
Chapter 1. Welcome to Applications	1	Web modules.	43
Chapter 2. Using Web applications	7	Assembling Web Modules	43
Web applications	7	Context parameters.	45
web.xml file	8	Servlet mappings	45
Migrating Web application components	9	Servlet caching	45
Default Application.	12	Web components	45
Snoop	12	Web property extensions	45
HelloHTML	13	Web resource collections	46
HitCount	13	Welcome files.	46
Servlets.	13	Context parameter assembly settings	46
Developing servlets with WebSphere Application		Initialization parameter assembly settings	46
Server extensions	14	JavaServer Pages attribute assembly settings	47
Application lifecycle listeners and events	14	Multipurpose Internet Mail Extensions (MIME)	
Listener classes for servlet context and session		filter assembly settings	49
changes.	15	Page list assembly settings	51
Example:		Security constraints.	51
com.ibm.websphere.DBConnectionListener.java	15	Security constraint assembly settings	51
Servlet filtering	16	Servlet mapping assembly settings	53
Filter, FilterChain, FilterConfig classes for servlet		Tag library assembly settings	53
filtering.	16	Welcome file assembly settings	54
Example: com.ibm.websphere.LoggingFilter.java	16	Servlet caching configuration assembly settings	54
Configuring page list servlet client configurations	17	Web components assembly settings	56
autoRequestEncoding and autoResponseEncoding	21	Web modules assembly settings	58
autoRequestEncoding and autoResponseEncoding		Assembly property extensions	61
encoding examples	21	File serving attribute assembly settings	61
JavaServer Pages files	22	Invoker attribute assembly settings	61
Developing JavaServer Pages files with WebSphere		Error page assembly settings	62
extensions	22	Web resource collections security constraint	
Tag libraries	23	properties	63
tsx:dbconnect tag JavaServer Pages syntax	23	Troubleshooting tips for Web application	
dbquery tag JavaServer Pages syntax	24	deployment	63
dbmodify tag JavaServer Pages syntax	25	Modifying the default Web container configuration	65
tsx:getProperty tag JavaServer Pages syntax and		Web container	65
examples	26	Web container settings.	65
tsx:userid and tsx:passwd tag JavaServer Pages		Web Container Services settings	66
syntax	27	Web module settings	67
tsx:repeat tag JavaServer Pages syntax	27	Web Module Deployment settings	68
Example: Combining tsx:repeat and		Web applications: Resources for learning	69
tsx:getProperty tag JavaServer Pages tags.	28	Chapter 3. Managing HTTP sessions	71
Example: tsx:dbmodify tag syntax	28	Sessions	71
Example: Using tsx:repeat JavaServer Pages tag		Migrating HTTP sessions	72
to iterate over a results set	28	Developing session management in servlets	73
JspBatchCompiler tool.	31	SessionSample.java	74
Bean Scripting Framework	32	Assembling so that session data can be shared	75
Example: Converting JavaScript source to the Bean		Servlet API Behavior	75
Scripting Framework	32	Session security support	75
Scenario: Creating a Bean Scripting Framework		Security integration rules for HTTP sessions	76
application	33	Programmatic details and scenarios	76
Scenario description	33	Session management support	77
Developing the BSF application.	34	Configuring session management by level	78
Deploying the BSF application	38	Session tracking options	78
Example: Bean Scripting Framework code example	39	Session tracking with cookies	79
		Session tracking with URL rewriting	79

Session tracking with SSL information	80
Configuring session tracking	80
Serializing access to session data	81
Session Management settings	81
Cookie settings	83
Distributed sessions	84
Session recovery support	84
Distributed Environment settings	84
Configuring for database session persistence	85
Switching to a multirow schema	85
Configuring tablespace and page sizes for DB2 session databases	86
Database settings	86
Multirow schema considerations	87
Clustered session support	88
Tuning session management	89
Configuring scheduled invalidation	89
Configuring write contents	90
Configuring write frequency	91
Base in-memory session pool size	91
Controlling write operations	92
Tuning parameter settings	92
Best practices for using HTTP Sessions	94
Managing HTTP sessions: Resources for learning:	97

Chapter 4. Using enterprise beans in applications 99

Enterprise beans	99
Developing enterprise beans	100
Migrating enterprise bean code to the supported specification	101
WebSphere extensions to the Enterprise JavaBeans specification	104
Best practices for developing enterprise beans	105
Using access intent policies	106
Access intent policies	106
Applying access intent policies to methods	108
Access intent exceptions	109
Access intent assembly settings	110
Access intent best practices	112
Frequently asked questions: Access intent	112
EJB modules	114
Assembling EJB modules	114
CMP field assembly settings	116
Container transactions	116
Container transaction assembly settings	116
EJB module assembly settings	118
Entity bean assembly settings	119
EJB local-reference assembly settings	128
Message-driven bean assembly settings	128
Method extensions	131
Method extension assembly settings	131
Method permissions	134
Method permission assembly settings	134
Query assembly settings	135
References	135
EJB reference assembly settings	136
EJB relation assembly settings	137
Exclude list assembly settings	137
Security role assembly settings	138
Session bean assembly properties	139

EJB containers	141
Managing EJB containers	142
EJB container settings	142
EJB container system properties	143
EJB cache settings	144
Container interoperability	144
Deploying EJB modules	148
EJB module collection	149
EJB module settings	149
Enterprise beans: Resources for learning	149

Chapter 5. Using message-driven beans in applications 151

Message-driven beans - an overview	151
Message-driven beans - components	152
Message-driven beans - transaction support	154
Designing an enterprise application to use message-driven beans	154
Developing an enterprise application to use message-driven beans	156
Migrating a JMS listener application to use message-driven beans	158
Deploying an enterprise application to use message-driven beans	159
Configuring deployment attributes for a message-driven bean	160
Configuring message listener resources for message-driven beans	162
Configuring the message listener service	162
Adding a new listener port	166
Configuring a listener port	166
Deleting a listener port	167
Configuring security for message-driven beans	167
Administering listener ports	168
Important files for message-driven beans and extended messaging	169
Troubleshooting message-driven beans	170
Message-driven beans samples	171

Chapter 6. Using application clients 173

Application clients	173
Application client functions	175
ActiveX application clients	176
Applet clients	177
J2EE application clients	178
Pluggable application clients	179
Thin application clients	181
Example: Migrating application clients	182
Migration tips for application clients	183
Installing application clients	184
Developing ActiveX application client code	185
Starting an ActiveX application	185
JClassProxy and JObjectProxy classes	188
Java virtual machine initialization tips	192
Example: Developing ActiveX to enterprise bean bridge, using Java proxy objects	192
Example: Calling Java methods in the ActiveX to enterprise bean bridge	193
Java field programming tips	195

ActiveX to Java primitive data type conversion values	195
Array tips for ActiveX application clients	197
Error handling codes for ActiveX application clients	197
Threading tips	198
Example: Viewing System.out message	199
Example: Enabling logging and tracing for application clients	200
ActiveX client programming best practices	201
Developing applet client code	204
Accessing secure resources using the TCP/IP protocol for applet clients	205
Applet client tag requirements.	206
Applet client code requirements	206
Developing J2EE application client code	207
J2EE application client class loading.	210
Developing pluggable application client code.	212
Developing thin application client code.	213
Assembling Application Client Modules	213
Application client assembly settings	214
Deploying application clients	215
Starting the Application Client Resource Configuration Tool and opening an EAR file	216
JDBC providers for application clients	216
Data sources for application clients	216
Configuring new data source providers (JDBC providers) for application clients	217
Configuring new data sources for application clients	220
Mail providers and mail sessions for the Application Client Assembly Tool	221
Configuring mail providers and sessions for application clients	221
Configuring new mail sessions for application clients	224
URLs for application clients	224
URL providers for the Application Client Resource Configuration Tool	225
Configuring new URL providers for application clients	225
Configuring new URLs with the Application Client Resource Configuration Tool	227
WebSphere asynchronous messaging using the Java Message Service API for the Application Client Resource Configuration Tool	228
Configuring Java messaging client resources	228
Configuring new connection factories for application clients	260
Configuring new Java Message Service destinations for application clients	260
Example: Configuring MQ Queue and Topic connection factories and destination factories for application clients	261
Example: Configuring WAS Queue and Topic connection factories and destination factories for application clients	262
Configuring new resource environment providers for application clients	264
Configuring new resource environment entries for application clients.	265

Managing application clients	266
Updating data source and data source provider configurations with the Application Client Resource Configuration Tool	266
Updating URLs and URL provider configurations for application clients	266
Updating mail session configurations for application clients	267
Updating Java Message Service provider, connection factories, and destination configurations for application clients	267
Updating MQ Java Message Service provider, MQ connection factories, and MQ destination configurations for application clients	268
Updating Resource Environment Entry and Resource Environment Provider configurations for application clients.	268
Removing application client resources	270
Running application clients.	270
launchClient tool	271
Application client troubleshooting tips	274

Chapter 7. Developing and managing Web services 281

Web services	281
Developing a Simple Object Access Protocol client	281
Administering enterprise beans through Simple Object Access Protocol	283
Assembling Web services	283
Simple Object Access Protocol deployment descriptor	284
Simple Object Access Protocol-enabling the services in an Enterprise Application Archive file	287
Deploying Web services applications	289
Administering deployed Web services (XML-SOAP administrative tool)	290
Securing Simple Object Access Protocol services	290
Migrating SOAP security	291
Securing Simple Object Access Protocol services with HTTP basic authentication	291
Securing Simple Object Access Protocol services on Secured Socket Layer.	291
Securing SOAP services on SSL with SOAP Signature.	292
UDDI4J specifications	295
Web services: Resources for learning.	295

Chapter 8. Enabling Web services to use the Web Services Invocation Framework 299

Goals of WSIF	299
WSIF - Web services are not just SOAP services	300
WSIF - tying client code to a particular protocol implementation is restricting	300
WSIF - incorporating new bindings into client code is hard	300
WSIF - multiple bindings can be used in flexible ways	300

WSIF - a freer Web services environment	
enables intermediaries	301
An overview of WSIF	301
WSIF architecture	301
Using WSIF with Web services that offer	
multiple bindings	302
WSIF and WSDL	302
WSIF usage scenarios.	303
Dynamic invocation	304
Using WSIF to invoke Web services	304
Using the WSIF providers	304
Developing a WSIF service	317
Using complex types	326
Using JNDI	327
Interacting with the WebSphere J2EE container	329
Running WSIF as a client	329
WSIF system management and administration	329
Maintaining the WSIF properties file	329
Enabling security for WSIF	330
WSIF troubleshooting tips	331
WSIF API	335
WSIF API reference: Creating a message for	
sending to a port	336
WSIF API reference: Finding a port factory or	
service.	337
WSIF API reference: Using ports	338
WSIF: Resources for learning	342

Chapter 9. Classloading 343

Classloaders	344
Classloader collection.	347
Classloader ID	348
Classloader Mode	348
Classloader settings	348
Migrating the classloader Module Visibility Mode	
setting.	348
Classloading: Resources for learning.	349

Chapter 10. Using EJB query 351

EJB query language	351
Example: EJB queries	352
FROM clause	354
Inheritance in EJB query.	355
Path expressions	355
WHERE clause	356
Scalar functions	364
Aggregation functions	367
SELECT clause	368
ORDER BY clause	369
Subqueries	369
EJB query restrictions.	370
EJB Query: Reserved words	371
EJB query: BNF syntax	371
Comparison of EJB 2.0 specification and	
WebSphere query language.	373

Chapter 11. Internationalizing applications 375

Internationalization	375
Identifying localizable text	376

Creating message catalogs	377
Composing language-specific strings	377
Localization API support	378
LocalizableTextFormatter class.	379
Creating a formatter instance	381
Setting optional localization values	382
Composing complex strings	383
Generating localized text	385
Preparing the localizable-text package for	
deployment	386
LocalizableTextEJBDeploy command	386
Internationalization: Resources for learning	387

Chapter 12. Using the transaction service 389

Transaction support in WebSphere Application	
Server	389
Resource manager local transaction (RMLT)	390
Global transactions	391
Local transaction containment (LTC).	391
Local and global transaction considerations	395
Developing components to use transactions	395
Setting transactional attributes in the	
deployment descriptor	396
Using bean-managed transactions	397
Configuring transaction properties for an	
application server	398
Transaction service settings.	400
Managing active transactions	400
Managing transaction logging for optimum server	
availability	401
Configuring transaction aspects of servers for	
optimum availability	402
Moving a transaction log from one server to	
another	403
Restarting an application server on a different	
host	404
Transactional interoperation with non-WebSphere	
application servers	405
Troubleshooting transactions	405
Transaction service exceptions.	406
Standard exceptions	406
Heuristic exceptions	406
UserTransaction interface - methods available	407

Chapter 13. Using naming 409

Naming	410
New features for name space support	410
Name space logical view	411
Name space partitions	412
Initial context support	414
Initial contexts registered with the ORB as initial	
references	414
Default initial contexts	415
Lookup names support in deployment descriptors	
and thin clients.	415
Relative names	415
Qualified names	416
JNDI support in WebSphere Application Server	418
Developing applications that use JNDI	418

Example: Getting the default initial context	421
Example: Getting an initial context by setting the provider URL property	424
Example: Setting the provider URL property to select a different root context as the initial context	426
Example: Looking up an EJB home with JNDI	428
Example: Looking up a JavaMail session with JNDI	430
JNDI interoperability considerations	430
JNDI caching	432
JNDI cache settings	433
Example: Controlling JNDI cache behavior from a program	434
JNDI name syntax	435
INS name syntax	436
JNDI to CORBA name mapping considerations	436
Example: Setting the syntax used to parse name strings	436
Developing applications that use CosNaming (CORBA Naming interface)	437
Example: Getting an initial context with CosNaming	437
Example: Looking up an EJB home with CosNaming	440
Configured name bindings	442
Configured binding types	443
Name space federation	444
Name space bindings	446
Configuring and viewing name space bindings	446
String binding settings	446
CORBA object binding settings	447
Indirect lookup binding settings	448
EJB binding settings	448
Name space binding collection	449
Configuring name servers	449
Name server settings	450
Troubleshooting name space problems	450
dumpNameSpace tool	450
Example: Invoking the name space dump utility	452
Name space dump utility for java: and local:name spaces	453
Example: Invoking the name space dump utility for java: and local: name spaces	455
Name space dump sample output	455
Naming and directories: Resources for learning	457

Chapter 14. Improving performance through the dynamic cache service 459

Dynamic cache	459
Configuring globally the dynamic cache service	459
DynamicCache service settings	460
Configuring servlet caching	460
Configuring the dynamic cache disk offload	461
Configuring Edge Side Include caching	461
Configuring external cache groups	463
Displaying cache information	466
Configuring cacheable objects with the cachespec.xml file	466
Verifying the cacheable page	468
Cachespec.xml file	468

Configuring command caching	474
Command class	474
CacheableCommandImpl class	475
Example: Caching a command object	475
Example: Caching Web services	476
Example: Configuring the dynamic cache	479

Chapter 15. Managing user profiles 483

User profile	483
UserProfileManager class	484
User profile development options	484
Extending the data represented in user profiles	484
Adding columns to the base user profile implementation	484
Extending the User Profile enterprise bean and importing legacy databases	485
UPServletExample.java	485
UserProfileExtendedSample.java	487
UPServletExampleExtended.java	488
UserProfileExtended.java	490
UPServletExtended.java	491
userprofile.xml	493

Chapter 16. Assembling applications 495

Application assembly and J2EE applications	496
Archive support in Version 5.0	497
Starting the Application Assembly Tool (AAT)	497
Migrating application modules from J2EE 1.2 to J2EE 1.3	498
earconvert tool	499
Assembling new or modifying existing modules	499
Adding files to assembled modules	502
Enterprise application assembly properties	503
Resource environment reference assembly settings	504
Resource Adapter Archive file assembly settings	504
Saving applications after assembly	508
Verifying archive files	508
Generating code for deployment	509
ejbdeploy tool	510
ejbdeploy syntax -- relationship to Application Assembly Tool options	511

Chapter 17. Deploying and managing applications 513

Enterprise applications	513
Installing a new application	513
Preparing for application install settings	518
Example: Installing an EAR file using the default bindings	522
Enterprise application collection	522
Name	522
Status	522
Enterprise application settings	523
Starting and stopping applications	525
Exporting applications	526
Exporting DDL files	526
Updating applications	526
Hot deployment and dynamic reloading	527
Uninstalling applications	536

Deploying and managing applications: Resources
for learning 536

Trademarks and service marks

The following terms are trademarks of IBM Corporation in the United States, other countries, or both:

- Everyplace
- iSeries
- IBM
- Redbooks
- ViaVoice
- WebSphere
- 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 trademarks of Microsoft 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.

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

Chapter 1. Welcome to Applications

The following items comprise the application programming model, including numerous services available to support deployed applications.

Web modules

Use Web components such as servlets and JavaServer Pages files to develop dynamic Web sites. Product extensions to the open source servlet and JSP APIs enhance standard features, and provide additional functionality.

Web modules consist of the following application components, each performing a different function:

- HTML and JSP pages provide the user interface and program logic
- Servlets coordinate work between other components of the application
- Bean Scripting Framework enables users to implement JavaScript code in JSP applications

HTTP sessions are a key area of product support for Web modules. By managing (**HTTP sessions**) for your Web applications, you can personalize a Web site for individual customers. A session is a series of requests to a servlet, originating from the same user at the same browser. Managing HTTP sessions allows servlets running in a Web container to keep track of individual users. For example, a servlet might use sessions to provide "shopping carts" to on-line shoppers. Suppose the servlet is designed to record the items each shopper indicates he or she will purchase from the Web site. It is important that the servlet be able to associate incoming requests with particular shoppers. Otherwise, the servlet might mistakenly add choices of Shopper 1 to the cart of Shopper 2.

EJB modules

IBM WebSphere Application Server provides broad support for enterprise beans, including the Enterprise JavaBeans (EJB) 2.0 specification. The EJB 2.0 specification introduces a container-managed persistence (CMP) 2.0 component model, which provides a number of improvements to aid developer productivity and application performance. In addition, this product continues to fully support enterprise beans written to the CMP 1.1 programming model and deployed in previous versions of this product; applications can use CMP 1.1 beans, CMP 2.0 beans, or a mixture of both. CMP 1.1 beans can be directly carried forward in an EJB 1.1 `ejb-jar` module or may be repackaged and combined with CMP 2.0 beans in an EJB 2.0 module.

For EJB 2.0 modules, a feature introduced in Version 5 of this product, called **access intent** policies, eases the management of interactions between CMP beans and their underlying datastores. Each policy sets such data access characteristics such as access type (read or update) and transaction isolation that affect the locking of resources, letting you choose the level of data integrity and performance for your application.

Several excellent trade books that cover EJB 2.0 and the CMP 2.0 persistence model are already available. A good way to locate some of these is to visit your favorite online bookstore and search on the term *Enterprise JavaBeans*. For a more basic orientation, see ("Enterprise beans: Resources for learning").

Your application development might include **asynchronous messaging**, which the product supports as a method of communication based on the Java Message Service (JMS) programming interface.

The base JMS support enables IBM WebSphere Application Server applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An application can explicitly poll for messages on a destination.

The product also provides a message listener service that applications can use to automatically retrieve messages from JMS destinations for processing by message-driven beans, without the application having to explicitly poll JMS destinations.

Refer to:

- "Asynchronous messaging with WebSphere - an overview" (not in this document)
- "Using JMS and messaging in applications" (not in this document)
- ("Using message-driven beans in applications")

Client modules

The product provides a CD-ROM and installation program for installing application clients without installing the entire application server. This smaller footprint is useful if you want to run client applications on multiple client machines. For more information, see ("Installing application clients").

Application clients follow several programming models with unique requirements and suitability for different types of applications. Most of these models are only available when you install the product. Models include:

ActiveX application client

Supported Windows platforms only; Client only

Applet client

Supported Windows platforms only; Client only

J2EE application client

All supported server platforms; All supported client platforms

Pluggable application client

Supported Windows platforms only, Sun JRE 1.3.1_03 or later (but not 1.4); Thin application client programming model only (no J2EE programming model); Client only

Thin application client

All supported client platforms

Web services

The Web services components included with this product version build upon the Apache Simple Object Access Protocol (SOAP) 2.3-based capabilities delivered with Version 4.0.x of the product.

New in this release is an open source implementation for a Web Services Invocation Framework (WSIF).

An additional Web services component, "IBM WebSphere Web Services for J2EE Technology Preview", is available for use with Version 5.0. It is a separate download available at:

<http://www7b.boulder.ibm.com/wsdd/downloads/techpreviews.html>

and includes additional documentation. The Web services technology preview supports emerging Java Web services standards like JAX-RPC and Web services for J2EE. It is recommended that new development efforts use the Web Services Technology Preview and follow these standards.

Additional features are available, such as UDDI Registry and Web Services Gateway. See their description in "Welcome to Servers" (not in this document).

Application services

IBM WebSphere Application Server provides essential services to ease the building of dynamic and flexible e-business applications. These services support and extend the open standards of J2EE and Web services, with a focus on application reuse and integration.

- **Classloading**

The WebSphere Application Server product provides several classloading modes, policies, and features to enable you to deploy and run your applications successfully. An application server provides an Application Classloader Policy that enables you to control the isolation of applications in a server. If you want applications to share classes, choose the SINGLE policy; otherwise choose the MULTIPLE policy, which isolates the classloaders for each application.

Similarly, at the application level, you can choose a (WAR Classloader Policy) that configures the isolation of Web modules within an application. If you choose the policy APPLICATION, then each Web module in your application can see the other Web module's classes. A policy of MODULE creates a separate classloader for each Web module resulting in isolation for each Web module's classes.

The classloader mode setting, which you can configure at the server, application, or Web module level depending on your classloader policy, enables you to control whether application classloaders override classes contained in base run-time classloaders. By default, the WebSphere Application Server classloaders have a classloader mode of PARENT_FIRST, which is the standard JDK mode and does not allow the application classloader to override classes. You must take care when using the PARENT_LAST classloader mode to make all dependent classes available within the application or you might get LinkageErrors or other classloader exceptions. For example, if you provide a newer version of the Xerces.jar file and your application is using XSLT, you must provide a compatible version of Xalan.jar within the application.

Finally, Version 5.0 of WebSphere Application Server introduces the concept of classloader "Managing shared libraries" (not in this document). A shared library is a CLASSPATH and a symbolic name for the classpath. You define shared libraries at the cell, node, or server level and then associate the shared libraries either with an application server (making the classes available to all applications in the server) or with individual applications (making the classes available only to the referencing application). This mechanism provides a convenient way to make libraries of classes available to your applications outside of a standard J2EE enterprise application (EAR) file for easier version management and space efficiency.

- **Internationalization**

If your application component must support multiple locales, the **localizable-text** API can help both developers and administrators through central management of displayed strings. The developer separates strings into a message catalog, which is then translated into the other languages required. All message catalogs are then deployed with the application component. From then on, the administrator can add or update message catalogs centrally as required.

- **Transactions**

IBM WebSphere Application Server applications can use transactions to coordinate multiple updates to resources as atomic units (as indivisible units of work) such that all or none of the updates are made permanent. The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself)
- Entity beans use container-managed transactions
- Web components (servlets) use bean-managed transactions

The product is a transaction manager that supports the coordination of resource managers through their XAResource interface and participates in distributed global transactions with other OTS 1.2 compliant transaction managers (for example J2EE 1.3 application servers). Applications can also be configured to interact with databases, JMS queues, and JCA connectors through their local transaction support when distributed transaction coordination is not required.

Resource managers that offer transaction support can be categorized into those that support 2-phase coordination (by offering an XAResource interface) and those that support only 1-phase coordination (for example through a LocalTransaction interface). The IBM WebSphere Application Server transaction support provides coordination, within a transaction, for any number of 2-phase capable resource managers. It also enables a single 1-phase capable resource manager to be used within a transaction in the absence of any other resource managers, although a WebSphere transaction is not necessary in this case. With the Last Participant Support of Enterprise Extensions, you can coordinate the use of a single 1-phase commit (1PC) capable resource with any number of 2-phase commit (2PC) capable resources in the same global transaction. At transaction commit, the 2-phase commit resources are prepared first using the 2-phase commit protocol, and if this is successful the 1-phase commit-resource is then called to commit(one_phase). The 2-phase commit resources are then committed or rolled back depending on the response of the 1-phase commit resource.

The ActivitySession service of Enterprise Extensions provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. It is a distributed context that can be used to coordinate multiple 1-phase resource managers. The product EJB container and deployment tooling support ActivitySessions as an extension to the J2EE programming model. Enterprise beans can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An application can then interact with a resource manager through its LocalTransaction interface for the period of a client-scoped ActivitySession rather than just the duration of an EJB method.

- **Naming**

Naming clients use (**Naming Services**) primarily to access objects, such as EJB homes, associated with applications installed on IBM WebSphere Application Server. Objects are made available to clients by being bound into a name space.

A name space is under the control of a name server. In this product, there are potentially many name servers, and the name spaces controlled by the various name servers are federated together to form the view of a single name space. Each name server presents the same logical view of the federated name spaces.

Name servers provided by this product are a CORBA CosNaming implementation. IBM WebSphere Application Server provides a CosNaming JNDI plug-in which enables clients to access the name servers through the JNDI interface. Clients to EJB applications typically use JNDI to perform Naming operations. Clients may access the name servers directly through the CORBA programming model. The CosNaming interface is part of the CORBA programming model. CORBA clients which need to access EJB homes or some other objects bound to the name space would typically use the CORBA CosNaming interface to perform Naming operations.

- **Dynamic cache**

Dynamic Cache improves application performance by caching outputs and contents of outputs of Servlets, Java Server Pages files, Web Services and Commands. On subsequent client requests to the same applications, Dynamic cache intercepts these calls and responds by serving the output or the contents of output from the cache.

Dynamic Cache in this product version includes:

- **Servlet/JSP caching**

- This caches output of dynamic servlets and JSP files by working with Java virtual machine of the application server by intercepting calls to service methods and serving Web pages from the cache. This improves server response time, throughput and scalability.

- **Command caching**

- Commands that are written to the Command Architecture encapsulate business logic tasks and provide a standard way to invoke the business logic request. Command objects need to implement CacheableCommand interface instead of TargetableCommand interface to cache. Like in servlets and JSP caching, requests to execute business logic in the command is intercepted by the cache. If a command with the same request attributes are available in cache, output properties are copied from the cached instance to the requested instance and returned without executing the business logic again.

- **Web Services caching**

- Web service responses can be cached just like servlet and JSP results. These requests are intercepted and cache ID computed based on how the cache ID rules are specified in the cache policy. Hash of the whole SOAPEnvelope can be used as a cache ID or it can be parsed and service, operation and parameters to these operations used as cache ID. If a cache entry is not found for the computed cache id, the request is forwarded to the SOAP engine and the result is cached.

- **Edge Side Include caching**

- This provides the ability to cache, assemble and deliver dynamic web pages at the edge of the enterprise network. Edge Side Includes (ESI) is a simple markup language which enables dynamic web pages (which by themselves are not so cache efficient) to be broken down into cacheable fragments. These fragments are then cached on the edge of the network and assembled into a single page upon user requests.

- **Distributed caching**

- Cache contents can be shared and replicated among servers by Dynamic

Caching using an underlying JMS based message broker system, DRS (Data Replication Service). Sharing characteristics of individual cache entry is configured using the cache policy specification.

- **User profiles**

Managing (**user profiles**) allows a company to maintain database tables containing fields for demographic data of individual customers or other users on the company system. For example, when a user repeatedly logs onto a Web site that supports user profiles, the Web site can display headlines and advertising tailored to the shopping preferences of that user. The site can address the user by his or her logon name. User profile API is deprecated in the current release.

Assembly tools

The Application Assembly Tool and a command line deployment tool are provided for packaging your application code components into the needed modules for deployment onto the server.

See ("Assembling applications").

EAR files are comprised of the following archives:

- Enterprise bean (JAR) files (known as ("EJB modules"))
- Web application (WAR) files (known as ("Web modules"))
- Application client (JAR) files (known as application client modules)
- Resource adapter (RAR) files (known as resource adapter modules)
- Optionally, additional JAR files containing dependent classes or other components required by the application

The standard file extension of an Enterprise application file is .ear.

For a discussion of archives and Web components supported by the Application Assembly Tool in Version 5, see ("Archive support in Version 5.0").

See also, "Assembling or packaging" (not in this document).

Deployment

Tools, such as the WebSphere Administrative Console, are provided for installing your modules onto the application server, then managing the installed modules in their respective containers.

Chapter 2. Using Web applications

A developer creates the files comprising a Web application, and then assembles the Web application components into a Web module. Next, the deployer (typically the developer in a unit-testing environment or the administrator in a production environment) installs the Web application on the server.

Therefore the development process generally consists of the following steps:

Steps for this task

1. **(Optional)** Migrate existing Web applications to run in the new version of WebSphere.
2. Design the Web application and develop its code artifacts: Servlets, JavaServer Pages (JSP) files, and static files, as for example, images and Hyper Text Markup Language (HTML) files.
See the "Resources for learning" article for links to design documentation.
3. **(Optional)** Implement JavaScript within JSP tags using the Bean Scripting Framework (BSF).
4. Develop the Web application, using WebSphere Application Server extensions to enhance its functionality.
5. Assemble the Web application into a Web module.

Web module assembly properties might include the ability to:

- Configure servlet page lists
 - Configure servlet filters
 - Serve servlets by class name
 - Enable file serving
6. (Deploy the Web module or application module) that contains the Web application.
Following deployment, you might find it handy to use the tool that enables batch compiling of the JSP files for quicker initial response times.
 7. **(Optional)** Troubleshoot your Web application.
 8. **(Optional)** Modify the default Web container configuration in the application server in which you deployed the Web module or application module containing the Web application.
 9. **(Optional)** (Manage the deployed Web application.).

Web applications

A Web application is comprised of one or more related servlets, JavaServer Pages technology (JSP files), and Hyper Text Markup Language (HTML) files that you can manage as a unit.

The files in a Web application are related in that they work together to perform a business logic function.

For example, one of the WebSphere Application Server samples is a Simple Greeting Web application. This application, comprised of a servlet and Web pages, greets new users when the application is accessed.

The Web application is a concept supported by the Java Servlet Specification. Web applications are typically packaged as .war files.

web.xml file

The web.xml file provides configuration and deployment information for the Web components that comprise a Web application. Examples of Web components are servlet parameters, servlet and JavaServer Pages (JSP) definitions, and Uniform Resource Locators (URL) mappings.

The servlet 2.3 specification dictates the format of the web.xml file, which makes this file portable among Java Two Enterprise Edition (J2EE) compliant products.

Location

The web.xml file must reside in the WEB-INF directory under the context of the hierarchy of directories that exist for a Web application. For example, if the application is client.war, then the web.xml file is placed in the *install_root/client.war*/WEB-INF directory.

Usage notes

- Is this file read-only?
No
- Is this file updated by a product component?
This file is updated by the Application Assembly Tool (AAT).
- If so, what triggers its update?
The AAT updates the web.xml file when you assemble Web components into a Web module, or when you modify the properties of the Web components or the Web module.
- How and when are the contents of this file used?
WebSphere Application Server functions use information in this file during the configuration and deployment phases of Web application development.

Sample file entry

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE web-app PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application
2.2//EN" "http://java.sun.com/j2ee/dtds/web-app_2_2.dtd">

  <web-app id="WebApp_1">
    <display-name>Persistence Manager Web Client</display-name>
    <description>Peristence Manager Web Client</description>
    <servlet id="Servlet_1">
      <servlet-name>CustomerLocalServlet</servlet-name>
      <description>Local Customer Servlet</description>
      <servlet-class>CustomerLocalServlet</servlet-class>
    </servlet>
    <servlet id="Servlet_2">
      <servlet-name>CustomerServlet</servlet-name>
      <description>Remote Customer Servlet</description>
      <servlet-class>CustomerServlet</servlet-class>
    </servlet>
    <servlet id="Servlet_3">
      <servlet-name>CreditCardServlet</servlet-name>
      <description>Credit Card Servlet - PM Verification</description>
      <servlet-class>CreditCardServlet</servlet-class>
    </servlet>
    <servlet-mapping id="ServletMapping_1">
      <servlet-name>CustomerLocalServlet</servlet-name>
```

```

        <url-pattern>/CustomerLocal</url-pattern>
    </servlet-mapping>
    <servlet-mapping id="ServletMapping_2">
        <servlet-name>CustomerServlet</servlet-name>
        <url-pattern>/Customer</url-pattern>
    </servlet-mapping>
    <servlet-mapping id="ServletMapping_3">
        <servlet-name>CreditCardServlet</servlet-name>
        <url-pattern>/CreditCard</url-pattern>
    </servlet-mapping>
    <welcome-file-list id="WelcomeFileList_1">
        <welcome-file>index.html</welcome-file>
    </welcome-file-list>
    <security-role id="SecurityRole_1">
        <description>Everyone role</description>
        <role-name>Everyone Role</role-name>
    </security-role>
    <security-role id="SecurityRole_2">
        <description>AllAuthenticated role</description>
        <role-name>All Role</role-name>
    </security-role>
    <security-role id="SecurityRole_3">
        <description>Deny all access role</description>
        <role-name>DenyAllRole</role-name>
    </security-role>
</web-app>

```

Migrating Web application components

Supported open specification levels in WebSphere Application Server Version 5 are documented in article, Migrating APIs and specifications.

Migration of Web applications deployed in WebSphere Application Server Version 4.x is not necessary; version 2.2 of the servlet specification and version 1.1 of the JavaServerPages (JSP) specification are still supported. However, where there are behavioral differences between the Java Two Enterprise Edition (J2EE) 1.2 and J2EE 1.3 specifications, bear in mind that J2EE 1.3 specifications are implemented in WebSphere Application Server Version 5 and will override any J2EE 1.2 behaviors.

Servlet migration might be a concern if your application:

- implements a WebSphere internal servlet to bypass a WebSphere Application Server Version 4.x single application path restriction.
- extends a PageListServlet that relies on configuration information in the servlet configuration XML file.
- uses a servlet to generate Hyper Text Markup Language (HTML) output.
- calls the `response.sendRedirect()` method for a servlet using the `encodeRedirectURL` function or executing within a non-context root.

JSP migration might be a concern if your application references JSP page implementation classes in unnamed packages, or if you install WebSphere Application Server Version 4.x EAR files (deployed in Version 4.x with the JSP Precompile option), in Version 5.

Follow these steps if migration issues apply to your Web application:

Steps for this task

1. Use WebSphere Application Server Version 5 package names for any WebSphere Application Server Version 4.x internal servlets, which are implemented in your application.

In WebSphere Application Server Version 4.x, Web modules with a context root setting of / are not supported. Accessing Web modules with this root context results in HTTP 404 - File not Found errors.

To bypass the errors, and to enable the serving of static files from the root context, WebSphere Application Server Version 4.x users are advised to add the servlet class, `com.ibm.servlet.engine.webapp.SimpleFileServlet`, to their Web module.

The Version 4.x single path limitation does not exist in Version 5. However, users who choose to use the `com.ibm.servlet.engine.webapp.SimpleFileServlet` in Version 5 must do one of the following:

- Rename `com.ibm.servlet.engine.webapp.SimpleFileServlet` to `com.ibm.ws.webcontainer.servlet.SimpleFileServlet`.
- Open the EAR file in the Application Assembly Tool (AAT) and enable the `SimpleFileServlet` static file setting.

The following list identifies the other internal servlets affected by the Version 5 package name change:

- `DefaultErrorReporter`
- `AutoInvoker`

Use the Version 5 package name, `com.ibm.ws.webcontainer.servlet.<servlet class name>` for these servlets.

2. Use the WASPostUpgrade tool to migrate servlets that extend `PageListServlet` and rely on configuration information in the associated XML servlet configuration file.

In Version 4.x, the XML servlet configuration file provides configuration data for page lists and augments servlet configuration information. This file is named as either `<servlet class name>.servlet` or `<servlet class name>.servlet`, and is stored in the same directory as the servlet class file.

The XML servlet configuration file is not supported in WebSphere Application Server Version 5.

3. Set a content type if your servlet generates Hyper Text Markup Language (HTML) output.

The default behavior of the Web container changed in WebSphere Application Server Version 5. If the servlet developer does not specify a content type in the servlet then the container is forbidden to set one automatically. Without an explicit content type setting, the content type is set to null. The Netscape browser displays HTML source as plain text with a null content type setting.

To resolve this problem, do one of the following:

- Explicitly set a content type in your servlet.
- Open the WAR file in the Application Assembly Tool (AAT) and enable the `autoResponseEncoding` static file setting.

4. Set the Java environment variable, `com.ibm.websphere.sendredirect.compatibility`, to **true** if you want your URLs interpreted relative to the application root.

The default value of the Java environment variable `com.ibm.websphere.sendredirect.compatibility` changed in WebSphere Application Server Version 5. In Version 4, the default setting of this variable is true. In Version 5, the setting is false.

When this variable is set to false, if a URL has a leading slash, the URL is interpreted relative to the Web module/application root. However, if the URL

does not have a leading slash, it is interpreted relative to the Web container root (also known as the Web server document root). Therefore, if an application has a WAR file that has a context root of `myPledge_app` and a servlet that has a servlet mapping of `/Intranet/`, a JSP file in the WAR file cannot access the servlet when its `encodeRedirectURL` is set to `/Intranet/myPledge`. The JSP file can access the servlet if the `encodeRedirectURL` is set to `myPledge_app/Intranet/myPledge`, or if the `com.ibm.websphere.sendredirect.compatibility` variable is set to `true`.

See the [Setting the sendredirect variable](#) article for more information.

5. Use the `WASPostUpgrade` tool to migrate WebSphere Version 4.x enterprise applications to Version 5.

Note: The WebSphere Application Server Version 4.x JSP page implementation class files are not compatible with the WebSphere Application Server Version 5 JSP container.

The `WASPostUpgrade` tool automatically precompiles JSP files, which ensures the JSP page implementation class files are compatible with Version 5.

If you install Version 4.x EAR files, deployed with the `JSP Precompile` option, in Version 5, and you choose not to follow the migration path, do one of the following:

- Select the `Pre-compile JSP` option in the administrative console `Install New Application` window.
See article [\(Installing a new application\)](#) for more information.
- Specify the `-preCompileJSPs` option when using the `Wsadmin` tool.

6. Import your classes if your application uses unnamed packages.

Section 8.2 of the JSP 1.2 specification states:

The JSP container creates a JSP page implementation class for each JSP page. The name of the JSP page implementation class is implementation dependent.

The JSP page implementation object belongs to an implementation-dependent named package. The package used may vary between one JSP and another, so minimal assumptions should be made. The unnamed package should not be used without an explicit `import` of the class.

For example, if `myBeanClass` is in the unnamed package, and you reference it in a `jsp:useBean` tag, then you must explicitly import `myBeanClass` with the page directive `import` attribute, as shown in the following example:

```
<%@page import="myBeanClass" %>
.
.
.
<jsp:useBean id="myBean" class="myBeanClass" scope="session"/>
```

In WebSphere Application Server Version 5, the JSP engine creates JSP page implementation classes in the `org.apache.jsp` package. If a class in the unnamed package is not explicitly imported, then the `javac` compiler assumes the class is in package `org.apache.jsp`, and the compilation fails.

Note: Avoid using the unnamed package altogether because of a change made in JDK 1.4 that will affect the JSP 2.0 specification. WebSphere Application Server Version 5 ships with JDK 1.3.1, so this is not an issue with the Version 5 JSP engine, but it will become an issue in future releases.

The *Incompatibilities* section of the version 1.4 Java 2 Platform, Standard Edition (J2SE) documentation states:

The compiler now rejects `import` statements that import a type from the unnamed namespace. Previous versions of the compiler would accept such `import` declarations,

even though they were arguably not allowed by the language (because the type name appearing in the import clause is not in scope). The specification is being clarified to state clearly that you cannot have a simple name in an import statement, nor can you import from the unnamed namespace.

To summarize, the syntax:

```
import SimpleName;
```

is no longer legal. Nor is the syntax:

```
import ClassInUnnamedNamespace.Nested;
```

which would import a nested class from the unnamed namespace.

To fix such problems in your code, move all of the classes from the unnamed namespace into a named namespace.

What to do next

See "Resources for learning" for links to the J2SE, JSP, and Servlet specification documentation.

Default Application

The IBM WebSphere Application Server provides a default configuration that allows administrators to easily verify that the Application Server is running. When the product is installed, it includes an application server called *server1* and an enterprise application called *Default Application*.

Default Application contains a Web Module called *DefaultWebApplication* and an enterprise bean JAR file called *Increment*. The *Default Application* provides a number of servlets, described below. These servlets are available in the product.

For additional code examples, visit the Samples Gallery. Learn how to locate and install the Samples Gallery by viewing the Samples Gallery reference page.

The URL for accessing Samples is: <http://localhost:9080/WSamples/>

Snoop

Use the Snoop servlet to retrieve information about a servlet request. This servlet returns the following information:

- Servlet initialization parameters
- Servlet context initialization parameters
- URL invocation request parameters
- Preferred client locale
- Context path
- User principal
- Request headers and their values
- Request parameter names and their values
- HTTPS protocol information
- Servlet request attributes and their values
- HTTP session information
- Session attributes and their values

The Snoop servlet includes security configuration so that when WebSphere Security is enabled, clients must supply a user ID and password to execute the servlet.

The URL for the Snoop servlet is: <http://localhost:9080/snoop/>.

HelloHTML

Use the HelloHTML pervasive servlet to exercise the PageList support provided by the WebSphere Web container. This servlet extends the PageListServlet, which provides APIs that allow servlets to call other Web resources by name or, when using the *Client Type detection* support, by type.

You can invoke the Hello servlet from an HTML browser, speech client, or most Wireless Application Protocol (WAP) enabled browsers using the URL: <http://localhost:9080/HelloHTML.jsp/>.

HitCount

Use the HitCount Demonstration application to demonstrate incrementing a counter using a variety of methods, including:

- A servlet instance variable
- An HTTP session
- An enterprise bean

You can instruct the servlet to execute any of these methods within a transaction that you can omit or roll back. If the transaction is committed, the counter is incremented. If the transaction is rolled back, the counter is not incremented.

The enterprise bean method uses a Container- Managed Persistence enterprise bean that persists the counter value to a Cloudscape database. This enterprise bean is configured to use the Default Datasource, which is set to the DefaultDB database.

When using the enterprise bean method, you can instruct the servlet to look up the enterprise bean, either in the WebSphere global namespace, or in the namespace local to the application.

The URL for the HitCount application is: <http://localhost:9080/HitCount.jsp/>.

Servlets

Servlets are Java programs that use the Java Servlet Application Programming Interface (API). You must package servlets in a Web ARchive (WAR) file or Web module for deployment to the application server.

Servlets run on a Java-enabled Web server and extend the capabilities of a Web server, similar to the way applets run on a browser and extend the capabilities of a browser.

Servlets can support dynamic Web page content, provide database access, serve multiple clients at one time, and filter data.

For the purposes of IBM WebSphere Application Server, discussions of servlets focus on Hyper Text Transfer Protocol (HTTP) servlets, which serve Web-based clients.

Developing servlets with WebSphere Application Server extensions

Several WebSphere Application Server extensions are provided for enhancing your servlets. This task provides a summary of the extensions that you can utilize.

Steps for this task

1. Review the supported specifications.

Create Java components, referring to the Servlet specifications from Sun Microsystems.

See Resources for learning for links to coding specifications and examples.

The application server includes its own packages that extend and add to the Java Servlet Application Programming Interface (API). These extensions and additions make it easier to manage session states, create personalized Web pages, generate better servlet error reports, and access databases. Locate the Javadoc for the application server APIs in the product `install_root\web\apidocs` directory.

All the public WebSphere Application Server APIs are located in the `com.ibm.websphere...` packages.

2. Use your favorite integrated development environment (IDE), or a text editor, to develop or migrate code artifacts that meet the specifications.
3. Test the code artifacts.

What to do next

Assemble your code artifacts into a Web module as a prerequisite to deploying the code to the application server.

Application lifecycle listeners and events

Application lifecycle listeners and events, now part of the Servlet API, enable you to notify interested listeners when servlet contexts and sessions change. For example, you can notify users when attributes change and if sessions or servlet contexts are created or destroyed.

The lifecycle listeners give the application developer greater control over interactions with `ServletContext` and `HttpSession` objects. Servlet context listeners manage resources at an application level. Session listeners manage resources associated with a series of requests from a single client. Listeners are available for lifecycle events and for attribute modification events. The listener developer creates a class that implements the `javax` listener interface, corresponding to the desired listener functionality.

At application startup time, the container uses introspection to create an instance of your listener class and registers it with the appropriate event generator.

When a servlet context is created, the `contextInitialized` method of your listener class is invoked, which creates the database connection for the servlets in your application to use, if this context is for your application.

When the servlet context is destroyed, your `contextDestroyed` method is invoked, which releases the database connection, if this context is for your application.

Listener classes for servlet context and session changes

The following methods are defined as part of the `javax.servlet.ServletContextListener` interface:

- `void contextInitialized(ServletContextEvent)` - Notification that the Web application is ready to process requests.

Place code in this method to see if the created context is for your Web application and if it is, allocate a database connection and store the connection in the servlet context.

- `void contextDestroyed(ServletContextEvent)` -Notification that the servlet context is about to shut down.

Place code in this method to see if the created context is for your Web application and if it is, close the database connection stored in the servlet context.

Two new listener interfaces are defined as part of the `javax.servlet` package:

- `ServletContextListener`
- `ServletContextAttributeListener`

One new filter interface is defined as part of the `javax.servlet` package:

- `FilterChain` interface - methods: `doFilter()`

Two new event classes are defined as part of the `javax.servlet` package:

- `ServletContextEvent`
- `ServletContextAttributeEvent`

Three new listener interfaces are defined as part of the `javax.servlet.http` package:

- `HttpSessionListener`
- `HttpSessionAttributeListener`
- `HttpSessionActivationListener`

One new event class is defined as part of the `javax.servlet.http` package:

- `HttpSessionEvent`

Example: `com.ibm.websphere.DBConnectionListener.java`

The following example shows how to create a servlet context listener:

```
package com.ibm.websphere;

import java.io.*;
import javax.servlet.*;

public class DBConnectionListener implements ServletContextListener
{
    // implement the required context init method
    void contextInitialized(ServletContextEvent sce)
    {
    }

    // implement the required context init method
    void contextDestroyed(ServletContextEvent sce)
    {
    }
}
```

Servlet filtering

Servlet filtering is an integral part of the Servlet 2.3 API. Servlet filtering provides a new type of object called a *filter* that can transform a request or modify a response.

You can chain filters together so that a group of filters can act on the input and output of a specified resource or group of resources.

Filters typically include logging filters, image conversion filters, encryption filters, and Multipurpose Internet Mail Extensions (MIME) type filters (functionally equivalent to the servlet chaining). Although filters are not servlets, their lifecycle is very similar.

Filters are handled in the following manner:

- The Web container determines whether it needs to construct a `FilterChain` containing the `LoggingFilter` for the requested resource.
The `FilterChain` begins with the invocation of the `LoggingFilter` and ends with the invocation of the requested resource.
- If other filters need to go in the chain, the Web container places them after the `LoggingFilter` and before the requested resource.
- The Web container then instantiates and initializes the `LoggingFilter` (if it was not done previously) and invokes its `doFilter(FilterConfig)` method to start the chain.
- The `LoggingFilter` preprocesses the request and response objects and then invokes the filter chain `doFilter(ServletRequest, ServletResponse)` method. This method passes the processing to the next resource in the chain (in this case, the requested resource).
- Upon return from the filter chain `doFilter(ServletRequest, ServletResponse)` method, the `LoggingFilter` performs post-processing on the request and response object before sending the response back to the client.

Filter, FilterChain, FilterConfig classes for servlet filtering

The following interfaces are defined as part of the `javax.servlet` package:

- `Filter` interface - methods: `doFilter()`, `getFilterConfig()`, `setFilterConfig()`
- `FilterChain` interface - methods: `doFilter()`
- `FilterConfig` interface - methods: `getFilterName()`, `getInitParameter()`, `getInitParameterNames()`, `getServletContext()`

The following classes are defined as part of the `javax.servlet.http` package:

- `HttpServletRequestWrapper` - methods: See the Servlet 2.3 Specification
- `HttpServletResponseWrapper` - methods: See the Servlet 2.3 Specification

Example: `com.ibm.websphere.LoggingFilter.java`

The following example shows how to implement a filter:

```
package com.ibm.websphere;

import java.io.*;
import javax.servlet.*;

public class LoggingFilter implements Filter
{
    File _loggingFile = null;
```

```

// implement the required init method
public void init(FilterConfig fc)
{
    // create the logging file
    xxx;
}

// implement the required doFilter method...this is where
// most of the work is done
public void doFilter(ServletRequest request, ServletResponse response,
FilterChain chain)
{
    try
    {
        // add request info to the log file
        synchronized(_loggingFile)
        {
            xxx;
        }

        // pass the request on to the next resource in the chain
        chain.doFilter(request, response);
    }
    catch (Throwable t)
    {
        // handle problem...
    }
}

// implement the required destroy method
public void destroy()
{
    // make sure logging file is closed
    _loggingFile.close();
}
}

```

Configuring page list servlet client configurations

You can define PageListServlet configuration information in the IBM Web Extensions file. The IBM Web Extensions file is created and stored in the Web Applications archive (WAR) file by the IBM WebSphere Application Assembly Tool (AAT)

To configure and implement page lists:

Steps for this task

1. Use the *PageList Extensions* tab in the Application Assembly Tool (AAT) to configure page list information.
2. Add the `callPage()` method to your servlet to invoke a JavaServer Page (JSP) file in response to a client request.

The PageListServlet has a `callPage()` method that invokes a JSP file in response to the HTTP request for a page in a page list. The `callPage()` method can be invoked in one of the following ways:

- `callPage(String pageName, HttpServletRequest request, HttpServletResponse response)`

where the method arguments are:

- `pageName` - a page name defined in the PageListServlet configuration
- `request` - the `HttpServletRequest` object
- `response` - the `HttpServletResponse` object

- `callPage(String mlName, String pageName, HttpServletRequest request, HttpServletResponse response)`

where the method arguments are:

- `mlName` - a markup language type
 - `pageName` - a page name defined in the `PageListServlet` configuration
 - `request` - the `HttpServletRequest` object
 - `response` - the `HttpServletResponse` object
3. Use the `PageListServlet` client type detection support to determine the markup language type a calling client requires for the response.

Page lists

Page lists allow you to avoid hardcoding URLs in servlets and JSP files. A page list specifies the location where a request is to be forwarded, but automatically tailors that location depending on the MIME type of the servlet. These properties allow you to specify a markup language and an associated MIME type. For the given MIME type, you also specify a set of pages to invoke.

WebSphere Application Server supplies the `PageListServlet`, which you can use to call a JavaServer Pages (JSP) file by name based on the configuration data in the `client_types.xml` file. This file maps a JSP file to a Uniform Resource Identifier (URI). When the URI is invoked, it specifies another JSP file in a Web module. This support allows you to access multiple Uniform Resource Locators (URLs) without hard-coding them in your servlets.

You can also logically group page lists according to the markup language type, as for example, Hypertext Markup Language (HTML) or Wireless Markup Language (WML). This allows applications, using servlets that extend the `PageListServlet`, to call JSP files that return the proper markup-language type for the client request. For example, if a request originates from a PDA device that requires WML data and is sent to a servlet that extends the `PageListServlet`, the servlet can call a JSP file that returns a WML response.

Client type detection support

In addition to providing the page list mapping capability, the `PageListServlet` also provides *Client Type Detection* support. A servlet determines the markup language type that a calling client needs in the response, using the configuration information in the `client_types.xml` file.

Client type detection support allows a servlet, extending the `PageListServlet`, to call an appropriate JavaServer Pages (JSP) file. The servlet invokes the `callPage()` method, which calls a JSP file based on the markup-language type of the request.

`client_types.xml`

The `client_types.xml` file provides client type detection support for servlets extending `PageListServlet`. Using the configuration data in the `client_types.xml` file, servlets can determine the language type that calling clients require for the response.

The client type detection support allows servlets to call appropriate JavaServer Pages (JSP) files with the `callPage()` method. Servlets select JSP files based on the markup-language type of the request.

Servlets must use the following version of the `callPage()` method to determine the markup language type required by the client:

```
callPage(String mlName, String pageName, HttpServletRequest request, HttpServletResponse response)
```

where the arguments are:

- mlName - a markup language type
- pageName - a page name defined in the PageListServlet configuration
- request - the HttpServletRequest object
- response - the HttpServletResponse object

Review the Extending PageListServlet code example to see how the callPage() method is invoked by a servlet.

In the example, the client type detection method, getMLTypeFromRequest(HttpServletRequest request), provided by the PageListServlet, inspects the HttpServletRequest object request headers, and searches for a match in the client_types.xml file.

The client type detection method does the following:

- Uses the input HttpServletRequest and the client_types.xml file, to check for a matching HTTP request name and value.
- Returns the markup-language value configured for the <client-type> element, if a match is found.
If multiple matches are found, this method returns the markup-language for the first <client-type> element for which a match is found.
- If no match is found, returns the value of the markup-language for the default page defined in the PageListServlet configuration.

Location

The client_types.xml file is located in the <install_root>/properties directory.

Usage notes

- Is this file read-only?
No
- Is this file updated by a product component?
No
- If so, what triggers its update?
This file is created and updated manually by users.
- How and when are the contents of this file used?
Servlets, extending PageListServlet, use this file to determine the language type that calling clients require for the response.

Sample file entry

```
<?xml version="1.0" >
<!DOCTYPE clients [
<!ELEMENT client-type (description, markup-language, request-header+)>
<!ELEMENT description (#PCDATA)>
<!ELEMENT markup-language (#PCDATA)>
<!ELEMENT request-header (name, value)>
<!ELEMENT clients (client-type+)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT value (#PCDATA)>]>
<clients>
  <client-type>
    <description>IBM Speech Client</description>
    <markup-language>VXML</markup-language>
    <request-header>
```



```

        <name>user-agent</name>
        <value>IBM VoiceXML pre-release version 000303</value>
</request-header>
<request-header>
    <name>accept</name>
    <value>text/vxml</value>
</request-header>
</client-type>
<client-type>
    <description>WML Browser</description>
    <markup-language>WML</markup-language>
<request-header>
    <name>accept</name>
    <value>text/x-wap.wml</value>
</request-header>
<request-header>
    <name>accept</name>
    <value>text/vnd.wap.xml</value>
</request-header>
</client-type>
</clients>

```

Example: Extending PageListServlet

The following example shows how a servlet extends the PageListServlet class and determines the markup-language type required by the client. The servlet then uses the callPage() method to call an appropriate JavaServer Pages (JSP) file. In this example, the JSP file that provides the the correct markup-language for the response is *Hello.page*.

```

public class HelloPervasiveServlet extends PageListServlet implements Serializable
{
    /*
    * doGet -- Process incoming HTTP GET requests
    */
    public void doGet(HttpServletRequest request, HttpServletResponse response)
    throws IOException, ServletException
    {
        // This is the name of the page to be called:
        String pageName = "Hello.page";

        // First check if the servlet was invoked with a
        // queryString that contains
        // a markup-language value.
        // For example, if this is how the servlet is invoked:
        // http://localhost/servlets/HelloPervasive?mlname=VXML
        // then use the following method:
        String mlname= getMLNameFromRequest(request);

        // If no markup language type is provided in the
        //queryString, then try to determine
        // the client type from the request, and use the markup-language
        // name configured in the client_types.xml file.
        if (mlName == null)
        {
            mlName = getMLTypeFromRequest(request);
        }
        try
        {
            // Serve the request page.
            callPage(mlName, pageName, request, response);
        }
        catch (Exception e)
        {
            handleError(mlName, request, response, e);
        }
    }
}

```

autoRequestEncoding and autoResponseEncoding

Two new WebSphere Application Server extensions are available in Version 5, `autoRequestEncoding` and `autoResponseEncoding`.

In WebSphere Application Server Version 5, the Web container no longer automatically sets request and response encodings, and response content types. Programmers are expected to set these values using available methods in the Servlet 2.3 Specification. If programmers choose not to use the character encoding methods, they can specify the `autoRequestEncoding` and `autoResponseEncoding` extensions, which enable the application server to set the encoding values and content type.

The values of the `autoRequestEncoding` and `autoResponseEncoding` extensions are either `true` or `false`. The default value for both extensions is `false`. If the value is `false` for both `autoRequestEncoding` and `autoResponseEncoding`, then the request and response character encoding is set to the Servlet 2.3 Specification default, which is ISO—8859-1. Also, If the value is set to `false` for a response, the Web container cannot set a response content type.

Use the Application Assembly Tool (AAT) to change the default values for the `autoRequestEncoding` and `autoResponseEncoding` extensions.

Review the `autoRequestEncoding` and `autoResponseEncoding` encoding examples for a description of Web container behavior when these values are set to `true`.

autoRequestEncoding and autoResponseEncoding encoding examples

The default value of the `autoRequestEncoding` and `autoResponseEncoding` extensions is `false`, which means that both the request and response character encoding is set to the Servlet 2.3 Specification default of ISO-8859-1. Different character encodings are possible if the client defines character encoding in the request header, or if the code includes the `setCharacterEncoding(String encoding)` method. Also, If the value is set to `false` for a response, the Web container cannot set a response content type.

If the `autoRequestEncoding` value is set to `true`, and the client did not specify character encoding in the request header, and the code does not include the `setCharacterEncoding(String encoding)` method, the Web container tries to determine the correct character encoding for the request parameters and data.

The Web container performs each step in the following list until a match is found:

- Looks at the character set (`charset`) in the *Content-Type* header.
- Attempts to map the servers locale to a character set using defined properties.
- Attempts to use the `DEFAULT_CLIENT_ENCODING` system property, if one is set.
- Uses the ISO-8859-1 character encoding as the default.

If the `autoResponseEncoding` value is set to `true`, and the client did not specify character encoding in the request header, and the code does not include the `setCharacterEncoding(String encoding)` method, the Web container does the following:

- Attempts to determine the response content type and character encoding from information in the request header.
- Uses the ISO-8859-1 character encoding as the default.

JavaServer Pages files

JavaServer Pages (JSP) files are application building blocks coded to the Sun Microsystems JavaServer Pages (JSP) Specification.

JSP files enable the separation of the Hypertext Markup Language (HTML) code from the business logic in Web pages so that HTML programmers and Java programmers can more easily collaborate in creating and maintaining pages.

The IBM extensions to the JSP Specification include JSP tags that resemble HTML tags making it easy for HTML authors to add the power of Java technology to Web pages, without being experts in Java programming.

JSP files support a division of roles:

HTML authors

Develop JSP files that access databases and reusable Java components, such as servlets and beans.

Java programmers

Create the reusable Java components and provide the HTML authors with the component names and attributes.

Database administrators

Provide the HTML authors with the name of the database access and table information.

Developing JavaServer Pages files with WebSphere extensions

Several IBM WebSphere extensions are provided for enhancing your JavaServer Pages (JSP) files. This task provides a summary of the extensions that you can utilize.

Steps for this task

1. Review the supported specifications.

Create Java components, referring to the JSP specifications from Sun Microsystems.

See Resources for learning for links to coding specifications and examples.

WebSphere Application Server Version 3.5 added IBM extensions to the base Application Programming Interfaces (APIs). Since the JavaServer Pages (JSP) 1.1 and JSP 1.2 Specifications are backward compatible to the JSP 1.0 Specifications, you can invoke the APIs with the IBM extensions without modification.

The extensions belong to these categories:

Syntax for variable data

Put variable fields in JSP files and have servlets and beans dynamically replace the variables with values from a database when the JSP output is returned to the browser.

Syntax for database access

Add a database connection to a Web page and then use that connection to query or update the database. You can provide the user ID and password for the database connection at request time, or you can hard code the user ID and password within the JSP file.

2. Use your favorite integrated development environment (IDE), or a text editor, to develop or migrate code artifacts that meet the specifications.
3. Test the code artifacts.

4. (Optional) Batch compile your JSP files if necessary.

Tag libraries

Java ServerPages (JSP) tag libraries contain classes for common tasks such as processing forms and accessing databases from JSP files.

Tag libraries encapsulate, as simple tags, core functionality common to many Web applications. The Java Standard Tag Library (JSTL) supports common programming tasks such as iteration and conditional processing, and provides tags for:

- manipulating XML documents
- supporting internationalization
- using Structured Query Language (SQL)

Tag libraries also introduce the concept of an expression language to simplify page development, and include a version of the JSP expression language.

A tag library has two parts - a Tag Library Descriptor (TLD) file and a JAR file.

tsx:dbconnect tag JavaServer Pages syntax

Use the `<tsx:dbconnect>` tag to specify information needed to make a connection to a Java Database Connectivity (JDBC) or an Open Database Connectivity (ODBC) database.

The `<tsx:dbconnect>` syntax does not establish the connection. Use the `<tsx:dbquery>` and `<tsx:dbmodify>` syntax instead to reference a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file to establish the connection.

When the JSP file compiles into a servlet, the Java processor adds the Java coding for the `<tsx:dbconnect>` syntax to the servlet service() method, which means a new database connection is created for each request for the JSP file.

This section describes the syntax of the `<tsx:dbconnect>` tag.

```
<tsx:dbconnect id="connection_id"
  userid="db_user" passwd="user_password"
  url="jdbc:subprotocol:database"
  driver="database_driver_name"
  jndiname="JNDI_context/logical_name">
</tsx:dbconnect>
```

where:

- **id**

Represents a required identifier. The scope is the JSP file. This identifier is referenced by the connection attribute of a `<tsx:dbquery>` tag.

- **userid**

Represents an optional attribute that specifies a valid user ID for the database that you want to access. Specify this attribute to add the attribute and its value to the request object.

Although the `userid` attribute is optional, you must provide the user ID. See `<tsx:userid>` and `<tsx:passwd>` for an alternative to hard coding this information in the JSP file.

- **passwd**

Represents an optional attribute that specifies the user password for the `userid` attribute. (This attribute is not optional if the `userid` attribute is specified.) If you specify this attribute, the attribute and its value are added to the request object. Although the `passwd` attribute is optional, you must provide the password. See `<tsx:userid>` and `<tsx:passwd>` for an alternative to hard coding this attribute in the JSP file.

- **url and driver**

Represents a required attribute if you want to establish a database connection. You must provide the URL and driver.

The application server supports connection to JDBC databases and ODBC databases.

- For a JDBC database, the URL consists of the following colon-separated elements: `jdbc`, the subprotocol name, and the name of the database to access. An example for a connection to the Sample database included with IBM DB2 is:

```
url="jdbc:db2:sample"  
driver="COM.ibm.db2.jdbc.app.DB2Driver"
```

- For an ODBC database, use the Sun JDBC-to-ODBC bridge driver included in their Java2 Software Developers Kit (SDK) or another vendor's ODBC driver. The `url` attribute specifies the location of the database. The `driver` attribute specifies the name of the driver to use in establishing the database connection.

If the database is an ODBC database, you can use an ODBC driver or the Sun JDBC-to-ODBC bridge. If you want to use an ODBC driver, refer to the driver documentation for instructions on specifying the database location with the `url` attribute and the driver name.

If you use the bridge, the `url` syntax is `jdbc:odbc:<database>`. An example follows:

```
url="jdbc:odbc:autos"  
driver="sun.jdbc.odbc.JdbcOdbcDriver"
```

Note: To enable the application server to access the ODBC database, use the ODBC Data Source Administrator to add the ODBC data source to the System DSN configuration. To access the ODBC Administrator, click the ODBC icon on the Windows NT Control Panel.

- **jndiname**

Represents an optional attribute that identifies a valid context in the application server Java Naming and Directory Interface (JNDI) naming context and the logical name of the data source in that context. The Web administrator configures the context using an administrative client such as the WebSphere Administrative Console.

If you specify the `jndiname` attribute, the JSP processor ignores the `driver` and `url` attributes on the `<tsx:dbconnect>` tag.

An empty element (such as `<url></url>`) is valid.

dbquery tag JavaServer Pages syntax

Use the `<tsx:dbquery>` tag to establish a connection to a database, submit database queries, and return the results set.

The `<tsx:dbquery>` tag does the following:

1. References a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file and uses the information the tag provides to determine the database URL and driver. You can also obtain the user ID and password from the `<tsx:dbconnect>` tag if those values are provided in the `<tsx:dbconnect>` tag.
2. Establishes a new connection
3. Retrieves and caches data in the results object.
4. Closes the connection and releases the connection resource.

This section describes the syntax of the `<tsx:dbquery>` tag.

```
<%-- SELECT commands and (optional) JSP syntax can be placed
  within the tsx:dbquery. --%>
<%-- Any other syntax, including HTML comments, are not valid. --%>
<tsx:dbquery id="query_id"
connection="connection_id"
limit="value" >
</tsx:dbquery>
```

where:

- **id**

Represents the identifier of this query. The scope is the JSP file. Use `id` to reference the query. For example, from the `<tsx:getProperty>` tag, use `id` to display the query results.

The `id` becomes the name of a bean that contains the results set. The bean properties are dynamic and the property names are the names of the columns in the results set. If you want different column names, use the SQL keyword for specifying an alias on the `SELECT` command. In the following example, the database table contains columns named `FNAME` and `LNAME`, but the `SELECT` statement uses the `AS` keyword to map those column names to `FirstName` and `LastName` in the results set:

```
Select FNAME, LNAME AS FirstName, LastName from Employee where FNAME='Jim'
```

- **connection**

Represents the identifier of a `<tsx:dbconnect>` tag in this JSP file. The `<tsx:dbconnect>` tag provides the database URL, driver name, and optionally, the user ID and password for the connection.

- **limit**

Represents an optional attribute that constrains the maximum number of records returned by a query. If this attribute is not specified, no limit is used. In such a case, the effective limit is determined by the number of records and the system caching capability.

- **SELECT command and JSP syntax**

Represents the only valid SQL command, `SELECT`. The `<tsx:dbquery>` tag must return a results set. Refer to your database documentation for information about the `SELECT` command. See other articles in this section for a description of JSP syntax for variable data and inline Java code.

dbmodify tag JavaServer Pages syntax

The `<tsx:dbmodify>` tag establishes a connection to a database and then adds records to a database table.

The `<tsx:dbmodify>` tag does the following:

1. References a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file and uses the information provided by that tag to determine the database URL and driver.

Note: You can also obtain the user ID and password from the <tsx:dbconnect> tag if those values are provided in the <tsx:dbconnect> tag.

2. Establishes a new connection.
3. Updates a table in the database.
4. Closes the connection and releases the connection resource.

This section describes the syntax of the <tsx:dbmodify> tag.

```
<%-- Any valid database update commands can be placed within the DBMODIFY tag. -->
<%-- Any other syntax, including HTML comments, are not valid. -->
<tsx:dbmodify connection="connection_id">
</tsx:dbmodify>
```

where:

- **connection**
Represents the identifier of a <DBCONNECT> tag in this JSP file. The <DBCONNECT> tag provides the database URL, driver name, and (optionally) the user ID and password for the connection.
- Database commands
Represents valid database commands. Refer to your database documentation for details

tsx:getProperty tag JavaServer Pages syntax and examples

The <tsx:getProperty> tag gets the value of a bean to display in a JavaServer Pages (JSP) file.

This IBM extension of the Sun JSP <jsp:getProperty> tag implements all of the <jsp:getProperty> function and adds the ability to introspect a database bean created using the IBM extension <tsx:dbquery> or <tsx:dbmodify>.

Note: You cannot assign the value from this tag to a variable. The value, generated as output from this tag, displays in the browser window.

This section describes the syntax of the <tsx:getProperty> tag:

```
<tsx:getProperty name="bean_name"
  property="property_name" />
```

where:

- **name**
Represents the name of the bean declared by the id attribute of a <tsx:dbquery> syntax within the JSP file. See <tsx:dbquery> for an explanation. The value of this attribute is case-sensitive.
- **property**
Represents the property of the bean to access for substitution. The value of the attribute is case-sensitive and is the locale-independent name of the property.

Tag example:

```
<tsx:getProperty name="userProfile" property="username" />
<tsx:getProperty name="request" property=request.getParameter("corporation") />
```

In most cases, the value of the property attribute is just the property name. However, to access the request bean or to access a property of a property (sub property), specify the full form of the property attribute. The full form also gives you the option to specify an index for indexed properties. You can specify the

optional index as a constant (such as 2), or an index like the one described in the `<tsx:repeat>` tag. Some examples using the full form of the property attribute follow:

```
<tsx:getProperty name="staffQuery" property=address(currentAddressIndex) />
<tsx:getProperty name="shoppingCart" property=items(4).price />
<tsx:getProperty name="fooBean" property=foo(2).bat(3).boo.far />
```

tsx:userid and tsx:passwd tag JavaServer Pages syntax

With the `<tsx:userid>` and `<tsx:passwd>` tags, you do not have to hard code a user ID and password in the `<tsx:dbconnect>` tag.

Use the `<tsx:userid>` and `<tsx:passwd>` tags to accept user input for the values and then add that data to the request object. You can access the request object with a JavaServer Pages (JSP) file, such as the *JSPEmployee.jsp* example that requests the database connection.

You must use `<tsx:userid>` and `<tsx:passwd>` tags within a `<tsx:dbconnect>` tag.

This section describes the syntax of the `<tsx:userid>` and `<tsx:passwd>` tags.

```
<tsx:dbconnect id="connection_id"
  <font color="red"><userid></font>
  <tsx:getProperty name="request"
    property=request.getParameter("userid") />
  <font color="red"></userid></font>
  <font color="red"><passwd></font>
  <tsx:getProperty name="request"
    property=request.getParameter("passwd") />
  <font color="red"></passwd></font>
  url="protocol:database_name:database_table"
  driver="JDBC_driver_name">
</tsx:dbconnect>
```

where:

- **<tsx:getProperty>**

Represents the syntax as a mechanism for embedding variable data.

- **userid**

Represents a reference to the request parameter that contains the user ID. You must add the parameter to the request object that passes to this JSP file. You can set the attribute and its value in the request object, using an HTML form or a URL query string to pass the user-specified request parameters.

- **passwd**

Represents a reference to the request parameter that contains the password. Add the parameter to the request object that passes to this JSP file. You can set the attribute and its value in the request object, using an HTML form or a URL query string, to pass user-specified values.

tsx:repeat tag JavaServer Pages syntax

The `<tsx:getProperty>` tag repeats a block of HTML tagging.

Use the `<tsx:repeat>` syntax to iterate over a database query results set. The `<tsx:repeat>` syntax iterates from the start value to the end value until one of the following conditions is met:

- The end value is reached.
- An exception is thrown.

The output of a `<tsx:repeat>` block is buffered until the block completes. If an exception is thrown before a block completes, no output is written for that block.

This section describes the syntax of the `<tsx:repeat>` tag:

```
<tsx:repeat index=name
  start="starting_index"
  end="ending_index">
</tsx:repeat>
```

where:

- **index**
Represents an optional name used to identify the index of this repeat block. The value is case-sensitive and its scope is the JSP file.
- **start**
Represents an optional starting index value for this repeat block. The default is 0.
- **end**
Represents an optional ending index value for this repeat block. The maximum value is 2,147,483,647.
If the value of the end attribute is less than the value of the start attribute, the end attribute is ignored.

Example: Combining `tsx:repeat` and `tsx:getProperty` JavaServer Pages tags

The following code snippet shows you how to code these tags:

```
<tsx:repeat>
<tr>
  <td><tsx:getProperty name="empqs" property="EMPNO" />
  <tsx:getProperty name="empqs" property="FIRSTNAME" />
  <tsx:getProperty name="empqs" property="WORKDEPT" />
  <tsx:getProperty name="empqs" property="EDLEVEL" />
</td>
</tr>
</tsx:repeat>
```

Example: `tsx:dbmodify` tag syntax

In the following example, a new employee record is added to a database. The values of the fields are based on user input from this JavaServer Pages (JSP) file and referenced in the database commands using the `<tsx:getProperty>` tag.

```
<tsx:dbmodify connection="conn" >
insert into EMPLOYEE
(EMPNO, FIRSTNAME, MIDINIT, LASTNAME, WORKDEPT, EDLEVEL) values
('<tsx:getProperty name="request" property=request.getParameter("EMPNO") />',
'<tsx:getProperty name="request" property=request.getParameter("FIRSTNAME") />',
'<tsx:getProperty name="request" property=request.getParameter("MIDINIT") />',
'<tsx:getProperty name="request" property=request.getParameter("LASTNAME") />',
'<tsx:getProperty name="request" property=request.getParameter("WORKDEPT") />',
<tsx:getProperty name="request" property=request.getParameter("EDLEVEL") />)
</tsx:dbmodify>
```

Example: Using `tsx:repeat` JavaServer Pages tag to iterate over a results set

The `<tsx:repeat>` tag iterates over a results set. The results set is contained within a bean. The bean can be a static bean, for example, a bean created by using the IBM WebSphere Studio database wizard, or a dynamically generated bean, for example,

a bean generated by the `<tsx:dbquery>` syntax. The following table is a graphic representation of the contents of a bean called, *myBean*:

	col1	col2	col3
row0	friends	Romans	countrymen
row1	bacon	lettuce	tomato
row2	May	June	July

Some observations about the bean:

- The column names in the database table become the property names of the bean. The `<tsx:dbquery>` section describes a technique for mapping the column names to different property names.
- The bean properties are indexed. For example, `myBean.get(Col1(row2))` returns May.
- The query results are in the rows. The `<tsx:repeat>` tag iterates over the rows, beginning at the start row.

The following table compares using the `<tsx:repeat>` tag to iterate over a static bean, versus a dynamically generated bean:

Static Bean Example	<code><tsx:repeat></code> Bean Example
<p>myBean.class</p> <pre>// Code to get a connection // Code to get the data Select * from myTable; // Code to close the connection</pre> <p>JSP file</p> <pre><tsx:repeat index=abc> <tsx:getProperty name="myBean" property="col1(abc)" /> </tsx:repeat></pre> <p>Notes:</p> <ul style="list-style-type: none"> • The bean (<code>myBean.class</code>) is a static bean. • The method to access the bean properties is <code>myBean.get(property (index))</code>. • You can omit the property index, in which case the index of the enclosing <code><tsx:repeat></code> tag is used. You can also omit the index on the <code><tsx:repeat></code> tag. • The <code><tsx:repeat></code> tag iterates over the bean properties row by row, beginning with the start row. 	<p>JSP file</p> <pre><tsx:dbconnect id="conn" userid="alice"passwd="test" url="jdbc:db2:sample" driver="COM.ibm.db2.jdbc.app.DB2Driver"> </tsx:dbconnect > <tsx:dbquery id="dynamic" connection="conn" > Select * from myTable; </tsx:dbquery> <tsx:repeat index=abc> <tsx:getProperty name="dynamic" property="col1(abc)" /> </tsx:repeat></pre> <p>Notes:</p> <ul style="list-style-type: none"> • The bean (dynamic) is generated by the <code><tsx:dbquery></code> tag and does not exist until the syntax executes. • The method to access the bean properties is <code>dynamic.getValue("property", index)</code>. • You can omit the property index, in which case the index of the enclosing <code><tsx:repeat></code> tag is used. You can also omit the index on the <code><tsx:repeat></code> tag. • The <code><tsx:repeat></code> tag syntax iterates over the bean properties row by row, beginning with the start row.

Implicit and explicit indexing

Examples 1, 2, and 3 show how to use the `<tsx:repeat>` tag. The examples produce the same output if all indexed properties have 300 or fewer elements. If there are more than 300 elements, Examples 1 and 2 display all elements, while Example 3 shows only the first 300 elements.

Example 1 shows *implicit indexing* with the default start and default end index. The bean with the smallest number of indexed properties restricts the number of times the loop repeats.

```
<table>
<tsx:repeat>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property="city" /></tr></td>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property="address" /></tr></td>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property="telephone" /></tr></td>
</tsx:repeat>
</table>
```

Example 2 shows indexing, starting index, and ending index:

```
<table>
<tsx:repeat index=myIndex start=0 end=2147483647>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property=city(myIndex) /></tr></td>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property=address(myIndex) /></tr></td>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property=telephone(myIndex) /></tr></td>
</tsx:repeat>
</table>
```

Example 3 shows *explicit indexing* and ending index with implicit starting index. Although the index attribute is specified, you can still implicitly index the indexed property city because the (myIndex) tag is not required.

```
<table>
<tsx:repeat index=myIndex end=299>
<tr><td>
<tsx:getProperty name="serviceLocationsQuery" property="city" />
</tr></td>
<tr><td>
<tsx:getProperty name="serviceLocationsQuery"
property="address(myIndex)" />
</tr></td>
<tr><td><tsx:getProperty name="serviceLocationsQuery"
property="telephone(myIndex)" />
</tr></td>
</tsx:repeat>
</table>
```

Nesting `<tsx:repeat>` blocks

You can nest `<tsx:repeat>` blocks. Each block is separately indexed. This capability is useful for interleaving properties on two beans, or properties that have subproperties. In the example, two `<tsx:repeat>` blocks are nested to display the list of songs on each compact disc in the user's shopping cart.

```
<tsx:repeat index=cdindex>
<h1><tsx:getProperty name="shoppingCart" property=cds.title /></h1>
<table>
<tsx:repeat>
<tr><td><tsx:getProperty name="shoppingCart" property=cds(cdindex).playlist />
```

```

</td></tr>
</table>
</tsx:repeat>
</tsx:repeat>

```

JspBatchCompiler tool

As an IBM enhancement to JavaServer Pages support, IBM WebSphere Application Server provides a batch JSP compiler. Use this function to batch compile your JSP files and thereby enable faster responses to the initial client requests for the JSP files on your production Web server.

Batch compiling makes the first request for a JSP file much faster because the JSP file is translated and compiled into a servlet. Batch compiling is also useful as a fast way to resynchronize all of the JSP files for an application.

To use the JSP batch compiler for JSP files, enter the following command on a single line at an operating system command prompt:

```

JspBatchCompiler -enterpriseapp.name <name>
                    [ -webmodule.name <name>]
                    [ -cell.name <name>]
                    [ -node.name <name>]
                    [ -server.name <name>]
                    [ -filename <jsp name>]
                    [ -keepgenerated <true|false>]
                    [ -verbose <true|false>]
                    [ -deprecation <true|false>]

```

If the names specified for these arguments are comprised of two or more words separated by spaces, you must add quotation marks around the names.

where:

- **enterpriseapp.name**
Represents the name of the enterprise application you want to compile.
- **webmodule.name**
Represents the name of the specific Web module that you want to compile. If this argument is not set, all Web modules in the enterprise application are compiled.
- **cell.name**
Represents the name of the cell in which the application is deployed. The default is `BaseApplicationServerCell`.
- **node.name**
Represents the name of the node in which the application is deployed. The default is `DefaultNode`.
- **server.name**
Represents the name of the server in which the application is deployed. The default is `server1`.
- **filename**
Represents the name of a single JSP file that you want to compile. If this argument is not set, all files in the Web module are compiled. Alternatively, if *filename* is set to the name of a directory, only the JSP files in that directory are compiled.
- **keepgenerated**
Represents the option to save or erase the generated files.

If set to yes, WebSphere Application Server saves the generated .java files used for compilation on your server. By default, this argument is set to no and the .java files are erased after the class files have compiled.

- **verbose**

Indicates the compiler should generate verbose output while compiling the generated sources.

- **deprecation**

Indicates the compiler should generate deprecation warnings while compiling the generated sources.

Bean Scripting Framework

The Bean Scripting Framework (BSF) enables you to use scripting language functions in your Java server-side applications. This framework also extends scripting languages so that you can use existing Java classes and Java beans in the JavaScript language.

With BSF, you can write scripts that create, manipulate and access values from Java objects, or you can write Java programs that evaluate and access results from scripts.

WebSphere Application Server provides the Bean Scripting Framework, which consists of a BSF manager, a BSF engine, and a scripting engine.

BSF provides an access mechanism to Java objects for the scripting languages it supports, so that both the scripting language and the Java code can access code exclusive functions. The access mechanism is implemented through a registry of objects maintained by BSF.

BSF in WebSphere Application Server supports the Rhino ECMAScript.

The "Resources for Learning" article provides external BSF links that document future supported languages.

Example: Converting JavaScript source to the Bean Scripting Framework

JavaScript code is one of the most popular languages of Web developers. This language supports the following base objects, plus additional objects from the Document Object Model:

- array
- date
- math
- number
- string

Server-side JavaScript code supports the same base objects, and additional objects that support user access to databases, file systems and e-mail systems.

Like client-side JavaScript code, server-side JavaScript code is also platform, browser, and language independent.

You can convert server-side JavaScript applications to the Bean Scripting Framework. This article describes how to perform this conversion.

Server-side JavaScript source code

Suppose you have the following server-side JavaScript application:

```
<html>
<head>
<title>Hello World server-side JavaScript example</title>
</head>
<body>
<br><br>
</body>
</html>

<server>
function writePage()
  write("<center><font size='6'>Hello World</font></center>");
</server>
```

Converting server-side JavaScript source code to the Bean Scripting Framework (BSF)

Make the following changes to the JavaScript source code to enable BSF:

```
<%@ page language="javascript" %>
<html>
<head>
<title>Hello World server-side BSF/JavaScript example</title>
</head>
<body>
<br><br>
</body>
</html>

<%
  out.println("<center><font size='6'>Hello World</font></center>");
%>
```

Review the other BSF reference articles for deployment information and additional programming examples.

Scenario: Creating a Bean Scripting Framework application

Scenario description

Programming skills in JavaScript code are more prevalent than programming skills using JavaServer Pages (JSP) tags. Using the Bean Scripting Framework, JavaScript programmers can gradually introduce JSP tags in their JavaScript applications without completely rewriting the source code. The BSF method not only reduces the potential of programming errors, but also provides a painless way to learn a new technology.

The following scenario illustrates how to implement a BSF application using JavaScript within JSP tags.

Developing the BSF application

At ABC elementary school, John Doe teaches third grade mathematics. He wants to help his students memorize their multiplication tables, and thinks a small Web-based quiz could help meet his objective. However, John Doe only knows JavaScript.

Using the Bean Scripting Framework to help leverage his JavaScript skills, John Doe creates two JSP files, `multiplication_test.jsp` and `multiplication_scoring.jsp`.

In the `multiplication_test.jsp` file, John Doe uses both client-side and server-side JavaScript code to generate a test of 100 random multiplication questions, displayed using a three minute timer. He then writes the `multiplication_scoring.jsp` file to read the data submitted by the `multiplication_test.jsp` file and to generate the scoring results.

John Doe creates the following two files:

multiplication_test.jsp:

```
<html>
<head>
<title>Multiplication Practice Test</title>
<script language="javascript">
var countMin=3;
var countSec=0;
function updateDisplay (min, sec) {
  var disp;
  if (min <= 9) disp = " 0";
  else disp = " ";
  disp += (min + ":");
  if (sec <= 9) disp += ("0" + sec);
  else disp += sec;
  return(disp);
}
function countDown() {
  countSec--;
  if (countSec == -1) {
    countSec = 59;
    countMin--;
  }
  document.multtest.counter.value = updateDisplay(countMin, countSec);
  if((countMin == 0) &&(countSec == 0)) document.multtest.submit();
  else var down = setTimeout("countDown();", 1000);
}
</script>
</head>
<body bgcolor="#ffffff" onLoad="countDown();">
<%@ page language="javascript" %>
<h1>Three Minute Multiplication Drill</h1>
<hr>
<h2>Remember: this is an opportunity to excel!</h2>
<p>
<form method="POST" name="multtest" action="multiplication_scoring.jsp">
<div align="center">
<table>
<tr>
<td>
<h3>Time left:
<input type="text" name="counter" size="9" value="03:00" readonly>
</h3>
</td>
<td>
<input type="submit" value="Submit for scoring!">
</td>
</tr>
</table>
</div>
</body>
```

```

</tr>
</table>
<table border="1">
<%
var newrow = 0;
var q_num = 0;
function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");
    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"" + q_num + "|" + num1 + ":" + num2 + "\" ");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");
    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;
    q_num++;
}
for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);
    addQuestion(rand1, rand2);
}
%>
</table>
</div>
</form>
</body>
</html>

```

multiplication_scoring.jsp:

```

<html>
<head>
<title>Multiplication Practice Test Results</title>
</head>
<body bgcolor="#ffffff">
<%@ page language="javascript" %>
<h1>Multiplication Drill Score</h1>
<hr>
<div align="center">
<table border="1">
<tr><th>Problem</th>
<th>Correct Answer</th><th>Your Answer
</th></tr>
<%
var total_score = 0;
function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;
    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");
    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}
var equations = request.getParameterNames();

```

```

while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");
    if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
}
%>
</table>
<h2>Total Score: <%= total_score %></h2>
<h3><a href="multiplication_test.jsp">Try again?</a></h3>
</div>
</body>
</html>

```

Follow these steps to see how John Doe uses BSF to implement JavaScript in a JSP application:

Steps for this task

1. Give your files a .jsp extension.
2. Use server-side JavaScript code in your application.

The `multiplication_test.jsp` file incorporates both client-side and server-side JavaScript. Server-side JavaScript is similar to client-side JavaScript; the primary difference consists of using a different set of objects. Whereas client-side Javascript programmers invoke `document` and `window` objects, server-side JavaScript programmers, using the Bean Scripting Framework, invoke a set of objects provided by the JSP technology. Also, client-side scripts are enclosed in `<script>` tags, but server-side scripts use JSP scriptlet and expression tags.

Examine the following blocks of code:

```

<script language="javascript">
var countMin=3;
var countSec=0;
function updateDisplay (min, sec) {
    var disp;
    if (min <= 9) disp = " 0";
    else disp = " ";
    disp += (min + ":");
    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;
    return(disp);
}
function countDown() {
    countSec--;
    if (countSec == -1) {
        countSec = 59;
        countMin--;
    }
    document.multtest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) && (countSec == 0)) document.multtest.submit();
    else var down = setTimeout("countDown();", 1000);
}
</script>
....
<body bgcolor="#ffffff" onLoad="countDown();">
...
<form method="POST" name="multtest" action="multiplication_scoring.jsp">
...
<input type="text" name="counter" size="9" value="03:00" readonly>
...

```

The JavaScript code contained in the `<script>` block implements a timer set within the `<input>` field named `counter`. The `onLoad` event handler in the `<body>` tag causes the browser to load and execute the code when the page is loaded.

The `document.mulpttest.submit()` statement causes the form named `mulptest` to be submitted when the timer expires.

3. Identify the code to the BSF function.

The following code example, from the `multiplication_test.jsp` file, displays the use of a JSP directive. This directive tells the WebSphere Application Server BSF function that this file is using the JavaScript language, and that the JavaScript code is enclosed by the `<% ... %>` scriptlet tags. The `out` implicit JSP object in this code example, creates the body section of a table from 100 randomly generated questions.

```
...
<%@ page language="javascript" %>
...
<%
var newrow = 0;
var q_num = 0;

function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");

    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"" + q_num + "|" + num1 + ":" + num2 + "\" ");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");

    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;

    q_num++;
}

for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);

    addQuestion(rand1, rand2);
}

%>
...

```

4. Read the results.

To score the results of the practice drill, John Doe uses the `request` implicit JSP object in the `multiplication_scoring.jsp` file to obtain the `POST` data created within the `<form>` tags in the `multiplication_test.jsp` file.

The `multiplication_scoring.jsp` file uses the `POST` data to build an output file containing the original question, the student's answer, and the correct answer, and then prints the text in a table format using the `out` implicit object.

The following code example from the `multiplication_scoring.jsp` file illustrates the use of the `request` and `out` JSP objects:

```

...
<%@ page language="javascript" %>
...
<%
var total_score = 0;
function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;
    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");
    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}
var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");
    if (splitPos1 >=0 && splitPos2 >= 0)
        score(currElt, splitPos1, splitPos2);
}
%>
...
<h2>Total Score: <%= total_score %></h2>
...

```

Note: Although using separate scriptlet blocks of code for different portions of a conditional expression is common in JSP files implemented in Java, it is invalid for JSP files implemented using JavaScript through the Bean Scripting Framework. The JavaScript code must be entirely contained within the scriptlet tags.

The following code example illustrates invalid usage:

```

<% if (pass == 0) %>
    <i>pass is true</i>
<% else %>
    <i>pass is not true</i>

```

Deploying the BSF application

You assemble and deploy BSF applications in the same manner as JSP applications. Review the (Assembling applications) article for more information.

Deploy the BSF code examples in WebSphere Application Server to view this applications processing and output. Use the following quick steps to deploy the application.

Note: The intent of these "quick steps" is to provide you with instant application output. However, the supported method for deployment is the same as for standard JSP files.

Steps for this task

1. Use the DefaultApplication to add your BSF files.
Copy your .jsp files to the DefaultApplication directory: *<application server install directory> /installedApps/<node name>/DefaultApplication.ear/DefaultApplication.war*
2. Start the application server.
3. Open a browser and request your BSF application
Use the following URL to request your application:
`http://hostName:9080/<JSP file name>.jsp`

Example: Bean Scripting Framework code example

The following code examples show how to implement JavaScript using the Bean Scripting Framework (BSF).

For a quick demonstration of the BSF function, copy these code examples into 2 separate files, and deploy them in WebSphere Application Server using the instructions in the BSF scenario article.

Multiplication practice test

```
<html>
<head>
<title>Multiplication Practice Test</title>
<!--
This file and its companion, multiplication_score.jsp, illustrate the
use of ECMAScript within the BSF framework. The task is a simple
timed math quiz, which is 3 minutes in duration. When the quiz ends,
the score is computed and displayed.
Users are then asked if they wish to try
the quiz again.
-->

<!--
This code fragment displays and updates the quiz
countdown in client side JavaScript code.
-->
<script language="javascript">
var countMin=3;
var countSec=0;

// This code computes the current countdown time.
function updateDisplay (min, sec) {
    var disp;

    if (min <= 9) disp = " 0";
    else disp = " ";

    disp += (min + ":");

    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;

    return(disp);
}

//This code fragment displays the current countdown time in the
user's browser window,
//and submits the results for scoring when the countdown ends.

function countDown() {
    countSec--;
    if (countSec == -1) {
```

```

        countSec = 59;
        countMin--;
    }
    document.milttest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) && (countSec == 0)) document.milttest.submit();
    else var down = setTimeout("countDown();", 1000);
}
</script>
</head>
<body bgcolor="#ffffff" onLoad="countDown();">

<!--
The body of the quiz runs as JavaServer Pages (JSP) code using BSF.
The code outputs the problems in table format using the POST method
and invokes the scoring module when the user chooses to
end the quiz or when the countdown ends.
-->
<%@ page language="javascript" %>

<h1>Three Minute Multiplication Drill</h1>
<hr>

<h2>Remember: this is an opportunity to excel!</h2>
<p>

<form method="POST" name="milttest" action="multiplication_scoring.jsp">
<div align="center">
<table>
<tr>
<td>
<h3>Time left:
<input type="text" name="counter" size="9" value="03:00" readonly>
</h3>
</td>
<td>
<input type="submit" value="Submit for scoring!">
</td>
</tr>
</table>
<table border="1">
<%
var newrow = 0;
var q_num = 0;

// This code generates a new random multiplication problem up to
// the number twelve, and
// enters it into the table of problems.

function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");

    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"" + q_num + "\"" + num1 + ":" + num2 + "\" ");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");

    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;

    q_num++;
}
}

```



```

//This code obtains two random operands and formats 100 quiz problems.

for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);

    addQuestion(rand1, rand2);
}

%>
</table>
</div>
</form>

</body>
</html>

```

Multiplication practice test results

```

<html>
<head>
<title>Multiplication Practice Test Results</title>
</head>
<body bgcolor="#ffffff">

<!--
This JSP code is invoked when the user submits a math quiz for scoring,
or when the quiz countdown expires. The JSP code tabulates the problem
list, the correct answer,
the user's answer, and scores the test. It then offers the user an
opportunity to try the quiz again.
-->
<%@ page language="javascript" %>

<h1>Multiplication Drill Score</h1>
<hr>

<div align="center">
<table border="1">
<tr><th>Problem</th><th>Correct Answer</th><th>Your Answer</th></tr>
<%
var total_score = 0;

// This code parses the submitted form, extracts the a problem generated by the
// multiplication_test.jsp file, outputs it, computes the correct answer,
// and displays this information and the user answer. The code scores
// the quiz using a running sum of correct answers.

function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;

    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");

    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
}

```

```

        out.println(your_product + "</td>");
        out.println("</tr>");
    }

    // This is the main body of the scoring application. It parses the posted quiz,
    // and calls the score() function to score remaining problems.

    var equations = request.getParameterNames();
    while(equations.hasMoreElements()) {
        var currElt = equations.nextElement();
        var splitPos1 = currElt.indexOf("|");
        var splitPos2 = currElt.indexOf(":");

        if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
    }

    %>
</table>

<h2>Total Score: <%= total_score %></h2>
<h3><a href="/multiplication_test.jsp">Try again?</a></h3>
</div>

</body>
</html>

```

Developing Web applications

Before you begin

Design a Web application and the components that it needs.

For general Web application design information, see "Resources for learning."

There are two basic approaches to selecting tools for developing Web applications:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the servlet and JavaServer Pages (JSP) code, and Hypertext Markup Language (HTML) files. They also contain integrated tools for packaging and testing the Web application components. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you decide to develop Web components without an IDE, you need at least an ASCII text editor. You can also use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the Web application components.

The following steps support the second approach, development without an IDE.

Steps for this task

1. If necessary, migrate any pre-existing code to the required version of the servlet and JSP specification.
2. Write and compile the components of the Web application.

What to do next

Assemble the application components in one or more Web modules.

Web modules

A Web module represents a Web application. A Web module is created by assembling servlets, JavaServer Pages (JSP) files, and static content such as HyperText Markup Language (HTML) pages into a single deployable unit. Web modules are stored in Web archive (WAR) files, which are standard Java archive files.

A Web module contains:

- One or more servlets, JSP files, and HTML files.
- A deployment descriptor, stored in an Extensible Markup Language (XML) file. The file, named `web.xml`, declares the contents of the module. It contains information about the structure and external dependencies of Web components in the module and describes how the components are used at run time.

You can create Web modules as standalone applications, or you can combine Web modules with other modules to create J2EE applications. You install and run a Web module in the Web container of an application server.

Assembling Web Modules

Before you begin

If you want to use existing J2EE 1.2 Web modules in your J2EE 1.3 application, migrate them to J2EE 1.3 first.

Assemble a Web module to contain servlets, JSP files, and related code artifacts. (Group enterprise beans, client code, and resource adapter code in separate modules).

A Web module can be installed as a standalone application or can be combined with other modules into an enterprise application.

The Application Assembly Tool (AAT) provides flexibility in assembling Web modules. Options described below include:

- Importing an existing Web module (WAR file)
- Creating a new Web module
- Copying code artifacts (such as servlets) from one Web module into a new Web module

Although you can input various properties for Web archives, available properties are specific to the Servlet, JSP, and J2EE specification level.

Steps for this task

1. ("Starting the Application Assembly Tool (AAT)")
2. Select **File>New>Web Module**.

The navigation tree now displays various sets of properties for configuring the new Web module.

3. (**Optional**) Use the property dialog shown in the AAT workspace to change the default file name and location.
 - a. It is recommended that you change the display name so that it differs from the file name.

- b. If you like, change the temporary location of the Web module from the default location, `install_root/bin`.
4. Add at least one Web component (servlet or JSP file) to the module. You must add at least one Web component, using one of the following methods.
 - Import an existing WAR file containing Web components.
 - a. In the navigation tree, right-click the **Web Components** folder.
 - b. Select **Import** from its right-click menu.
 - c. Use the file browser to locate and select the archive file for the module.
 - d. Click **Open**. The Web applications in the selected archive are displayed.
 - e. Select a Web application. Its Web components are displayed in the workspace.
 - f. Select the servlets or JSP files to be added and click **Add**. The components are displayed in the **Selected Components** window.
 - g. Click **OK**. The properties associated with the archive are also imported. The property dialog boxes in the workspace are populated automatically with values.
 - h. Double-click the Web Components icon to verify that the servlets or JSP files are included in the module.
 - i. Double-click the **Web Components** icon to verify that the servlets or JSP files are included in the module.
 - j. ("Saving applications after assembly")
 - Copy and paste archive files from an existing module.
 - Create a new Web component.
 - a. In the navigation tree, right-click the **Web Components** folder.
 - b. Select **New** from its right-click menu.
 - c. When the new module is displayed, enter a component name and choose a component type.
 - d. Use the file browser to locate and select the archive file for the module.
 - e. Click the plus sign (+) to verify its contents and enter assembly properties.
 - f. In the **New Web Component** property dialog box, click **OK**.
 - g. Verify that the Web component has been added to the module by double-clicking the **Web components** icon in the navigation tree.
 - h. Click the component to view its corresponding property dialog box in the bottom portion of the pane.
5. Enter assembly properties for each Web component.
 - a. Click the plus sign (+) next to the each component to reveal its property groups.
 - b. Right-click each property group's icon and click **New** to display properties in the workspace.
6. Specify additional properties for the Web module.

Right-click each property group's icon. Choose **New** to add new values, or edit existing values in the property pane. (Click **Help** for descriptions of the settings).

Note that if you add a security constraint, you must add at least one Web resource collection.
7. Add any other files needed by the application.

- In the navigation tree, click the plus sign (+) next to the Files icon. Right-click **Add Class Files**, **Add JAR Files**, or **Add Resource Files**. Select **Add Files**.
 - Add files, using the Add Files dialog.
8. ("Saving applications after assembly")

What to do next

Assemble zero or more other new modules of your choice:

- ("Assembling EJB modules")
- ("Assembling Application Client Modules")
- "Assembling Resource Adapter modules" (not in this document)

You can also ("Migrating application modules from J2EE 1.2 to J2EE 1.3 ").

Another option is to proceed directly to ("Assembling new or modifying existing modules"). While assembling an application module, you can create any new modules that you need.

Context parameters

A servlet context defines a server's view of the Web application within which the servlet is running. The context also allows a servlet to access resources available to it.

Using the context, a servlet can log events, obtain URL references to resources, and set and store attributes that other servlets in the context can use. These properties declare a Web application's parameters for its context. They convey setup information, such as a webmaster's e-mail address or the name of a system that holds critical data.

Servlet mappings

A servlet mapping is a correspondence between a client request and a servlet.

Servlet containers use URL paths to map client requests to servlets, and follow the URL path-mapping rules as specified in the Java Servlet specification. The container uses the URI from the request, minus the context path, as the path to map to a servlet. The container chooses the longest matching available context path from the list of Web applications that it hosts.

Servlet caching

Dynamic caching can be used to improve the performance of servlet and JavaServer Pages (JSP) files by serving requests from an in-memory cache. Cache entries contain the servlet's output, results of the servlet's execution, and metadata.

Web components

A web component is a servlet, Java Server Page (JSP), or HTML file. One or more web components make up a web module.

Web property extensions

Web property extensions are IBM extensions to the standard deployment descriptors for Web applications. These extensions include mime filtering and servlet caching.

Web resource collections

A Web resource collection defines a set of URL patterns (resources) and HTTP methods belonging to the resource.

HTTP methods handle HTTP-based requests, such as GET, POST, PUT, and DELETE. A URL pattern is a partial Uniform Resource Locator that acts as a template for matching the pattern with existing full URLs in an attempt to find a valid file.

Welcome files

A Welcome file is an entry point file (for example, `index.html`) for a group of related HTML files.

Welcome files are located by using a group of partial URIs. The Web container uses the partial URIs to find a valid file when the initial URI is not found.

Context parameter assembly settings

A servlet context defines the server view of the Web application within which the servlet is running. The context also allows a servlet to access resources available to it. Using the context, a servlet can log events, obtain URL references to resources, and set and store attributes that other servlets in the context can use.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance*

Parameter name (Required, String)

Specifies the name of a parameter, for example, `dataSourceName`.

Data type String

Parameter value (Required, String)

Specifies the value of a parameter, for example, `jdbc/sample`.

Data type String

Description

Contains a description of the context parameter.

Data type String

Initialization parameter assembly settings

Use this page to specify the initialization parameters that are sent to a servlet in its `HttpConfig` object when the servlet is first started.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules >
component_instance >
Web Components

Parameter name (Required, String)

Specifies the name of an initialization parameter.

Data type String

Parameter value (Required, String)

Specifies the value of the initialization parameter.

Data type String

Description

Contains text describing the use of the parameter.

Data type String

JavaServer Pages attribute assembly settings

Use the JavaServer Pages (JSP) attributes page to set JSP attributes that are used by servlets that implement JSP processing behavior.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance* > **Assembly Property Extensions**

JSP Attribute (Name)

Specifies the name of an attribute.

Data type String

JSP Attribute (Value)

Specifies the value of an attribute.

Data type String

The WebSphere JSP container supports the following JSP attributes:

classdebuginfo

Indicates the compiler should include debugging information in the generated classfile.

classdebuginfo true or false

Default is false.

classpath

Specifies an additional classpath for compiling the generated servlets.

classpath *classpath* or *null*

Default is *null*.

deprecation

Indicates the compiler should generate deprecation warnings when compiling the generated Java source.

deprecation true or false

Default is false.

ieClassID

Indicates the Java plugin COM class ID for Internet Explorer. The <jsp:plugin> tags use this value.

ieClassID *classid*

Default is clsid:8AD9C840-044E-11D1-B3E9-00805F499D93.

javaEncoding

Indicates the Java platform encoding to use to generate the JSP page servlet.

javaEncoding *encoding value*

Default is UTF-8.

jspCompilerPath

Indicates the path of the compiler to use for compiling JSP pages.

jspCompilerPath *path name or null*

Default is *null*.

keepgenerated

Indicates the Java files generated by the JSP compiler during the translation phase of the processing should be kept.

keepgenerated *true or false*

Default is false.

largefile

Specifies support for large files. When the Java code is generated, the HTML data in a JSP file is stored separately instead of being saved as constant string data in the generated servlet.

largefile *true or false*

Default is false.

mappedfile

Indicates the compiler should generate Java source that includes a print statement for every line in the JSP file. Use this option for debugging purposes only. It is not recommended for production environments because the mappedfile option generates too many out.print() statements.

mappedfile *true or false*

Default is false.

scratchdir

Specifies the directory where the generated classfiles are created.

scratchdir *directory name.*

Default is [WAS_INSTALL_ROOT]/temp.

Note: The system property `com.ibm.websphere.servlet.temp.dir` can be used to set the `scratchdir` option on a server-wide basis. This setting, if it is present, overrides the system property.

usePageTagPool

Enables or disables the reuse of custom tag handlers on an individual JavaServer Page basis.

usePageTagPool true or false

Default is false.

The note in the `useThreadTagPool` attribute description also applies to the `usePageTagPool` attribute.

useThreadTagPool

Enables or disables the reuse of custom tag handlers on a per request thread basis.

useThreadTagPool true or false

Default is false.

The note in the `useThreadTagPool` attribute description also applies to the `usePageTagPool` attribute.

Note: Enabling custom tag handler reuse might reveal problems in your tag handler code regarding the tags ability to be reused. A custom tag handler should always do two things:

1. The `release()` method of the tag handler should reset its state and release any private resources that it might have used. The JSP engine guarantees the `release()` method will be called before the tag handler is garbage collected.
2. In the `doEndTag()` method, all instance states associated with this instance must be reset.

verbose

Indicates the compiler should generate verbose output when compiling the generated Java source code.

verbose true or false

Default is false.

Multipurpose Internet Mail Extensions (MIME) filter assembly settings

Use this page to configure Multipurpose Internet Mail Extensions (MIME) filters.

Access this page by traversing the following path in the Application Assembly Tool:

Component name (Required, String)

Specifies the name of the servlet or JavaServer Pages(TM) (JSP) file. This name must be unique within the Web module.

Data type String

Display name

Specifies a short name that is intended for display by GUIs.

Data type String

Description

Contains a description of the servlet or JSP file.

Data type String

Component type

Specifies the type of Web component. Valid values are `servlet` or `JSP file`.

Data type String

Class name (Required, String)

Specifies the full path name for the servlet class.

Data type String

JSP file (Required, String)

Specifies the full path name for the JSP file.

Data type String

Load on startup

Indicates whether this servlet loads at the startup of the Web application.

The default is `false` (the check box is not selected). Also specifies a positive integer indicating the order in which to load the servlet. Lower integers are loaded before higher integers. If no value is specified, or if the value specified is not a positive integer, the container is free to load the servlet at any time in the startup sequence.

Data type String

Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type JPEG, GIF

Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type JPEG, GIF

Page list assembly settings

Page lists allow you to avoid hardcoding URLs in servlets and JSP files.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance* > **Web Components**

Name

Specifies the name of the markup language—for example, Hypertext Markup Language (HTML), Wireless Markup Language (WML), and Voice Extensible Markup Language (VXML).

Data type String

MIME Type

Specifies the Multi-Purpose Internet Mail Extensions (MIME) type of the markup language, for example, *text/html* and *text/x-vxml*.

Data type String

Error Page

Specifies the name of an error page.

Data type String

Default Page

Specifies the name of a default page.

Data type String

Pages - Name

Specifies the name of the page to serve, for example, *StockQuoteRequest.page*.

Data type String

Pages - URI

Specifies the URI of the page to serve, for example, *examples/StockQuoteHTMLRequest.jsp*.

Data type String

Security constraints

Security constraints declare how Web content is to be protected.

These properties associate security constraints with one or more Web resource collections. A constraint consists of a Web resource collection, an authorization constraint, and a user data constraint.

Security constraint assembly settings

Use the Security constraints panel to configure security constraints.

To view this AAT panel, open an existing or create a new Web module. Right-click **Security Constraints** from the left navigation menu. Click **New**.

Security constraints declare how to protect Web content. These properties associate security constraints with one or more Web resource collections. A constraint consists of a Web resource collection, an authorization constraint, and a user data constraint.

- A Web resource collection is a set of resources (URL patterns) and HTTP methods on those resources. All requests that contain a request path that matches the URL pattern described in the Web resource collection is subject to the constraint. If no HTTP methods are specified, then the security constraint applies to all HTTP methods.
- An authorization constraint is a set of roles that users must be granted in order to access the resources described by the Web resource collection. If a user who requests access to a specified URI is not granted at least one of the roles specified in the authorization constraint, the user is denied access to that resource.
- A user data constraint indicates that the transport layer of the client/server communications process must satisfy the requirement of either guaranteeing content integrity (preventing tampering in transit) or guaranteeing confidentiality (preventing reading while in transit).

If multiple security constraints are specified, the container uses the "first match wins" rule when processing a request to determine what authentication method to use, or what authorization to allow

Security constraint name

Specifies the name of the security constraint.

Data type String

Authorization Constraints - Roles

Specifies the user roles that are permitted access to this resource collection.

Data type String

Authorization Constraints - Description

Contains a description of the authorization constraints

Data type String

User Data Constraints - Transport guarantee

Indicates how data communicated between the client and the server is to be protected.

Specifies that the protection for communications between the client and server is None, Integral, or Confidential.

- None means that the application does not require any transport guarantees.
- Integral means that the application requires that the data sent between the client and the server must be sent in such a way that it cannot be changed in transit.
- Confidential means that the application requires that the data must be transmitted in a way that prevents other entities from observing the contents of the transmission.

In most cases, Integral or Confidential indicates that the use of SSL is required.

Data type String

User Data Constraints - Description

Contains a description of the user data constraints.

Data type String

Servlet mapping assembly settings

A servlet mapping is a correspondence between a client request and a servlet.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance*

URL pattern (Required, String)

Specifies the URL pattern of the mapping.

The URL pattern must conform to the Servlet specification. Use the following syntax:

- A string beginning with a slash character (/) and ending with the slash and asterisk characters (/*) represents a path mapping.
- A string beginning with the characters *. represents an extension mapping.
- All other strings are used as exact matches only.
- A string containing only the slash character (/) indicates that the servlet specified by the mapping becomes the default servlet of the application. In this case, the servlet path is the request Uniform Resource Identifier (URI) minus the context path, and the path information is null.

Data type String

Servlet (Required, String)

Specifies the name of the servlet associated with the URL pattern.

Data type String

Tag library assembly settings

Use this page to define the tag library parameters.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance*

Java ServerPages (JSP) tag libraries contain classes for common tasks such as processing forms and accessing databases from JSP files.

Tag library file name (Required, String)

Specifies a file name relative to the location of the web.xml document, identifying a tag library used in the Web application.

Data type String

Tag library location (Required, String)

Contains the location, as a resource relative to the root of the Web application, where you can find the Tag Library Definition file for the tag library.

Data type String

Welcome file assembly settings

Use this page to configure your welcome page.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance*

Welcome file (Required, String)

The Welcome file list is an ordered list of partial URLs with no trailing or leading slash characters (/).

The Web server appends each file in the order specified and checks whether a resource in the Web archive (WAR) file is mapped to that request Uniform Resource Identifier (URI). The container forwards the request to the first resource in the WAR file that matches.

Data type String

Servlet caching configuration assembly settings

Use this page to configure your cache groups.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance* > **Assembly Property Extensions**

The properties on the General tab define a cache group and govern how long an entry remains in the cache. The properties on the ID Generation tab define how cache IDs are built and the criteria used to cache or invalidate entries. The properties on the Advanced tab define external cache groups and specify custom interfaces for handling servlet caching.

Caching group name

Specifies a name for the group of servlets or JavaServer Pages (JSP) files to cache.

Priority

Defines the default priority for cached servlets. Specify as an integer. The default value is 1.

Priority is an extension of the Least Recently Used (LRU) caching algorithm. It represents the number of cycles through the LRU algorithm that an entry is guaranteed to stay in the cache. The priority represents the length of time that an entry remains in the cache before becoming eligible for removal. On each cycle of the algorithm, the priority of an entry is decremented. When the priority reaches zero, the entry is eligible for invalidation. If an entry is requested while in the cache, its priority is reset to the priority value. Regardless of the priority value and the number of requests, an entry is invalidated when its timeout occurs. Consider increasing the priority of a servlet or JSP file when it is difficult to calculate the output of the servlet or JSP file or when the servlet or JSP file is executed more often than average. Priority values should be low. Higher values do not yield

much improvement but use extra LRU cycles. Use timeout to guarantee the validity of an entry. Use priority to rank the relative importance of one entry to other entries. Giving all entries equal priority results in a standard LRU cache that increases performance significantly.

Timeout

Specifies the length of time, in seconds, that a created cache entry remains in the cache.

When this time elapses, the entry is removed from the cache. If the timeout is zero or a negative number, the entry does not time out. It is removed when the cache is full or programmatically, from within an application.

Invalidate only

Specifies that invalidations for a servlet take place, but that no servlet caching is performed.

For example, you can use this property to prevent caching of control servlets. Control servlets treat HTTP requests as commands and execute those commands. By default, this check box is not selected.

Caching group members

Specifies the names of the servlets or JSP files to cache. The URIs are determined from the servlet mappings.

Use URIs for cache ID building

Specifies whether or not to use the URI of the requested servlet to create a cache ID. By default, URIs are used.

Use specified string

Specifies a string representing a combination of request and session variables to use for creating cache IDs. This property defines request and session variables, and the cache uses the values of these variables to create IDs for the entries.

Variables - ID

Specifies the name of a request parameter, request attribute, session parameter, or cookie.

Variables - Type

Specifies the type of variable indicated in the ID field. The valid values are Request parameter, Request attribute, Session parameter, or Cookie.

Variables - Method

Specifies the name of a method in the request attribute or session parameter. The output of this method is used to generate cache entry IDs. If this value is not specified, the toString method is used by default.

Variables - Data ID

Specifies a string that, combined with the value of the variable, generates a group name for the cache entry. The cache entry is placed in this group. You can invalidate this group.

Variables - Invalidate ID

Specifies a string that is combined with the value of the variable on the request or session to form a group name. The cache invalidates the group name.

Required

Specifies whether a value must exist in the request. If this check box is selected, and either the request parameter, request attribute, session parameter, or the method is not specified, the request is not cached.

External cache groups - Group name

Specifies the name of the external cache group to which this servlet is published.

ID generator

Specifies a user-written interface for handling parameters, attributes, and sessions.

The value must represent a full package and class name of a class extending `com.ibm.websphere.servlet.cache.IdGenerator`. The properties specified in the Application Assembly Tool are used and passed to the `IdGenerator` in the `initialize` method inside a `com.ibm.websphere.servlet.cache.CacheConfig` object.

Data type String

Meta data generator

Specifies a user-written interface for handling invalidation, priority levels, and external cache groups.

The value must represent the full package and class name of a class extending `com.ibm.websphere.servlet.cache.MetaDataGenerator`. The properties specified in the Application Assembly Tool are used and passed to the `MetaDataGenerator` in the `initialize` method inside a `com.ibm.websphere.servlet.cache.CacheConfig` object.

Data type String

Web components assembly settings

Use this page to set the assembly properties for the components that make up a Web module.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance* > **Web Components**

Component name

Specifies the name of the servlet or JavaServer Pages(TM) (JSP) file. This name must be unique within the Web module.

Data type String

Display name

Specifies a short name that is intended for display by GUIs.

Data type String

Description

Contains a description of the servlet or JSP file.

Data type String

Component type

Specifies the type of Web component. Valid values are servlet or JSP file.

Data type String

Class name

Specifies the full path name for the servlet class.

Data type String

JSP file

Specifies the full path name for the JSP file.

Data type String

Load on startup

Indicates whether this servlet loads at the startup of the Web application. The default is false (the check box is not selected).

This field also specifies a positive integer indicating the order in which the servlet is to load. Lower integers are loaded before higher integers. If no value is specified, or if the value specified is not a positive integer, the container is free to load the servlet at any time in the startup sequence.

Data type Boolean
Default False

Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type String

Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type String

Run as role name

Enter a role name that represents the user account under which the servlet executes. The default role name is blank, which indicates the servlet runs under the user that logged into the application server. The role name of "all role" indicates the servlet can execute under different users.

Data type String

Description

In this optional field, enter a description that explains the importance of the role, and where and how the role can be used.

Data type String

Run as role mode

Indicates a security role that is defined in the enterprise application.

Data type

String

Unresolved action

This output field defines a behavior for a transaction that does not complete successfully. For example, a servlet does not have to specify rollback in its code. Rollback occurs automatically if the servlet initiates a transaction and the transaction does not complete but rollback is specified as the unresolved action.

Data type

String

Web modules assembly settings

Use this page to set the assembly properties for web modules. Web modules are composed of one or more web components.

File name

Specifies the file name of the Web module, relative to the top level of the application package.

Alternative DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module's JAR file.

This file is the postassembly version of the deployment descriptor file. (The original deployment descriptor file can be edited to resolve dependencies and security information. Directing the use of the alternative deployment descriptor allows you to keep the original deployment descriptor file intact). The value of the Alternative DD property must be the full path name of the deployment descriptor file relative to the module's root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read directly from the module's JAR file.

Context root

Specifies the context root of the Web application. The context root is combined with the defined servlet mapping (from the WAR file) to compose the full URL that users type to access the servlet.

For example, if the context root is /gettingstarted and the servlet mapping is MySession, then the URL is http://host:port/gettingstarted/MySession.

Classpath

Specifies the class path for resources used by the Web application, relative to the ear file..

If your Web application requires access to classes within an ear file, specify the relative path of the classes in this field.

Display name

Specifies a short name that is intended to be displayed by GUIs.

Description

Contains a description of the Web module.

Distributable

Specifies that this Web application is programmed appropriately to be deployed into a distributed servlet container.

Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels). The image is used as an icon to represent the module in a GUI.

Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels). The image is used as an icon to represent the module in a GUI.

Session configuration

Indicates that session configuration information is present. Checking this box makes the Session timeout property editable.

Session timeout

Specifies a time period, in seconds, after which a client is considered inactive. The default value is zero, indicating that the session timeout never expires.

Login configuration -- Authentication method

Specifies an authentication method to use. As a prerequisite to gaining access to any Web resources protected by an authorization constraint, a user must authenticate by using the configured mechanism.

A Web application can authenticate a user to a Web server by using one of the following mechanisms: HTTP basic authentication, HTTP digest authentication, HTTPS client authentication, and form-based authentication.

- HTTP basic authentication is not a secure protocol because the user password is transmitted with a simple Base64 encoding and the target server is not authenticated. In basic authentication, the Web server requests a Web client to authenticate the user and passes a string called the realm of the request in which the user is to be authenticated.
- HTTP digest authentication transmits the password in encrypted form.
- HTTPS client authentication uses HTTPS (HTTP over SSL) and requires the user to possess a public key certificate.
- Form-based authentication allows the developer to control the appearance of login screens.

The Login configuration properties are used to configure the authentication method that should be used, the realm name that should be used for HTTP basic authentication, and the attributes that are needed by the form-based login mechanism. Valid values for this property are Unspecified, Basic, Digest, Form, and Client certification.

Note:

HTTP digest authentication is not supported as a login configuration in this product. Also, not all login configurations are supported in all of the product's global security authentication mechanisms (Local Operating system, LTPA, and custom pluggable user registry). HTTP basic authentication and form-based login authentication are the only authentication methods supported by the Local Operating system user registry. LTPA and the custom pluggable user registry are capable of supporting HTTP basic authentication, form-based login, and HTTPS client authentication.

Login configuration -- Realm name

Specifies the realm name to use in HTTP basic authorization. It is based on a user name and password, sent as a string (with a simple Base64 encoding).

An HTTP realm is a string that allows URIs to be grouped together. For example, if a user accesses a secured resource on a Web server within the "finance realm," subsequent access to the same or different resource within the same realm does not result in a repeat prompt for a user ID and password.

Login configuration -- Login page

Specifies the location of the login form. If form-based authentication is not used, this property is disabled.

Form Login Config -- Error page

Specifies the location of the error page. If form-based authentication is not used, this property is disabled.

Reload interval

Specifies a time interval, in seconds, in which the Web application's file system is scanned for updated files. The default is 0 (zero).

Reloading enabled

Specifies whether file reloading is enabled. The default is false.

Default error page

Specifies a file name for the default error page. If no other error page is specified in the application, this error page is used.

Additional classpath

Specifies the full class path that will be used to reference classes outside of those specified in the archive.

If your Web application requires access to classes not contained in the archive file, specify the full path for those classes in this field.

File serving enabled

Specifies whether file serving is enabled. File serving allows the application to serve static file types, such as HTML and GIF. File serving can be disabled if the application contains only dynamic components. The default value is true.

Directory browsing enabled

Specifies whether directory browsing is enabled. Directory browsing allows the application to browse disk directories. Directory browsing can be disabled if, you want to protect data. The default value is true.

Serve servlets by classname

Specifies whether a servlet can be served by requesting its class name.

Usually, servlets are served only through a URI reference. The class name is the actual name of the servlet on disk. For example, a file named SnoopServlet.java compiles into SnoopServlet.class. (This is the class name.) SnoopServlet.class is normally invoked by specifying snoop in the URI. However, if Serve Servlets by Classname is enabled, the servlet is invoked by specifying SnoopServlet. The default value is true.

Virtual hostname

Specifies a virtual host name. A virtual host is a configuration enabling a single host machine to resemble multiple host machines. This property allows you to bind the application to a virtual host in order to enable execution on that virtual host.

Filter mappings

Specifies the filter mapping declarations in this application. The container uses the filter mapping declarations to decide on the type and order of filters to apply to a request.

After the container matches the request URI to a servlet, for each filter mapping element, it determines what filters to apply based on the servlet name or the URL pattern, depending on the style specified. Filters are invoked in the same order as the one specified in the list of filter mapping elements. The value that you specify for the filter name must be the same value as that specified in the `<filter><filtername>` sub-element declarations in the deployment descriptor.

Assembly property extensions

Use this panel to configure WebSphere Application Server specific Web module extensions, or also referred to as assembly property extensions.

This panel lists the extensions that can be configured through the tool.

Reach the applicable extension panel by clicking on the panel name in the navigation at the left, or by double-clicking the attribute name in the list provided.

Access this page by traversing the following path in the Application AssemblyTool:
Web Modules > component_instance > Assembly Property Extensions

Assembly properties for a Web module include:

- File serving attributes
- Invoker attributes
- JavaServer Pages (JSP) attributes
- Multipurpose Internet Mail Extensions (MIME) filters
- Servlet caching configurations

File serving attribute assembly settings

File serving allows a Web application to serve static file types, such as HTML. File-serving attributes are used by the servlet that implements file-serving behavior.

Access this page by traversing the following path in the Application AssemblyTool:
Web Modules > component_instance > Assembly Property Extensions

File Serving Attribute (Name)

Specifies the name of an attribute.

Data type String

File Serving Attribute (Value)

Specifies the value of an attribute.

Data type String

Invoker attribute assembly settings

Invoker attributes are used by the servlet that implements the invocation behavior.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance* > **Assembly Property Extensions**

Invoker Attribute (Name)

Specifies the name of an attribute.

Data type String

Invoker Attribute (Value)

Specifies the value of an attribute.

Data type String

Error page assembly settings

Error page locations allow a servlet to find and serve a URI to a client based on a specified error status code or exception type. These properties are used if the error handler is another servlet or JSP file.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance*

The error page properties specify a mapping between an error code or exception type and the path of a resource in the Web application. The container examines the list in the order that it is defined, and attempts to match the error condition by status code or by exception class. On the first successful match of the error condition, the container serves back the resource defined in the Location property.

Error code

Indicates that the error condition is a status code.

Data type Integer

Error Code (Required, String)

Specifies an HTTP error code, for example, 404.

Data type String

Exception

Indicates that the error condition is an exception type.

Data type String

Exception type name (Required, String)

Specifies an exception type.

Data type String

Location (Required, String)

Contains the location of the error-handling resource in the Web application.

Data type String

Web resource collections security constraint properties

A Web resource collection defines a set of URL patterns or resources and HTTP methods belonging to the resource, which define the security constraints for a Web component.

Access this page by traversing the following path in the Application Assembly Tool:

Web Modules > *component_instance*

HTTP methods handle HTTP-based requests, such as GET, POST, PUT, and DELETE. A URL pattern is a partial Uniform Resource Locator that acts as a template for matching the pattern with existing full URLs in an attempt to find a valid file.

Web resource name

Specifies the name of a Web resource collection.

Data type String

Web resource description

Contains a description of the Web resource collection.

HTTP methods

Specifies the HTTP methods to which the security constraints apply. If no HTTP methods are specified, then the security constraint applies to all HTTP methods. The valid values are GET, POST, PUT, DELETE, HEAD, OPTIONS, and TRACE.

Data type String

URL pattern

Specifies URL patterns for resources in a Web application. All requests that contain a request path that matches the URL pattern are subject to the security constraint.

Data type String

Troubleshooting tips for Web application deployment

Deployment of a Web application is successful if you can access the application by typing a Uniform Resource Locator (URL) in a browser, or if you can access the application by following a link.

If you cannot access your application, follow these steps to eliminate some common errors that can occur during migration or deployment.

Web module does not run in WebSphere Application Server Version 5.

Symptom	Your Web module does not run when you migrate it to Version 5
Problem	In Version 4.x, the classpath setting that affected visibility was <i>Module Visibility Mode</i> . In Version 5, you must use classloader policies to set visibility.

Recommended response

Reassemble an existing module, or change the visibility settings in the classloader policies.

See article (Migration of module visibility modes from Version 4.x) for more information and examples.

Welcome page is not visible.

Symptom

You cannot access an application with a Web path of:

```
/webapp/myapp
```

Problem

The default welcome page for a Web application is assumed to be *index.html*. You cannot access the default page of the *myapp* application unless it is named *index.html*.

Recommended response

To identify a different welcome page, modify the properties of the Web module during assembly. See article Assembling Web modules for more information.

HTML files are not found.

Symptom

Your Web application ran successfully on prior versions, but now you encounter errors that the welcome page (typically *index.html*), or referenced HTML files are not found:

```
Error 404: File not found: Banner.html  
Error 404: File not found: HomeContent.html
```

Problem

For security and consistency reasons, Web application URLs are now case-sensitive on all operating systems.

Suppose the content of the index page is as follows:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 5.0 Frameset//EN">  
<HTML>  
<TITLE>  
Insurance Home Page  
</TITLE>  
<frameset rows="18,80">  
  <frame src="Banner.html" name="BannerFr  
  <frame src="HomeContent.html" name="HomeCont  
</frameset>  
</HTML>
```

However the actual file names in the `\WebSphere\AppServer\installedApps\...` directory where the application is deployed are:

```
banner.html  
homecontent.html
```

Recommended response

To correct this problem, modify the *index.html* file to change the names *Banner.html* and *HomeContent.html* to *banner.html* and *homecontent.html* to match the names of the files in the deployed application.

Modifying the default Web container configuration

Before you begin

If you choose not to use the default application server, create an application server instance.

An application server instance contains a single instance of a Web container, which is created automatically when an application server is created.

The Web container is created initially with default properties values suitable for simple Web applications. However, these values might not be appropriate for more complex Web applications.

Your application is complex if it requires any of the following features:

- virtual host
- servlet caching
- special client request loads
- persistent HTTP session support
- special HTTP transport settings

Modify the following properties if you have a complex application:

Steps for this task

1. **(Optional)** Modify the Web container General Properties if your Web application requires a virtual host, other than the default_host, or requires servlet caching.
2. **(Optional)** Modify the Web Container Additional Properties Thread Pool setting if your application handles special client request loads.
3. **(Optional)** Modify the Web Container Additional Properties Session Management setting if your application requires persistent HTTP session support.
4. **(Optional)** Modify the Web Container Additional Properties HTTP transports setting if your application requires one of the following HTTP transport settings:
 - unique hostname and port for client access
 - SSL enablement

Web container

A Web container handles requests for servlets, JavaServer Pages (JSP) files, and other types of files that include server-side code. The Web container creates servlet instances, loads and unloads servlets, creates and manages request and response objects, and performs other servlet management tasks.

The Web server plug-ins, provided by the WebSphere Application Server, help supported Web servers pass servlet requests to Web containers.

Web container settings

Use this page to configure the container settings.

Access this page by traversing the following path in the Web console:

Servers > Application Servers > *server_instance* > Web container

Configuration - General Properties

Default virtual host

Specifies a virtual host is a configuration enabling a single host machine to resemble multiple host machines. Resources associated with one virtual host cannot share data with resources associated with another virtual host, even if the virtual hosts share the same physical machine.

Select a virtual host option:

Default Host

The product provides a default virtual host with some common aliases, such as the machine IP address, short host name, and fully qualified host name. The alias comprises the first part of the path for accessing a resource such as a servlet. For example, it is `localhost:9080` in the request `http://localhost:9080/myServlet`.

Admin Host

This is another name for the application server; also known as *server1* in the base installation. This process supports the use of the administrative console.

Servlet caching

Specifies that if a servlet is invoked once and it generates output to be cached, a cache entry is created containing not only the output, but also side effects of the invocation. These side effects can include calls to other servlets or Java Server Pages (JSP) files, as well as metadata about the entry, including timeout and entry priority information.

Enable servlet caching

Check this box to enable servlet caching.

Web Container Services settings

Use this page to configure container services settings.

Access this page by traversing the following path in the Web console:

Servers > Application Servers > *server_instance* > Web container > Additional Properties

Can Be Grown

Specifies the number of threads to increase beyond the maximum size configured for the pool.

Data type	Boolean
Default	False

Default Priority

Specifies the default priority for servlets that you can cache. This setting determines how long an entry stays in a full cache. The recommended value is 1.

Data type	Integer
Default	Not applicable
Range	1-3

Dynamic Properties

Specifies a set of name-value pairs for configuring properties beyond those displayed in the interface.

Data type String

Enable servlet caching

Enable the servlet and JavaServer Pages (JSP) dynamic Java Naming and Directory Interface (JNDI) caching feature.

Data type Boolean
Default True

HTTP Transport

Specifies the HTTP transports associated with this Web container.

Data type String

Inactivity Timeout

Specifies the time after which a thread is reclaimed due to inactivity.

Data type Integer
Units seconds

Installed Web Modules

Specifies the Web modules installed into the Web container of this server.

Data type String

Maximum Size

Specifies the maximum number of threads to allow in the pool.

Data type Integer

Minimum Size

Specifies the minimum number of threads to allow in the pool.

Data type Integer
Units Threads

Session Manager

Specifies the Session Manager associated with this Web container.

Data type String

Thread Pool

Specifies the thread pool settings for the Web container.

Data type String

Web module settings

Use this page to configure Web module settings.

Access this page by traversing the following path in the Web console:

Enterprise Application > *application_instance* > **Web Module**

URI

Specifies a URI that, when resolved relative to the application URL, specifies the location of the module archive contents on a file system. The URI must match the ModuleRef URI in the deployment descriptor of an application if the module was packaged as part of a deployed application or enterprise archive (EAR) file.

Name

Specifies the unique display name for the module.

Alternate DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module JAR file.

This file is the *post-assembly* version of the deployment descriptor file. You can edit the original deployment descriptor file to resolve dependencies and security information. Specifying the use of the alternative deployment descriptor keeps the original deployment descriptor file intact.

The value of the *Alternate DD* property must be the full path name of the deployment descriptor file, relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read from the module JAR file.

Starting weight

Specifies the order in which modules are started. Lower weighted modules are started before higher weighted modules.

Prefer WEB-INF Classes

Specifies classes to load in WEB-INF before any other classes. Implementing the application class loader is recommended so that classes and resources packaged within the WAR file load before classes and resources residing in container-wide library JAR files.

Initial State

Specifies the default state of this application at server startup.

Web Module Deployment settings

Use this page to configure an instance of Web module deployment.

Access this page by traversing the following path in the Web console:

Enterprise Application > *application_instance* > **Web Module** > *Web Module_instance*

URI

Specifies a URI that, when resolved relative to the application URL, specifies the location of the module archive contents on a file system. The URI must match the ModuleRef URI in the deployment descriptor of an application if the module was packaged as part of a deployed application or enterprise archive (EAR) file.

Alternate DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module JAR file.

This file is the *post-assembly* version of the deployment descriptor file. You can edit the original deployment descriptor file to resolve dependencies and security information. Specifying the use of the alternative deployment descriptor keeps the original deployment descriptor file intact.

The value of the *Alternate DD* property must be the full path name of the deployment descriptor file, relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read from the module JAR file.

Starting weight

Specifies the order in which modules are started. Lower weighted modules are started before higher weighted modules.

Classloader Mode

Specifies whether the classloader should search in the parent classloader or in the application classloader first to load a class. The standard for JDK classloaders and WebSphere classloaders is PARENT_FIRST. By specifying PARENT_LAST, your application can override classes contained in the parent classloader, but this action can potentially result in ClassCastException or LinkageErrors if you have mixed use of overridden classes and non-overridden classes.

The options are PARENT_FIRST and PARENT_LAST. The default is to search in the parent classloader before searching in the application classloader to load a class.





Data type	String
Default	PARENT_FIRST









Web applications: Resources for learning

Use the following links to find relevant supplemental information about Web applications. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Programming specifications
-  **J2EE BluePrints for Web applications**
http://java.sun.com/blueprints/guidelines/designing_enterprise_applications/web_tier/index.html
-  **Redbook on the design and implementation of Servlets, JSP files, and enterprise beans**
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg245754.html?OpenDocument>
-  **Redbook on Servlet and JSP file Programming**
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg245755.html?OpenDocument>
-  **Sun's Java™ Tutorial on Servlets**
<http://java.sun.com/docs/books/tutorial/servlets/index.html>

-  **Introduction to JavaServer Pages - Tutorial**
<http://www-4.ibm.com/software/webservers/appserv/education.html#online>
-  **Bean Scripting Framework description**
<http://www.mozilla.org/rhino/bsf.html>
-  **Web delivered samples in the Samples Gallery**
<http://www7b.software.ibm.com/wsdd/library/samples/AppServer.html>
-  **Java 2 Software Development Kit (SDK)**
<http://java.sun.com/j2se/1.3/>
-  **Servlet 2.3 Specification**
<http://jcp.org/aboutJava/communityprocess/first/jsr053/index.html>
-  **JavaServer Pages 1.2 Specification**
<http://jcp.org/aboutJava/communityprocess/first/jsr053/index.html>
-  **Differences between JavaScript and ECMAScript**
<http://www.webstandards.org/learn/resources/javascript/index.html>
-  **ISO 8859 Specifications**
<http://www.iso.org/iso/en/CatalogueListPage.CatalogueList>

Chapter 3. Managing HTTP sessions

IBM WebSphere Application Server provides a service for managing HTTP sessions: Session Manager. The key activities for session management are summarized below.

Before you begin these steps, make sure you are familiar with the programming model for accessing HTTP session support in the applications following the Servlet 2.3 API.

Steps for this task

1. Plan your approach to session management, which could include session tracking, session recovery, and session clustering.
2. Create or modify your own applications to use session support to maintain sessions on behalf of Web applications.
3. (Assemble your application.)
4. (Deploy your application.)
5. Ensure the administrator appropriately configures session management in the administrative domain.
6. Adjust configuration settings and perform other tuning activities for optimal use of sessions in your environment.

Sessions

A session is a series of requests to a servlet, originating from the same user at the same browser.

Sessions allow applications running in a Web container to keep track of individual users.

For example, a servlet might use sessions to provide "shopping carts" to online shoppers. Suppose the servlet is designed to record the items each shopper indicates he or she wants to purchase from the Web site. It is important that the servlet be able to associate incoming requests with particular shoppers. Otherwise, the servlet might mistakenly add Shopper_1's choices to the cart of Shopper_2.

A servlet distinguishes users by their unique session IDs. The session ID arrives with each request. If the user's browser is cookie-enabled, the session ID is stored as a cookie. As an alternative, the session ID can be conveyed to the servlet by URL rewriting, in which the session ID is appended to the URL of the servlet or JavaServer Pages (JSP) file from which the user is making requests. For requests over HTTPS or Secure Sockets Layer (SSL), Another alternative is to use SSL information to identify the session.

Migrating HTTP sessions

Note: In Version 5 default write frequency mode is `TIME_BASED_WRITES`, which is different from Version 4.0 and 3.5 default mode of `END_OF_SERVICE`.

Migrating from Version 4.0

No programmatic changes are required to migrate from version 4.0 to version 5.

Migrating from Version 3.5

If you have Version 3.5 applications running in Servlet 2.1 mode, some of the following Version 5 differences might influence how you choose to track and manage sessions.

Steps for this task

1. During application development, modify session-related APIs as needed.
Some API changes are required in order to redeploy existing applications on Version 5. These include changes to the `HttpSession` API itself as well as issues associated with moving to support for the Servlet 2.3 specification. Certain Servlet 2.1 API methods have been deprecated in Servlet 2.3 API . These deprecated APIs still work in Version 5.0, but they may be removed in a future version of the API. Changes are summarized in the following list:
 - Replace instances of `getValue()` with `getAttribute()`
 - Replace instances of `getValueNames()` with `getAttributeNames()`
 - Replace instances of `removeValue()` with `removeAttribute()`
 - Replace instances of `putValue()` with `setAttribute()`
2. During application development, modify Web application behavior as needed.
In accordance with the Servlet 2.3 specification, `HttpSession` objects must be scoped within a single Web application context; they may not be shared between contexts. This means that a session can no longer span Web applications. Objects added to a session by a servlet or JSP in one Web application cannot be accessed from another Web application. The same session ID may be shared (because the same cookie is in use), but each Web application will have a unique session associated with the session ID. Version 5 provides a feature that can be used to extend scope of a session to enterprise application.
3. Use administrative tools to configure Session Manager security settings as needed.
Relative to session security, the default Session Manager setting for Integrate Security is now `false`. This is different from the default setting in some earlier releases.
4. Use administrative tools to configure the JSP enabler and application server as needed.
In Version 3.5 of the product, JSP files that contained the usebean tag with scope set to `session` did not always work properly when session persistence was enabled. Specifically, the JSP writer needed to write a scriplet to explicitly set the attribute (that is, to call `setAttribute()`) if it was changed as part of JSP processing.
Two new features in Version 5.0 help address this problem:
 - You can set `dosetattribute` to `true` on the JSP `InitParameter`.
 - You can set the `Write Contents` option to `Write all`.

The differences between the two solutions are summarized in the following table:

	Applies to	Configured at	Action
dosetattribute set to <i>true</i>	JSP	JSP enabler	Assures that JSP session-scoped beans always call <code>setAttribute()</code>
Write Contents option set to <i>Write all</i>	servlet or JSP	application server	All session data (changed or unchanged) is written to the external location

If session persistence is enabled and a class reload for the Web application occurs, the sessions associated with the Web application are maintained in the persistent store and will be available after the reload.

Developing session management in servlets

This information, combined with the coding example `SessionSample.java`, provides a programming model for implementing sessions in your own servlets.

Steps for this task

1. Get the `HttpSession` object.

To obtain a session, use the `getSession()` method of the `javax.servlet.http.HttpServletRequest` object in the Java Servlet 2.3 API.

When you first obtain the `HttpSession` object, the Session Management facility uses one of three ways to establish tracking of the session: cookies, URL rewriting, or Secure Sockets Layer (SSL) information.

Assume the Session Management facility uses cookies. In such a case, the Session Management facility creates a unique session ID and typically sends it back to the browser as a *cookie*. Each subsequent request from this user (at the same browser) passes the cookie containing the session ID, and the Session Management facility uses this ID to find the user's existing `HttpSession` object.

In Step 1 of the code sample, the `Boolean(create)` is set to `true` so that the `HttpSession` object is created if it does not already exist. (With the Servlet 2.3 API, the `javax.servlet.http.HttpServletRequest.getSession()` method with no boolean defaults to `true` and creates a session if one does not already exist for this user.)

2. Store and retrieve user-defined data in the session.

After a session is established, you can add and retrieve user-defined data to the session. The `HttpSession` object has methods similar to those in `java.util.Dictionary` for adding, retrieving, and removing arbitrary Java objects.

In Step 2 of the code sample, the servlet reads an integer object from the `HttpSession`, increments it, and writes it back. You can use any name to identify values in the `HttpSession` object. The code sample uses the name `sessiontest.counter`.

Because the `HttpSession` object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

3. (Optional) Output an HTML response page containing data from the `HttpSession` object.

4. Provide feedback to the user that an action has taken place during the session. You may want to pass HTML code to the client browser indicating that an action has occurred.

For example, in step 3 of the code sample, the servlet generates a Web page that is returned to the user and displays the value of the `sessiontest.counter` each time the user visits that Web page during the session.

5. (Optional) Notify Listeners.

Objects stored in a session that implement the `javax.servlet.http.HttpSessionBindingListener` interface are notified when the session is preparing to end and become invalidated. This notice enables you to perform post-session processing, including permanently saving the data changes made during the session to a database.

6. End the session.

You can end a session:

- Automatically with the Session Management facility if a session is inactive for a specified time. The administrators provide a way to specify the amount of time after which to invalidate a session.
- By coding the servlet to call the `invalidate()` method on the session object.

SessionSample.java

```
import java.io.*;
import java.util.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class SessionSample extends HttpServlet {
    public void doGet (HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {

        // Step 1: Get the Session object

        boolean create = true;
        HttpSession session = request.getSession(create);

        // Step 2: Get the session data value

        Integer ival = (Integer)
            session.getAttribute ("sessiontest.counter");
        if (ival == null) ival = new Integer (1);
        else ival = new Integer (ival.intValue () + 1);
        session.setAttribute ("sessiontest.counter", ival);

        // Step 3: Output the page

        response.setContentType("text/html");
        PrintWriter out = response.getWriter();
        out.println("<html>");
        out.println("<head><title>Session Tracking Test<
            /title></head>");
        out.println("<body>");
        out.println("<h1>Session Tracking Test</h1>");
        out.println ("You have hit this page " + ival + " times" +
            "<br>");
        out.println ("Your " + request.getHeader("Cookie"));
        out.println("</body></html>");
    }
}
```

Assembling so that session data can be shared

In accordance with the Servlet 2.3 API specification, by default the Session Management facility supports session scoping by Web module only. Only servlets in the same Web module can access the data associated with a particular session. WebSphere Application Server provides an option that you can use to extend the scope of the session attributes to an enterprise application. Therefore, you can share session attributes across all the Web modules in an enterprise application. This option is provided as an IBM extension.

Restriction: To use this option, you must install all the Web modules in the enterprise application on a given server. You cannot split up Web modules in the enterprise application by servers. For example, with an enterprise application containing two Web modules, you cannot use this option when one Web module is installed on one server and second Web module is installed on a different server. In such split installations, applications might share session attributes across Web modules using distributed sessions, but session data integrity is lost when concurrent access to a session is made in different Web modules. It also severely restricts use of some Session Management features, like `TIME_BASED_WRITES`. For enterprise applications on which this option is enabled, the Session Management configuration on the Web module inside the enterprise application is ignored. Then Session Management configuration defined on enterprise application is used if Session Management is overwritten at the enterprise application level. Otherwise, the Session Management configuration on the Web container is used.

Do the following to share session data across Web modules in an enterprise application:

Steps for this task

1. Launch the (Application Assembly Tool) (AAT).
2. Click the application (EAR file) you want to share.
3. Click the **IBM extension tab**.
4. Click **Shared httpsession context**.
5. Click **Apply**.

Make sure the class definition of attributes put into session are available to all the Web modules in the enterprise application.

6. Save the application (EAR) file.

Servlet API Behavior

If shared HttpSession context is turned on in an enterprise application, HttpSession listeners defined in all the Web modules inside the enterprise application are invoked for session events. The order of listener invocation is not guaranteed.

Session security support

You can integrate HTTP sessions and security in IBM WebSphere Application Server. When security integration is enabled in the Session Management facility and a session is accessed in a protected resource, you can access that session only in protected resources from then on. You cannot mix secured and unsecured resources accessing sessions when security integration is turned on. Security integration in the Session Management facility is not supported in form-based login with SWAM.

Security integration rules for HTTP sessions

Only authenticated users can access sessions created in secured pages and are created under the identity of the authenticated user. Only this authenticated user can access these sessions in other secured pages. To protect these sessions from unauthorized users, you cannot access them from an unsecured page.

Programmatic details and scenarios

IBM WebSphere Application Server maintains the security of individual sessions.

An identity or user name, readable by the `com.ibm.websphere.servlet.session.IBMSession` interface, is associated with a session. An unauthenticated identity is denoted by the user name `anonymous`. IBM WebSphere Application Server includes the `com.ibm.websphere.servlet.session.UnauthorizedSessionRequestException` class, which is used when a session is requested without the necessary credentials.

The Session Management facility uses the WebSphere Application Server security infrastructure to determine the authenticated identity associated with a client HTTP request that either retrieves or creates a session. WebSphere Application Server security determines identity using certificates, LPTA, and other methods.

After obtaining the identity of the current request, the Session Management facility determines whether to return the session requested using a `getSession()` call or not.

The following table lists possible scenarios in which security integration is enabled with outcomes dependent on whether the HTTP request is authenticated and whether a valid session ID and user name was passed to the Session Management facility.

	Unauthenticated HTTP request is used to retrieve a session	HTTP request is authenticated, with an identity of "FRED" used to retrieve a session
No session ID was passed in for this request, or the ID is for a session that is no longer valid	A new session is created. The user name is <code>anonymous</code>	A new session is created. The user name is <code>FRED</code>
A session ID for a valid session is passed in. The current session user name is <code>"anonymous"</code>	The session is returned.	The session is returned. Session Management changes the user name to <code>FRED</code>
A session ID for a valid session is passed in. The current session user name is <code>FRED</code>	The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown*	The session is returned.
A session ID for a valid session is passed in. The current session user name is <code>BOB</code>	The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown*	The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown*

* A `com.ibm.websphere.servlet.session.UnauthorizedSessionRequestException` error is thrown to the servlet.

Session management support

WebSphere Application Server provides facilities, grouped under the heading *Session Management*, that support the `javax.servlet.http.HttpSession` interface described in the Servlet API specification.

In accordance with the Servlet 2.3 API specification, the Session Management facility supports session scoping by Web module only. Only servlets in the same Web module can access the data associated with a particular session. Multiple requests from the same browser, each specifying a unique Web application, result in multiple sessions with a shared session ID. You can invalidate any of the sessions that share a session ID without affecting the other sessions.

You can configure a session timeout for each Web application. A Web application timeout value of 0 (the default value) means that the invalidation timeout value from the Session Management facility is used.

When an HTTP client interacts with a servlet, the state information associated with a series of client requests is represented as an HTTP session and identified by a session ID. Session Management is responsible for managing HTTP sessions, providing storage for session data, allocating session IDs, and tracking the session ID associated with each client request through the use of cookies or URL rewriting techniques. Session Management can store session-related information in several ways:

- In application server memory (the default). This information cannot be shared with other application servers.
- In a database. This storage option is known as *database persistent sessions*.

The last two options are referred to as *distributed sessions*. Distributed sessions are essential for using HTTP sessions for failover facility. When an application server receives a request associated with a session ID that it currently does not have in memory, it can obtain the required session state by accessing the external store (database or memory-to-memory). If distributed session support is not enabled, an application server cannot access session information for HTTP requests that are sent to servers other than the one where the session was originally created. Session Management implements caching optimizations to minimize the overhead of accessing the external store, especially when consecutive requests are routed to the same application server.

Storing session states in an external store also provides a degree of fault tolerance. If an application server goes offline, the state of its current sessions is still available in the external store. This availability enables other application servers to continue processing subsequent client requests associated with that session.

Saving session states to an external location does not completely guarantee their preservation in case of a server failure. For example, if a server fails while it is modifying the state of a session, some information is lost and subsequent processing using that session can be affected. However, this situation represents a very small period of time when there is a risk of losing session information.

The drawback to saving session states in an external store is that accessing the session state in an external location can use valuable system resources. Session Management can improve system performance by caching the session data at the server level. Multiple consecutive requests that are directed to the same server can find the required state data in the cache, reducing the number of times that the

actual session state is accessed in external store and consequently reducing the overhead associated with external location access.

Configuring session management by level

When you configure session management at the Web container level, all applications and the respective Web modules in the Web container normally inherit that configuration, setting up a basic default configuration for the applications and Web modules below it.

However, you can set up different configurations individually for specific applications and Web modules that vary from the Web container default. These different configurations override the default for these applications and Web modules only.

Note: When you overwrite the default session management settings on the application level, all the Web modules below that application inherit this new setting unless they too are set to overwrite these settings.

Steps for this task

1. Open the Administrative console.
2. Select the level that this configuration applies to:
 - For the web container level:
 - a. Click **Servers > Application Servers**.
 - b. Select a server from the list of application servers.
 - c. Under Additional Properties, click **Web Container**.
 - For the enterprise application level:
 - a. Click **Applications > Applications**.
 - b. Select an applications from the list of applications.
 - For the Web module level:
 - a. Click **Applications > Applications**.
 - b. Select an applications from the list of applications.
 - c. Under Related Items, click **Web Modules defined for this Application**.
 - d. Click a Web module from the list of Web modules.
3. Under Additional Properties, click **Session Management**.
4. Make whatever changes you need to manage sessions
5. **(Optional)** If you are working on the Web module or application level and want these settings to override the inherited Session Management settings, select **Overwrite**.
6. Click **Apply** and **Save**.

Session tracking options

There are several options for session tracking, depending on what sort of tracking method you want to use:

- Session tracking with cookies
- Session tracking with URL rewriting
- Session tracking with SSL information

Session tracking with cookies

Tracking sessions with cookies is the default. No special programming is required to track sessions with cookies.

Session tracking with URL rewriting

An application that uses URL rewriting to track sessions must adhere to certain programming guidelines. The application developer needs to do the following:

- Program servlets to encode URLs
- Supply a servlet or Java Server Pages (JSP) file as an entry point to the application

Using URL rewriting also requires that you enable URL rewriting in the Session Management facility.

Note: In certain cases, clients cannot accept cookies. Therefore, you cannot use cookies as a session tracking mechanism. Applications can use URL rewriting as a substitute.

Program session servlets to encode URLs

Depending on whether the servlet is returning URLs to the browser or redirecting them, include either `encodeURL()` or `encodeRedirectURL()` in the servlet code. Examples demonstrating what to replace in your current servlet code follow.

Rewrite URLs to return to the browser

Suppose you currently have this statement:

```
out.println("<a href=\"/store/catalog\">catalog<a>");
```

Change the servlet to call the `encodeURL` method before sending the URL to the output stream:

```
out.println("<a href=\"\"");  
out.println(response.encodeURL ("/store/catalog"));  
out.println(">catalog</a>");
```

Rewrite URLs to redirect

Suppose you currently have the following statement:

```
response.sendRedirect ("http://myhost/store/catalog");
```

Change the servlet to call the `encodeRedirectURL` method before sending the URL to the output stream:

```
response.sendRedirect (response.encodeRedirectURL  
("http://myhost/store/catalog"));
```

The `encodeURL()` and `encodeRedirectURL()` methods are part of the `HttpServletResponse` object. These calls check to see if URL rewriting is configured before encoding the URL. If it is not configured, the calls return the original URL.

If both cookies and URL rewriting are enabled and `response.encodeURL()` or `encodeRedirectURL()` is called, the URL is encoded, even if the browser making the HTTP request processed the session cookie.

You can also configure session support to enable protocol switch rewriting. When this option is enabled, the product encodes the URL with the session ID for switching between HTTP and HTTPS protocols.

Supply a servlet or JSP file as an entry point

The entry point to an application (such as the initial screen presented) may not require the use of sessions. However, if the application in general requires session support (meaning some part of it, such as a servlet, requires session support), then after a session is created, all URLs are encoded in perpetuate the session ID for the servlet (or other application component) requiring the session support.

The following example shows how you can embed Java code within a JSP file:

```
<%  
response.encodeURL ("/store/catalog");  
%>
```

Session tracking with SSL information

No special programming is required to track sessions with Secure Sockets Layer (SSL) information.

To use SSL information, turn on **Enable SSL Tracking** in the Session Management property sheet. Because the SSL session ID is negotiated between the Web browser and HTTP server, this ID cannot survive an HTTP server failure. However, the failure of an application server does not affect the SSL session ID if an external HTTP Server is present between WebSphere Application Server and the browser.

SSL tracking is supported for the IBM HTTP Server and iPlanet Web servers only. You can control the lifetime of an SSL session ID by configuring options in the Web server. For example, in the IBM HTTP Server, set the configuration variable SSLV3TIMEOUT to provide an adequate lifetime for the SSL session ID. An interval that is too short can cause a premature termination of a session. Also, some Web browsers might have their own timers that affect the lifetime of the SSL session ID. These Web browsers may not leave the SSL session ID active long enough to serve as a useful mechanism for session tracking. Internal Http Server of WebSphere also supports SSL Tracking.

When using the SSL session ID as the session tracking mechanism in a cloned environment, use either cookies or URL rewriting to maintain session affinity. The cookie or rewritten URL contains session affinity information that enables the Web server to properly route a session back to the same server for each request.

Configuring session tracking

To configure session tracking, complete the following:

Steps for this task

1. Go to the appropriate level of Session Management.
2. Specify which session tracking mechanism you want to pass the session ID between the browser and the servlet:
 - To track sessions with cookies, click **Enable Cookies**.
To change the cookie settings, click **Modify**.
 - To track sessions with URL rewriting, click **Enable URL Rewriting**.

If you want to enable protocol switch rewriting, click **Enable protocol switch rewriting**.

- To track sessions with SSL information, click **Enable SSL ID tracking**.

3. Click **Apply**.
4. Click **Save**.
5. Define the session recovery characteristics.

Serializing access to session data

The Servlet API supports concurrent access to a session in a given server instance. WebSphere Application Server provides an option to prevent the concurrent access to a session in a given server instance so that concurrent modification of a session does not occur in a given server instance. This prevention is achieved by synchronizing the requests based on session. When this feature is turned on, a session is obtained for the request before invoking the servlet and requests are synchronized by locking the session for the servlet execution time. Note that synchronization is based on the memory copy of session. So this feature cannot serialize requests across servers based on session when session affinity fails.

Restriction: Use this feature only when concurrent modification of the same session data is possible and is not desirable by the application. This feature has overhead of serializing the requests based on a session.

Do the following to synchronize session access:

Steps for this task

1. Select the level of Session Management on which you want to serialize session access.
2. Under Serialize Session access, click **Allow serial access**.
3. In the Maximum wait time box, type the amount of time, in milliseconds, a servlet waits on a session before continuing execution. The default is 120000 milliseconds or two minutes.
4. **(Optional)** Select **Allow access on timeout** if you want the servlet execution to abort when the session request times out. This selection creates error logs. If you do not select this box, the request is executed normally.
5. Click **Apply**.
6. Click **Save**.

Session Management settings

Use this page to manage HTTP session support. This support includes specifying a session tracking mechanism, setting maximum in-memory session count, controlling overflow, and configuring session timeout.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Web Container > Session Management**.

Overwrite Session Management

Specifies whether or not these Session Management settings take precedence over those normally inherited from a higher level for the current application or Web module.

By default, Web modules inherit Session Management settings from the application level above it, and applications inherit Session Management settings from the Web container level above it.

Session tracking mechanism

Specifies a mechanism for HTTP session management.

Mechanism	Function
Enable SSL ID Tracking	Specifies that session tracking uses Secure Sockets Layer (SSL) information as a session ID. Enabling SSL tracking takes precedence over cookie-based session tracking and URL rewriting.
Enable Cookies	<p>Specifies that session tracking uses cookies to carry session IDs. If cookies are enabled, session tracking recognizes session IDs that arrive as cookies and tries to use cookies for sending session IDs. If cookies are not enabled, session tracking uses URL rewriting instead of cookies (if URL rewriting is enabled).</p> <p>Enabling cookies takes precedence over URL rewriting. Do not disable cookies in the Session Management facility of the application server that is running the administrative application because this action causes the administrative application not to function after a restart of the server. As an alternative, run the administrative application in a separate process from your applications.</p>
Enable URL Rewriting	<p>Click Modify to change these settings.</p> <p>Specifies that the Session Management facility uses rewritten URLs to carry the session IDs. If URL rewriting is enabled, the Session Management facility recognizes session IDs that arrive in the URL if the encodeURL method is called in the servlet.</p>
Enable Protocol Switch Rewriting	Specifies that the session ID is added to a URL when the URL requires a switch from HTTP to HTTPS or from HTTPS to HTTP. If rewriting is enabled, the session ID is required to go between HTTP and HTTPS.

Maximum in-memory session count

Specifies the maximum number of sessions to maintain in memory.

The meaning differs depending on whether you are using in-memory or distributed sessions. For in-memory sessions, this value specifies the number of sessions in the base session table. Use the Allow Overflow property to specify whether to limit sessions to this number for the entire Session Management facility or to allow additional sessions to be stored in secondary tables. For distributed sessions, this value specifies the size of the memory cache for sessions. When the session cache has reached its maximum size and a new session is requested, the Session Management facility removes the least recently used session from the cache to make room for the new one.

Overflow

Specifies that the number of sessions in memory can exceed the value specified by the Max In Memory Session Count property. This option is valid only in nondistributed sessions mode.

Session timeout

Specifies how long a session can go unused before it is no longer valid. Specify either Set timeout or No timeout. Specify the value in minutes greater than or equal to two.

The value of this setting is used as a default when the session timeout is not specified in a Web module deployment descriptor. Note that to preserve performance, the invalidation timer is not accurate to the second. When the Write Frequency is time based, ensure that this value is least twice as large as the write interval.

Security integration

Specifies that when security integration is enabled, the Session Management facility associates the identity of users with their HTTP sessions

Serialize session access

Specifies that concurrent session access in a given server is not allowed.

Maximum wait time

Specifies the maximum amount of time a servlet request waits on an HTTP session before continuing execution. This parameter is optional and expressed in milliseconds. The default is 120000, or 2 minutes. Under normal conditions, a servlet request waiting for access to an HTTP session gets notified by the request that currently owns the given HTTP session when the request finishes.

Session access on timeout

Specifies whether the request executes the servlet or aborts servlet execution in the event of a timeout and creates error logs. This parameter is optional and expressed as the boolean true or false. If the value is false, multiple servlet requests that have timed out concurrently, execute concurrently. The default value is true, servlet execution aborts.

Cookie settings

Use this page to configure cookie settings for session management.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Web Container > Session Management > Enable Cookies.**

Cookie name

Specifies a unique name for the session management cookie. The servlet specification requires the name JSESSIONID. However, for flexibility this value can be configured.

Secure cookies

Specifies that the session cookies include the secure field. Enabling the feature restricts the exchange of cookies to HTTPS sessions only.

Cookie domain

Specifies the domain field of a session tracking cookie. This value controls whether or not a browser sends a cookie to particular servers. For example, if you specify a particular domain, session cookies are sent to hosts in that domain. The default domain is the server.

Cookie path

Specifies that a cookie is sent to the URL designated in the path. Specify any string representing a path on the server. "/" indicates root directory. Specify a value to restrict the paths to which the cookie will be sent. By restricting paths, you prevent the cookie from going to certain URLs on the server. If you specify the root directory, the cookie is sent no matter which path on the given server is accessed.

Cookie maximum age

Specifies the amount of time that the cookie lives on the client browser. Specify that the cookie lives only as long as the current browser session, or to a maximum age. If you choose the maximum age option, specify the age in seconds. This value corresponds to the Time to Live (TTL) value described in the Cookie specification.

Default is the current browser session which is equivalent to setting the value to -1.

Distributed sessions

WebSphere Application Server provides the following session mechanisms in a distributed environment:

- **Database Session persistence**, where sessions are stored in the database specified.

When a session contains attributes that implement `HttpSessionActivationListener`, notification occurs anytime the session is activated (that is, session is read to the memory cache) or passivated (that is, session leaves the memory cache). Passivation can occur because of a server shutdown or when the session memory cache is full and an older session is removed from the memory cache to make room for a newer session. It is not guaranteed that a session is passivated in one application server prior to being activated in another.

Session recovery support

For session recovery support, WebSphere Application Server provides distributed session support in the form of database sessions. Use session recovery support under the following conditions:

- When the user's session data must be maintained across a server restart
- When the user's session data is too valuable to lose through an unexpected server failure

All the attributes set in a session must implement `java.io.Serializable` if the session requires external storage. In general, consider making all objects held by a session serialized, even if immediate plans do not call for session recovery support. If the Web site grows, and session recovery support becomes necessary, the transition occurs transparently to the application if the sessions only hold serialized objects. If not, a switch to session recovery support requires coding changes to make the session contents serialized.

Distributed Environment settings

Use this page to specify a type for saving a session in a distributed environment.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Web Container > Session Management > Distributed Environment Settings**.

Distributed Environment

Specifies a type of distributed environment for saving a session in.

None	Specifies that the Session Management facility discards the session data when the server shuts down.
Database	Specifies that the Session Management facility stores session information in the data source specified by the data source connection settings. Click Modify to change these data source settings.

Configuring for database session persistence

To configure the Session Management facility for database session persistence, complete the following:

Steps for this task

1. Create a JDBC Driver.
2. Create a data source pointing to an existing database, using the JDBC driver that you created. Note the JNDI name of the data source.
3. Go to the appropriate level of Session Management.
4. Click **Distributed Environment Settings**.
5. Select and click **Database**.
6. Specify the Data Source JNDI name from step 2.
7. Specify the database user ID and password for accessing the database.
8. Retype the password for confirmation.
9. **(Optional)** Configure a table space and page sizes for DB2 session databases.
10. **(Optional)** Switch to a multirow schema.
11. Click **OK**.
12. **(Optional)** If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.
13. Click **Apply**.
14. Click **Save**.

Switching to a multirow schema

By default, a single session maps to a single row in the database table used to hold sessions. With this setup, there are hard limits to the amount of user-defined, application-specific data that WebSphere Application Server can access.

Steps for this task

1. Modify the Session Management facility properties to switch from single to multirow schema.
2. Manually drop the database table or delete all the rows in the database table that the product uses to maintain HttpSession objects.

To drop the table:

- a. Determine which data source configuration Session Management is using.
- b. In the data source configuration, look up the database name.
- c. Use the database facilities to connect to the database.
- d. Drop the SESSIONS table.

Configuring tablespace and page sizes for DB2 session databases

If you are using DB2 for session persistence, you can increase the page size to optimize performance for writing large amounts of data to the database. Page sizes of 8K, 16K, and 32K are supported.

To use a page size other than the default (4K), do the following:

Steps for this task

1. If the SESSIONS table already exists, drop it from the DB2 database.
2. Create a new DB2 buffer pool and table space, specifying the same page size (8K, 16K or 32K) for both, and assign the new buffer pool to this table space.

```
DB2 Connect to session
DB2 CREATE BUFFERPOOL sessionBP SIZE 1000 PAGESIZE 8K
DB2 Connect reset
DB2 Connect to session
DB2 CREATE TABLESPACE sessionTS PAGESIZE 8K MANAGED BY SYSTEM
      USING ('D:\DB2\NODE0000\SQL00005\sessionTS.0') BUFFERPOOL sessionBP
DB2 Connect reset
```

Refer to DB2 product documentation for details.

3. Configure the correct table space name and page size in the Session Management facility.

Page size is referred to as *row size* on the Session Management page.)

Results

When the product is restarted, the Session Management facility creates a new SESSIONS table in the specified tablespace based on the indicated page size.

Database settings

Use this page to specify the settings for database session support.

To view this administrative console page, click **Servers > Application Servers > server_name > Web Container > Session Management > Distributed Environment Settings > Database**.

Datasource JNDI Name

Specifies the datasource description

The JNDI name of the non-XA enabled data source from which Session Management obtains database connections. For example, if the JNDI name of the datasource is "jdbc/sessions", specify "jdbc/sessions." The data source represents a pool of database connections and a configuration for that pool (such as the pool size). The data source must already exist as a configured resource in the environment.

User ID

Specifies the user ID for database access

Password

Specifies the password for database access

Confirm Password

Specifies the password a second time to ensure it recorded correctly.

DB2 Row Size

Specifies the tablespace page size configured for the sessions table, if using a DB2 database. Possible values are 4, 8, 16, and 32 kilobytes (K). The default row size is 4K.

The default row size is 4K. In DB2, it can be updated to a larger value. This can help database performance in some environments. When this value is other than 4, you must specify Table Space Name to use. For 4K pages, the Table Space Name is optional.

Table Space Name

Specifies that tablespace to be used for the sessions table.

This value is required when the DB2 Page Size is other than 4K.

Use Multirow Sessions

Specifies that each instance of application data be placed in a separate row in the database, allowing larger amounts of data to be stored for each session. This action can yield better performance in certain usage scenarios. If using multirow schema is not enabled, instances of application data can be placed in the same row.

Multirow schema considerations

IBM WebSphere Application Server supports the use of a multirow schema option in which each piece of application specific data is stored in a separate row of the database. With this setup, the total amount of data you can place in a session is now bound only by the database capacities. The only practical limit that remains is the size of the session attribute object.

The multirow schema potentially has performance benefits in certain usage scenarios, such as when larger amounts of data are stored in the session but only small amounts are specifically accessed during a given servlet processing of an HTTP request. In such a scenario, avoiding unneeded Java object serialization is beneficial to performance.

Understand that switching between multirow and single row is not a trivial proposition.

In addition to allowing larger session records, using multirow schema can yield performance benefits. However, it requires a little work to switch from single-row to multirow schema, as shown in the instructions below.

Coding considerations and test environment

Consider configuring direct single-row usage to one database and multirow usage to another database while you verify which option suits your application needs. (Do this in code by switching the data source used; then monitor performance.)

Programming issue	Application scenario
Reasons to use single-row	<ul style="list-style-type: none">You can read or write all values with just one record read and write.This takes up less space in a database because you are guaranteed that each session is only one record long.
Reasons not to use single-row	2-megabyte limit of stored data per session.

Programming issue	Application scenario
Reasons to use multirow	<ul style="list-style-type: none"> • The application can store an unlimited amount of data; that is, you are limited only by the size of the database and a 2-megabyte-per-record limit. • The application can read individual fields instead of the whole record. When large amounts of data are stored in the session but only small amounts are specifically accessed during servlet processing of an HTTP request, multirow sessions can improve performance by avoiding unneeded Java object serialization.
Reasons not to use multirow	If data is small in size, you probably do not want the extra overhead of multiple row reads when you can store everything in one row.

In the case of multirow usage, design your application data objects not to have references to each other, to prevent circular references. For example, suppose you are storing two objects A and B in the session using `HttpSession.put(..)` method, and A contains a reference to B. In the multirow case, because objects are stored in different rows of the database, when objects A and B are retrieved later, the object graph between A and B is different than stored. A and B behave as independent objects.

Clustered session support

A clustered environment supports load balancing, where the workload is distributed among the application servers that compose the cluster. In a cluster environment, the same Web application must exist on each of the servers that can access the session. You can accomplish this setup by installing an application onto a cluster definition. Each of the servers in the group can then access the Web application

In a clustered environment, the Session Management facility requires an affinity mechanism so that all requests for a particular session are directed to the same application server instance in the cluster. This requirement conforms to the Servlet 2.3 specification in that multiple requests for a session cannot coexist in multiple application servers. One such solution provided by IBM WebSphere Application Server is *session affinity* in a cluster; this solution is available as part of the WebSphere Application Server plug-ins for Web servers. It also provides for better performance because the sessions are cached in memory. In clustered environments other than WebSphere Application Server clusters, you must use an affinity mechanism (for example, IBM WebSphere Edge Server affinity).

If one of the servers in the cluster fails, it is possible for the request to reroute to another server in the cluster. If distributed sessions support is enabled, the new server can access session data from the database or another WebSphere Application Server instance. You can retrieve the session data only if a new server has access to an external location from which it can retrieve the session.

Tuning session management

IBM WebSphere Application Server session support has features for tuning session performance and operating characteristics, particularly when sessions are configured in a distributed environment. These options support the administrator flexibility in determining the performance and failover characteristics for their environment.

The table summarizes the features, including whether they apply to sessions tracked in memory, in a database, with memory-to-memory replication, or all. Click a feature for details about the feature. Some features are easily manipulated using administrative settings; others require code or database changes.

Feature or option	Goal	Applies to sessions in memory, database, or memory-to-memory
Write frequency	Minimize database write operations.	Database and Memory-to-Memory
Session affinity	Access the session in the same application server instance.	All
Multirow schema	Fully utilize database capacities.	Database
Base in-memory session pool size	Fully utilize system capacity without overburdening system.	All
Write contents	Allow flexibility in determining what session data to write	Database and Memory-to-Memory
Scheduled invalidation	Minimize contention between session requests and invalidation of sessions by the Session Management facility. Minimize write operations to database for updates to last access time only.	Database and Memory-to-Memory
Tablespace and row size	Increase efficiency of write operations to database.	Database (DB2 only)

Configuring scheduled invalidation

You can set specific times for the Session Management facility to scan for invalidated sessions in a distributed environment. When used with distributed sessions, this feature has the following benefits:

- You can schedule the scan for invalidated sessions for times of low application server activity, avoiding contention between invalidation scans of database or another WebSphere Application Server instance and read and write operations to service HTTP session requests.
- Significantly fewer external write operations can occur when running with the End of Service Method write mode because the last access time of the server does not require writing out each HTTP session request. (Manual Update and Time Based Write options already minimize the writing of the last access time.)

Usage considerations

- With scheduled invalidation configured, HttpSession timeouts are not strictly enforced. Instead, all invalidation processing is handled at the configured invalidation times.
- HttpSessionBindingListener processing is handled at the configured invalidation times unless the HttpSession.invalidate() method is explicitly called.
- The HttpSession.invalidate() method immediately invalidates the session from both the session cache and the external store.

Configuring write contents

In Session Management, you can configure which session data is written to the database or to another WebSphere instance, depending on whether you are using database persistent sessions or memory to memory replication. This flexibility allows for fewer code changes for the JSP writer when the application will be operating in a clustered environment. The following options are available in Session Management for tuning what is to be written back:

- Write changed (the default) - Write only session data properties that have been updated through setAttribute() and removeAttribute() method calls.
- Write all - Write all session data properties.

The **Write all** setting might benefit servlet and JSP writers who change Java objects' states that reside as attributes in HttpSession and do not call HttpSession.setAttribute().

However, the use of **Write all** could result in more data being written back than is necessary. If this situation applies to you, consider combining the use of **Write all** with **Time-based write** to boost performance overall. As always, be sure to evaluate the advantages and disadvantages for your installation.

With either Write Contents setting, when a session is first created, complete session information is written, including all of the objects bound to the session. When using database session persistence, in subsequent session requests, what is written to the database depends on whether a single-row or multirow schema has been set for the session database, as follows:

Write Contents setting	Behavior with single-row schema	Behavior with multirow schema
Write changed	If any session attribute is updated, all objects bound to the session are written.	Only the session data modified through setAttribute() or removeAttribute() calls is written.
Write all	All bound session attributes are written.	All session attributes that currently reside in the cache are written. If the session has never left the cache, all session attributes are written.

Steps for this task

1. Go to the appropriate level of Session Management.
2. Click Distributed Environment Settings
3. Click Custom Tuning Parameters.
4. Select Custom Settings, and click Modify.

5. Select the appropriate write contents setting.

Configuring write frequency

In the Session Management facility, you can configure the frequency for writing session data to the database or to a WebSphere instance, depending on whether you use database distributed sessions or memory-to-memory replication. This flexibility enables you to weigh session performance gains against varying degrees of failover support. The following options are available in the Session Management facility for tuning write frequency:

- `END_OF_SERVICE` - Write session data at the end of the servlet `service()` method call.
- `MANUAL_UPDATE` - Write session data only when the servlet calls the `IBMSession.sync()` method.
- `TIME_BASED_WRITE` (the default) - Write session data at periodic intervals, in seconds (called the *write interval*).

When a session is first created, session information is always written at the end of the `service()` call.

Base in-memory session pool size

The base in-memory session pool size number has different meanings, depending on session support configuration:

- With in-memory sessions, session access is optimized for up to this number of sessions.
- With distributed sessions (meaning, when sessions are stored in a database or in another WebSphere Application Server instance); it also specifies the cache size and the number of last access time updates saved in manual update mode.

For distributed sessions, when the session cache has reached its maximum size and a new session is requested, the Session Management facility removes the least recently used session from the cache to make room for the new one.

General memory requirements for the hardware system, and the usage characteristics of the e-business site, determines the optimum value.

Note that increasing the base in-memory session pool size can necessitate increasing the heap sizes of the Java processes for the corresponding WebSphere Application Servers.

Overflow in nondistributed sessions

By default, the number of sessions maintained in memory is specified by base in-memory session pool size. If you do not wish to place a limit on the number of sessions maintained in memory and allow overflow, set `overflow` to *true*.

Allowing an unlimited amount of sessions can potentially exhaust system memory and even allow for system sabotage. Someone could write a malicious program that continually hits your site and creates sessions, but ignores any cookies or encoded URLs and never utilizes the same session from one HTTP request to the next.

When overflow is disallowed, the Session Management facility still returns a session with the `HttpServletRequest.getSession(true)` method when the memory limit is reached, and this is an invalid session that is not saved.

With the WebSphere Application Server extension to HttpSession, com.ibm.websphere.servlet.session.IBMSession, an isOverflow() method returns true if the session is such an invalid session. An application can check this status and react accordingly.

Controlling write operations

You can manually control when modified session data is written out to the database or to another WebSphere Application Server instance by using the sync() method in the com.ibm.websphere.servlet.session.IBMSession interface, which extends the javax.servlet.http.HttpSession interface. By calling the sync() method from the service() method of a servlet, you send any changes in the session to the external location. When MANUAL_UPDATE is selected as the write frequency mode, session data changes are written to an external location only if the application calls the sync() method. If the sync() method is not called, session data changes are lost when a session object leaves the server cache. With END_OF_SERVICE or TIME_BASE_WRITE is the write frequency mode, the session data changes are written out whenever the sync() method is called. If the sync() method is not called, changes are written out at the end of service method or on a time interval basis based on the write frequency mode selected.

```
IBMSession iSession = (IBMSession) request.getSession();
iSession.setAttribute("name", "Bob");

//force write to external store
iSession.sync( )
```

Tuning parameter settings

Use this page to set tuning parameters for distributed sessions.

To view this administrative console page, click **Servers > Application Servers > server_name > Web Container > Session Management > Distributed Environment Settings > Custom Tuning Parameters**.

Tuning Level

Specifies that the Session Management facility provides certain predefined settings that affect performance.

Select one of these predefined settings or customize a setting.

To customize a setting, select one of the predefined settings that comes closest to the setting you want, click **Custom settings**, make your changes, and then click **OK**.

Predefined options and their settings follow:

Very high (optimize for performance):

Write frequency	TIME_BASED_WRITES
Write interval	300 sec
Write Contents	Only updated attributes
Schedule Sessions Clean up	true
First Hour	00

High:

Write frequency	TIME_BASED_WRITES
-----------------	-------------------

Write interval	300 sec
Write Contents	All session attributes

Medium:

Write frequency	END_OF_SERVICE
Write Contents	Only updated attributes

Low (optimize for failover)

Write frequency	END_OF_SERVICE
Write Contents	All session attributes

Write frequency

Specifies when the session writes to the database or another WebSphere Application Server instance happens.

End of servlet service	A session writes to a database or another WebSphere Application Server instance after the servlet completes execution.
Manual update	A programmatic sync on the IBMSession object is required to write the session data to the database or another WebSphere Application Server instance.
Time based	Session data writes to the database or another WebSphere Application Server instance based on the specified Write Interval value.

Write contents

Specifies whether updated attributes are only written to the external location or all of the session attributes are written to the external location, regardless of whether or not they changed. The external location can be either a database or another application server instance.

Schedule sessions cleanup

Specifies when to clean the invalid sessions from a database or another application server instance.

Specify sessions cleanup schedule

Enables the scheduled invalidation process for cleaning up the invalidated HTTP sessions from the external location. Enable this option to reduce the number of updates to a database or another application server instance required to keep the HTTP sessions alive. When this option is not enabled, the invalidator process runs every few minutes to remove invalidated HTTP sessions.

When this option is enabled, specify the two hours of a day for the process to clean up the invalidated sessions in the external location. Specify the times when there is the least activity in the application servers. An external location can be either a database or another application server instance.

First Time of Day

Indicates the first hour during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

Second Time of Day

Indicates the second hour during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

Best practices for using HTTP Sessions

- **Release HttpSession objects using `javax.servlet.http.HttpSession.invalidate()` when finished.**

HttpSession objects live inside the Web container until:

- The application explicitly and programmatically releases it using the `javax.servlet.http.HttpSession.invalidate()` method; quite often, programmatic invalidation is part of an application logout function.
- WebSphere Application Server destroys the allocated HttpSession when it expires (default = 1800 seconds or 30 minutes). The WebSphere Application Server can only maintain a certain number of HTTP sessions in memory based on Session Management settings. In case of distributed sessions, when maximum cache limit is reached in memory, the Session Management facility removes the least recently used (LRU) one from cache to make room for a session.

- **Avoid trying to save and reuse the HttpSession object outside of each servlet or JSP file.**

The HttpSession object is a function of the HttpRequest (you can get it only through the `req.getSession()` method), and a copy of it is valid only for the life of the `service()` method of the servlet or JSP file. You *cannot* cache the HttpSession object and refer to it outside the scope of a servlet or JSP file.

- **Implement the Serializable class when developing new objects to be stored in the HTTP session.**

This action allows the object to properly serialize when using distributed sessions. Without this extension, the object cannot serialize correctly and throws an error. An example of this follows:

```
public class MyObject implements java.io.Serializable {...}
```

Make sure all instance variable objects that are not marked transient are serializable.

- **The HttpSession API does not dictate transactional behavior for sessions.**

Distributed HttpSession support does not guarantee transactional integrity of an attribute in a failover scenario or when session affinity is broken. Use transactionally aware resources like enterprise Java beans to guarantee the transaction integrity required by your application.

- **Ensure the Java objects you add to a session are in the correct class path.**

If you add Java objects to a session, place the class files for those objects in the correct classpath (the Application Classpath if utilizing sharing across Web modules in an enterprise application, or the WebModule Classpath if using the Servlet 2.2-complaint session sharing) or in the directory containing other servlets used in WebSphere Application Server. In the case of session clustering, this action applies to every node in the cluster.

Because the HttpSession object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

- **Avoid storing large object graphs in the HttpSession object.**

In most applications each servlet only requires a fraction of the total session data. However, by storing the data in the HttpSession object as one large object, an application forces WebSphere Application Server to process all of it each time.

- **Utilize Session Affinity to help achieve higher cache hits in the WebSphere Application Server.**

WebSphere Application Server has functionality in the HTTP Server plug-in to help with session affinity. The plug-in will read the cookie data (or encoded URL) from the browser and helps direct the request to the appropriate application or clone based on the assigned session key. This functionality increases use of the in-memory cache and reduces hits to the database or another WebSphere Application Server instance

- **Maximize use of session affinity and avoid breaking affinity.**

Using session affinity properly can enhance the performance of the WebSphere Application Server. Session affinity in the WebSphere Application Server environment is a way to maximize the in-memory cache of session objects and reduce the amount of reads to the database or another WebSphere Application Server instance. Session affinity works by caching the session objects in the server instance of the application with which a user is interacting. If the application is deployed in multiple servers of a server group, the application can direct the user to any one of the servers. If the user starts on server1 and then comes in on server2 a little later, the server must write all of the session information to the external location so that the server instance in which server2 is running can read the database. You can avoid this database read using session affinity. With session affinity, the user starts on server1 for the first request; then for every successive request, the user is directed back to server1. Server1 has to look only at the cache to get the session information; server1 never has to make a call to the session database to get the information.

You can improve performance by not breaking session affinity. Some suggestions to help avoid breaking session affinity are:

- Combine all Web applications into a single application server instance, if possible, and use modeling or cloning to provide failover support.
- Create the session for the frame page, but do not create sessions for the pages within the frame when using multiframe JSP files. (See discussion later in this topic.)

- **When using multi-framed pages, follow these guidelines:**

- Create a session in only one frame or before accessing any frame sets. For example, assuming there is no session already associated with the browser and a user accesses a multi-framed JSP file, the browser issues concurrent requests for the JSP files. Because the requests are not part of any session, the JSP files end up creating multiple sessions and all of the cookies are sent back to the browser. The browser honors only the last cookie that arrives. Therefore, only the client can retrieve the session associated with the last cookie. Creating a session before accessing multi-framed pages that utilize JSP files is recommended.
- By default, JSPs get a HttpSession using `request.getSession(true)` method. So by default JSPs create a new session if none exists for the client. Each JSP page in the browser is requesting a new session, but only one session is used per browser instance. A developer can use `<% @ page session="false" %>` to turn off the automatic session creation from the JSP files that will not access the session. Then if the page needs access to the session information, the

developer can use `<%HttpSession session = javax.servlet.http.HttpServletRequest.getSession(false); %>` to get the already existing session that was created by the original session creating JSP file. This action helps prevent breaking session affinity on the initial loading of the frame pages.

- Update session data using only one frame. When using framesets, requests come into the HTTP server concurrently. Modifying session data within only one frame so that session changes are not overwritten by session changes in concurrent frameset is recommended.
- Avoid using multi-framed JSP files where the frames point to different Web applications. This action results in losing the session created by another Web application because the JSESSIONID cookie from the first Web application gets overwritten by the JSESSIONID created by the second Web application.
- **Secure all of the pages (not just some) when applying security to servlets or JSP files that use sessions with security integration enabled, .**

When it comes to security and sessions, it is all or nothing. It does not make sense to protect access to session state only part of the time. When security integration is enabled in the Session Management facility, all resources from which a session is created or accessed must be either secured or unsecured. You cannot mix secured and unsecured resources.

The problem with securing only a couple of pages is that sessions created in secured pages are created under the identity of the authenticated user. Only the same user can access sessions in other secured pages. To protect these sessions from use by unauthorized users, you cannot access these sessions from an unsecured page. When a request from an unsecured page occurs, access is denied and an `UnauthorizedSessionRequestException` error is thrown.

(`UnauthorizedSessionRequestException` is a runtime exception; it is logged for you.)

- **Use manual update and either the `sync()` method or time-based write in applications that read session data, and update infrequently.**

With `END_OF_SERVICE` as write frequency, when an application uses sessions and anytime data is read from or written to that session, the `LastAccess` time field updates. If database sessions are used, a new write to the database is produced. This activity is a performance hit that you can avoid using the Manual Update option and having the record written back to the database only when data values update, not on every read or write of the record.

To use manual update, turn it on in the Session Management Service. (See the tables above for location information.) Additionally, the application code must use the `com.ibm.websphere.servlet.session.IBMSession` class instead of the generic `HttpSession`. Within the `IBMSession` object there is a method called `sync()`. This method tells the WebSphere Application Server to write the data in the session object to the database. This activity helps the developer to improve overall performance by having the session information persist only when necessary.

Note: An alternative to using the manual updates is to utilize the timed updates to persist data at different time intervals. This action provides similar results as the manual update scheme.

- Implement the following suggestions to achieve high performance:
 - If your applications do not change the session data frequently, use Manual Update and the `sync()` function (or timed interval update) to efficiently persist session information.








- Keep the amount of data stored in the session as small as possible. With the ease of using sessions to hold data, sometimes too much data is stored in the session objects. Determine a proper balance of data storage and performance to effectively use sessions.
- Use a dedicated database for the session database. Avoid using the application database. This helps to avoid contention for JDBC connections and allows for better database performance.
- Verify that you have the latest e-fixes for the WebSphere Application Server.
- Utilize the following tools to help monitor session performance.
 - Run the `com.ibm.servlet.personalization.sessiontracking.IBMTrackerDebug` servlet. - To run this servlet, you must have the servlet invoker running in the Web application you want to run this from. Or, you can explicitly configure this servlet in the application you want to run.
 - Use the WebSphere Application Server Resource Analyzer which comes with WebSphere Application Server to monitor active sessions and statistics for the WebSphere Application Server environment.
 - Use database tracking tools such as "Monitoring" in DB2. (See the respective documentation for the database system used.)

Managing HTTP sessions: Resources for learning:

Use the following links to find relevant supplemental information about HTTP sessions. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Programming specifications
-  **Best practices** (<http://www7b.software.ibm.com/wsdd/zones/bp/>)
-  **HTTP Session Persistence Best Practices**
(http://www7b.software.ibm.com/wsdd/library/techarticles/0209_draeger/draeger.html)
-  **Improving session persistence performance with DB2**
(<http://www7b.software.ibm.com/dmdd/library/techarticle/0203kitchlu/0203kitchlu.html>)
-  **Persistent client state HTTP cookies specification**
(http://www.netscape.com/newsref/std/cookie_spec.html)
-  **Java Servlet documentation, tutorials, and examples site**
(<http://java.sun.com/products/servlet/docs.html>)
-  **Java Servlet 2.3 API specification download site**
(<http://java.sun.com/products/servlet/download.html>)
-  **J2EE 1.3 specification download site**
(<http://java.sun.com/j2ee/download.html>)

Chapter 4. Using enterprise beans in applications

Steps for this task

1. Design a J2EE application and the enterprise beans that it needs.
See “Enterprise beans: Resources for learning” on page 149 for links to design information that is specific to enterprise beans.
2. Develop any enterprise beans that your application will use.
3. Prepare for assembly. For your EJB 2.x-compliant entity beans, decide on an appropriate access intent policy.
4. Assemble the beans into one or more EJB modules. This includes “Securing EJB applications” (not in this document).
5. Assemble the modules into a J2EE application.
6. For a given application server, update the EJB container configuration if needed for the application to be deployed.
7. Deploy the application in an application server.
8. Test the modules.
 - As needed, debug problems with the container.
 - Debug access and deployment problems.
9. Assemble the production application.
10. Deploy the application to a production environment.
11. Manage the application:
 - a. Manage installed EJB modules.
After an application has been installed, you can manage its EJB modules individually through administrative console settings.
 - b. Manage other aspects of the J2EE application.
12. Update the module and redeploy it.
13. Tune the application.

Enterprise beans

An enterprise bean is a Java component that can be combined with other resources to create J2EE applications. There are three types of enterprise beans, *entity* beans, *session* beans, and *message-driven* beans.

All beans reside in EJB containers, which provide an interface between the beans and the application server on which they reside.

Entity beans store permanent data. Entity beans with container-managed persistence (CMP) require connections to a form of persistent storage. This storage might be a database, an existing legacy application, a file, or other types of persistent storage. Entity beans with bean-managed persistence manage permanent data in whichever manner is defined in the bean code. This can include writing to databases or XML files, for example.

Session beans do not require database access, although they can obtain it indirectly as needed through entity beans. Session beans can also obtain direct access to databases (and other resources) through the use of resource references. Session beans can be either *stateful* or *stateless*.

New in the Enterprise JavaBeans (EJB) specification, version 2.0, message-driven beans enable asynchronous message servicing. The EJB container and a Java Message Service (JMS) provider work together to process messages. When a message arrives from another application component through JMS, the EJB container forwards it through an `onMessage()` call to a message-driven bean instance, which then processes the message. In other respects, message-driven beans are similar to stateless session beans.

Beans that require data access use *data sources*, which are administrative resources that define pools of connections to persistent storage mechanisms.

For more information about enterprise beans, see “Enterprise beans: Resources for learning” on page 149.

Developing enterprise beans

Before you begin

Design a J2EE application and the enterprise beans that it needs.

- For general design information, see “Enterprise beans: Resources for learning” on page 149.
- Before developing entity beans with container-managed persistence (CMP), read “Concurrency control” on page 106.

There are two basic approaches to selecting tools for developing enterprise beans:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the enterprise bean code and contain integrated tools for packaging and testing enterprise beans. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you have decided to develop enterprise beans without an IDE, you need at least an ASCII text editor. You can also use a Java development tool that does not support enterprise bean development. You can then use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the beans.

The following steps primarily support the second approach, development without an IDE.

Steps for this task

1. If necessary, migrate any pre-existing code to the required version of the Enterprise JavaBeans specification.
2. Write and compile the components of the enterprise bean.
 - At a minimum, an EJB 1.1 session bean requires a bean class, a home interface, and a remote interface. An EJB 1.1 entity bean requires a bean class, a primary-key class, a home interface, and a remote interface.
 - At a minimum, an EJB 2.0 session bean requires a bean class, a home or local home interface, and a remote or local interface. An EJB 2.0 entity bean requires a bean class, a primary-key class, a remote home or local home interface, and a remote or local interface. The types of interfaces go together: If you implement a local interface, you must define a local home interface as well.
 - Available only through EJB 2.0, a message-driven bean requires only a bean class.

3. **(Optional) (CMP entity beans for EJB 1.1 only: an IBM extension)** Create a finder helper interface for each CMP entity bean that contains specialized finder methods (other than the `findByPrimaryKey` method).

The following logic is required for each finder method (other than the `findByPrimaryKey` method) contained in the home interface of an entity bean with CMP:

- The logic must be defined in a public interface named *NameBeanFinderHelper*, where *Name* is the name of the enterprise bean (for example, *AccountBeanFinderHelper*).
 - The logic must be contained in a String constant named *findMethodNameWhereClause*, where *findMethodName* is the name of the finder method. The String constant can contain zero or more question marks (?) that are replaced from left to right with the value of the finder method's arguments when that method is called.
4. **(CMP entity beans for EJB 2.0 only)** Define finder queries with EJB Query Language (EJB QL).
With EJB QL, you define finders in terms of CMP fields and container-managed relationships, as follows:
 - *Public* finders are visible in the bean's home interface. Implemented in the bean class, they return only remote interfaces and collection types.
 - *Private* finders, expressed as SELECT statements, are used only within the bean class. They can return both local and remote interfaces, dependent values, other CMP field types, and collection types.
 5. **(Entity beans only)** Create a database schema for the entity bean's persistent data.
 - For entity beans with CMP, you must store the bean's persistent data in one of the supported databases. The Application Assembly Tool automatically generates SQL code for creating database tables for CMP entity beans. If your CMP beans require complex database mappings, it is recommended that you use the WSAD product to generate code for the database tables.
 - For entity beans with bean-managed persistence (BMP), you can create the database and database table by using the database tools or use an existing database and database table.

For more information on creating databases and database tables, consult your database documentation.

What to do next

Assemble the beans in one or more EJB modules.

Migrating enterprise bean code to the supported specification

Support for Version 2.0 of the Enterprise JavaBeans (EJB) specification is new for Version 5 of this product. Migration of enterprise beans deployed in Version 4.0.x of this product is not generally necessary; Version 1.1 of the EJB specification is still supported. Follow these steps as appropriate for your application deployment.

Steps for this task

1. Modify enterprise bean code for changes in the specification.
 - For Version 1.0 beans, migrate at least to Version 1.1.

- As stated previously, migration from Version 1.1 to Version 2.0 of the EJB specification is not required for redeployment on this version of the product. However, if your application requires the capabilities of Version 2.0, migrate your Version 1.1-compliant code.

Note: The EJB Version 2.0 specification mandates that prior to the EJB container's executing a *findByMethod* query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

2. Modify enterprise bean code for changes in deployment requirements.

If the enterprise beans were previously deployed in Version 3.0.x of this product, modify import statements to match standard package names. In Version 3.0.2.x, the following standard packages were present under nonstandard names:

```
javax.sql.*
javax.transaction.*
```

Any code using WebSphere data sources, including BMP entity beans and session beans that access databases, must be modified.

3. You might have to modify code for some EJB 1.1-compliant modules that were not migrated to Version 2.0. Use the following information to help you decide.
 - Some stub classes for deployed enterprise beans have changed in the Java 2 SDK, Version 1.3.
 - The task of generating deployment code for enterprise beans changed significantly for EJB 1.1-compliant modules relative to EJB 1.0-compliant modules.
 - If the CMP beans write to databases with mixed-case table or column names and you used IBM VisualAge for Java, Version 3.5.x, to generate the original deployment code, you cannot simply reassemble the beans in this product. You must export the original EJB project from the VisualAge for Java product as an EJB 1.1 JAR. This preserves the metadata needed to generate the correct deployment code for mixed-case database tables and columns. For more information, see the documentation for the Deployment Tool for Enterprise JavaBeans.

For detailed information about source and binary compatibility between deployed versions, see "Enterprise beans: Resources for learning" on page 149.

4. Reassemble and redeploy all modules to incorporate migrated code.

Migrating enterprise bean code from Version 1.0 to Version 1.1

The following information generally applies to any enterprise bean that currently complies with Version 1.0 of the Enterprise JavaBeans (EJB) specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Enterprise beans: Resources for learning" on page 149.

Steps for this task

1. In session beans, replace all uses of `javax.jts.UserTransaction` with `javax.transaction.UserTransaction`. Entity beans may no longer use the `UserTransaction` interface at all.
2. In finder methods for entity beans, include `FinderException` in the `throws` clause.
3. Remove `throws` of `java.rmi.RemoteException`; throw `javax.ejb.EJBException` instead.

However, continue to include `RemoteException` in the `throws` clause of home and remote interfaces as required by the use of Remote Method Invocation (RMI).

4. Remove uses of the `finalize()` method.
5. Replace calls to `getCallerIdentity()` with calls to `getCallerPrincipal()`.
The use of `getCallerIdentity()` is deprecated.
6. Replace calls to `isCallerInRole(Identity)` with calls to `isCallerInRole(String)`.
The use of `isCallerInRole(Identity)` and `java.security.Identity` is deprecated.
7. Replace calls to `getEnvironment()` in favor of JNDI lookup.

As an example, change the following code:

```
String homeName =
    aLink.getEntityContext().getEnvironment().getProperty("TARGET_HOME_NAME");
if (homeName == null) homeName = "TARGET_HOME_NAME";
```

The updated code would look something like the following:

```
Context env = (Context)(new InitialContext()).lookup("java:comp/env");
String homeName = (String)env.lookup("ejb10-properties/TARGET_HOME_NAME");
```

8. In `ejbCreate` methods for an entity bean with container-managed persistence (CMP), return the bean's primary key class instead of `void`.
9. Add the `getHomeHandle()` method to home interfaces.
`public javax.ejb.HomeHandle getHomeHandle() {return null;}`

What to do next

Consider enhancements to match the following changes in the specification:

- Primary keys for entity beans can be of type `java.lang.String`.
- Finder methods for entity beans return `java.util.Collection`.
- Check the format of any JNDI names being used. Local name spaces are also supported.
- Security is defined by role, and isolation levels are defined at the method level rather than at the bean level.

Migrating enterprise bean code from Version 1.1 to Version 2.0

Enterprise JavaBeans (EJB) Version 2.0-compliant beans may be assembled only in an EJB 2.0-compliant module, although an EJB 2.0-compliant module can contain a mixture of Version 1.x and Version 2.0 beans.

The EJB Version 2.0 specification mandates that prior to the EJB container's executing a `findByMethod` query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

The following information generally applies to any enterprise bean that currently complies with Version 1.1 of the EJB specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see “Enterprise beans: Resources for learning” on page 149.

Steps for this task

1. In beans with container-managed persistence (CMP) version 1.x, replace each CMP field with abstract get and set methods.
In doing so, you must make each bean class abstract.
2. In beans with CMP version 1.x, change all occurrences of `this.field = value` to `setField(value)`.
3. In each CMP bean, create abstract get and set methods for the primary key.
4. In beans with CMP version 1.x, create an EJB Query Language statement for each finder method.
5. In finder methods for beans with CMP version 1.x, return `java.util.Collection` instead of `java.util.Enumeration`.
6. Update handling of non-application exceptions.
 - To report non-application exceptions, throw `javax.ejb.EJBException` instead of `java.rmi.RemoteException`.
 - Modify rollback behavior as needed: In EJB versions 1.1 and 2.0, all non-application exceptions thrown by the bean instance result in the rollback of the transaction in which the instance is running; the instance is discarded. In EJB 1.0, the container does not roll back the transaction or discard the instance if it throws `java.rmi.RemoteException`.
7. Update rollback behavior as the result of application exceptions.
 - In EJB versions 1.1 and 2.0, an application exception does not cause the EJB container to automatically roll back a transaction.
 - In EJB Version 1.1, the container performs the rollback only if the instance has called `setRollbackOnly()` on its `EJBContext` object.
 - In EJB Version 1.0, the container is required to roll back a transaction when an application exception is passed through a transaction boundary started by the container.

WebSphere extensions to the Enterprise JavaBeans specification

This article outlines extensions to the Enterprise JavaBeans (EJB) specification that IBM provides with this product:

Inheritance in enterprise beans

In the Java language, *inheritance* is the creation of a new class from an existing class or a new interface from an existing interface. This product supports two forms of inheritance: standard class inheritance and EJB inheritance.

In standard class inheritance, the home interface, remote interface, or enterprise bean class inherits properties and methods from base classes that are not themselves enterprise bean classes or interfaces.

By contrast in enterprise bean inheritance, an enterprise bean inherits properties (such as container-managed persistence (CMP) fields and container-managed relationship (CMR) fields), methods, and method-level control descriptor attributes from another enterprise bean.

For more information, see the documentation for the IBM WebSphere Studio Application Developer product.

Optimistic concurrency control for container-managed persistence

This product supports optimistic concurrency control of data access.

Access intents for EJB persistence

This product supports the application of named data-access policies at the method level.

Performance enhancements

Through the lifetime-in-cache settings, this product provides a way for you to improve performance for beans that are only occasionally updated. For more information, see “Entity bean assembly settings” on page 119.

Some enterprise beans created with the IBM WebSphere Studio Application Developer product can utilize *read-ahead* for loading a bean and its related beans in a single database operation. An entire object graph or any part of the graph can be preloaded by configuring a finder method to use read-ahead.

Assembly and deployment extensions

This product supports IBM extensions of assembly and deployment options. IBM extensions are clearly marked in reference topics for assembly settings.

Best practices for developing enterprise beans

Use the following guidelines when designing and developing enterprise beans:

- Use a stateless session bean to act as the entry point for business logic. For more information about using session facades, see “Enterprise beans: Resources for learning” on page 149.
- Entity beans should use container-managed persistence.
- In an Enterprise JavaBeans (EJB) Version 2.0 environment, use local interfaces to improve communication between enterprise beans in the same Java virtual machine.

Local calls avoid the overhead of RMI/IIOP and use pass-by-reference semantics instead of pass-by-value. For each call, the caller and callee beans share the state of arguments. EJB 2.0 beans can have both a local and remote interface but more typically have one or the other.

- For communicating with remote clients, provide remote and remote home interfaces. For communicating with local clients like servlets, entity beans, and message-driven beans, provide local and local home interfaces.

Using access intent policies

You can use access intent policies to help the product run-time environment manage various aspects of Enterprise JavaBeans (EJB) persistence. You apply access intent policies to methods of EJB Version 2.0 entity beans by using the Application Assembly Tool. This product provides a set of default access intent policies.

Steps for this task

1. Apply access intent policies to methods of CMP entity beans.

Access intent policies

An access intent policy is a named set of properties (access intents) that governs data access for Enterprise JavaBeans (EJB) persistence. You can assign a policy to individual methods on an entity bean's home, remote, or local interfaces during assembly. Access intents are settable only within EJB Version 2.x-compliant modules for entity beans with CMP Version 2.x.

This product supplies a number of access intent policies that specify permutations of read intent and concurrency control; the pessimistic/update policy can be qualified further. The selected policy determines the appropriate isolation level and locking strategy used by the run-time environment.

Access intent policies are specifically designed to supplant the use of isolation level and access intent method-level modifiers found in the extended deployment descriptor for EJB version 1.1 enterprise beans. You cannot specify isolation level and read-only modifiers for EJB version 2.0 enterprise beans.

Access intent policies are named and defined at the module level. A module can have one or many such policies. Policies are assigned, and apply, to individual methods of the declared interfaces of entity beans and their associated home interfaces. A policy is acted upon by the combination of the EJB container and persistence manager.

For entity beans that are backed by tables with nullable columns, use an optimistic policy with caution. Nullable columns are automatically excluded from overqualified updates at deployment time; concurrent changes to a nullable field might result in lost updates. When used with the IBM WebSphere Studio Application Developer product, this product provides support for selecting a subset of the nonnullable columns that are to be reflected in the overqualified update statement that is generated in the deployment code to support optimistic policies.

A method that is configured with a read-only policy that causes a bean to be activated can cause problems if updates are attempted within the same transaction. Those changes will not be committed, and an exception will be thrown because data integrity might be compromised.

Concurrency control

Concurrency control is the management of contention for data resources. A concurrency control scheme is considered *pessimistic* when it locks a given resource early in the data-access transaction and does not release it until the transaction is closed. A concurrency control scheme is considered *optimistic* when locks are acquired and released over a very short period of time at the end of a transaction.

The objective of optimistic concurrency is to minimize the time over which a given resource would be unavailable for use by other transactions. This is especially important with long-running transactions, which under a pessimistic scheme would lock up a resource for unacceptably long periods of time.

Under an optimistic scheme, locks are obtained immediately before a read operation and released immediately afterwards. Update locks are obtained immediately before an update operation and held until the end of the transaction.

To enable optimistic concurrency, this product uses an *overqualified update scheme* to test whether the underlying data source has been updated by another transaction since the beginning of the current transaction. With this scheme, the columns marked for update and their original values are added explicitly through a WHERE clause in the UPDATE statement so that the statement fails if the underlying column values have been changed. As a result, this scheme can provide column-level concurrency control; pessimistic schemes can control concurrency at the row level only.

Optimistic schemes typically perform this type of test only at the end of a transaction. If the underlying columns have not been updated since the beginning of the transaction, pending updates to container-managed persistence fields are committed and the locks are released. If locks cannot be acquired or if some other transaction has updated the columns since the beginning of the current transaction, the transaction is rolled back: All work performed within the transaction is lost.

Pessimistic and optimistic concurrency schemes require different transaction isolation levels. Enterprise beans that participate in the same transaction and require different concurrency control schemes cannot operate on the same underlying data connection.

Whether or not to use optimistic concurrency depends on the type of transaction. Transactions with a high penalty for failure might be better managed with a pessimistic scheme. (A high-penalty transaction is one for which recovery would be risky or resource-intensive.) For low-penalty transactions, it is often worth the risk of failure to gain efficiency through the use of an optimistic scheme. In general, optimistic concurrency is more efficient when update collisions are expected to be infrequent; pessimistic concurrency is more efficient when update collisions are expected to occur often.

Read-ahead hints

Read-ahead schemes enable applications to minimize the number of database roundtrips by retrieving a working set of container-managed persistence (CMP) beans for the transaction within one query. Read-ahead involves activating the requested CMP beans and caching the data for their related beans, which ensures that data is present for the beans that are most likely to be needed next by an application. A *read-ahead hint* is a canonical representation of the related beans that are to be read. It is associated with a finder method for the requested bean type, which must be an EJB 2.x-compliant CMP entity bean.

Read-ahead hints can be set only through the Add Access Intent wizard of the IBM WebSphere Studio Application Developer product. In the wizard, the **Read Ahead Hint** check box is enabled only with access intent policies with optimistic concurrency.

Read-ahead is limited to optimistic policies because locking persistent data store for all beans represented in the hint would be more likely to cause lock conflicts, and optimistic policies do not obtain locks until immediately before the database operation.

Currently, only `findByPrimaryKey` methods can have read-ahead hints. Only beans related to the requested beans by a container-managed relationship (CMR), either directly or indirectly through other beans, can be read ahead.

A read-ahead hint takes the form of a character string. You do not have to provide the string; the wizard generates it for you based on CMRs defined for the bean. The following example is provided as supplemental information only.

Suppose a CMP bean type A has a finder method that returns instances of bean A. A read-ahead hint for this method is specified using the following notation:
RelB.RelC; RelD

Interpret the preceding notation as follows:

- Bean type A has a CMR with bean types B and D.
- Bean type B has a CMR with bean type C.

For each bean of type A that is retrieved from the database, its directly-related B and D beans and its indirectly-related C beans are also retrieved. The order of the retrieved bean data columns in each row of the result set is the same as their order in the read-ahead hint: an A bean, a B bean (or null), a C bean (or null), a D bean (or null). For hints in which the same relationship is mentioned more than once (for example, *RelB.RelC;RelB.RelE*), a bean's data columns appear only once, at the position it first appears in the hint.

The tokens shown in the notation (*RelB* and so on) must be CMR field names for the relationships as defined in the deployment descriptor for the bean. In indirect relationships such as *RelB.RelC*, *RelC* is a CMR field name defined in the deployment descriptor for bean type B.

For more information about how to set read-ahead hints, see the documentation for the WebSphere Studio Application Developer product.

Applying access intent policies to methods

You apply an access intent policy to a method, or set of methods, in an application's entity beans through the Application Assembly Tool (AAT).

Steps for this task

1. Start the AAT.
2. Create or edit the application EAR file.
For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.
3. Select **EJB Modules > *moduleName* > Access Intent**.
4. To configure a new access intent policy, right-click and select **New**.
5. On the **New Access Intent** panel, specify a name and a description.
These attributes are provided as a convenience to the developer and are not used at run time.
6. To select the methods to which the access intent policy should apply, click **Add** beside the Methods table.

7. From the **Applied access intent** list, select an access intent policy.
8. **(Optional)** To override an attribute defined in the applied policy, click **Add** beside the Access intent attribute overrides table.
9. Click **OK** to exit the New Access Intent panel.
10. Save your configuration by selecting **File > Save**.

Access intent exceptions

The following exceptions are thrown in response to the application of access intent policies:

com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException

If the method that drives the `ejbLoad()` method is configured to be read-only but updates are then made within the transaction that loaded the bean's state, an exception is thrown during invocation of the `ejbStore()` method, and the transaction is rolled back. Likewise, the `ejbRemove()` method cannot succeed in a transaction that is set as read-only. If an update hint is applied to methods of entity beans with bean-managed persistence, the same behavior and exception results. The forwarded exception object contains the message string `PMGR1103E: update instance level read only bean beanName`

This exception is also thrown if the applied access intent policy cannot be honored because a finder, `ejbSelect`, or container-managed relationship (CMR) accessor method returns an inherently read-only result. The forwarded exception object contains the message string `PMGR1001: No such DataAccessSpec - methodName`

The most common occurrence of this error is when a custom finder that contains a read-only EJB Query Language (EJB QL) statement is called with an applied access intent of `wsPessimisticUpdate` or `wsPessimisticUpdate-Exclusive`. These policies require the use of a FOR UPDATE clause on the SQL SELECT statement to be executed, but a read-only query cannot support FOR UPDATE. Other examples of read-only queries include joins; the use of ORDER BY, GROUP BY, and DISTINCT keywords.

To eliminate the exception, edit the EJB query so that it does not return an inherently read-only result or change the access intent policy being applied.

- If an update access is required, change the applied access intent setting to `wsPessimisticUpdate-WeakestLockAtLoad` or `wsOptimisticUpdate`.
- If update access is not truly required, use `wsPessimisticRead` or `wsOptimisticRead`.
- If connection sharing between entity beans is required, use `wsPessimisticUpdate-WeakestLockAtLoad` or `wsPessimisticRead`.

com.ibm.websphere.ejb.container.CollectionCannotBeFurtherAccessed

If a lazy collection is driven after it is no longer in scope, and beyond what has already been locally buffered, a `CollectionCannotBeFurtherAccessed` exception is thrown.

com.ibm.ws.exception.RuntimeWarning

If an application is configured incorrectly, a run-time warning exception is thrown as the application starts; startup is ended. You can validate an application's configuration by choosing the verify function in the WebSphere Application Assembly Tool. Some examples of misconfiguration include:

- A method configured with two different access intent policies

- A method configured with an undefined access intent policy

javax.ejb.NoSuchEntityException

If an update fails under optimistic concurrency because fields changed within another transaction between load and store requests, a `NoSuchEntityException` is raised and the commit fails.

Depending on the access intent policy used, an exception might also be thrown on an entity bean with container-managed persistence (CMP) version 2.0 that has one or more CMR fields. If this occurs, the default setting (`wsPessimisticUpdateWeakestLockAtLoad`) can be used.

If application needs require that another policy be used, the application code must be changed so that all CMR accessor methods are called from business logic within the CMP entity bean rather than from a session bean. Furthermore, an entity bean cannot be added or removed from a CMR collection unless the add or remove method of the CMR collection is called from business logic in the entity bean rather than from a session bean.

Access intent assembly settings

Access intent policies contain data-access settings for use by the persistence manager. Specify one or more methods and associate an access intent policy with each method.

These settings are applicable only for EJB 2.x-compliant entity beans that are packaged in EJB 2.x-compliant modules. Connection sharing between beans with bean-managed persistence and those with container-managed persistence is possible if they all use the same access intent policy.

Name

Specifies a name for the mapping between an access intent policy and one or more methods.

Description

Contains text that describes the mapping.

Methods - Name

Specifies the name of an enterprise bean method, or the asterisk character (*). The asterisk is used to denote all of the methods of an enterprise bean's remote and home interfaces.

Methods - Enterprise bean

Specifies which enterprise bean contains the methods indicated in the Name setting.

Methods - Type

Used to distinguish between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if an access intent policy applies to all methods of the bean.

Data type	String
Range	Valid values are Home, Remote, Local, LocalHome or Unspecified

Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

Applied access intent

Specifies how the container must manage data access for persistence.

Data type	String
Default	wsPessimisticUpdate-WeakestLockAtLoad. However, this policy cannot be used with Oracle; see the table that follows.
Range	Valid settings are wsPessimisticUpdate, wsPessimisticUpdate-NoCollision, wsPessimisticUpdate-Exclusive, wsPessimisticUpdate-WeakestLockAtLoad, wsPessimisticRead, wsOptimisticUpdate, or wsOptimisticRead. Only wsPessimisticRead and wsOptimisticRead are valid when class-level caching is enabled in the EJB container.

This product supports lazy collections. For each segment of a collection, iterating through the collection (*next()*) does not trigger a remote method call to retrieve the next remote reference. Two policies (wsPessimisticUpdate and wsPessimisticUpdate-Exclusive) are extremely lazy; the collection increment size is set to 1 to avoid overlocking the application. The other policies have a collection increment size of 25.

If a method is not configured with an access intent policy, the run-time environment typically uses wsPessimisticUpdate-WeakestLockAtLoad by default. If, however, the Lifetime in cache property is set on the bean, the default value of **Applied access intent** is wsOptimisticRead; updates are not permitted. If a method of a **Lifetime in cache**-configured bean is configured with an access intent policy that permits updates, the application will not run until the method or bean is reconfigured.

Additional information about valid settings follows:

Profile name	Concurrency control	Access type	Transaction isolation
wsPessimisticRead (Note 1)	pessimistic	read	For Oracle, read committed. Otherwise, repeatable read
wsPessimisticUpdate (Note 2)	pessimistic	update	For Oracle, read committed. Otherwise, repeatable read
wsPessimisticUpdate-Exclusive (Note 3)	pessimistic	update	serializable
wsPessimisticUpdate-NoCollision (Note 4)	pessimistic	update	read committed
wsPessimisticUpdate-WeakestLockAtLoad (Note 5)	pessimistic	update	Repeatable read
wsOptimisticRead	optimistic	read	read committed
wsOptimisticUpdate (Note 6)	optimistic	update	read committed

Profile name	Concurrency control	Access type	Transaction isolation
<p>Notes:</p> <ol style="list-style-type: none"> 1. Read locks are held for the duration of the transaction. 2. The generated SELECT FOR UPDATE query grabs locks at the beginning of the transaction. 3. SELECT FOR UPDATE is generated; locks are held for the duration of the transaction. 4. A plain SELECT query is generated. No locks are held, but updates are permitted. Relative to <code>wsPessimisticUpdate</code>, this difference results in generally better transaction throughput. 5. The generated SELECT query does not include FOR UPDATE; locks are escalated by the persistent store at storage time if updates were made. Do not use this policy with Oracle; doing so results in a <code>NoSuchDataAccessSpec</code> exception. Comparable alternatives are <code>wsPessimisticUpdate-NoCollision</code> or <code>wsOptimisticUpdate</code>. If you choose <code>wsOptimisticUpdate</code>, be sure to review the rules for forming overqualified-update query predicates. Certain column types (for example, BLOB) are ineligible for inclusion in the overqualified-update query predicate and might affect your design. 6. Generated overqualified-update query forces failure if CMP column values have changed since the beginning of the transaction. 			

Access intent best practices

This topic outlines issues to consider when applying access intent policies to Enterprise JavaBeans (EJB) methods.

- **Start with defaults.** The default access intent policy (`wsPessimisticUpdate-WeakestLockAtLoad`) loads persistent data with the weakest lock that is supported by the persistent store (typically a read lock). Updates are allowed, and the database is permitted to undertake lock escalation when necessary. This option generally works best for most EJB application patterns. After your application is built and running, you can more finely tune certain access paths in your application.
- **Don't mix access types.** Avoid using both pessimistic and optimistic policies in the same transaction. For most databases, pessimistic and optimistic policies use different isolation levels. This results in multiple database connections, which prevents you from taking advantage of the performance benefits possible through connection sharing.
- **Access intent for the `ejbSelect` method must be applied indirectly.** Because `ejbSelect` methods are not exposed through a home, remote, or local interface, you cannot apply a policy to them directly. An `ejbSelect` method is called by a home or business method, so apply the appropriate policy to the home or business method that governs the behavior of the `ejbSelect` method.
- **Take care when applying `wsPessimisticUpdate-NoCollision`.** This policy does not ensure data integrity. No database locks are held, so concurrent transactions can overwrite each other's updates. Use this policy only if you can be sure that only one transaction will attempt to update persistent store at any given time.

Frequently asked questions: Access intent

I have not applied any access intent policies at all. My application runs just fine with a DB2 database, but it fails with an Oracle database with the following message: `com.ibm.ws.ejbpersistence.utilpm.PersistanceManagerException: PMGR1001E: No such DataAccessSpec :FindAllCustomers. The backend datastore does not support the SQLStatement needed by this AccessIntent: (pessimistic`

update-weakestLockAtLoad)(collections: transaction/25) (resource manager prefetch: 0) (AccessIntentImpl@d23690a). Why?

If you have not configured access intent, all of your data is accessed under the default access intent policy (`wsPessimisticUpdate-WeakestLockAtLoad`). On DB2 databases, the weakest lock is a shared one, and the query runs without a FOR UPDATE clause. On Oracle databases, however, the weakest lock is an update lock; this means that the SQL query must contain a FOR UPDATE clause. However, not every SQL statement necessarily supports FOR UPDATE; for example, if the query is being run against multiple tables in a join, FOR UPDATE is not supported.

To avoid this problem, try either of the following:

- Modify your SQL query or reconfigure your application so that an update lock is supported
- Apply an access intent policy that supports optimistic concurrency

I am calling a finder method from one entity bean to another, and I get an InconsistentAccessIntentException at run time. Why?

This behavior suggests that the second entity bean was previously loaded in the same transaction. This could happen if you called a multifinder method that returned the second bean instance with access intent policy X applied; you are now trying to load the second bean again by calling its `findByPrimaryKey` method with access intent Y applied. Both methods must have the same access intent policy applied.

Likewise, you might have called a container-managed relationship (CMR) accessor method that included the second entity bean. Like a multifinder method, this would have caused the second bean to be loaded under the access intent policy applied to the business method of the first entity bean.

To avoid this problem, ensure that your code does not load the same bean instance twice within the same transaction with different access intent policies applied.

I have two beans in a container-managed relationship. I call `findByPrimaryKey()` on the first bean and then call `getBean2()`, a CMR accessor method, on the returned instance. At that point, I get an `InconsistentAccessIntentException`.

Why? You are probably using read-ahead. When you loaded the first bean, you caused the second bean to be loaded under the access intent policy applied to the finder method for the first bean. However, your CMR accessor method from the first bean to the second had a different access intent policy applied. CMR accessor methods are really finder methods in disguise; the run-time environment behaves as if you were trying to change the access intent for an instance you have already read from persistent store.

To avoid this problem, ensure that the `findByPrimaryKey()` and CMR accessor methods have the same access intent policy applied when you are using read-ahead.

I have a bean with a one-to-many relationship to a second bean. The first bean has a pessimistic-update intent policy applied. When I try to add an instance of the second bean to the first bean's collection, I get an `UpdateCannotProceedWithIntegrityException`. Why?

The second bean probably has a read intent policy applied. When you add the second bean to the first bean's collection, you are not updating the first bean's state, you are implicitly modifying the second bean's state. (The second bean contains a foreign key to the first bean, which is modified.)

To avoid this problem, ensure that both ends of the relationship have an update intent policy applied if you expect to change the relationship at run time.

EJB modules

An EJB module is used to assemble one or more enterprise beans into a single deployable unit. An EJB module is stored in a standard Java archive (JAR) file.

An EJB module contains the following:

- One or more deployable enterprise beans.
- A deployment descriptor, stored in an Extensible Markup Language (XML) file. This file declares the contents of the module, defines the structure and external dependencies of the beans in the module, and describes how the beans are to be used at run time.

An EJB module can be used as a standalone application, or it can be combined with other EJB modules, or with Web modules, to create a J2EE application. An EJB module is installed and run in an enterprise bean container.

For more information about EJB modules, see “Enterprise beans: Resources for learning” on page 149.

Assembling EJB modules

Before you begin

If you want to use existing Java 2 Platform, Enterprise Edition (J2EE) Version 1.2 modules in your J2EE Version 1.3 application, migrate them to the Version 1.3 specification first.

Assemble an Enterprise JavaBeans (EJB) module to contain enterprise beans and related code artifacts. Group Web components, client code, and resource adapter code in separate modules.

An EJB module can be installed as a standalone application or can be combined with other modules into an enterprise application.

The Application Assembly Tool (AAT) provides flexibility in assembling EJB modules. Options described below include:

- Importing an existing EJB module (EJB JAR file)
- Creating a new EJB module
- Copying code artifacts (such as entity beans) from one EJB module into a new EJB module

Steps for this task

1. Start the AAT.
2. Select **File>New>EJB Module**.
The navigation tree displays various sets of properties for configuring the new EJB module.
3. **(Optional)** Use the property dialog shown in the AAT workspace to change the default file name and location.

- a. It is recommended that you change the display name so that it differs from the file name.
 - b. If you like, change the temporary location of the EJB module from the default location, *install_root/bin*.
4. Add at least one EJB component to the module.
 - Add at least one enterprise bean to the EJB component.
 - Import an existing JAR or EAR file containing EJB components.
 - a. In the Navigation pane, right-click the **EJB Components** icon.
 - b. Select **Import** from the pop-up menu.
 - c. Click **Browse** to locate the archive file to import.
 - d. Click **Open** to display the contents of the archive file.
 - e. Select an EJB application from the archive file.
 - f. Select the servlets or JSP files to be added and click **Add** to display the components in the Selected Components window.
 - g. Click **OK** to add the selected components.
 - Copy and paste values from an existing module.
 - Create a new EJB component.
 - a. In the Navigation pane, right-click the **EJB Components** icon.
 - b. Select **New** from the pop-up menu.
 - c. Enter the component name and archive type.
 - d. Select class files.
 - e. Click **OK** in the **New EJB Component** property dialog.
 - f. Enter properties for the EJB component as needed.
 5. Enter assembly properties for each bean.
 - a. Click the plus sign (+) next to the component instance to show property groups.
 - b. Right-click the icon for a property group.
 - c. Select **New** from the pop-up menu to add new values, or edit existing values in the property pane.
 6. Add any other files needed by the application.
 - a. Right-click the **Files** icon.
 - b. Select **Add Files** from the pop-up menu.
 - c. Select **Browse** to navigate the directory structure.
 - d. Click **Select** to open an archive.
 - e. Select the files to add and click **Add**.
 - f. In the Selected Files window, click **OK** to add the files.

What to do next

Assemble any other new modules of your choice:

- EJB modules
- Application client modules
- Resource adapter modules

You can also migrate existing modules.

Another option is assembling a new application module. While assembling (module). While assembling an application module, you can create any new modules that you need.

CMP field assembly settings

In Enterprise JavaBeans (EJB) Version 1.1-compliant beans, container-managed persistence (CMP) fields define the variables in the bean class for which the container must handle persistence management. In EJB Version 2.0-compliant beans, these are replaced by abstract get and set methods; generated code provides the implementation of these abstract methods.

Name

Specifies a subset of public variables in the enterprise bean's implementation class.

Container transactions

Container transaction properties specify how an EJB container is to manage transaction scopes for the enterprise bean's method invocations. A transaction attribute is mapped to one or more methods.

Container transaction assembly settings

Container transaction settings specify how an EJB container is to manage transaction scopes for the enterprise bean's method invocations. Specify one or more methods and associate a transaction attribute with each method.

Name

Specifies a name for the mapping between a transaction attribute and one or more methods.

Description

Contains text that describes the mapping.

Transaction attribute

Specifies how the container must manage the transaction boundaries when delegating a method invocation to an enterprise bean's business method.

Data type	String
Default	Required
Range	For all but message-driven beans, valid values are Mandatory, Never, Not Supported, Required, Requires New, Supports. For session beans, Bean Managed is also valid. For message-driven beans, only Bean Managed, Not Supported, and Required are valid.

Additional information about valid values follows:

Bean Managed

Notifies the container that the bean class directly handles transaction demarcation. This setting can be specified for session beans and (in EJB 2.0 implementations only) for message-driven beans, and it cannot be specified for individual bean methods.

Mandatory

Directs the container to always call the bean method within the transaction context associated with the client. If the client attempts to invoke the bean method without a transaction context, the container throws the

javax.jts.TransactionRequiredException exception to the client. The transaction context is passed to any EJB object or resource accessed by an enterprise bean method.

EJB clients that access these entity beans must do so within an existing transaction. For other enterprise beans, the enterprise bean or bean method must implement the Bean Managed value or use the Required or Requires New value. For non-enterprise bean EJB clients, the client must access a transaction by using the javax.transaction.UserTransaction interface.

Never Directs the container to invoke bean methods without a transaction context.

- If the client calls a bean method from within a transaction context, the container throws the java.rmi.RemoteException exception.
- If the client calls a bean method from outside a transaction context, the container behaves in the same way as if the Not Supported transaction attribute was set. The client must call the method without a transaction context.

Not Supported

Directs the container to call the bean method without a transaction context. If a client calls a bean method from within a transaction context, the container suspends the association between the transaction and the current thread before invoking the method on the enterprise bean instance. The container then resumes the suspended association when the method invocation returns. The suspended transaction context is not passed to any enterprise bean objects or resources that are used by this bean method.

Required

Directs the container to call the bean method within a transaction context. If a client calls a bean method from within a transaction context, the container calls the bean method within the client transaction context. If a client calls a bean method outside a transaction context, the container creates a new transaction context and calls the bean method from within that context. The transaction context is passed to any enterprise bean objects or resources that are used by this bean method.

Requires New

Directs the container to always call the bean method within a new transaction context, regardless of whether the client calls the method within or outside a transaction context. The transaction context is passed to any enterprise bean objects or resources that are used by this bean method.

Supports

Directs the container to call the bean method within a transaction context if the client calls the bean method within a transaction. If the client calls the bean method without a transaction context, the container calls the bean method without a transaction context. The transaction context is passed to any enterprise bean objects or resources that are used by this bean method.

Methods - Name

Specifies the name of an enterprise bean method, or the asterisk character (*). The asterisk is used to denote all methods of an enterprise bean's remote and home interfaces.

Methods - Enterprise bean

Specifies which enterprise bean contains the methods indicated in the **Name** setting.

Methods - Type

Used to distinguish between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if a transaction attribute applies to all methods of the bean.

Data type	String
Range	Valid values for EJB 1.1 implementations are <code>Home</code> , <code>Remote</code> , or <code>Unspecified</code> . For EJB 2.0 implementations, <code>Local</code> and <code>LocalHome</code> are also valid.

Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

EJB module assembly settings

An EJB module is used to assemble enterprise beans into a single deployable unit. An EJB module contains one or more enterprise beans and a deployment descriptor.

File name

Specifies the file name of the EJB module, relative to the top level of the application package.

Alternate DD

Specifies a deployment descriptor to be used at run time instead of the one installed in the module.

Classpath

The path that contains additional classes required by the application that are not contained in the module's archive file. The class loader uses this path. Specify the values relative to the root of the EAR file and separate the values with spaces. Absolute values that refer to files or directories on the hard drive are ignored.

To specify classes that are not in JAR files but are in the root of the EAR file, use a period and forward slash (`./`). Consider the following example directory structure in which the file `myapp.ear` contains an EJB module named `myejb.jar`. Additional classes reside in `class1.jar` and `class2.zip`. A class named `xyz.class` is not packaged in a JAR file but is in the root of the EAR file.

```
myapp.ear/myejb.jar
myapp.ear/class1.jar
myapp.ear/class2.zip
myapp.ear/xyz.class
```

Specify `class1.jar class2.zip ./` as the value of the `Classpath` setting. (Name only the directory for `.class` files.)

Display name

Specifies a short name that is intended to be displayed by GUIs.

Description

Contains text that describes the module.

EJB client JAR

Specifies the location of a JAR file that contains a subset of deployed classes needed by the client.

Small icon

Specifies the name of a JPEG or GIF file that contains a small image (16x16 pixels). The image is used as an icon to represent the module in a GUI.

Large icon

Specifies the name of a JPEG or GIF file that contains a large image (32x32 pixels). The image is used as an icon to represent the module in a GUI.

Generalizations - Subtype

Information about this property is not available.

This property is an IBM extension to the standard J2EE deployment descriptor.

Generalizations - Supertype

Information about this property is not available.

This property is an IBM extension to the standard J2EE deployment descriptor.

EJB relationships - Name

The logical name for a container-managed relationship between EJB 2.0-compliant entity beans.

Default data source - JNDI name

Specifies the default JNDI name for the data source. This default is used if binding information is not specified in the deployment descriptor for an individual enterprise bean.

Default CMP connection factory

Specifies the JNDI name for a CMP connection factory. This setting is applicable only for EJB 2.x-compliant CMP beans.

Default authorization - User ID

Specifies the default user ID for connecting to an enterprise bean's data store.

Default authorization - Password

Specifies the default password for connecting to an enterprise bean's data store.

Entity bean assembly settings

An entity bean encapsulates persistent data, which is stored in a data source, and associated methods to manipulate that data.

EJB name

Specifies a logical name for the enterprise bean. This name must be unique within the EJB module. There is no relationship between this name and the JNDI name.

Display name

Specifies a short name that is intended to be displayed by GUIs.

Description

Contains text that describes the entity bean.

EJB class

Specifies the full name of the enterprise bean class (for example, `com.ibm.ejs.doc.account.AccountBean`).

Remote - Home

(Required for EJB 1.x) Specifies the full name of the enterprise bean's home interface class (for example, `com.ibm.ejs.doc.account.AccountHome`).

Remote - Interface

(Required for EJB 1.x) Specifies the full name of the enterprise bean's remote interface class (for example, com.ibm.ejs.doc.account.Account).

Local interface - Home

(Required for EJB 1.x) Specifies the full name of the enterprise bean's local home interface class (for example, com.ibm.ejs.doc.account.AccountLocalHome).

Local interface - Interface

(Required for EJB 1.x) Specifies the full name of the enterprise bean's local interface class (for example, com.ibm.ejs.doc.account.AccountLocal).

Persistence type

Specifies whether an entity bean manages its own persistent storage or whether storage is managed by the container.

Data type	String
Range	Valid values are Bean managed and Container managed.

Reentrant

Specifies whether the entity bean is reentrant. If an enterprise bean is reentrant, it can call methods on itself or call another bean that calls a method on the calling bean. Only entity beans can be reentrant.

If an entity bean is not reentrant and a bean instance is executing a client request in a transaction context and another client using the same transaction context makes a request on the same bean instance, the EJB container throws the `java.rmi.RemoteException` exception to the second client. If a bean is reentrant, the container cannot distinguish this type of illegal loopback call from a legal concurrent call, so the bean must be coded to detect illegal loopback calls.

Primary key class

Specifies the full name of the bean's primary key class (for example, com.ibm.ejs.doc.account.AccountKey). Composite primary keys map to multiple fields in the entity bean class (or to data structures built from the primitive Java data types) and must be encapsulated in a primary key class.

More complicated enterprise beans are likely to have composite primary keys, with multiple instance variables representing the primary key. A subset of the container-managed fields is used to define the primary key class associated with each instance of an enterprise bean.

Primary key field

Specifies the name of a simple primary key. Simple primary keys map to a single field in the entity bean class and are made up of primitive Java data types (such as integer or long). If exactly one CMP field is the primary key, it can be specified here.

Data type	String
Range	Valid values are the name of any one CMP field or Compound key, which appears when the primary key class is set

Version

Specifies the version of EJB specification with which a container-managed persistence (CMP) entity bean complies.

Data type	String
-----------	--------

Range Valid values are 1.x or 2.x

Abstract schema name

Specifies the name of the abstract schema type of an EJB Version 2.x CMP entity bean. It is used in EJB Query Language (QL) queries.

For example, the abstract schema name might be Order for an entity bean whose local interface is com.acme.commerce.Order.

Small icon

Specifies the name of a JPEG or GIF file that contains a small image (16x16 pixels). The image is used as an icon to represent the entity bean in a GUI.

Large icon

Specifies the name of a JPEG or GIF file that contains a large image (32x32 pixels). The image is used as an icon to represent the entity bean in a GUI.

Security identity

Specifies that a principal's credential properties are to be handled as indicated in the Run-As mode property. If this setting is enabled, the Run-As mode property can be edited.

Run-As mode

Specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources.

At appropriate points, the security service determines whether the principal is authorized to use a particular resource based on the principal's permissions. If the method call is authorized, the security service acts on the principal's credential properties according to the Run-As mode setting of the enterprise bean.

Data type	Enumerated integer
Range	Valid values are Use identity of caller and Use identity assigned to specified role

Additional information about valid settings follows:

Use identity of caller

The security service makes no changes to the principal's credential properties.

Use identity assigned to specified role

A principal that has been assigned to the specified security role is used for the execution of the bean's methods. This association is part of the application binding in which the role is associated with a user ID and password of a user who is granted that role.

Role name

Specifies the name of a security role. If **Run-As mode** is set to Use identity assigned to specified role, a principal that has been granted this role is used.

Description

Contains further information about the security role.

Concurrency control

Specifies how the bean is to handle concurrent access to its data. This setting is applicable only for EJB 1.x-compliant entity beans.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type String
Range Valid values are `Optimistic` or `Pessimistic`

Inheritance root

Specifies whether the enterprise bean is at the root of an inheritance hierarchy.

This property is an IBM extension to the standard J2EE deployment descriptor.

Bean Cache - Activate at

Specifies the point at which an enterprise bean is activated and placed in the cache. Removal from the cache and passivation is also governed by this setting.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type String
Default Transaction
Range Valid values are `Once`, `Transaction`, and `Activity session`

More information about valid values follows:

Once Indicates that the bean is activated when it is first accessed in the server process, and passivated (and removed from the cache) at the discretion of the container, for example, when the cache becomes full.

Transaction

Indicates that the bean is activated at the start of a transaction and passivated (and removed from the cache) at the end of the transaction.

Activity session

Indicates that the bean is activated and passivated as follows:

- On an `ActivitySession` boundary, if an `ActivitySession` context is present on activation
- On a transaction boundary, if a transaction context (but no `ActivitySession` context) is present on activation
- Otherwise, on an invocation boundary

The values of the **Activate at** and **Load at** settings govern which commit options are used, as follows:

- For Commit Option A (implies exclusive DB access), use `Activate at = Once` and `Load at = Activation`.

This option reduces database I/O (avoids calls to the `ejbLoad` function) but serializes all transactions accessing the bean instance. Option A can increase memory usage by maintaining more objects in the cache, but could provide better response time if bean instances are not generally accessed concurrently by multiple transactions. To use Option A successfully, you must also set `Concurrency control` to `Pessimistic`.

- For Commit Option B (implies shared DB access), use `Activate at = Once`, `Load at = Transaction`.

Option B can increase memory usage by maintaining more objects in the cache. However, because each transaction creates its own copy of an object, there can be multiple copies of an instance in memory at any given time (one per transaction), requiring that the database be accessed at each transaction. If an enterprise bean contains a significant number of calls to the `ejbActivate` function,

using Option B can be beneficial because the required object is already in the cache. Otherwise, this option does not provide significant benefit over Option A.

- For Commit Option C (implies shared DB access), use Activate at = Transaction and Load at = Transaction.

This option can reduce memory usage by maintaining fewer objects in the cache; however, there can be multiple copies of an instance in memory at any given time (one per transaction). This option can reduce transaction contention for enterprise bean instances that are accessed concurrently but not updated.

Bean Cache - Load at

Specifies when the bean loads its state from the database. The value of this setting implies whether the container has exclusive or shared access to the database.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type	String
Default	Transaction
Range	Valid values are Activation and Transaction

Additional information about valid values follows:

Activation

Indicates that the bean is loaded when it is activated (regardless of Activate at setting) and implies that the container has exclusive access to the database.

Transaction

Indicates that the bean is loaded at the start of a transaction and implies that the container has shared access to the database.

The Activate at and Load at settings govern which commit options are used. The commit options themselves are described in the Enterprise JavaBeans specification. For more information about setting this setting to achieve a given commit behavior, see Bean Cache - Activate at.

Commit option

Specifies which commit option is used as a result of bean cache settings. The commit options themselves are described in the Enterprise JavaBeans specification.

Data type	String
Range	Valid values are A, B, and C

Local Transactions - Unresolved action

Specifies the action that the EJB container must take if resources are uncommitted by an application in a local transaction.

This property is an IBM extension to the standard J2EE deployment descriptor. This setting is applicable only when **Resolution control** is set to Application. A local transaction context is created when a method runs in what the EJB specification refers to as an unspecified transaction context.

Data type	String
Default	Rollback
Range	Valid values are Commit and Rollback

Additional information about these settings follows:

Commit

At end of the local transaction context, the container instructs all unresolved local transactions to commit.

Rollback

(Default) At end of the local transaction context, the container instructs all unresolved local transactions to roll back.

Local Transactions - Resolution control

Specifies how the local transaction is to be resolved before the local transaction context ends: by the application through user code or by the EJB container.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type String
Range Valid values are Application and ContainerAtBoundary

Additional information about these settings follows:

Application

When this setting is used, your code must either commit or roll back the local transaction. If this does not occur, the runtime environment logs a warning and automatically commits or rolls back the connection as specified by the **Unresolved action** setting.

ContainerAtBoundary

When this setting is used, the container takes responsibility for resolving each local transaction. This provides you with a programming model similar to global transactions in which your code simply gets a connection and performs work within it. User code does not have to handle local transactions.

- - If the **Boundary** attribute is set to `ActivitySession`, then the local transactions are enlisted as `ActivitySession` resources and directed to complete by the `ActivitySession`.
 - If the the **Boundary** attribute is set to `BeanMethod`, then the local transactions are committed at method end by the container.

Connections are never committed automatically by the resource adapter when this value is configured for the bean **Unresolved action** is not used. An application cannot call `Connection.LocalTransaction.begin()` when using this policy and receives an exception from the resource adapter if it does so.

When using a **Resolution control** of `ContainerAtBoundary`, applications must get connection handles *after* the local transaction context boundary has been started by the container. The application should close the connection before the end of the boundary, although any work performed on the connection is not committed or rolled back until the local transaction context ends. This model of connection usage is sometimes referred to as the "get-use-close" model.

This value is supported only for EJB components that use container-managed transactions or ActivitySessions. It is not supported for web components or for enterprise beans that use bean-managed transactions or ActivitySessions.

Local Transactions - Boundary

Specifies the duration of a local transaction context.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type	String
Default	BeanMethod
Range	Valid values are BeanMethod and ActivitySession

Additional information about valid settings follows:

BeanMethod

When this setting is used, the local transaction begins when the method begins and ends when the method ends.

ActivitySession

When this setting is used, the local transaction must be resolved within the scope of any ActivitySession in which it was started or, if no ActivitySession context is present, within the same bean method in which it was started.

This property can be changed on WAS Enterprise only.

Local Relationship Roles - Name

Within a local relationship between EJB 1.x-compliant entity beans, the logical name for the view an entity bean presents to other beans in the relationship.

For example, in a relationship between Account and Customer beans, the role of the Account instance relative to the Customer instance might be *savingsAccount*.

This property is an IBM extension to the standard J2EE deployment descriptor. This is separate from the container-managed relationships defined in the Enterprise JavaBeans specification, Version 2.0.

Local Relationship Roles - Source EJB Name

The name of the entity bean for which the role is defined.

This property is an IBM extension to the standard J2EE deployment descriptor.

Local Relationship Roles - is Forward

Specifies how deployment code for navigating the relationship is generated. This setting is applicable only for navigable relationships.

If **isForward** is enabled (set to true), deployment code is generated in the source bean. That is, navigation of the relationship proceeds *forward* from the source to the target.

Otherwise, deployment code is generated in the target bean. That is, navigation of the relationship proceeds from the target to the source.

This property is an IBM extension to the standard J2EE deployment descriptor. For more information, see the documentation for the Deployment Tool for Enterprise JavaBeans.

Local Relationship Roles - is Navigable

Specifies whether data in related beans may be retrieved through queries to the source bean.

This property is an IBM extension to the standard J2EE deployment descriptor.

Lifetime in cache

The lifetime, in seconds, of cached data for an instance of this bean type.

This value indicates how long the cached data is to exist beyond the end of the transaction in which the data was retrieved. This might avoid another retrieval from persistent storage if the same bean instance were to be used in later transactions. How this value is interpreted depends on the value of **Lifetime in cache usage**.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type	Long
Units	Seconds
Default	0
Range	0 to $2^{61} - 1$

Lifetime in cache usage

Indicates how the lifetime-in-cache setting is to be used by the caching mechanism.

This property is an IBM extension to the standard J2EE deployment descriptor.

If your application uses CMP beans in which the underlying data changes infrequently, you might gain significantly better performance by using this setting with **Lifetime in cache**. Typically, data read from persistent storage is held temporarily in an internal cache until the state of the instance is restored. Cached data normally does not persist beyond state restoration or the end of the transaction in which the finder method was called. By setting **Lifetime in cache usage** to a value other than Off, you indicate that the cached data is to be held for a longer time, potentially hours or days, before invalidating the version of the data in the cache and fetching a new version. Avoiding a trip to persistent storage greatly speeds up access to such beans by applications.

In addition, the use of a value other than Off requires that finders on the bean have an access type of Read, because EJB applications are not permitted to update such CMP beans.

- For EJB 1.x-compliant beans, see Access intent - access type.
- For EJB 2.x-compliant beans, see Applied access intent.

Setting Bean Cache - Activate at to activation and Bean Cache - Load at to Once also minimizes retrievals from persistent storage. However, this settings combination might not be supported by certain CMP beans because it results in the `ejbLoad()` method being called once instead of at the beginning of each transaction in which they are used. The lifetime-in-cache settings combination is independent of CMP bean implementation, though it does incur the modest overhead of calling `ejbLoad()` on each use.

Data type	Enumerated int
Units	Not applicable
Default	0 (Off)
Range	Valid values are Clock Time, Elapsed Time, Week Time, or Off

Additional information about valid values follows:

Off When this value is used, the value of **Lifetime in cache** is ignored. Beans of this type are cached only in a transaction-scoped cache. The cached data for this instance expires after the transaction in which it was retrieved is completed.

Elapsed Time

When this value is used, the value of **Lifetime in cache** is added to the time at which the transaction in which the bean instance was retrieved is completed. The resulting value becomes the time at which the cached data expires. The value of **Lifetime in cache** can add up to minutes, hours, days, and so on.

Clock Time

When this value is used, the value of **Lifetime in cache** represents a particular time of day. The value is added to the immediately preceding or following midnight to calculate a future time value, which is then treated as for Elapsed Time. Using Clock Time enables you to specify that all instances of this bean type are to have their cached data invalidated at, for example, 3 AM, no matter when they were retrieved. This is important if, for example, the data underlying this bean type is batch-updated at 3 AM every day.

The selection of midnight (preceding or following) depends on the value of **Lifetime in cache**. If **Lifetime in cache** plus the value that represents the preceding midnight is earlier than the current time, the following midnight is used.

When you use Clock Time, the value of **Lifetime in cache** is not supposed to represent more than 24 hours. If it does, the cache manager subtracts 24-hour increments from it until a value less than or equal to 24 hours is achieved. To invalidate data at midnight, set **Lifetime in cache** to 0.

Week Time

Usage of this value is the same as for Clock Time, except that the value of **Lifetime in cache** is added to the preceding or following Sunday midnight (11:59 PM Saturday plus 1 minute). When Week Time is used, the value of **Lifetime in cache** can represent more than 24 hours but not more than 7 days.

JNDI name

Specifies the JNDI name of the bean's home interface. This is the name under which the enterprise bean's home interface is registered and therefore, is the name that must be specified when an EJB client does a lookup of the home interface.

Data source - JNDI name

Specifies the JNDI name for the bean's data source.

Default Authorization - User ID

Specifies the default user ID for connecting to a data source.

Default Authorization - Password

Specifies the default password for connecting to a data source.

CMP Resource - JNDI name

Specifies the JNDI name for the resource by which CMP data is stored.

CMP Resource - Resource authentication

Specifies the scope at which resources are to be authenticated: by the container or by the resource.

EJB local-reference assembly settings

For EJB 2.0-compliant beans, the EJB local reference element declares a reference to another enterprise bean's local home interface.

Name

Specifies the name of an EJB reference.

This is the JNDI name that the servlet code uses to get a reference to the enterprise bean. The following example illustrates how this element is specified in the deployment descriptor:

```
<ejb-ref-name>ejb/Payroll</ejb-ref-name>
```

Description

Contains a description of the parent element.

This can include any information that the EJB archive-file producer wants to provide to the consumer of the EJB archive file.

Link

Used in the `ejb-ref` element to specify that an EJB reference is linked to an enterprise bean in the encompassing web-application package.

The value of the link element must be the EJB name of an enterprise bean in the same web-application package. The following example illustrates how this element is specified in the deployment descriptor:

```
<ejb-link>EmployeeRecord</ejb-link>
```

Local interface

Specifies the fully-qualified name of the enterprise bean's local interface.

Local home

Specifies the fully-qualified name of the enterprise bean's local home interface.


Type

Specifies the expected type of the referenced enterprise bean.

Message-driven bean assembly settings

Use this page to configure the assembly properties of message-driven beans

For more information about the effect of JMS properties, such as message selectors and message acknowledgement, see the WebSphere MQ *Using Java* book,

SC34-5456 or  "Messaging Systems and the Java Message Service" (<http://developer.java.sun.com/developer/technicalArticles/Networking/messaging/>).

The following notebook pages are available:

General properties

Specify general assembly properties for the message bean.

- EJB name
- Display name
- Description
- EJB class
- Transaction type

Advanced properties

Specify advanced assembly properties for the message bean.

- Message selector
- Acknowledge mode
- Destination type

Bindings properties

Specify bindings assembly properties for the message bean.

EJB name

The logical name for the message bean (as an enterprise bean)

The logical name for the message bean (as an enterprise bean). This name must be unique within the EJB module. There is no relationship between this name and the JNDI name.

Data type	String
Units	Not applicable
Default	Null
Range	Up to 30 ASCII characters

Display name

A short name that is intended to be displayed by graphical user interfaces

Data type	String
Units	Not applicable
Default	Null
Range	Up to 30 ASCII characters

Description

A description of the message bean, for administrative use

Data type	String
Units	Not applicable
Default	Null
Range	Up to 30 ASCII characters

EJB class

The full package name of the message bean class

Specify the full package name of the message bean class, for example, `com.ibm.ejs.doc.account.MessageBean`. You can either type the class name or click **Browse** to locate an existing class file.

Data type	String
Units	Not applicable
Default	Null
Range	Up to 30 ASCII characters

Transaction type

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean. All messages retrieved from a specific destination have the same transactional behavior. To enable the transactional behavior that you want, you must configure the JMS destination with the same transactional behavior as you configure for the message bean.

Data type	Enum
Units	Not applicable
Default	Bean
Range	Bean The message bean manages its own transactions Container The container manages transactions on behalf of the bean

Message selector

The JMS message selector to be used to determine which messages the message bean receives

The JMS message selector to be used to determine which messages the message bean receives; for example:

```
JMSType='car' AND color='blue' AND weight>2500
```

The selector string can refer to fields in the JMS message header and fields in the message properties. Message selectors cannot reference message body values.

Data type	String
Units	Not applicable
Default	Null
Range	A String, up to 30 ASCII characters, whose syntax is based on a subset of the SQL92 conditional expression syntax.

Acknowledge mode

How the session acknowledges any messages it receives.

This property applies only to message-driven beans that uses bean-managed transaction demarcation (**Transaction type** is set to Bean).

Data type	Enum
Units	Not applicable
Default	Auto Acknowledge
Range	Auto Acknowledge The session automatically acknowledges a message when it has either successfully returned from a call to receive, or the message listener it has called to process the message successfully returns. Dups OK Acknowledge The session lazily acknowledges the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

Destination type

Whether the message bean uses a queue or topic destination.

Data type	Enum
Units	Not applicable
Default	Null
Range	Queue The message bean uses a queue destination. Topic The message bean uses a topic destination.

Listener port name

The name of the listener port for this message bean.

The name of the listener port for this message bean (as defined on the WebSphere administrative console).

Data type	String
Units	Not applicable
Default	Null
Range	Up to 30 ASCII characters

Method extensions

Method extensions are IBM extensions to the standard deployment descriptors for enterprise beans.

Method extension properties are used to define transaction isolation levels for methods, to control the delegation of a principal's credentials, and to define custom finder methods.

Method extension assembly settings

Method extensions are IBM extensions to the standard J2EE deployment descriptors for Enterprise JavaBeans (EJB) Version 1.x-compliant beans. Method extension settings define transaction isolation levels for methods and control the delegation of a principal's credentials.

Method type

Specifies the type of the enterprise bean method.

Data type	String
Range	Valid values are Home, Remote, and Unspecified.

Name

Specifies the name of an enterprise bean method, or the asterisk character (*). The asterisk is used to denote all methods of an enterprise bean's remote and home interfaces.

Parameters

Contains a list of fully qualified Java type names of the method parameters. Used to identify a single method among multiple methods with an overloaded method name.

Isolation level attributes

The transaction isolation level determines how isolated one transaction is from another. This can be set for individual methods in an enterprise bean or for all

methods in the enterprise bean. An asterisk is used to indicate all methods in the bean. This setting is not applicable for EJB 2.x-compliant beans.

Within a transactional context, the isolation level associated with the first method call becomes the required isolation level for all methods called within that transaction. If a method is called with a different isolation level from that of the first method, the `java.rmi.RemoteException` exception is thrown.

Isolation level

Specifies the level of transactional isolation.

The container uses the transaction isolation level attribute as follows:

- Session beans and entity beans with bean-managed persistence (BMP): For each database connection used by the bean, the container sets the transaction isolation level at the start of each transaction unless the bean explicitly sets the isolation level on the connection.
- Entity beans with container-managed persistence (CMP): The container generates database access code that implements the specified isolation level.

Data type	String
Range	Valid values are <code>Serializable</code> , <code>Repeatable read</code> , <code>Read committed</code> , and <code>Read uncommitted</code>

Serializable

This level prohibits the following types of reads:

- *Dirty* reads, in which a transaction reads a database row containing uncommitted changes from a second transaction.
- *Nonrepeatable* reads, in which one transaction reads a row, a second transaction changes the same row, and the first transaction rereads the row and gets a different value.
- *Phantom* reads, in which one transaction reads all rows that satisfy an SQL WHERE condition, a second transaction inserts a row that also satisfies the WHERE condition, and the first transaction applies the same WHERE condition and gets the row inserted by the second transaction.

Repeatable read

This level prohibits dirty reads and nonrepeatable reads, but it allows phantom reads.

Read committed

This level prohibits dirty reads but allows nonrepeatable reads and phantom reads.

Read uncommitted

This level allows dirty reads, nonrepeatable reads, and phantom reads.

Access intent - Intent type

Specifies whether to load the enterprise bean as read-only or for update. This setting is applicable only for EJB 1.x-compliant beans.

This setting is applicable for the following types of beans:

- EJB 1.x-compliant entity beans
- Enterprise beans with CMP version 1.x that are packaged in EJB 2.x-compliant modules

To specify the access intent for EJB 2.x-compliant beans, select an access intent policy.

Data type	String
Range	Valid values are Read or Update

Finder descriptor - User

Specifies that the user has provided a finder helper class in the entity bean's home interface. The class contains specialized finder methods. This setting is applicable only for EJB 1.x-compliant entity beans.

Finder descriptor - EJB QL

Describes the semantics of a finder method that uses EJB QL (Enterprise JavaBeans query language). This setting is applicable only for EJB 1.x-compliant entity beans.

EJB QL is a declarative, SQL-like language that is intended to be compiled to the target language of the persistent datastore used by a persistence manager. The language is independent of the bean's mapping to a relational datastore and is therefore portable. The EJB query specifies a search based on the persistent attributes and relationships of the bean. An EJB query can contain the following clauses:

- SELECT (optional), which specifies the EJB objects to return
- FROM (required), which specifies the collections of objects to which the query is to be applied
- WHERE (optional), which contains search predicates over the collections
- ORDER BY (optional), which specifies the ordering of the resulting collection

Finder descriptor - Full SELECT

Describes the semantics of a finder method that uses an SQL SELECT clause. For information on restrictions, see the documentation for the Deployment Tool for Enterprise JavaBeans.

Finder descriptor - WHERE clause

Describes the semantics of a finder method that uses an SQL WHERE clause. This clause restricts the results that are returned by the query. For information on restrictions, see the documentation for the Deployment Tool for Enterprise JavaBeans.

Security identity

Specifies whether a principal's credential settings are to be handled as indicated in the **Run-As mode** setting. If this is enabled, the **Run-As mode** setting can be edited.

Description

Contains further information about the security instructions.

Run-As mode

Specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources.

At appropriate points, the security service determines whether the principal is authorized to use a particular resource based on the principal's permissions. If the method call is authorized, the security service acts on the principal's credential settings according to the **Run-As mode** setting of the enterprise bean.

Data type	Enumerated integer
-----------	--------------------

Range Valid values are Use identity of caller, Use identity of EJB server, and Use identity assigned to specified role

Additional information about valid values for this setting follows:

Use identity of caller

The security service makes no changes to the principal's credential settings.

Use identity of EJB server

The security service alters the principal's credential settings to match the credential settings associated with the EJB server.

Use identity assigned to specified role

A principal that has been assigned to the specified security role is used for the execution of the bean's methods. This association is part of the application binding in which the role is associated with a user ID and password of a user who is granted that role.

Role name

Specifies the name of a security role. If **Run-As mode** is set to Use identity assigned to specified role, a principal that has been granted this role is used.

Description

Contains further information about the security role.

Method permissions

A method permission is a mapping between one or more security roles and one or more methods that a member of the role can call.

Method permission assembly settings

A method permission is a mapping between one or more security roles and one or more methods that a member of the role can call. Assembly settings for method permissions include an optional description, a list of security role names, and a list of methods. The security roles must be defined, and the methods must be defined in the enterprise bean's remote or home interfaces.

Method permission name

Specifies a name for the mapping between method permissions and security roles.

Description

Contains text that describes the mapping between method permissions and security roles.

Methods - Name

Specifies the name of an enterprise bean method, or the asterisk (*) character. The asterisk is used to denote all the methods of an enterprise bean's remote and home interfaces.

Methods - Enterprise bean

Specifies the name of the enterprise bean that contains the method.

Methods - Type

Distinguishes between a method with the same signature that is defined in both the home and remote interface. Use Unspecified if a method permission applies to all methods of a bean.

Data type String

Range Valid values are Unspecified, Remote, or Home.

Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

Unchecked

Specifies whether the method permission is checked before the method is run.

Roles - Role name

Specifies the name of the security role that must be granted in order to call the method.

Query assembly settings

Use these to specify a finder or SELECT query.

A query element contains the following:

- Optional description of the query
- Name of the finder or SELECT method that uses the query
- The return type of mapping, if it is used
- Whether the query is for a SELECT method
- EJB query language (EJB QL) query string that defines the query

Queries that are expressed in EJB QL must use the `ejb-ql` element to specify the query. If a query cannot be expressed in EJB QL, describe the semantics of the query by using the description element and leave the `ejb-ql` element empty.

Name

Contains the name of an enterprise bean method or the asterisk (*) character. An asterisk in the method-name element denotes all methods of an enterprise bean's remote and home interfaces.

Parameters

Contains a list of the fully-qualified Java names of the method parameters.

Result type

Used in the query element to indicate whether a returned abstract schema type for a SELECT method should be mapped to an `EJBLocalObject` or `EJBObject` type.

References

References are logical names used to locate external resources for enterprise applications. References are defined in the application's deployment descriptor file. At deployment, the references are bound to the physical location (global JNDI name) of the resource in the target operational environment.

This product supports the following types of references:

- An EJB reference is a logical name used to locate the home interface of an enterprise bean.
- A resource reference is a logical name used to locate a connection factory object.

These objects define connections to external resources such as databases and messaging systems. The container makes references available in a JNDI naming subcontext. By convention, references are organized as follows:

- EJB references are made available in the `java:comp/env/ejb` subcontext.
- Resource references are made available as follows:
 - JDBC DataSource references are declared in the `java:comp/env/jdbc` subcontext.
 - JMS connection factories are declared in the `java:comp/env/jms` subcontext.
 - JavaMail connection factories are declared in the `java:comp/env/mail` subcontext.
 - URL connection factories are declared in the `java:comp/env/url` subcontext.

EJB reference assembly settings

An EJB reference is a logical name used to locate the home interface of an enterprise bean used by an application.

At deployment, the EJB reference is bound to the enterprise bean's home in the target operational environment. The container makes the application's EJB references available in a JNDI naming context. It is recommended that references to enterprise beans be organized in the `ejb` subcontext of the application's environment (in `java:comp/env/ejb`).

Name

Specifies the JNDI name of the enterprise bean's home interface relative to the `java:comp/env` context.

For example, if `ejb/EmplRecord` is specified, the referring code looks up the enterprise bean's home interface at `java:comp/env/ejb/EmplRecord`. This JNDI name is an alias used by the code (the actual JNDI name is specified on the Binding tab).

Description

Contains text that describes the EJB reference.

Link

Used to link an EJB reference to an enterprise bean in the current module (the same module as the one making the reference) or in another module within the same J2EE application. This setting specifies the name of the target enterprise bean.

The target enterprise bean can be in any EJB module in the same J2EE application as the referring module. To avoid having to rename enterprise beans to have unique names within an J2EE application, specify the path name of the EJB archive file that contains the referenced enterprise bean and append the target bean's name, separated by a # symbol (for example, `...products/product.jar#ProductEJB`). The path name is relative to the referring module's archive file specification. If a link is not specified, the reference must be resolved to a JNDI name during installation.

Home

Specifies the fully qualified name of the enterprise bean's home interface (for example, `com.ibm.ejbs.EmplRecordHome`).

Remote

Specifies the fully qualified name of the enterprise bean's remote interface (for example, `com.ibm.ejbs.EmplRecord`).

Type

Specifies the expected type of the referenced enterprise bean.

Data type	String
Default	None; must be set
Range	Entity or Session

JNDI name

Binding information that is used by the run-time environment to resolve the location of a resource.

For EJB references, the value of this setting must match the JNDI name of the enterprise bean as it was specified on the **Binding** tab for the EJB module that contains the bean.

EJB relation assembly settings

An EJB relation describes a relationship between two entity beans with container-managed persistence.

The name of the relationship, if specified, is unique within an EJB archive file.

Description

Contains text to describe the EJB relationship role.

Source EJB

Specifies the source of the role that participates in a relationship.

Multiplicity

Specifies the multiplicity of the role that participates in a relation.

Cascade delete

Within a particular relationship, specifies that the lifetime of one or more entity beans is dependent on the lifetime of another entity bean.

Cascade delete can be specified only for an EJB relationship role contained in an EJB relation in which the other EJB relationship role specifies a multiplicity of one.

CMR field

Enables the declaration of a container-managed relationship (CMR) field.

The CMR field describes the bean provider's view of a relationship. It consists of an optional description and the name and class type of the source enterprise bean's role in a relationship.

Exclude list assembly settings

The exclude list indicates which methods in the enterprise beans may not be called. You should also configure security for the enterprise bean so that access to the listed methods is not permitted.

This capability applies only to Enterprise JavaBeans (EJB) Version 2.x-compliant beans. For more information about exclude lists, see the EJB specification.

Description

Provides additional information about this exclude list.

Methods - Name

Specifies the name of an enterprise bean method, or the asterisk (*) character. The asterisk is used to denote all the methods of an enterprise bean's remote and home interfaces.

Methods - Enterprise bean

Specifies the name of the enterprise bean that contains the method.

Methods - Type

Distinguishes between a method with the same signature that is defined in both the home and remote interface. Use Unspecified if the exclusion applies to all methods of a bean.

Data type	String
Range	Valid values are Unspecified, Remote, or Home

Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

Security role assembly settings

A security role is a logical grouping of principals. Access to operations (such as enterprise-bean methods) is controlled by granting access to a role.

Role name

Specifies the name of a security role that is unique to an application. This setting applies only when you are specifying security roles at the application level (EAR file).

Description

Contains text that describes the application-specific security role. This setting applies only when you are specifying security roles at the application level (EAR file).

Binding - Groups - Name

Specifies the user groups that are granted the application-specific security role. This setting applies only when you are specifying security roles at the application level (EAR file).

Binding - Users - Name

Specifies the users that are granted the application-specific security role. This setting applies only when you are specifying security roles at the application level (EAR file).

Binding - Special Subjects - Name

Specifies one of two special categories of authenticate users to which application-specific security roles can be granted: Everyone or All. This setting applies only when you are specifying security roles at the application level (EAR file).

If the special subject All is granted a role, any user who can authenticate by using a valid user ID and password is considered to be granted that role.

If the special subject Everyone is granted a role, all users, including those who did not authenticate, are granted the role. In other words, a method on an enterprise bean or a URI is unprotected if any of the required roles for that method are granted to the special subject Everyone.

Data type	String
Range	Valid values are All or Everyone

Session bean assembly properties

A session bean encapsulates transient data that is associated with a particular EJB client. Unlike data in an entity bean, the data in a session bean is not stored in a persistent data source.

EJB name

Specifies a logical name for the enterprise bean. This name must be unique within the EJB module. There is no relationship between this name and the JNDI name.

Display name

Specifies a short name that is intended to be displayed by GUIs.

Description

Contains text that describes the session bean.

EJB class

Specifies the full name of the enterprise bean class (for example, `com.ibm.ejs.doc.account.AccountBean`).

Remote - Home

Specifies the full name of the enterprise bean's home interface class (for example, `com.ibm.ejs.doc.account.AccountHome`).

Remote - Interface

Specifies the full name of the enterprise bean's remote interface class (for example, `com.ibm.ejs.doc.account.Account`).

Local interface - Home

Specifies the full name of the enterprise bean's home interface class (for example, `com.ibm.ejs.doc.account.AccountLocalHome`).

Local interface - Interface

Specifies the full name of the enterprise bean's local interface class (for example, `com.ibm.ejs.doc.account.AccountLocal`).

Session type

Specifies whether the enterprise bean maintains a conversational state (is stateful) or does not (is stateless).

Data type	String
Range	Valid values are Stateful and Stateless

Transaction type

Specifies whether the enterprise bean manages its own transactions or whether the container manages transactions on behalf of the bean.

Data type	String
Range	Valid values are Container or Bean

Small icon

Specifies the name of a JPEG or GIF file that contains a small image (16x16 pixels). The image is used as an icon to represent the session bean in a GUI.

Large icon

Specifies the name of a JPEG or GIF file that contains a large image (32x32 pixels). The image is used as an icon to represent the session bean in a GUI.

Security identity

Specifies whether a principal's credential properties are to be handled as indicated in the **Run-As mode** property. If this setting is enabled (that is, set to true), the **Run-As mode** setting can be edited.

Description

Contains further information about the security instructions.

Run-As mode

Specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources.

At appropriate points, the security service determines whether the principal is authorized to use a particular resource based on the principal's permissions. If the method call is authorized, the security service acts on the principal's credential properties according to the **Run-As mode** setting of the enterprise bean.

Data type	Enumerated integer
Range	Valid values are Use identity of caller and Use identity assigned to specified role

Additional information about valid values for this setting follows:

Use identity of caller

The security service makes no changes to the principal's credential properties.

Use identity assigned to specified role

A principal that has been assigned to the specified security role is used for the execution of the bean's methods. This association is part of the application binding in which the role is associated with a user ID and password of a user who is granted that role.

Role name

Specifies the name of a security role. If **Run-As mode** is set to Use identity assigned to specified role, a principal that has been granted this role is used.

Description

Contains further information about the security role.

Timeout

This property applies only to stateful session beans.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type	Integer
Units	Seconds

Inheritance root

Specifies whether the enterprise bean is at the root of an inheritance hierarchy.

This property is an IBM extension to the standard J2EE deployment descriptor.

Bean Cache - Activate at

Specifies the point at which an enterprise bean is activated and placed in the cache. Removal from the cache and passivation is also governed by this setting. This setting applies to stateful session beans only (not to stateless beans).

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type	String
Default	Once
Range	Valid values are Once and Transaction

Additional information about valid values follows:

Once Indicates that the bean is activated when it is first accessed in the server process, and passivated (and removed from the cache) at the discretion of the container, for example, when the cache becomes full.

Transaction

Indicates that the bean is activated at the start of a transaction and passivated (and removed from the cache) at the end of the transaction.

Local Transactions - Unresolved action

Specifies the action the container must take if resources are uncommitted by an application in a local transaction. A local transaction context is created when a method runs in what the EJB specification refers to as an unspecified transaction context.

This property is an IBM extension to the standard J2EE deployment descriptor.

Data type	String
Default	Rollback
Range	Valid values are Commit and Rollback

Local Transactions - Boundary

This property does not apply to session beans.

This property is an IBM extension to the standard J2EE deployment descriptor.

JNDI name

Specifies the JNDI name of the bean's home interface. This is the name under which the enterprise bean's home interface is registered and therefore, is the name that must be specified when an EJB client does a lookup of the home interface.

EJB containers

An Enterprise JavaBeans (EJB) container provides a run-time environment for enterprise beans within the application server. The container handles all aspects of an enterprise bean's operation within the application server and acts as an intermediary between the user-written business logic within the bean and the rest of the application server environment.

One or more EJB modules, each containing one or more enterprise beans, can be installed in a single container.

The EJB container provides many services to the enterprise bean, including the following:

- Beginning, committing, and rolling back transactions as necessary.
- Maintaining pools of enterprise bean instances ready for incoming requests and moving these instances between the inactive pools and an active state, ensuring that threading conditions within the bean are satisfied.

- Most importantly, automatically synchronizing data in an entity bean's instance variables with corresponding data items stored in persistent storage.

By dynamically maintaining a set of active bean instances and synchronizing bean state with persistent storage when beans are moved into and out of active state, the container makes it possible for an application to manage many more bean instances than could otherwise simultaneously be held in the application server's memory. In this respect, an EJB container provides services similar to virtual memory within an operating system.

Between transactions, the state of an entity bean can be cached. The EJB container supports option A, B, and C caching.

For more information about EJB containers, see "Enterprise beans: Resources for learning" on page 149.

Managing EJB containers

Each application server can have a single EJB container; one is created automatically for you when the application server is created. The following steps are to be performed only as needed to improve performance after the EJB application has been deployed.

Steps for this task

1. **(Optional)** Adjust EJB container settings.
2. **(Optional)** Adjust EJB cache settings.

What to do next

If adjustments do not improve performance, consider adjusting access intent policies for entity beans, reassembling the module, and redeploying the module in the application.

EJB container settings

Use this page to configure and manage a specific EJB container.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container**.

Passivation directory

Specifies the directory into which the container saves the persistent state of passivated stateful session beans.

Beans are passivated when the number of active bean instances becomes greater than the cache size specified in the container configuration. When a stateful bean is passivated, the container serializes the bean instance to a file in the passivation directory and discards the instance from the bean cache. If, at a later time, a request arrives for the passivated bean instance, the container retrieves it from the passivation directory, deserializes it, returns it to the cache, and dispatches the request to it. If any step fails (for example, if the bean instance is no longer in the passivation directory), the method invocation fails.

Inactive pool cleanup interval

Specifies the interval at which the container examines the pools of available bean instances to determine if some instances can be deleted to reduce memory usage.

Data type	Integer
Units	Milliseconds
Range	Greater than 0

Default datasource JNDI name

Specifies the JNDI name of a data source to use if no data source is specified during application deployment. This setting is not applicable for EJB 2.x-compliant CMP beans.

Servlets and enterprise beans use *data sources* to obtain these connections. When configuring a container, you can specify a default data source for the container. This data source becomes the default data source used by any entity beans installed in the container that use container-managed persistence (CMP).

The default data source for a container is secure. When specifying it, you must provide a user ID and password for accessing the data source.

Specifying a default data source is optional if each CMP entity bean in the container has a data source specified in its configuration. If a default data source is not specified and a CMP entity bean is installed in the container without specifying a data source for that bean, applications cannot use that CMP entity bean.

Initial state

Specifies the execution state requested when the server first starts.

Data type	String
Default	Started
Range	Valid values are Started and Stopped

EJB container system properties

In addition to the settings accessible from the administrative console, you can set the following system property by command-line scripting:

com.ibm.websphere.ejbcontainer.poolSize

Specifies the size of entity-bean pools. If you do not specify a default value, the container defaults of 50 and 500 are used.

Set the pool size for a given entity bean as follows:

```
beantype = min, max[:beantype = min, max...]
```

beantype is the J2EE name of the bean, formed by concatenating the application name, the # character, the module name, the # character, and the name of the bean class. *min* and *max* are the minimum and maximum pool sizes, respectively, for that bean type. Do not specify the square brackets shown in the previous prototype; they denote optional additional bean types that you can specify after the first. Each bean-type specification is delimited by a colon (:).

Use an asterisk (*) as the value of *beantype* to indicate that all bean types are to use those values unless overridden by an exact bean-type specification somewhere else in the string, as follows:

```
*=30,100
```

To specify that a default value be used, omit either *min* or *max* but retain the comma (,) between the two values, as follows (split for publication):

```
SMAApp#PerfModule#TunerBean=54,  
:SMAApp#SModule#TypeBean=100,200
```

You can specify the bean types in any order within the string.

EJB cache settings

Use this page to configure and manage the cache for a specific EJB container.

To view this administrative console page, click **Servers > Application Servers > serverName > EJB Container > EJB Cache Settings**.

Cleanup interval

Specifies the interval at which the container attempts to remove unused items from the cache in order to reduce the total number of items to the value of the cache size.

The cache manager tries to maintain some unallocated entries that can be allocated quickly as needed. A background thread attempts to free some entries while maintaining some unallocated entries. If the thread runs while the application server is idle, then when the application server needs to allocate new cache entries, it does not pay the performance cost of removing entries from the cache. In general, increase this parameter as the cache size increases.

Data type	Integer
Units	Milliseconds
Range	Greater than 0

Cache size

Specifies the number of buckets in the active instance list within the EJB container.

A bucket can contain more than one active enterprise bean instance, but performance is maximized if each bucket in the table has a minimum number of instances assigned to it. When the number of active instances within the container exceeds the number of buckets (that is, the cache size), the container periodically attempts to reduce the number of active instances in the table by passivating some of the active instances. For the best balance of performance and memory usage, set this value to the maximum number of active instances expected during a typical workload.

Data type	Integer
Units	Buckets in the hash table
Range	Greater than 0. The container selects the next largest prime number equal to or greater than the specified value.

Container interoperability

Container interoperability describes the ability of WebSphere Application Server clients and servers at different versions to successfully negotiate differences in native Enterprise JavaBeans (EJB) Version 1.1 finder methods support and Java 2 Platform, Enterprise Edition (J2EE) Version 1.3 compliance.

At one time, there were significant interoperability problems among WebSphere Application Server, versions 4.0.x and 3.5.x distributed, and Version 4.0.x for

zSeries. The introduction of interoperable versions of some class types solved these problems for distributed versions 3.5.6, 4.0.3, and 5 as well as for zSeries Version 4.0.x.

Older 4.0.x and 3.5.x client and application server versions do not support the interoperability classes, which makes them uninteroperable with versions that use the classes. The system property *com.ibm.websphere.container.portable* remedies this situation by enabling newer versions of the application server to turn off the interoperability classes. This lets a more recent application server return class types that are interoperable with an older client.

Depending on the value of *com.ibm.websphere.container.portable*, application servers at versions 5, 4.0.3 and later, and 3.5.6 and later, return different classes for the following:

- Enumerations and collections returned by EJB 1.1 finder methods
- EJBMetaData
- Handles to:
 - Entity beans
 - Session beans
 - Home interfaces

If the property is set to *false*, application servers return the old class types, to enable interoperability with versions 3.5.5 and earlier, and 4.0.2 and earlier. If the property is set to *true*, application servers return the new classes.

Instructions for setting the *com.ibm.websphere.container.portable* property are in the release notes for versions 3.5.6 and later, and 4.0.3 and later. The following tables show interoperability characteristics for various version combinations of application servers and clients as well as default property values for each combination.

Interoperability of Version 3.5.x client with Version 5 application server

Clients at Version 3.5.5 and earlier are not interoperable with Version 5 servers when using:

- EJBMetaData
- Enumerations returned by EJB 1.x finder methods
- Handles to entity beans

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier) installed, set the system property *com.ibm.websphere.container.portable.finder* to *false*. With this setting in place, the Version 5 application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

To interoperate with Version 5 application servers, you must upgrade all Version 3.5.x clients to Version 3.5.6 or later.

Interoperability of Version 5 client with Version 3.5.x application server

Client at Version 5, using this function	Application server at Version 3.5.6, property true	Application server at Version 3.5.6, property false (default)	Application server at Version 3.5.5 and earlier
EJBMetaData	Does not work across domains	Works	Does not work
Handle to session bean	Works	Works	Does not work
Handle to entity bean	Does not work across domains	Does not work across domains	Does not work across domains
Enumeration returned by EJB 1.x finder method	Works	Works	Works

Interoperability of Version 4.0.x client with Version 5 application server

Ideally, all 4.0.x clients that use Version 5 application servers should be at Version 4.0.3 or later.

Version 5 application servers return the interoperability class types by default (true). This can cause interoperability problems for distributed clients at versions 4.0.1 or 4.0.2. In particular, problems can occur with collections and enumerations returned by EJB 1.1 finder methods.

Although it is strongly discouraged, you can set `com.ibm.websphere.container.portable` to `false` on a Version 5 application server. This causes the application server to return the old class types, providing interoperability with clients at Version 4.0.2 and earlier. This is discouraged because:

- The Version 5 application server instance would become non-J2EE 1.3 compliant with regard to handles, home interface handles, and EJBMetaData.
- EJB 1.x finder methods return collection and enumeration objects that do not originate from `ejbportable.jar`.
- Interoperability restrictions still exist with the property set to `false`.
- Version 5 client handles to entity beans and home interfaces do not work across domains for the server you set to `false`.

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 4.0.2 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

Interoperability of client at Version 4.0.2 and earlier with Version 5 application server

Client at Version 4.0.2 and earlier, using this function	Application server at Version 5, property true (default)	Application server at Version 5, property false
EJBMetaData	Does not work	Works for 4.0.2 client
Handle to session bean	Does not work	Works
Handle to entity bean	Does not work	Does not work across cells

Client at Version 4.0.2 and earlier, using this function	Application server at Version 5, property true (default)	Application server at Version 5, property false
Enumeration returned by EJB 1.x finder method	Does not work	Works
Collection returned by EJB 1.x finder method	Does not work	Works
Handle to home interface	Does not work	Does not work across cells

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier, and 4.0.2 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 server uses the new Handle classes but returns the older enumeration and collection classes.

Interoperability of client at Version 4.0.3 and later with Version 5 application server

Clients at Version 4.0.3 and later work well with Version 5 application servers. However, if you set the `com.ibm.websphere.container.portable` to `false`, client handles to entity beans and home interfaces do not work across domains for the server you set to `false`.

Client at Version 4.0.3 and later, using this function	Application server at Version 5, property true (default)	Application server at Version 5, property false
EJBMetaData	Works	Works
Handle to session bean	Works	Works
Handle to entity bean	Works	Does not work across cells
Enumeration returned by EJB 1.x finder method	Works	Works
Collection returned by EJB 1.x finder method	Works	Works
Handle to home interface	Works	Does not work across cells

Interoperability of Version 5 client with Version 4.0.x application server

Clients at Version 5 work well with Version 4.0.3 application servers if you set `com.ibm.websphere.container.portable` to `true`. Client handles to entity beans and home interfaces do not work across domains for any Version 4.0.3 server with `com.ibm.websphere.container.portable` at the default value, `false`. Version 5 client handles to application servers at Version 4.0.2 and earlier also have restrictions.

Client at Version 5, using this function	Application server at Version 4.0.3, property true	Application server at Version 4.0.3, property false (default)	Application server at Version 4.0.2 or earlier
EJBMetaData	Works	Works	Works for 4.0.2 server only
Handle to session bean	Works	Works	Works

Client at Version 5, using this function	Application server at Version 4.0.3, property true	Application server at Version 4.0.3, property false (default)	Application server at Version 4.0.2 or earlier
Handle to entity bean	Works	Does not work across domains	Does not work across domains
Enumeration returned by EJB 1.x finder method	Works	Works	Works
Collection returned by EJB 1.x finder method	Works	Works	Works
Handle to home interface	Works	Does not work across domains	Does not work across domains

Interoperability of zSeries Version 4.0.x client with Version 5 application server

The only valid configuration for container interoperability with zSeries Version 4.0.x clients is the default configuration for the Version 5 application server.

Interoperability of Version 5 client with zSeries Version 4.0.x application server

Version 5 clients should work with a zSeries Version 4.0.x application server with the correct interoperability fixes described in the zSeries documentation. The interoperability characteristics should be the same as for a Version 4.0.3 distributed application server with the property set to true.

Client at Version 5, using this function	zSeries application server at Version 4.0.x
EJBMetaData	Works
Handle to session bean	Works
Handle to entity bean	Works
Enumeration returned by EJB 1.x finder method	Works
Collection returned by EJB 1.x finder method	Works
Handle to home interface	Works

Deploying EJB modules

Before you begin

Assemble one or more EJB modules, assemble one or more Web modules, and assemble them into a J2EE application.

Steps for this task

1. Prepare the deployment environment.
2. Deploy the application.
3. Update the configuration for each EJB module as needed for the deployment environment.

What to do next

The next step is to test and debug the module.

EJB module collection

Use this page to manage the EJB modules deployed in a specific application.

To view this administrative console page, click **Applications > Applications > *applicationName* > EJB modules**. Click the check boxes to select one or more of the EJB modules in your collection.

URI

When resolved relative to the application URL, this specifies the location of the module's archive contents on a file system. The URI matches the <ejb> or <web> tag in the <module> tag of the application deployment descriptor.

EJB module settings

Use this page to configure and manage a specific deployed EJB module.

To view this administrative console page, click **Applications > Applications > *applicationName* > EJB modules > *moduleName***.

URI

When resolved relative to the application URL, this specifies the location of the module archive contents on a file system. The URI must match the URI of a ModuleRef URI in the deployment descriptor of the deployed application (EAR).

Alternate DD

Specifies a deployment descriptor to be used at run time instead of the one installed in the module.

Starting weight

Specifies the order in which modules are started when the server starts. The module with the lowest starting weight is started first.

Data type	Integer
Default	1
Range	0 to 100

Enterprise beans: Resources for learning


Use the following links to find relevant supplemental information about enterprise beans. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.



View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming model and decisions
- Programming instructions and examples
- Programming specifications




Planning, business scenarios, and IT architecture

-  **Mastering Enterprise JavaBeans**
(<http://www.theserverside.com/books/masteringEJB/index.jsp>)
A comprehensive treatment of Enterprise JavaBeans (EJB) programming in nonprintable form (PDF). One must be registered to download the PDF, but registration is free. Information about purchasing a hardcopy is available on the Web site.
- *Enterprise JavaBeans* by Richard Monson-Haefel (O'Reilly and Associates, Inc.: Third Edition, 2001)



Programming model and decisions

-  **Read all about EJB 2.0** (<http://www.ibm.com/developerworks/library/jw-ejb20/index.html>)
A comprehensive overview of the specification.
-  **The J2EE Tutorial** (http://java.sun.com/j2ee/tutorial/1_3-fcs/index.html)
This set of articles by Sun Microsystems covers several EJB-related topics, including the basic programming models, persistence, and EJB Query Language.

Programming instructions and examples

-  **Rules and Patterns for Session Facades**
(http://www7b.boulder.ibm.com/wsdd/library/techarticles/0106_brown/sessionfacades.html)
EJB programming practice: Fronting entity beans with a session-bean facade.
-  **WebSphere Application Server Development Best Practices for Performance and Scalability** (http://www-4.ibm.com/software/webservers/appserv/ws_bestpractices.pdf)
Programming practice for enterprise beans and other types of J2EE components.
-  **Optimistic Locking in IBM WebSphere Application Server 4.0.2**
(<http://www7b.boulder.ibm.com/wsdd/>)
Examples of the effect of optimistic concurrency on application behavior. Although the paper is based on a previous version of this product, the data access issues discussed in it are current.
This paper does not seem to be available directly by URL. To view this paper, visit the specified URL and search on "optimistic locking"

Programming specifications

-  **What's new in the Enterprise JavaBeans 2.0 Specification?**
(<http://java.sun.com/products/ejb/2.0.html>)
You can also download the specification itself from this URL.
-  **Java™ 2 Platform: Compatibility with Previous Releases**
(<http://java.sun.com/products/jdk/1.3/compatibility.html>)
This Sun Microsystems article includes both source and binary compatibility issues.

Chapter 5. Using message-driven beans in applications

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface.

Message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) extend the base JMS support and the Enterprise JavaBean component model to provide automatic asynchronous messaging. When a message arrives on a destination, a listener passes the message to a new instance of a user-developed message-driven bean for processing.

You can use WebSphere Studio Application Developer to develop applications that use message-driven beans. You can use the WebSphere Application Server runtime tools, like the administrative console, to deploy and administer applications that use message-driven beans.

For more information about implementing WebSphere enterprise applications that use message-driven beans, see the following topics:

- An overview of message-driven beans
- Designing an enterprise application to use a message-driven bean
- Developing an enterprise application to use a message-driven bean
- Deploying an enterprise application to use a message-driven bean
- Configuring message listener resources for message-driven beans
- Troubleshooting problems with message-driven beans

Message-driven beans - an overview


WebSphere Application Server supports automatic asynchronous messaging with *message-driven beans* (a type of enterprise bean defined in the EJB 2.0 specification). Messaging with message-driven beans is shown in the figure Message-driven beans and the message listener service.

The support for message-driven beans is based on the message listener service, which comprises a *listener manager* that controls and monitors one or more *listeners*. Each listener monitors a JMS destination for incoming messages. When a message arrives on the destination, the listener passes the message to a new instance of a user-developed message-driven bean (an enterprise bean) for processing. The listener then looks for the next message without waiting for the bean to return.

Messages arriving at a destination being processed by a listener have no client credentials associated with them; the messages are anonymous. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component. For more information about EJB security, see "EJB component security" (not in this document).

You are recommended to develop a message-driven bean to delegate the business processing of incoming messages to another enterprise bean, to provide clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client.


Message-driven beans and the message listener service. This figure shows an incoming message being passed by a JMS listener to a message-driven bean, which passes the message on to a business logic bean for business processing. This messaging is controlled by the listener manager. For more information, see the text that accompanies this figure.



Message-driven beans - components

The WebSphere Application Server support for message-driven beans is based on JMS message listeners and the message listener service, and builds on the base support for JMS. The main components of WebSphere Application Server support for message-driven beans are shown in the following figure and described after the figure:

The main components for message-driven beans. This figure shows the main components of WebSphere support for message-driven beans, from JMS provider through a connection to a destination, listener port, then deployed message-driven bean that processes the message retrieved from the destination. Each listener port defines the association between a connection factory, destination, and a deployed message-driven bean. The other main components are the message listener service, which comprises a listener for each listener port, all controlled by the same listener manager. For more information, see the text that accompanies this figure.



The *message listener service* is an extension to the JMS functions of the JMS provider and provides a *listener manager*, which controls and monitors one or more JMS *listeners*.

Each listener monitors either a JMS queue destination (for point-to-point messaging) or a JMS topic destination (for publish/subscribe messaging).

A *connection factory* is used to create connections with the JMS provider for a specific JMS queue or topic destination. Each connection factory encapsulates the configuration parameters needed to create a connection to a JMS destination.

A *listener port* defines the association between a connection factory, a destination, and a deployed *message-driven bean*. Listener ports are used to simplify the administration of the associations between these resources.

When a deployed message-driven bean is installed, it is associated with a listener port and the listener for a destination. When a message arrives on the destination, the listener passes the message to a new instance of a message-driven bean for processing.

When an application server is started, it initializes the listener manager based on the configuration data. The listener manager creates a dynamic session thread pool for use by listeners, creates and starts listeners, and during server termination controls the cleanup of listener message service resources. Each listener completes several steps for the JMS destination that it is to monitor, including:

- Creating a JMS server session pool, and allocating JMS server sessions and session threads for incoming messages.

- Interfacing with JMS ASF to create JMS connection consumers to listen for incoming messages.
- If specified, starting a transaction and requesting that it is committed (or rolled back) when the EJB method has completed.
- Processing incoming messages by invoking the `onMessage()` method of the specified enterprise bean.

Message-driven beans - transaction support

Message-driven beans can handle messages read from JMS destinations within the scope of a transaction. If transaction handling is specified for a JMS destination, the JMS listener starts a global transaction *before* it reads any incoming message from that destination. When the message-driven bean processing has finished, the JMS listener commits or rolls back the transaction (using JTA transaction control).

Note:

- All messages retrieved from a specific destination have the same transactional behavior.

If messages are queued to be sent within a global transaction they are sent when the transaction is committed. If the processing of a message causes the transaction to be rolled back, then the message that caused the bean instance to be invoked is left on the JMS destination.

You can configure the **Maximum retries** property of the listener port to define the maximum number of times the listener attempts to read a message from a destination. When the Max retries limit is reached, the listener for that destination is stopped. When you have resolved the problem, you must then restart the listener.

Designing an enterprise application to use message-driven beans

This topic describes things to consider when designing an enterprise application to use message-driven beans.

The considerations in this topic are based on a generic enterprise application that uses one message-driven bean to retrieve messages from a JMS queue destination and passes the messages on to another enterprise bean that implements the business logic.

To design an enterprise application to use message-driven beans, complete the following steps:

Steps for this task

1. Identify the JMS resources that the application is to use.

This helps to identify the properties of resources that need to be used within the application and configured as application deployment descriptors or within WebSphere Application Server.

JMS resource type	Properties
Queue connection factory	Name: SamplePtoPQueueConnectionFactory JNDI Name: Sample/JMS/QCF
Queue destination	Name: Q1 JNDI Name: Sample/JMS/Q1

JMS resource type	Properties
Listener port (for the destination)	Name: SamplePtoPListenerPort Connection Factory JNDI Name: Sample/JMS/QCF Destination JNDI Name: Sample/JMS/Q1 Maximum Sessions: 5 Maximum Retries: 10 Maximum Messages: 1
Message-driven bean (deployment properties)	Name: JMSppSampleMDBBean Transaction type: Container Destination type: Queue Listener port name: SamplePtoPListenerPort
Business logic bean	Name: MyLogicBean

Ensure that you use consistent values where needed; for example, the JNDI names for the connection factory and destination must be the same for both those resources and the equivalent properties of the listener port.

2. Separation of business logic.

You are recommended to develop a message-driven bean to delegate the business processing of incoming messages to another enterprise bean. This provides clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client.

3. Security considerations.

Messages arriving at a destination being processed by a listener have no client credentials associated with them; the messages are anonymous. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component. For more information about EJB security, see "EJB component security" (not in this document).

4. General JMS considerations

For Publish/Subscribe messaging, choose the JMS server port to be used depending on your needs for transactions or performance:

Queued port

The TCP/IP port number of the listener port used for all point-to-point and Publish/Subscribe support.

Direct port

The TCP/IP port number of the listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for Publish/Subscribe support.

Note: Message-driven beans cannot use the direct listener port for Publish/Subscribe support. Therefore, any topic connection factory configured with **Portset** to **Direct** cannot be used with message-driven beans.

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see "The effect of transaction context on non-durable subscribers" (not in this document).

Developing an enterprise application to use message-driven beans

Use this task to develop an enterprise application to use a message-driven bean. The message-driven bean is invoked by a JMS listener when a message arrives on the input queue that the listener is monitoring.

You are recommended to develop the message-driven bean to delegate the business processing of incoming messages to another enterprise bean, to provide clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client. Responses can be handled by another enterprise bean acting as a sender bean, or handled in the message-driven bean.

You develop an enterprise application to use a message-driven bean like any other enterprise bean, except that a message-driven bean does not have a home interface or a remote interface.

This topic describes how to develop a completely new message-driven bean class. If you have a WAS 4.0 enterprise application that uses the JMS listener, you can migrate that application to use message-driven beans, as described in *Migrating a WAS 4.0 JMS listener application to use message-driven beans*.

For more information about writing the message-driven bean class, see *Creating a message-driven bean* in the WebSphere Studio help bookshelf.

To develop an enterprise application to use a message-driven bean, complete the following steps:

Steps for this task

1. Creating the Enterprise Application project, as described in the WebSphere Studio article .
2. Creating the message-driven bean class.

You can use the New Enterprise Bean wizard of WebSphere Studio Application Developer to create an enterprise bean with a bean type of Message-driven bean. The wizard creates appropriate methods for the type of bean.

By convention, the message bean class is named *nameBean*, where *name* is the name you assign to the message bean; for example:

```
public class MyJMSppMDBBean implements MessageDrivenBean, MessageListener
```

The message-driven bean class must define and implement the following methods:

- `onMessage(message)`, which must meet the following requirements:
 - The method must have a single argument of type `javax.jms.Message`.
 - The throws clause must *not* define any application exceptions.
 - If the message-driven bean is configured to use bean-managed transactions, it must call the `javax.transaction.UserTransaction` interface to scope the transactions. Because these calls occur inside the `onMessage()` method, the transaction scope does not include the initial message receipt. This means the application server is given one attempt to process the message.

To handle the message within the `onMessage()` method (for example, to pass the message on to another enterprise bean), you use standard JMS. (This is known as bean-managed messaging.)

- `ejbCreate()`

You must define and implement an `ejbCreate` method for each way in which you want a new instance of an enterprise bean to be created.

- `ejbRemove()`.

This method is invoked by the container when a client invokes the `remove` method inherited by the enterprise bean's home interface from the `javax.ejb.EJBHome` interface. This method must contain any code that you want to execute before an enterprise bean instance is removed from the container (and the associated data is removed from the data source).

For example, the following code extract shows how to access the text and the JMS MessageID, from a JMS message of type `TextMessage`:

Code example: The `onMessage()` method of a message bean. This figure shows a code extract for a basic `onMessage()` method of a sample message-driven bean. The method unpacks the incoming text message to extract the text and message identifier and calls a private `putMessage` method (defined within the same message bean class) to put the message onto another queue.

```
public void onMessage(javax.jms.Message msg)
{
    String text      = null;
    String messageID = null;

    try
    {
        text = ((TextMessage)msg).getText();

        System.out.println("senderBean.onMessage(), msg text2: "+text);

        //
        // store the message id to use as the Correlator value
        //
        messageID = msg.getJMSMessageID();

        // Call a private method to put the message onto another queue
        putMessage(messageID, text);
    }
    catch (Exception err)
    {
        err.printStackTrace();
    }
    return;
}
```

The result of this step is a message-driven bean that can be assembled into an `.EAR` file for deployment.

3. Assembling and packaging the application for deployment.

You can use WebSphere Studio to assemble and package the application for deployment.

Results

The result of this task is an `.EAR` file, containing an application message-driven bean, that can be deployed in WebSphere Application Server.

What to do next

After you have developed an enterprise application to use message-driven beans, configure and deploy the application; for example, define the listener ports for the message-driven beans and, optionally, change the deployment descriptor attributes for the application. For more information about configuring and deploying an application that uses message-driven beans, see "Deploying an enterprise application to use message-driven beans"

Migrating a JMS listener application to use message-driven beans

Use this task to migrate an enterprise application that uses message beans with the JMS Listener from WebSphere Application Server 4.0 to use EJB 2.0 message-driven beans.

This task uses a command line utility, `mb2mdb`, that takes as its input either a deployed `MessageBean.jar` module or a deployed Enterprise Application (`.ear`) that contains a message bean, along with the JMS listener configuration XML file that defines the WebSphere Application Server 4.0 message beans. The result is a new `.jar/.ear` module that can then be deployed directly into a WebSphere Application Server 5.0 application server.

You can display the usage help for the migration utility, by typing the command `mb2mdb` at a command line.

To migrate a WebSphere Application Server 4.0 enterprise application that uses message beans to use EJB 2.0 message-driven beans, type the following command at an operating system command line:

```
mb2mdb inputMB.jar-ear jmsListenerConfig.xml workingDirectory outputMDB.jar-ear options
```

Where:

inputMB.jar-ear

The name of the deployed WebSphere Application Server 4.0 jar or ear file containing a stateless session message bean.

jmsListenerConfig.xml

The name of the XML configuration file used to configure the WebSphere Application Server 4.0 JMS listeners.

workingDirectory

The name of a new or existing directory that is used to generate the new message-driven bean and package the `outputMDB.jar` or `.ear` file.

Note: By default, the tool clears the working directory after it has completed. If you want to preserve the contents of the working directory, you must specify the `-keep` option.

outputMDB.jar-ear

The name of the output `.jar` or `.ear` file for the migrated message-driven bean application.

options

An optional set of parameters that you can use to control the `mb2mdb` utility.

-keep This prevents the tool from clearing out the working directory after completion.

-verbose

This causes the tool to display informational messages as to the progress of the migration and its parameters.

-map listenerHome=bindingHome

This option provides a mechanism to map between the JNDIHomeName specified for a listener in the JMS listener configuration XML file and the default binding home name specified in the *inputMB.jar-ear* file.

If the *jmsListenerConfig.xml* file contains a deployed EJB home JNDI name that is different to the default binding within the *inputMB.jar-ear*, use this option to map between the two names.

This enables you to install the output .jar or .ear file for the message-driven bean into an application server and bind the bean with a different JNDIHomeName than is specified in the bean's bindings.xmi.

Results

The result of this task is a new .jar or .ear file for a message-driven bean that can then be deployed directly into a WebSphere Application Server 5.0 application server.

To successfully install the .jar or .ear file, you need to bind the message-driven bean against a listener port defined to the message listener service of the application server. You need to have used the WebSphere Application Server administrative console to define the listener port, which defines the JMS connection factory and destination that a message-driven bean bound to it listens on. For more information about installing and configuring a .jar or .ear file for a message-driven bean, see *Deploying an enterprise application to use message-driven beans*.

Deploying an enterprise application to use message-driven beans

Use this task to deploy an enterprise application to use message-driven beans.

This task description assumes that you have an .EAR file, which contains an application enterprise bean with code for message-driven beans, that can be deployed in WebSphere Application Server.

To deploy an enterprise application to use message-driven beans, complete the following steps:

Steps for this task

1. Use the WebSphere administrative console to define the listener ports for the application, as described in *Adding a new listener port*.
2. For each message-driven bean in the application, use the application assembly tool to configure the deployment attributes to match the listener port definitions, as described in *"Configuring deployment attributes for a message-driven bean"*.
3. Use the WebSphere administrative console to install the application

This stage is a standard WebSphere Application Server task, as described in *"Installing a new application"* on page 513.

When you install the application, you are prompted to specify the name of the listener port that the application is to use for late responses. Select the listener port, then click **OK**.

Configuring deployment attributes for a message-driven bean

Use this task to configure the message-driven beans deployment attributes for an enterprise bean, to override the deployment attributes defined within the application EAR file.

This task description assumes that you have an EAR file, which contains an application enterprise bean developed as a message-driven bean, that can be deployed in WebSphere Application Server.

Note: After deployment code has been generated for an application, the deployable archive is renamed with the prefix `Deployed_`. Any subsequent changes to the archive from within the Application Assembly Tool are applied to the version of the archive that existed prior to code generation. To see changes reflected in your application, you must regenerate deployment code and re-install the deployable archive.

To configure the message-driven beans deployment attributes for an enterprise bean, use the the application assembly tool to configure the deployment attributes of the application to match the listener port definitions:

Steps for this task

1. Launch the Application Assembly Tool.
2. Create or edit the application EAR file.

For example, to change attributes of an existing application, click **File-> Open** then select the the EAR file.

3. In the navigation pane, select the message-driven bean instance; for example, expand `ejb_module_instance-> Message-driven beans` then select the bean instance.

A property dialog notebook for the message-driven bean is displayed in the property pane.

4. Specify general deployment properties.
 - a. In the property pane, select the General tab.
 - b. Specify the following properties:

Transaction type

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean. All messages retrieved from a specific destination have the same transactional behavior. To enable the transactional behavior that you want, you must configure the JMS destination with the same transactional behavior as you configure for the message bean.

Bean The message bean manages its own transactions

Container

The container manages transactions on behalf of the bean

5. Specify advanced deployment properties.
 - a. In the property pane, select the Advanced tab.
 - b. Specify the following properties:

Message selector

The JMS message selector to be used to determine which messages the message bean receives; for example:

```
JMSType='car' AND color='blue' AND weight>2500
```

The selector string can refer to fields in the JMS message header and fields in the message properties. Message selectors cannot reference message body values.

Acknowledge mode

How the session acknowledges any messages it receives.

This property applies only to message-driven beans that uses bean-managed transaction demarcation (**Transaction type** is set to Bean).

Auto Acknowledge

The session automatically acknowledges a message when it has either successfully returned from a call to receive, or the message listener it has called to process the message successfully returns.

Dups OK Acknowledge

The session lazily acknowledges the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

As defined in the EJB specification, clients cannot use using Message.acknowledge() to acknowledge messages. If a value of CLIENT_ACKNOWLEDGE is passed on the createxxxSession call, then messages are automatically acknowledged by the application server and Message.acknowledge() is not used.

Destination type

Whether the message bean uses a queue or topic destination.

Queue

The message bean uses a queue destination.

Topic The message bean uses a topic destination.

Subscription durability

Whether a JMS topic subscription is durable or non-durable.

Durable

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

Nondurable

Non-durable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages

published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see "The effect of transaction context on non-durable subscribers" (not in this document).

6. Specify bindings deployment properties.
 - a. In the property pane, select the Bindings tab.
 - b. Specify the following property:

Listener port name

The name of the listener port for this message-driven bean.

7. To apply the changes and close the Application Assembly Tool, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
8. **(Optional)** To see changes reflected in your application, regenerate deployment code and reinstall the deployable archive.

Configuring message listener resources for message-driven beans

Use the following tasks to configure resources needed by the message listener service to support message-driven beans.

- Configuring the message listener service
- Adding a new listener port
- Configuring a listener port
- Configuring security for message-driven beans

Configuring the message listener service

Use this task to configure the properties of the message listener service for an application server.

To configure the properties of the message listener service for an application server, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**
This displays the Message Listener Service properties in the content pane.
4. Specify appropriate properties of the message listener service.
5. Click **OK**.
6. Save your configuration.
7. **(Optional)** To have the changed configuration take effect, stop then restart the Application Server.

Message listener service

The message listener service is an extension to the JMS functions of the JMS provider. It provides a listener manager that controls and monitors one or more JMS listeners, which each monitor a JMS destination on behalf of a deployed message-driven bean.

Name: The name by which the message listener service is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	MsgLService
Range	1 through 30 ASCII characters

Description: A description of the message listener service, for administrative purposes

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Thread pool: Select this link to display the service thread pool properties.

Data type	Not applicable
Units	Not applicable
Default	Not applicable
Range	Not applicable

Custom Properties: Custom properties of the message listener service.

MQJMS.POOLING.TIMEOUT: The number of milliseconds after which a connection in the pool is destroyed if it has not been used.

An MQSimpleConnectionManager allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if it has not been used for five minutes.

Data type	Integer
Units	Milliseconds
Default	5 minutes
Range	

MQJMS.POOLING.THRESHOLD: The maximum number of unused connections in the pool.

An MQSimpleConnectionManager allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if there are more than ten unused connections in the pool.

Data type	Integer
Units	Number of connections
Default	10
Range	

MAX.RECOVERY.RETRIES: The maximum number of times that the listener service tries to get a message from a listener port before the associated listener is stopped, in the range 0 through 2147483647.

Data type	Integer
Units	Retry attempts
Default	0 (no retries)
Range	0 (no retries) through 2147483647

RECOVERY.RETRY.INTERVAL: The time in seconds between retry attempts by the listener service to get a message from a listener port.

Data type	Integer
Units	Seconds
Default	10
Range	1 through 2147483647

Message listener port collection: The message listener ports configured in the administrative domain

This panel displays a list of the message listener ports configured in the administrative domain. Each listener port is used with a message-driven bean to automatically receive messages from an associated JMS destination. You can use this panel to add new listener ports or to change the properties of existing listener ports. For more information about the property fields for listener ports, see Listener port properties.

Listener port settings: A listener port is used to simplify administration of the association between a connection factory, destination, and deployed message-driven bean.

Use this panel to view or change the configuration properties of the selected listener port.

Initial state: The state that you want the listener port to have when the application server is next restarted

Data type	Enum
Units	Not applicable
Default	Started
Range	Started When the application server is next started, the listener port is started automatically. Stopped When the application server is next started, the listener port is not started automatically. If message-driven beans are to use this listener port on the application server, the system administrator must start the port manually or select the Started value of this property then restart the application server.

Description: A description of the listener port, for administrative purposes

Data type	String
Units	En_US ASCII characters

Default	Null
Range	1 through 30 ASCII characters

Listener port: The name by which the listener port is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Connection factory JNDI name: The JNDI name for the JMS connection factory to be used by the listener port; for example, `jms/connFactory1`.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Destination JNDI name: The JNDI name for the destination to be used by the listener port; for example, `jms/destn1`.

If the extended messaging service is to use this listener port to handle late responses, the value of this property must match the JMS response destination on the output port used by the sender bean. Also cannot use a temporary destination for late responses.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Maximum sessions: The maximum number of concurrent JMS server sessions used by a listener to process messages, in the range 1 through 2147483647.

The maximum number of concurrent JMS server sessions used by a listener to process messages for its associated message-driven bean.

Data type	Integer
Units	Sessions
Default	1
Range	1 through 2147483647

Maximum retries: The maximum number of times that the listener tries to deliver a message before the listener is stopped, in the range 0 through 2147483647.

The maximum number of times that the listener tries to deliver a message to a message-driven bean instance before the listener is stopped.

Data type	Integer
Units	Retry attempts
Default	0 (no retries)
Range	0 (no retries) through 2147483647

Maximum messages: The maximum number of messages that the listener can process in one JMS server session, in the range 0 through 2147483647.

The maximum number of messages that a listener can process in one JMS server session.

Data type	Integer
Units	Number of messages
Default	1
Range	1 through 2147483647

Adding a new listener port

Use this task to add a new listener port to the message listener service, so that message-driven beans can be associated with the port to retrieve messages.

To add a new listener port, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**
This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**.
This displays a list of the listener ports.
5. In the content pane, click **New**.
6. Specify appropriate properties for the listener port.
7. Click **OK**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.
9. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

Results

If enabled, the listener port is started automatically when a message-driven bean associated with that port is installed.

Configuring a listener port

Use this task to change the properties of an existing listener port, used by message-driven beans associated with the port to retrieve messages.

To configure the properties of a listener port, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**

- This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.
This displays the properties of the application server in the content pane.
 3. In the Additional Properties table, select **Message Listener Service**
This displays the Message Listener Service properties in the content pane.
 4. In the content pane, click **Listener Ports**.
This displays a list of the listener ports.
 5. Click the listener port that you want to modify.
This displays the properties of the listener port in the content pane.
 6. Specify appropriate properties for the listener port.
 7. Click **OK**.
 8. To save your configuration, click **Save** on the task bar of the Administrative console window.
 9. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

Deleting a listener port

Use this task to delete a listener port from the message listener service, to prevent message-driven beans associated with the port from retrieving messages.

To delete a listener port, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**
This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**.
This displays a list of the listener ports.
5. In the content pane, select the checkbox for the listener port that you want to delete.
6. Click **Delete**. This action stops the port (needed to allow the port to be deleted) then deletes the port.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

Configuring security for message-driven beans

Use this task to configure resource security and security permissions for message-driven beans.

Messages arriving at a listener port have no client credentials associated with them. The messages are anonymous.

To call secure enterprise beans from a message-driven bean, the message-driven bean needs to be configured with a RunAs Identity deployment descriptor. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component.

For more information about EJB security, see "EJB component security" (not in this document). For more information about configuring security for your application, see Assembling secured applications.

Administering listener ports

Use the following tasks to administer listener ports, which each define the association between a connection factory, a destination, and a message-driven bean.

You can use the WebSphere administrative console to administer listener ports, as described in the following tasks.

- Adding a new listener port
Use this task to create a new listener port, to specify a new association between a connection factory, a destination, and a message-driven bean. This enables deployed message-driven beans associated with the port to retrieve messages from the destination.
- Configuring a listener port
Use this task to view or change the configuration properties of a listener port.
- Starting a listener port
Use this task to start a listener port manually.
- Stopping a listener port
Use this task to stop a listener port manually.

Note: If configured as enabled, a listener port is started automatically when a message-driven bean associated with that port is installed. You do not normally need to start or stop a listener port manually.

Starting a listener port

Use this task to start a listener port on an application server, to enable the listeners for message-driven beans associated with the port to retrieve messages.

A listener is active, that is able to receive messages from a destination, if the deployed message-driven bean, listener port, and message listener service are all started. Although you can start these components in any order, they must all be in a started state before the listener can retrieve messages.

If configured as enabled, a listener port is started automatically when a message-driven bean associated with that port is installed. However, you can start a listener port manually, as described in this topic.

When a listener port is started, the listener manager tries to start the listeners for each message-driven bean associated with the port. If a message-driven bean is stopped, the port is started but the listener is not started, and remains stopped. If you start a message-driven bean, the related listener is started.

To start a listener port on an application server, use the administrative console to complete the following steps:

Steps for this task

1. **(Optional)** If you want the listener for a deployed message-driven bean to be able to receive messages at the port, check that the message-driven bean has been started.
2. In the navigation pane, select **Servers-> Application Servers**
This displays a table of the application servers in the administrative domain.
3. In the content pane, click the name of the application server.
This displays the properties of the application server in the content pane.
4. In the Additional Properties table, select **Message Listener Service**
This displays the Message Listener Service properties in the content pane.
5. In the content pane, select **Listener Ports**.
This displays a list of the listener ports.
6. **(Optional)** Select the checkbox for the listener port that you want to start.
7. Click **Start**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.

Stopping a listener port

Use this task to stop a listener port on an application server, to prevent the listeners for message-driven beans associated with the port from retrieving messages.

When you stop a listener port as described in this topic, the listener manager stops the listeners for all message-driven beans associated with the port.

To stop a listener port on an application server, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**
This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**.
This displays a list of the listener ports.
5. **(Optional)** In the content pane, select the listener port that you want to stop.
6. Click **Stop**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

Important files for message-driven beans and extended messaging

The following files in the WAS_HOME/temp directory are important for the operation of the WebSphere Application Server messaging service, so should not be deleted. If you do need to delete the WAS_HOME/temp directory or other files in it, ensure that you preserve the following files.

server_name-durableSubscriptions.ser

You should not delete this file, because the messaging service uses it to keep track of durable subscriptions for message-driven beans. If you uninstall an application that contains a message-driven bean, this file is used to unsubscribe the durable subscription.

server_name-AsyncMessageRequestLog.ser

You should not delete this file, because the messaging service uses it to keep track of late responses that need to be delivered to the late response message handler for the extended messaging provider.

Troubleshooting message-driven beans

Use this overview task to help resolve a problem that you think is related to message-driven beans.

Message-driven beans support uses the standard WebSphere Application Server troubleshooting facilities. If you encounter a problem that you think might be related to the message-driven beans, complete the following stages:

Steps for this task

1. Check for messages about message-driven beans in the application server's SystemOut log at *was_home\logs\server\SystemOut*.
Look in the SystemOut log for messages that indicate a problem with JMS resources for message-driven beans, such as listener ports.
2. Check for more messages in the application server's SystemOut log.
If the JMS server is running, but you have problems accessing JMS resources, check the SystemOut log file, which should contain more error messages and extra details about the problem.
3. Check the Release Notes for specific problems and workarounds
The section *Possible Problems and Suggested Fixes* of the Release Notes, available from the WebSphere Application Server library web site, is updated regularly to contain information about known defects and their workarounds. Check the latest version of the Release Notes for any information about your problem. If the Release Notes does not contain any information about your problem, you can also search the Technotes database on the WebSphere Application Server web site.
4. Check that message listener service has started.
The message listener service is an extension to the JMS functions of the JMS provider. It provides a listener manager that controls and monitors one or more JMS listeners, which each monitor a JMS destination on behalf of a deployed message-driven bean.
5. Check your JMS resource configurations
If the WebSphere Messaging functions seem to be running properly (the JMS server is running without problems), check that the JMS resources have been configured correctly. For example, check that the listener ports have been configured correctly and have been started.
6. Check for problems with the WebSphere Messaging functions
For more information about troubleshooting WebSphere Messaging, see the related topics.
7. **(Optional)** Get a detailed exception dump for messaging.

If the information obtained in the preceding steps is still inconclusive, you can enable the application server debug trace for the "Messaging" group to provide a detailed exception dump.

Message-driven beans samples

The following examples are provided to illustrate use of the message-driven beans support:

- Point-to-point samples:

- "Tutorial: Creating JMS message sample"

This tutorial is designed to help you develop and deploy a JMS message sample application that tests the WebSphere Application Server message-driven beans support in a point-to-point scenario. This sample illustrates how to develop and deploy an application that comprises the following components:

- A Java/JMS program that writes a message to a queue.
- A message-driven bean that is invoked by a JMS listener when a message arrives on a defined queue.

For more information about this sample, see the samples article "Tutorial: Creating JMS message sample" at `WAS_HOME/Enterprise/samples/messaging/doc/PtoP/wsblDptoP.xml` (if you have installed the samples option).

- "Sample: Message Listener (point-to-point)"

This sample is designed to demonstrate the use and behavior of message-driven beans for a simple point-to-point scenario. This sample uses the JMS message sample deployed in the sample above.

For more information about this sample, see the samples article "Sample: Message Listener (Point-to-Point)" at `WAS_HOME/Enterprise/samples/messaging/doc/PtoP/wsrnPtoP.xml` (if you have installed the samples option).

- Publish/subscribe samples

- "Tutorial: Creating JMS message publish/subscribe sample"

This tutorial is designed to help you develop and deploy a JMS message sample application that tests the WebSphere Application Server message-driven beans support in a publish/subscribe scenario. This sample illustrates how to develop and deploy an application that comprises the following components:

- A client program that starts the message sequence by publishing a message to a selected topic.
- A message-driven bean that is invoked by a JMS listener when the broker passes a message to the listener from a topic to which it has subscribed.

For more information about this sample, see the samples article "Tutorial: Creating JMS message publish/subscribe sample" at `WAS_HOME/Enterprise/samples/messaging/doc/PubSub/wsblDPubSub.xml` (if you have installed the samples option).

- "Sample: Message Listener (publish/subscribe)"

This sample is designed to demonstrate the use and behavior of message-driven beans for a simple publish/subscribe scenario. This sample uses the JMS message sample deployed in the publish/subscribe sample above.

For more information about this sample, see the samples article "Sample: Message Listener (publish/subscribe)" at [WAS_HOME/Enterprise/samples/messaging/doc/PubSub/wsruntimePubSub.xml](#) (if you have installed the samples option).

Chapter 6. Using application clients

An application client module is a JAR (Java ARchive) file containing a client for accessing a Java application.

Steps for this task

1. Decide on a type of application client.
2. Develop the application client code.
 - Develop ActiveX application client code.
 - Develop applet client code.
 - Develop J2EE application client code.
 - Develop pluggable application client code.
 - Develop thin application client code.

Usage scenario

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

Application clients

In a traditional client server environment, the client requests a service and the server fulfills the request. Multiple clients use a single server. Clients can also access several different servers. This model persists for Java clients except now these requests make use of a client run-time environment.

In this model, the client application requires a servlet to communicate with the enterprise bean, and the servlet must reside on the same machine as the WebSphere Application Server.

With WebSphere Application Server V5.0, application clients now consist of the following models:

- ActiveX application client
- Applet client
- J2EE application client
- Pluggable application client
- Thin application client

The *ActiveX application client* model, uses the Java Native Interface (JNI) architecture to programmatically access the Java virtual machine (JVM) API. Therefore the JVM code exists in the same process space as the ActiveX application (Visual Basic, VBScript, or Active Server Pages (ASP)) and remains attached to the process until that process terminates.

In the *Applet client* model, a Java applet embeds in a HyperText Markup Language (HTML) document residing on a remote client machine from the WebSphere Application Server. With this type of client, the user accesses an enterprise bean in the WebSphere Application Server through the Java applet in the HTML document.

The *J2EE application client* is a Java application program that accesses enterprise beans, Java Database Connectivity (JDBC), and Java Message Service message queues. The J2EE application client program runs on client machines. This program follows the same Java programming model as other Java programs; however, the J2EE application client depends on the application client run time to configure its execution environment, and uses the Java Naming and Directory Interface (JNDI) name space to access resources.

The *Pluggable and thin application clients* provide a lightweight Java client programming model. These clients are best suited in situations where a Java client application exists but the application needs enhancements to use enterprise beans, or where the client application requires a thinner, more lightweight environment than the one offered by the J2EE application client. The difference between the thin application client and the pluggable application client is that the thin application client includes a Java virtual machine (JVM) API, and the pluggable application client requires the user to provide this code. The pluggable application client uses the Sun Java Development Kit, and the thin application client uses the IBM Developer Kit For the Java Platform.

The J2EE application client programming model provides the benefits of the J2EE platform for the Java client application. The J2EE application client offers the ability to seamlessly develop, assemble, deploy and launch a client application. The tooling provided with the WebSphere platform supports the seamless integration of these stages to help the developer create a client application from start to finish.

When you develop a client application using and adhering to the J2EE platform, you can put the client application code from one J2EE platform implementation to another. The client application package can require redeployment using each J2EE platform deployment tool, but the code that comprises the client application does not change.

The application client run time supplies a container that provides access to system services for the client application code. The client application code must contain a main method. The application client run time invokes this main method after the environment initializes and runs until the Java virtual machine code terminates.

The J2EE platform allows the application client to use *nicknames* or *short names*, defined within the client application deployment descriptor. These deployment descriptors identify enterprise beans or local resources (JDBC, Java Message Service (JMS), JavaMail and URL APIs) for simplified resolution through JNDI use. This simplified resolution to the enterprise bean reference and local resource reference also eliminates changes to the client application code, when the underlying object or resource either changes or moves to a different server. When these changes occur, the application client can require redeployment.

The application client also provides initialization of the run-time environment for the client application. The deployment descriptor defines this unique initialization for each client application. The application client run time also provides support for security authentication to the enterprise beans and local resources.

The application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access enterprise bean references and to use CORBA services provided by the J2EE platform implementation. Use of the RMI-IIOP protocol and the accessibility of CORBA services assist users in developing a client application that requires access to both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments or programming models in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

Application client functions

Use the following table to identify the available functions in the different types of clients:

Available functions	ActiveX client	Applet client	J2EE client	Pluggable client	Thin client
Provides all the benefits of a J2EE platform	Yes	No	Yes	No	No
Portable across all J2EE platforms	No	No	Yes	No	No
Provides the necessary run-time to support communication between client and server	Yes	Yes	Yes	Yes	Yes
Allows the use of nicknames in the deployment descriptors	Yes	No	Yes	No	No
Supports use of the RMI-IIOP protocol	Yes	Yes	Yes	Yes	Yes
Browser based application	No	Yes	No	No	No
Enables development of client applications that can access enterprise bean references and CORBA object references	Yes	Yes	Yes	Yes	Yes
Enables the initialization of the client application run-time environment	Yes	No	Yes	No	No
Supports security authentication to enterprise beans	Yes	Limited	Yes	Yes	Yes
Supports security authentication to local resources	Yes	No	Yes	No	No

Requires distribution of application to client machines	Yes	No	Yes	Yes	Yes
Enables access to enterprise beans and other Java classes through Visual Basic, VBScript, and Active Server Pages (ASP) code	Yes	No	No	No	No
Provides a lightweight client suitable for download	No	Yes	No	Yes	Yes
Enables access to Java Naming and Directory Interface (JNDI) for enterprise bean resolution	Yes	Yes	Yes	Yes	Yes
Runs on client machines that use the Sun Java Runtime Environment	No	No	No	Yes	No
Supports CORBA services (using CORBA services can render the application client code nonportable)	No	No	Yes	No	No

ActiveX application clients

WebSphere Application Server provides an ActiveX to EJB bridge that enables ActiveX programs to access WebSphere Enterprise JavaBeans through a set of ActiveX automation objects.

The bridge accomplishes this by loading the Java virtual machine (JVM) into any ActiveX automation container such as Visual Basic, VBScript, and Active Server Pages (ASP).

There are two main environments in which the ActiveX to EJB bridge runs:

- **Client applications**, such as Visual Basic and VBScript, are programs that a user starts from the command line, desktop icon, or Start menu shortcut.
- **Client services**, such as Active Server Pages, are programs started by some automated means like the Services control panel applet.

The ActiveX to EJB bridge uses the Java Native Interface (JNI) architecture to programmatically access the JVM code. Therefore the JVM code exists in the same process space as the ActiveX application (Visual Basic, VBScript, or ASP) and remains attached to the process until that process terminates. To create JVM code, an ActiveX client program calls the `XJBInit()` method of the `XJB.JClassFactory`

object. For more information about creating JVM code for an ActiveX program, see [ActiveX to EJB bridge, initializing JVM code](#).

After an ActiveX client program has initialized the JVM code, the program calls several methods to create a proxy object for the Java class. When accessing a Java class or object, the real Java object exists in the JVM code; the automation container contains the proxy for that Java object. The ActiveX program can use the proxy object to access the Java class, object fields, and methods. For more information about using Java proxy objects, see [ActiveX to EJB bridge, using Java proxy objects](#). For more information about calling methods and access fields, see [ActiveX to EJB bridge, calling Java methods](#) and [ActiveX to EJB bridge, accessing Java fields](#).

The client program performs primitive data type conversion through the COM IDispatch interface (use of the IUnknown interface is not directly supported). Primitive data types are automatically converted between native Automation types and Java types. All other types are handled automatically by the Proxy Objects. For more information about data type conversion, see [ActiveX to EJB bridge, converting data types](#).

Any exceptions thrown in Java code are encapsulated and re-thrown as a COM error, from which the ActiveX program can determine the actual Java exceptions. For more information about handling exceptions, see [ActiveX to EJB bridge, handling errors](#).

The ActiveX to EJB bridge supports both free-threaded and apartment-threaded access and implements the Free Threaded Marshaler to work in a hybrid environment such as Active Server Pages. For more information about the support for threading, see [ActiveX to EJB bridge, using threading](#).

Applet clients

The applet client provides a browser-based Java run time capable of interacting with enterprise beans directly, instead of indirectly through a servlet.

This client is designed to support users who want a browser-based Java client application programming environment that provides a richer and more robust environment than the one offered by the Applet > Servlet > enterprise bean model.

The programming model for this client is a cross between the Java application thin client and a servlet client. When accessing enterprise beans from this client, the applet can consider the enterprise bean object references as CORBA object references.

No tooling support exists for this client to develop, assemble or deploy the applet. You are responsible for developing the applet, generating the necessary client bindings for the enterprise beans and CORBA objects, and bundling these pieces together to install or download to the client machine. The Java applet client provides the necessary run time to support communication between the client and the server.

Client side bindings generate using the (Application Assembly Tool). An applet can utilize these bindings, or you can generate client side bindings using the **rmic** command that is part of the IBM Developer Kit, Java edition, installed with the WebSphere Application Server.

The Applet client uses the RMI-IIOP protocol. Using this protocol enables the applet to access enterprise bean references and CORBA object references, but the applet is restricted in using some supported CORBA services.

If you combine the enterprise bean and CORBA environments in one applet, you must understand the differences between the two programming models, and you must use and manage each appropriately.

The applet client provides the run time to support the J2EE applet client. The applet client does not have tooling support for developing, assembling or deploying the applet. The applet client run time is provided through the Java applet browser plug-in that you install on the client machine using the WebSphere Application Server Client CD.

The applet environment restricts access to external resources from the browser run-time environment. You can make some of these resources available to the applet by setting the correct security policy settings in the WebSphere Application Server `client.policy` file. If given the correct set of permissions, the applet client must explicitly create the connection to the resource using the appropriate API. This client does not perform initialization of any service that the client applet can need. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or Java Naming and Directory Interface (JNDI) APIs.

J2EE application clients

The J2EE application client programming model provides the benefits of Java™ 2 Platform Enterprise Edition (J2EE).

The J2EE platform offers the ability to seamlessly develop, assemble, deploy and launch a client application. The tooling provided with the WebSphere platform supports the seamless integration of these stages to help the developer create a client application from start to finish.

When you develop a client application using and adhering to the J2EE platform, you can put the client application code from one J2EE platform implementation to another. The client application package can require redeployment using each J2EE platform deployment tool, but the code that comprises the client application does not change.

The J2EE application client run time supplies a container that provides access to system services for the application client code. The J2EE application client code must contain a main method. The J2EE application client run time invokes this main method after the environment initializes and runs until the Java virtual machine application terminates.

Application clients can use *nicknames* or *short names*, defined within the client application deployment descriptor with the J2EE platform. These deployment descriptors identify enterprise beans or local resources (Java Database Connectivity (JDBC), Java Message Service (JMS), JavaMail and URL APIs) for simplified resolution through JNDI use. This simplified resolution to the enterprise bean reference and local resource reference also eliminates changes to the application client code, when the underlying object or resource either changes or moves to a different server. When these changes occur, the application client can require redeployment.

The J2EE application client also provides initialization of the run-time environment for the client application. The deployment descriptor defines this unique initialization for each client application. The J2EE application client run time also provides support for security authentication to the enterprise beans and local resources.

The J2EE application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access enterprise bean references and to use CORBA services provided by the J2EE platform implementation. Use of the RMI-IIOP protocol and the accessibility of CORBA services assist users in developing a client application that requires access to both enterprise bean references and CORBA object references.

When you combine the J2EE and the CORBA enterprise edition environments or programming models in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

Pluggable application clients

The pluggable application client provides a lightweight, downloadable Java application run-time capable of interacting with enterprise beans.

The pluggable application client requires that you have previously installed the Sun Java Runtime Environment (JRE) files. In all other aspects, the pluggable application client, and the thin application client are similar.

This client is designed to support those users who want a lightweight Java client application programming environment, without the overhead of the J2EE platform on the client machine. The programming model for this client is heavily influenced by the CORBA programming model, but supports access to enterprise beans.

When accessing enterprise beans from this client, the client application can consider the enterprise beans object references as CORBA object references.

Tooling does not exist on the client, it exists on the server. You are responsible for developing the client application, generating the necessary client bindings for the enterprise bean and CORBA objects, and bundling these pieces together, install on the client machine.

The pluggable application client provides the necessary run-time to support the communication needs between the client and the server.

The pluggable application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access not only enterprise bean references and CORBA object references, but also allows the client application to use any supported CORBA services. Using the RMI-IIOP protocol along with the accessibility of CORBA services can assist a user in developing a client application that needs to access both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments in one client application, you must understand the differences between the two programming models, to use and manage each appropriately.

The pluggable application client run time provides the necessary support for the client application for object resolution, security, Reliability Availability and

Servicability (RAS), and other services. However, this client does not support a container that provides easy access to these services. For example, no support exists for using *nicknames* for enterprise beans or local resource resolution. When resolving to an enterprise bean (using either Java Naming and Directory Interface (JNDI) or CosNaming) sources, the client application must know the location of the name server and the fully qualified name used when the reference was bound into the name space. When resolving to a local resource, the client application cannot resolve to the resource through a JNDI lookup. Instead the client application must explicitly create the connection to the resource using the appropriate API (JDBC, Java Message Service (JMS), and so on). This client does not perform initialization of any of the services that the client application might require. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or JNDI APIs.

The pluggable application client offers access to most of the available client services in the J2EE application client. However, you cannot access the services in the pluggable client as easily as you can in the J2EE application client. The J2EE client has the advantage of performing a simple Java Naming and Directory Interface (JNDI) name space lookup to access the desired service or resource. The pluggable client must code explicitly for each resource in the client application. For example, looking up an enterprise bean Home requires the following code in a J2EE application client:

```
java.lang.Object ejbHome=initialContext.lookup("java:/comp/env/ejb/MyEJBHome");
MyEJBHome=(MyEJBHome) javax.rmi.PortableRemoteObject.narrow
(ejbHome,MyEJBHome.class);
```

However, you need more explicit code in a Java pluggable application client:

```
java.lang.Object ejbHome=initialContext.lookup("the/fully/qualified/path/to/
actual/home/in/namespace/MyEJBHome");
MyEJBHome=(MyEJBHome) javax.rmi.PortableRemoteObject.narrow
(ejbHome,MyEJBHome.class);
```

In this example, the J2EE application client accesses a logical name from the `java:/comp` name space. The J2EE client run time resolves that name to the physical location and returns the reference to the client application. The pluggable client, must know the fully qualified physical location of the enterprise bean Home in the name space. If this location changes, the pluggable client application must also change the value placed on the `lookup()` statement.

In the J2EE client, the client application is protected from these changes because it uses the logical name. A change can require a redeployment of the EAR file, but the actual client application code remains the same.

The pluggable application client is a traditional Java application that contains a *main* function. The WebSphere pluggable application client provides run-time support for accessing remote enterprise beans, and provides the implementation for various services (security, Workload Management (WLM), and others). This client can also access CORBA objects and CORBA based services. When using both environments in one client application, you need to understand the differences between the enterprise bean and CORBA programming models to manage both environments.

For instance, the CORBA programming model requires the CORBA CosNaming name service for object resolution in a name space. The enterprise beans programming model requires the JNDI name service. The client application must initialize and properly manage these two naming services.

Another difference applies to the enterprise bean model. Use the Java Naming and Directory Interface (JNDI) implementation in the enterprise bean model to initialize the Object Request Broker (ORB). The client application is unaware that an ORB is present. The CORBA model, however, requires the client application to explicitly initialize the ORB through the `ORB.init()` static method.

The pluggable application client provides a batch command that you can use to set the `CLASSPATH` and `JAVA_HOME` environment variables to enable the pluggable application client run time.

Thin application clients

The thin application client provides a lightweight, downloadable Java application run time capable of interacting with enterprise beans.

This client is designed to support those users who want a lightweight Java client application programming environment, without the overhead of the J2EE platform on the client machine. The programming model for this client is heavily influenced by the CORBA programming model, but supports access to enterprise beans.

When accessing enterprise beans from this client, the client application can consider the enterprise beans object references as CORBA object references.

Tooling does not exist on the client, it exists on the server. You are responsible for developing the client application, generating the necessary client bindings for the enterprise bean and CORBA objects, and bundling these pieces together to install on the client machine.

The thin application client provides the necessary run-time to support the communication needs between the client and the server.

The thin application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access not only enterprise bean references and CORBA object references, but also allows the client application to use any supported CORBA services. Using the RMI-IIOP protocol along with the accessibility of CORBA services can assist a user in developing a client application that needs to access both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments in one client application, you must understand the differences between the two programming models, to use and manage each appropriately.

The thin application client run time provides the necessary support for the client application for object resolution, security, Reliability Availability and Servicability (RAS), and other services. However, this client does not support a container that provides easy access to these services. For example, no support exists for using *nicknames* for enterprise beans or local resource resolution. When resolving to an enterprise bean (using either Java Naming and Directory Interface (JNDI) or CosNaming) sources, the client application must know the location of the name server and the fully qualified name used when the reference was bound into the name space. When resolving to a local resource, the client application cannot resolve to the resource through a JNDI lookup. Instead the client application must explicitly create the connection to the resource using the appropriate API (JDBC, Java Message Service (JMS), and so on). This client does not perform initialization

of any of the services that the client application might require. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or JNDI APIs.

The thin application client offers access to most of the available client services in the J2EE application client. However, you cannot access the services in the thin client as easily as you can in the J2EE application client. The J2EE client has the advantage of performing a simple Java Naming and Directory Interface (JNDI) name space lookup to access the desired service or resource. The thin client must code explicitly for each resource in the client application. For example, looking up an enterprise bean Home requires the following code in a J2EE application client:

```
java.lang.Object ejbHome=initialContext.lookup("java:/comp/env/ejb/MyEJBHome");
MyEJBHome=(MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome,MyEJBHome.class);
```

However, you need more explicit code in a Java thin application client:

```
java.lang.Object ejbHome=initialContext.lookup("the/fully/qualified/path/to/actual/
home/in/namespace/MyEJBHome");
MyEJBHome=(MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome,MyEJBHome.class);
```

In this example, the J2EE application client accesses a logical name from the `java:/comp` name space. The J2EE client run time resolves that name to the physical location and returns the reference to the client application. The thin client must know the fully qualified physical location of the enterprise bean Home in the name space. If this location changes, the thin client application must also change the value placed on the `lookup()` statement.

In the J2EE client, the client application is protected from these changes because it uses the logical name. A change might require a redeployment of the EAR file, but the actual client application code remains the same.

The thin application client is a traditional Java application that contains a *main* function. The WebSphere thin application client provides run-time support for accessing remote enterprise beans, and provides the implementation for various services (security, Workload Management (WLM), and others). This client can also access CORBA objects and CORBA based services. When using both environments in one client application, you need to understand the differences between the enterprise bean and CORBA programming models to manage both environments.

For instance, the CORBA programming model requires the CORBA CosNaming name service for object resolution in a name space. The enterprise beans programming model requires the JNDI name service. The client application must initialize and properly manage these two naming services.

Another difference applies to the enterprise bean model. Use the Java Naming and Directory Interface (JNDI) implementation in the enterprise bean model to initialize the Object Request Broker (ORB). The client application is unaware that an ORB is present. The CORBA model, however, requires the client application to explicitly initialize the ORB through the `ORB.init()` static method.

The thin application client provides a batch command that you can use to set the `CLASSPATH` and `JAVA_HOME` environment variables to enable the thin application client run time.

Example: Migrating application clients

Use the **ClientUpgrade** command to migrate application clients.

The command file is located in the bin subdirectory of the <WAS_install_root>, or the <ND_install_root> directory. By default, the *WAS_install_root* for WebSphere Application Server and WebSphere Application Server Enterprise is:

- Windows NT or Windows 2000 operating platforms - <drive>\WebSphere\AppServer directory
- AIX or UNIX-based operating platforms - /usr/WebSphere/AppServer directory

By default, the *ND_install_root* for WebSphere Application Server Network Deployment is:

- Windows NT or Windows 2000 operating platforms - <drive>\WebSphere\DeploymentManager directory
- AIX or UNIX-based operating platforms - /usr/WebSphere/DeploymentManager directory

The command uses the `com.ibm.websphere.migration.clientupgrade.ClientUpgrade` class.

Syntax

```
ClientUpgrade EAR_file [-clientJar client_jar ]  
                        [-traceString trace_spec [-traceFile file_name ]]
```

Parameters

Supported arguments include the following:

EAR_file

Use this parameter to specify the fully qualified path to the EAR file that contains client JAR files to process.

-clientJar

Use this optional parameter to specify a JAR file for processing. If not specified, the program transforms all client JAR files in the EAR file.

-traceString -traceFile

Use these optional parameters to gather trace information for IBM Service personnel. Specify a *trace_spec* of `"*=all=enabled"` (with quotation marks) to gather all trace information.

Examples

The following example demonstrates correct syntax:

```
ClientUpgrade EAR_file -clientJar ejbJarFile
```

Migration tips for application clients

Tips for migrating thin application client code:

- The Java invocation used to run thin application clients has changed in Version 5.0. You must specify `-Xbootclasspath/p:%WAS_BOOTCLASSPATH%` on Windows systems or `-Xbootclasspath/p:$WAS_BOOTCLASSPATH` on Unix systems when you invoke the Java command. Set the `WAS_BOOTCLASSPATH` environment variable in one of the following:
 - `setupClient.bat` for Windows systems or `setupClient.sh` for Unix systems
 - `setupCmdLine.bat` for Windows systems or `setupCmdLine.sh` for Unix systems

Tips for migrating J2EE application client code:

- If your J2EE application client uses resource references and you have configured those resources using the Application Client Resource Configuration Tool (ACRCT), you must run the **ClientUpgrade** command to migrate the resource configuration information in WebSphere Application Server V5.

Installing application clients

Before you begin

Application clients do not require the IBM WebSphere Application Server base. They are supported at the level of the operating system.

Before you install the pluggable application client, you must have one of the following Sun Java Runtime Environment (JRE) levels on your system: 1.3.1_03 and all levels up to but not including 1.4. Similar Sun Java Development Kit (JDK) levels are required to install the pluggable application client Samples.

Steps for this task

1. Issue the following command, according to platform, at a command line:
install.sh for AIX
Install.sh for Solaris
install.exe for Windows
The WebSphere Application Server client install wizard appears.
2. Click **next**, to continue.
3. Select a language from the drop-down menu and click **OK**.
If you have an earlier version of the WebSphere Application Server client installed on your machine, an information window appears.
4. Click **Yes to all** to overwrite this older version.
5. Read the license agreement, and select the radio button next to **I accept the terms in the license agreement**. Click **next**.
6. Choose a setup type and click **next**.
7. Choose the features you want to install and click **Next**.
8. Click **Yes**, if you want to install the Samples development environment.
9. Click **Next**, to install the WebSphere Application Server client to the default directory. Click **Browse**, to install the WebSphere Application Server client to a different directory.
10. Click **Next**.
11. Enter the host name of the server machine, and click **Next**.
Specify a port number if you are not using the default port.
12. Click **Next**, to install.
A progress window appears. If you chose to install the applet client, an information window appears.
13. Choose **yes** to overwrite the existing JRE registry entry. Choose **no** to return to the feature panel.
14. Click **Finish** to exit the wizard, after the WebSphere Application Server client installs.

Results

Application clients are installed on your machine.

Developing ActiveX application client code

Before you begin

This topic provides an outline for developing an ActiveX program, such as Visual Basic, VBScript, and Active Server Pages, to use the WebSphere ActiveX to EJB bridge to access enterprise beans.

This topic assumes that you are familiar with ActiveX programming. You should also consider the information given in ActiveX to EJB bridge, good programming guidelines.

To use the ActiveX to EJB bridge to access a Java class, develop your ActiveX program to complete the following steps:

Steps for this task

1. Create an instance of the XJB.JClassFactory object.
2. Create JVM code within the ActiveX program process, by calling the XJBInit() method of the XJB.JClassFactory object.
After the ActiveX program has created an XJB.JClassFactory object and called the XJBInit() method, the JVM code is initialized and ready for use.
3. Create a proxy object for the Java class, by using the XJB.JClassFactory FindClass() and NewInstance() methods.
The ActiveX program can use the proxy object to access the Java class, object fields, and methods.
4. Call methods on the Java class, using the Java method invocation syntax, and access Java fields as required.
5. Use the helper functions to do the conversion in cases where automatic conversion is not possible. You can convert between the following data types:
 - Java Byte and Visual Basic Byte
 - Visual Basic Currency types and Java 64-bit
6. Implement methods to handle any errors returned from the Java class. In Visual Basic or VBScript, use the Err.Number and Err.Description fields to determine the actual Java error.

What to do next

After you develop the ActiveX client code, start the ActiveX application.

Starting an ActiveX application

Before you begin

To run an ActiveX client application that is to use the ActiveX to EJB bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time. This initial configuration sets up the environment within which the ActiveX client application can run.

To perform the required configuration, complete one or more of the following subtasks:

Steps for this task

1. Starting an ActiveX application and configuring service programs

2. Starting an ActiveX application and configuring non-service programs

Starting an ActiveX application and configuring service programs

Before you begin

To run an ActiveX service program such as Active Server Page (ASP) that is to use the ActiveX to EJB bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time. This configuration sets up the environment within which the ActiveX service program can run.

The XJB.JClassFactory must find the Java run-time Dynamic Link Library (DLL) when initializing. In a service program such as Internet Information Server you cannot specify a path for its processes independently; you must set the process paths in the system PATH variable. This limitation means that you can only have a single JVM version available on a machine using ASP.

To add the JRE directories to your System path, complete one of the following subtasks:

Steps for this task

1. On Windows 2000, complete the following substeps:
 - a. Open the Control Panel, then double-click the System icon.
 - b. Click the Advanced tab on the System Properties window.
 - c. Click **Environment Variables**.
 - d. Edit the **Path** variable in the System Variables window.
 - e. Add the following to the beginning of the path displayed in the Variable Value input box:
`C:\WebSphere\AppClient\Java\jre\bin;C:\WebSphere\AppClient\Java\jre\bin\classic;`

where C:\WebSphere\AppClient is the directory in which you installed the WebSphere Java client
 - f. Click **OK** in the Edit System Variable window to apply the changes.
 - g. Click **OK** in the Environment Variables window.
 - h. Click **OK** in the System Properties window.
 - i. Restart Windows 2000.
2. On Windows NT, complete the following substeps:
 - a. Open the Control Panel, then double-click the System icon.
 - b. Click the Environment tab on the System Properties window.
 - c. In the System Variables window, edit the **Path** variable.
 - d. Add the following to the beginning of the path displayed in the Value input box:
`C:\WebSphere\AppClient\Java\jre\bin;C:\WebSphere\AppClient\Java\jre\bin\classic;`

Where C:\WebSphere\AppClient is the directory in which you installed the WebSphere Java client
 - e. Click **Set** to apply the changes.
 - f. Click **OK**.
 - g. Restart Windows NT.

What to do next

After you change the System PATH variable you must reboot the Internet Information Server machine so that Internet Information Server can see the change.

Starting an ActiveX application and configuring non-service programs

Before you begin

To run an ActiveX program initiated from an icon or command-line (a non-service program) that is to use the ActiveX to EJB bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java runtime. This uses a batch file to set up the environment within which the ActiveX program can run.

To perform the required configuration, complete the following steps:

Steps for this task

1. **(Optional)** Edit the setupCmdLineXJB.bat file to specify appropriate values for the environment variables required by the ActiveX to EJB bridge. For more information about these environment variables, see ActiveX to EJB bridge, environment and configuration.

For more information about creating a JVM for an ActiveX program, see ActiveX to EJB bridge, initializing the JVM.

After the ActiveX program has created an XJB.JClassFactory object and called the XJBInit() method, the JVM is initialized and ready for use.

2. Start the ActiveX client application by using one of the following methods:
 - Use the launchClientXJB.bat file to start the application; for example:
`launchClientXJB MyApplication.exe parm1 parm2`
 - or
`launchClientXJB MyApplication.vbp`
 - Use the setupCmdLineXJB.bat file to create an environment in which the application can be run, then start the application from within that environment.

setupCmdLineXJB.bat, launchClientXJB.bat, and other ActiveX batch files

This topic provides reference information about the aids that client applications and client services can use to access the ActiveX to EJB bridge. These enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time.

Location

The include file is located in the <was_client_home\aspIncludes directory. You can include the file into your ASP application with the following syntax in your ASP page:

```
<<-- #include virtual ="/WSASPIncludes/setupASPXJB.inc" -->
```

This assumes that you have created a virtual directory in Internet Information Server called **WSASPIncludes** that points to the <was_client_home\aspIncludes directory.

Usage notes

The following batch files are provided for client applications to use the ActiveX to EJB bridge:

- **setupCmdLineXJB.bat**

Sets the client environment variables.

- **launchClientXJB.bat**

Calls the setupCmdLineXJB.bat file and launches the application you specify as its arguments; for example:

```
launchClientXJB.bat myapp.exe parm1 parm2
```

or

```
launchClientXJB MyApplication.vbp
```

- **Active Server Pages (ASP) include file**

An include file is provided for ASP users to automatically set the following page-level (local) environment variables:

- **com_ibm_websphere_javahome** Path to the Java run-time directory installed with the WebSphere Advanced Server Client.
- **com_ibm_websphere_washome** Path to the WebSphere Advanced Server Client directory.
- **com_ibm_websphere_namingfactory** Sets the Java java.naming.factory.initial system property.
- **com_ibm_websphere_computername** (Optional) Name of the computer where the WebSphere Advanced Server Client is installed. If you intend to talk to a single specific computer, you are recommended to change this value to become the server name that you intend to access.

- **System Settings**

To enable the ActiveX to EJB bridge to access the Java run-time Dynamic Link Library (DLL), the following directories must exist in the system PATH environment variable:

```
was_client_home\java\jre\bin;was_client_home\java\jre\bin\classic
```

Where *was_client_home* is the name of the directory where you installed the WebSphere Application Server Client (for example, C:\WebSphere\AppClient).

Note: This technique enables only *one* Java run-time to activate on a machine, therefore all client services on that machine must use the same Java run-time. Client applications do not have this limitation because they each have their own private, non-system scope.

JClassProxy and JObjectProxy classes

This topic provides reference information about the object classes of the ActiveX to EJB bridge.

JClassFactory is the object used to access the majority of JVM features. It handles JVM initialization, accessing classes and creating class instances (objects). The majority of tasks for accessing your Java classes and objects are handled with the JClassProxy and JObjectProxy objects.

- XJBInit(String astrJavaParameterArray())

Initializes the JVM environment using an array of strings that represent the command line parameters you would normally send to the java.exe file.

If you have invalid parameters in the XJBInit() string array, the following error results:

Error: 0x6002 "XJBNI::Init() Failed to create VM" when calling XJBInit()

If you have C++ logging enabled, the activity log displays the invalid parameter.

- `JClassProxy FindClass(String strClassName)`

Uses the current thread class loader to load the specified fully qualified class name and returns a `JClassProxy` object representing the Java Class object.

- `JObjectProxy NewInstance()`

Creates a Class instance for the specified `JClassProxy` object using the parameters supplied to call the Class Constructor. For more information about using `JMethodArgs`, see *ActiveX to EJB bridge, calling Java methods*.

```
JObjectProxy NewInstance(JClassFactory obj, Variant vArg1, Variant vArg2,
    Variant vArg3, ...)
```

```
JObjectProxy NewInstance(JClassFactory obj, JMethodArgs args)
```

- `JMethodArgs GetArgsContainer()`

Returns a `JMethodArgs` object (Class instance).

You can create a `JClassProxy` object from the `JClassFactory.FindClass()` method and also from any Java method call that would normally return a Java Class object. You can use this object as if you had direct access to the Java Class object. All of the class static methods and fields are accessible as are the `java.lang.Class` methods. In case of a clash between static method names of the reflected user class and those of the `java.lang.Class` (for example, `getName()`), the reflected static methods would execute first.

For example, the following is a static method called `getName()`. The `java.lang.Class` object also has a method called `getName()`:

– In Java:

```
class foo{
    foo(){};
    public static String getName(){return "abcdef";}
    public static String getName2(){return "ghijkl";}
    public String toString2(){return "xyz";}
}
```

– In Visual Basic:

```
...
Dim clsFoo as Object
set clsFoo = oXJB.FindClass("foo")
clsFoo.getName() ' Returns "abcdef" from the static foo class
clsFoo.getName2() ' Returns "ghijkl" from the static foo class
clsFoo.toString() ' Returns "class foo" from the java.lang.Class object.
oFoo = oXJB.NewInstance(clsFoo)
oFoo.toString() ' Returns some text from the java.lang.Object's
                ' toString() method which foo inherits from.
oFoo.toString2() ' Returns "xyz" from the foo class instance
```

You can create a `JObjectProxy` object from the `JClassFactory.NewInstance()` method, and can be created from any Java method call that would normally return a Class instance object. You can use this object as if you had direct access to the Java object and can access all the static methods and fields of the object. All of object instance methods and fields are accessible (including those accessible through inheritance).

The `JMethodArgs` object is created from the `JClassFactory.GetArgsContainer()` method. Use this object as a container for method and constructor arguments. You must use this object when overriding the object type when calling a method (for example, when sending a `java.lang.String JProxyObject` to a constructor that normally takes a `java.lang.Object` type).

There are two groups of methods to add arguments to the collection: **Add** and **Set**. You can use **Add** to add arguments in the order that they are declared. Alternatively, you can use **Set** to set an argument based on its position in the argument list (where the first argument is in position 1).

For example, if you had a Java Object Foo that took a constructor of Foo(int, String, Object), you could use a JMethodArgs object as shown in the following code extract:

```
...
Dim oArgs as Object
set oArgs = oXJB.GetArgsContainer()

oArgs.AddInt(CLng(12345))
oArgs.AddString("Apples")
oArgs.AddObject("java.lang.Object", oSomeJObjectProxy)

Dim clsFoo as Object
Dim oFoo as Object
set clsFoo = oXJB.FindClass("com.mypackage.foo")
set oFoo = oXJB.NewInstance(clsFoo, oArgs)

' To reuse the oArgs object, just clear it and use the add method
' again, or alternatively, use the Set method to reset the parameters
' Here, we will use Set
oArgs.SetInt(1, CLng(22222))
oArgs.SetString(2, "Bananas")
oArgs.SetObject(3, "java.lang.Object", oSomeOtherJObjectProxy)

Dim oFoo2 as Object
set oFoo2 = oXJB.NewInstance(clsFoo, oArgs)
```

- **AddObject (String strObjectName, Object oArg)**

Adds an arbitrary object to the argument container in the next available position, casting the object to the class name specified in the first parameter. Arrays are specified using the traditional [] syntax; for example:

```
AddObject("java.lang.Object[] []", oMy2DArrayOfFooObjects)
```

or

```
AddObject("int[]", oMyArrayOfInts)
```

- **AddByte (Byte byteArg)**

Adds a primitive byte value to the argument container in the next available position.

- **AddBoolean (Boolean bArg)**

Adds a primitive boolean value to the argument container in the next available position.

- **AddShort (Integer iArg)**

Adds a primitive short value to the argument container in the next available position.

- **AddInt (Long lArg)**

Adds a primitive int value to the argument container in the next available position.

- **AddLong (Currency cyArg)**

Adds a primitive long value to the argument container in the next available position.

- **AddFloat (Single fArg)**

Adds a primitive float value to the argument container in the next available position.

- `AddDouble (Double dArg)`

Adds a primitive double value to the argument container in the next available position.

- `AddChar (String strArg)`

Adds a primitive char value to the argument container in the next available position.

- `AddString (String strArg)`

Adds the argument in string form to the argument container in the next available position.

- `SetObject (Integer iArgPosition, String strObjectName, Object oArg)`

Adds an arbitrary object to the argument container in the specified position casting it to the class name or primitive type name specified in the second parameter. Arrays are specified using the traditional `[]` syntax; for example:
`SetObject(1, "java.lang.Object[] []", oMy2DArrayOfFooObjects)`

or

`SetObject(2, "int[]", MyArrayOfInts)`

- `SetByte (Integer iArgPosition, Byte byteArg)`

Sets a primitive byte value to the argument container in the position specified.

- `SetBoolean (Integer iArgPosition, Boolean bArg)`

Sets a primitive boolean value to the argument container in the position specified.

- `SetShort (Integer iArgPosition, Integer iArg)`

Sets a primitive short value to the argument container in the position specified.

- `SetInt (Integer iArgPosition, Long lArg)`

Sets a primitive int value to the argument container in the position specified.

- `SetLong (Integer iArgPosition, Currency cyArg)`

Sets a primitive long value to the argument container in the position specified.

- `SetFloat (Integer iArgPosition, Single fArg)`

Sets a primitive float value to the argument container in the position specified.

- `SetDouble (Integer iArgPosition, Double dArg)`

Sets a primitive double value to the argument container in the position specified.

- `SetChar (Integer iArgPosition, String strArg)`

Sets a primitive char value to the argument container in the position specified.

- `SetString (Integer iArgPosition, String strArg)`

Sets a `java.lang.String` value to the argument container in the position specified.

- `Object Item(Integer iArgPosition)`

Returns the value of an argument at a specific argument position.

- `Clear()`

Removes all arguments from the container and resets the next available position to one.

- `Long Count()`

Returns the number of arguments in the container.

Java virtual machine initialization tips

Initialize the Java virtual machine (JVM) code with the ActiveX to enterprise JavaBeans bridge. For an ActiveX client program (Visual Basic, VBScript, or ASP) to access Java classes or objects, the first step that the program must do is to create JVM code within its process. To create JVM code, the ActiveX program calls the XJBInit() method of the XJB.JClassFactory object. When an XJB.JClassFactory object is created and the XJBInit() method called, the JVM is initialized and ready to use.

- To enable the XJB.JClassFactory to find the Java run-time Description Definition Language (DLL) when initializing, the JRE bin and bin\classic directories must exist in the system path environment variable.
- The XJBInit() method accepts only one parameter; an array of strings. Each string in the array represents a command line argument that for a Java program you would normally specify on the Java.exe command line. This string interface is used to set the classpath, stack size, heap size, and debug settings. You can get a listing of these parameters by typing `java -?` from the command line.
- If you set a parameter incorrectly, you receive a 0x6002 "Failed to initialize VM" error message.
- Due to the current limitations of Java Native Interface (JNI), you cannot unload or reinitialize JVM code after it has loaded. Therefore, after XJBInit() has been called once, subsequent calls have no effect other than to create a duplicate JClassFactory object for you to access. It is best to store your XJB.JClassFactory object globally and continue to reuse that object.
- The following Visual Basic extract shows an example of initializing JVM code:

```
Dim oXJB as Object
set oXJB = CreateObject("XJB.JClassFactory")
Dim astrJavaInitProps(0) as String
astrJavaInitProps(0) = _
    "-Djava.class.path=.;c:\myjavaclasses;c:\myjars\myjar.jar"
oXJB.XJBInit(astrJavaInitProps)
```

Example: Developing ActiveX to enterprise bean bridge, using Java proxy objects

To use Java proxy objects with the ActiveX to enterprise JavaBeans bridge:

- After an ActiveX client program (Visual Basic, VBScript, or ASP) has initialized the XJB.JClassFactory (and thereby the JVM), it can access Java classes and initialize Java objects. To do this, the client program uses the XJB.JClassFactory FindClass() and NewInstance() methods.
- In Java programming there are two ways to access Java classes: direct invocation through the Java compiler, and through the Java Reflection interface. Because the ActiveX to Java bridge needs no compilation and is a complete runtime interface to Java, it depends on the latter Reflection interface to access its classes, objects, methods, and fields. The XJB.JClassFactory FindClass() and NewInstance() methods behave very similarly to the Java Class.forName() and the Method.invoke() and Field.invoke() methods.
- XJB.JClassFactory.FindClass() takes the fully-qualified class name as its only parameter and returns a Proxy Object (JClassProxy). You can use the returned Proxy object like a normal Java Class object and call static methods and access static fields. You can also create a Class Instance (or object) from it, as described below. For example, the following Visual Basic code extract returns a Proxy object for the Java class `java.lang.Integer`:

```
...
Dim clsMyString as Object
Set clsMyString = oXJB.FindClass("java.lang.Integer")
```

- After the proxy is created, you can access its static information directly. For example, you can use the following code extract to convert a decimal integer to its hexadecimal representation.

```
...
Dim strHexValue as String
strHexValue = clsMyString.toHexString(CLng(255))
```

- The equivalent Java syntax is: `static String toHexString(int i)`. Because ints in Java programming are really 32-bits (which translates to Long in VB), the `CLng()` function converts the value from the default int to a long. Also, even though the `toHexString()` function returns a `java.lang.String`, the code extract does not return an Object Proxy. Instead, the returned `java.lang.String` is automatically converted to a native Visual Basic String.

To create an object from a class, you use the `JClassFactory.NewInstance()` method. This method creates an Object Instance and takes whatever parameters your Class Constructor needs. Once the object is created, you have access to all of its public instance methods and fields. For example, you can use the following Visual Basic code extract to create an instance of `java.lang.Integer`:

```
...
Dim oMyInteger as Object
set oMyInteger = oXJB.NewInstance(CLng(255))

Dim strMyInteger as String
strMyInteger = oMyInteger.toString
```

Example: Calling Java methods in the ActiveX to enterprise bean bridge

In the ActiveX to EJB bridge, methods are called using the native language method invocation syntax.

- The following are important differences between Java invocation and ActiveX Automation invocation:
 - Unlike Java methods, ActiveX does not support method (and constructor) polymorphism; that is, you cannot have two methods in the same class with the same name.
 - Java methods are case-sensitive, but ActiveX Automation is not case-sensitive.
- You should take care when invoking Java methods through ActiveX Automation. If you use the wrong case on a method call or use the wrong parameter type, you get an Automation Error 438 "Object doesn't support this property or method" thrown.
- To compensate for Java polymorphic behavior, give the exact parameter types to the method call. The parameter types determine the correct method to invoke. For a listing of correct types to use, see ActiveX to EJB bridge, converting data types.
- For example, the following Visual Basic code would fail if `CLng()` was not present or `toHexString` was incorrectly typed as `ToHexString`:

```
...
Dim strHexValue as String
strHexValue = clsMyString.toHexString(CLng(255))
```

- Sometimes it is difficult to force some development environments to leave the case of your method calls unchanged. For example, in Visual Basic if you want to call a method `close()` (uncapitalized), Visual Basic would try to capitalize it "Close()". In Visual Basic, the only way to effectively get around this behavior is to use the `CallByName()` method. For example:

```
o.Close(123) 'Incorrect...
CallByName(o, "close", vbMethod, 123) 'Correct...
```

or in VBScript, use the Eval function:

```
o.Close(123)                'Incorrect...
Eval("o.Close(123)")       'Correct...
```

- The return value of a function is always converted dynamically to the correct type. However, you must take care to use the **set** keyword in Visual Basic. If you expect a non-primitive data type to return, you must use **set**. (If you expect a primitive data type to return, you do not need to use **set**.) For example:

```
Set oMyObject = o.GetObject
iMyInt = o.GetInt
```

- In some cases, you might not know the type of object returning from a method call, because wrapper classes are converted automatically to primitives (for example, java.lang.Integer returns an ActiveX Automation Long). In such cases, you might need to use your language built-in exception handling techniques to try to coerce the returned type (for example, On Error and Err.Number in Visual Basic).

- Methods with Character Arguments

Because ActiveX automation does not natively support character types supported by Java methods, the ActiveX to EJB bridge uses strings (byte or VT_I1 do not work, because characters have multiple bytes in Java). If you try to call a method that takes a char or java.lang.Character type you must use the JMethodArgs argument container to pass character values to methods or constructors. For more information about how this argument container is used, see Methods with "Object" Type as Argument and Abstract Arguments.

- Methods with "Object" Type as Argument and Abstract Arguments

Because of the polymorphic nature of Java programming, the ActiveX to Java bridge uses direct argument type mapping to find a method. This works well in most cases, but sometimes methods are declared with a Parent or Abstract Class as an argument type (for example, java.lang.Object). You need the ability to send an object of arbitrary type to a method. To acquire this ability, you must use the XJB.JMethodArgs object to coerce your parameters to match the parameters on your method. You can get a JMethodArgs instance by using the JClassFactory.GetArgsContainer() method.

The JMethodArgs object is a container for method parameters or arguments. This container enables you to add parameters to it one-by-one and then you can send the JMethodArgs object to your method call. The JClassProxy and JObjectProxy objects recognize the JMethodArgs object and attempt to find the correct method and let the Java language coerce your parameters appropriately.

For example, to add an element to a Hashtable object the method syntax is Object.put(Object key, Object value). In Visual Basic, the method usage looks like this:

```
Dim oMyHashtable as Object
Set oMyHashtable = _
    oXJB.NewInstance(oXJB.FindClass("java.util.Hashtable"))
```

```
' This line will not work. The ActiveX to EJB bridge cannot find a method
' called "put" that has a short and String as a parameter:
```

```
oMyHashtable.put 100, "Dogs"
oMyHashtable.put 200, "Cats"
```

```
' You must use a XJB.JMethodArgs object instead:
```

```
Dim oMyHashtableArgs as Object
Set oMyHashtableArgs = oXJB.GetArgsContainer
oMyHashtableArgs.AddObject("java.lang.Object", 100)
oMyHashtableArgs.AddObject("java.lang.Object", "Dogs")
```

```
oMyHashtable.put oMyHashtableArgs
```

```
' Reuse the same JMethodArgs object by clearing it.
oMyHashtableArgs.Clear
oMyHashtableArgs.AddObject("java.lang.Object", 200)
oMyHashtableArgs.AddObject("java.lang.Object", "Cats")

oMyHashtable.put oMyHashTableArgs
```

Java field programming tips

Using the ActiveX to EJB bridge to access Java fields has the same case sensitivity issue that it has when invoking methods. Field names must use the same case as the Java field syntax.

- Visual Basic has the same problem with unsolicited case changing on fields as it does with methods. (For more information about this problem, see ActiveX to EJB bridge, calling Java methods). You might need to use the CallByName() function to set a field in the same way that you would call a method in some cases. For Fields, you use VBLet for primitive types and VBSet for Objects. For example:

```
o.MyField = 123 'Incorrect...
CallByName(o, "MyField", vbLet, 123) 'Correct...
```

or in VBScript:

```
o.MyField = 123 'Incorrect...
Eval("o.myField = 123") 'Correct...
```

ActiveX to Java primitive data type conversion values

All primitive Java data types are automatically converted to native ActiveX Automation types. However, not all Automation data types are converted to Java types (for example, VT_DATE). Variant data types are used for data conversion. Variant data types are a requirement of any Automation interface, and are used automatically by Visual Basic and VBScript. The tables below provide details about how primitive data types are converted between Automation types and Java types.

ActiveX to Java primitive data type conversion

Visual Basic Type	Variant Type	Java Type	Notes
Byte	VT_I1	byte	Byte in Visual Basic is unsigned, but is signed in Java data type.
Boolean	VT_BOOL	boolean	
Integer	VT_I2	short	
Long	VT_I4	int	
Currency	VT_CY	long	
Single	VT_R4	float	
Double	VT_R8	double	
String	VT_BSTR	java.lang.String	
String	VT_BSTR	char	
Date	VT_DATE	n/a	

Example: Using helper methods for data type conversion

Generally, data type conversion between ActiveX (Visual Basic and VBScript) and Java methods occurs automatically, as described in ActiveX to EJB bridge,

converting data types. However, the following helper functions are provided for cases where automatic conversion is not possible:

- Byte helper function
- Currency helper function
- Byte helper function

Because the Java Byte data type is signed (-127 through 128) and the Visual Basic Byte data type is unsigned (0 through 255), convert unsigned Bytes to a Visual Basic Integers, which look like the Java signed byte. To make this conversion, you can use the following helper function:

```
Private Function GetIntFromJavaByte(Byte jByte) As Integer
    GetIntFromJavaByte = (CInt(jByte) + 128) Mod 256 - 128
End Function
```

- Currency helper function

Visual Basic 6.0 cannot properly handle 64-bit integers like Java methods can (as the Long data type). Therefore, Visual Basic uses the Currency type, which is intrinsically a 64-bit data type. The only side effect of using the Currency type (the Variant type VT_CY) is that a decimal point is inserted into the type. To extract and manipulate the 64-bit Long value in Visual Basic, use code like the following example. For more details on this technique for converting Currency data types, see Q189862, "HOWTO: Do 64-bit Arithmetic in VBA", on the Microsoft Knowledge Base.

```
' Currency Helper Types
Private Type MungeCurr
    Value As Currency
End Type
Private Type Munge2Long
    LoValue As Long
    HiValue As Long
End Type

' Currency Helper Functions
Private Function CurrToText(ByVal Value As Currency) As String
    Dim Temp As String, L As Long
    Temp = Format$(Value, "#.0000")
    L = Len(Temp)
    Temp = Left$(Temp, L - 5) & Right$(Temp, 4)
    Do While Len(Temp) > 1 And Left$(Temp, 1) = "0"
        Temp = Mid$(Temp, 2)
    Loop
    Do While Len(Temp) > 2 And Left$(Temp, 2) = "-0"
        Temp = "-" & Mid$(Temp, 3)
    Loop
    CurrToText = Temp
End Function

Private Function TextToCurr(ByVal Value As String) As Currency
    Dim L As Long, Negative As Boolean
    Value = Trim$(Value)
    If Left$(Value, 1) = "-" Then
        Negative = True
        Value = Mid$(Value, 2)
    End If
    L = Len(Value)
    If L < 4 Then
        TextToCurr = CCur(IIf(Negative, "-0.", "0.") & _
            Right$("0000" & Value, 4))
    Else
        TextToCurr = CCur(IIf(Negative, "-", "") & _
            Left$(Value, L - 4) & "." & Right$(Value, 4))
    End If
End Function
```

```

' Java Long as Currency Usage Example
Dim LC As MungeCurr
Dim L2 As Munge2Long

' Assign a Currency Value (really a Java Long)
' to the MungeCurr type variable
LC.Value = cyTestIn

' Coerce the value to the Munge2Long type variable
LSet L2 = LC

' Perform some operation on the value, now that we
' have it available in two 32-bit chunks
L2.LoValue = L2.LoValue + 1

' Coerce the Munge value back into a currency value
LSet LC = L2
cyTestIn = LC.Value

```

Array tips for ActiveX application clients

Arrays are very similar between Java and Automation Containers like Visual Basic and VBScript. Here are some important points to consider when passing arrays back and forth between these containers:

- Java arrays cannot mix types. All Java arrays contain a single type, so when passing arrays of Variants to a Java Array, you must make sure that all of the elements in the Variant array are of the same base type. For example, in Visual Basic:

```

...
Dim VariantArray(1) as Variant
VariantArray(0) = CLng(123)
VariantArray(1) = CDb1(123.4)
oMyJavaObject.foo(VariantArray) ' Illegal!

VariantArray(0) = CLng(123)
VariantArray(1) = CLng(1234)
oMyJavaObject.foo(VariantArray) ' This works

```

- Arrays of Primitive Types are converted using the rules defined in Primitive Data Type Conversion.
- Arrays of Java Objects are handled through arrays of JObjectProxy objects.
- Arrays of JObjectProxy objects must be fully-initialized and of the correct associated Java type. When initializing an array in Visual Basic (for example, Dim oJavaObjects(1) as Object), you must set each object to a JObjectProxy before you send the array to Java. The bridge is unable to determine the type of null or empty Object values.
- When receiving an array from a Java method, the lower-bound is always zero. Java methods only support zero-based arrays.
- Nested or multi-dimensional arrays are treated as zero-based multi-dimensional arrays in Visual Basic and VBScript.
- Uninitialized arrays or Array Types are unsupported. When calling a Java method that takes an array of objects as a parameter, you must fully initialize the array of JObjectProxy objects.

Error handling codes for ActiveX application clients

All exceptions thrown in Java code are encapsulated and re-thrown as a COM error through the ISupportErrorInfo interface and the EXCEPINFO structure of IDispatch::Invoke(); the Err object in Visual Basic and VBScript. Because there are

no error numbers associated with Java exceptions, whenever a Java exception is thrown, the entire stack trace is stored in the error description text and the error number assigned is 0x6003.

In Visual Basic or VBScript, you need to use the Err.Number and Err.Description fields to determine the actual Java error. Non-Java errors are thrown as you would expect via the IDispatch interface; for example, if a method cannot be found, then error 438 "Object doesn't support this property or method" is thrown.

Error number	Description
0x6001	Java Native Interface (JNI) error
0x6002	Initialization error
0x6003	Java exception. Error description is the Java Stack Trace.
0x6FFF	General Internal Failure

Threading tips

The ActiveX to EJB bridge supports both free-threaded and apartment-threaded access and implements the Free Threaded Marshaler to work in a hybrid environment such as Active Server Pages. Each thread created in the ActiveX process is mirrored in the Java environment when the thread communicates through the ActiveX to EJB bridge. In addition, once all references to Java objects (there are no JObjectProxy or JClassProxy objects) are loaded in an ActiveX thread, the ActiveX to EJB bridge detaches the thread from the JVM code. Therefore, you must be careful that any Java code that you access from a multi-threaded Windows application is thread-safe. Visual Basic and VBScript applications are both essentially single-threaded. Therefore, Visual Basic and VBScript applications do not have threading issues in the Java programs they access. Active Server Pages and multi-threaded C and C++ programs can have issues.

Consider the following scenario:

1. A multi-threaded Windows Automation Container (our ActiveX Process) starts. It exists on Thread A.
2. The ActiveX Process initializes the ActiveX to EJB bridge, which starts the JVM code. The JVM attaches to the same thread and internally calls it Thread1.
3. The ActiveX Process starts two threads: B and C.
4. Thread B in the ActiveX Process uses the ActiveX to EJB bridge to access an object that was created in Thread A. The JVM attaches to thread B and calls it Thread 2.
5. Thread C in the ActiveX Process never talks to the JVM code, so the VM never needs to attach to it. This is a case where the JVM code does not have a one-to-one relationship between ActiveX threads and Java threads.
6. Thread B later releases all of the JObjectProxy and JClassProxy objects that it used. The Java Thread 2 is detached.
7. Thread B again uses the ActiveX to EJB bridge to access an object that was created in Thread A. The JVM code attaches again to the thread and calls it Thread 3.

ActiveX Process	JVM Access by ActiveX Process
Thread A - Created in 1	Thread 1 - Attached in 2

ActiveX Process	JVM Access by ActiveX Process
Thread B - Created in 4	Thread 2 - Attached in 4, detached in 6 Thread 3 - Attached in 7
Thread C - Created in 4	

Threads and Active Server Pages

Active Server Pages (ASP) in Microsoft's Internet Information Server is a multi-threaded environment. When you create the XJB.JClassFactory object, you can store it in the Application collection as an Application-global object. All threads within your ASP environment can now access the same ActiveX to EJB bridge object. Active Server Pages by default creates 10 Apartment Threads per ASP process per CPU. This means that when your ActiveX to EJB bridge object is initialized any of the 10 threads can call this object, not just the thread that created it.

If you need to simulate single-apartment behavior, you can create a Single-Apartment Threaded ActiveX Dynamic Link Library (DLL) in Visual Basic and encapsulate the ActiveX to EJB bridge object there. This encapsulation guarantees that all access to the JVM object is on the same thread. You need to use the <OBJECT> tag to assign the XJB.JClassFactory to an Application object and must be aware of the consequences of introducing single-threaded behavior to a Web application.

The Microsoft KnowledgeBase has several articles about ASP and threads, including:

- Q243543 INFO: Do Not Store STA Objects in Session or Application
- Q243544 INFO: Component Threading Model Summary Under Active Server Pages
- Q243548 INFO: Design Guidelines for VB Components Under ASP

Example: Viewing System.out message

The ActiveX to EJB bridge does not have a console available to view Java System.out messages. To view these messages when running a standalone client program (such as Visual Basic), you need to redirect the output to a file. For example:

```
launchClientXJB.bat MyProgram.exe > output.txt
```

- To view the System.out messages when running a Service program such as Active Server Pages, you need to override the Java System.out OutputStream object to FileOutputStream. For example, in VBScript:

```
'Redirect system.out to a file
' Assume that oXJB is an initialized XJB.JClassFactory object
Dim clsSystem
Dim oOS
Dim oPS
Dim oArgs

' Get the System class
Set clsSystem = oXJB.FindClass("java.lang.System")

' Create a FileOutputStream object
' Create a PrintStream object and assign to it our FileOutputStream
Set oArgs = oXJB.GetArgsContainer oArgs.AddObject "java.io.OutputStream", oOS
```

```

Set oPS = oXJB.NewInstance(oXJB.FindClass("java.io.PrintStream"), oArgs)

' Set our System OutputStream to our file
clsSystem.setOut oPS

```

Example: Enabling logging and tracing for application clients

The ActiveX to EJB bridge provides two logging and tracing formats: Windows Application Event Log and Java Trace Log.

- Windows Application Event Log

The Windows Application Event Log shows JNI errors, Java console error messages, and XJB initialization messages. This log is most useful for determining `XJBInit()` errors and any unusual exceptions that do not come from the Java environment. By default, critical error logging will be enabled and debug and event logging will be disabled.

To enable or disable logging certain types of events to the Windows Application Event Log, you need to specify one or more parameters to `XJBInit()`. If more than one parameter is set, they will be processed in the order in which they appear in the input string array to `XJBInit()`. Once `XJBInit()` is initialized, these parameters can no longer be set/reset for the life of the process. Using Java `java.lang.System.setProperty()` to set these values also will have no effect.

- `-Dcom.ibm.ws.client.xjb.native.logging.debug=enabled|disabled`
Enables or disables debug level messages from appearing in the Windows NT event log. This level of logging is most useful and shows most internal errors, user programming issues or configuration problems.
- `-Dcom.ibm.ws.client.xjb.native.logging.event=enabled|disabled`
Enables or disables event level messages from appearing in the Windows NT event log.
- `-Dcom.ibm.ws.client.xjb.native.logging.*=enabled|disabled`
Enables or disables both event and debug level messages from appearing in the Windows NT event log. It is not possible to disable some critical error messages from being displayed in the error log. Only debug and event level messages can be disabled.

Viewing the Windows application event log with the event viewer:

To open the event viewer in the Windows operating system, click Start > Settings > Control Panel. Double-click Administrative Tools, and then double-click Event Viewer. All ActiveX to EJB bridge events will have the text "WebSphere XJB" in the source column and will appear in the Application log. For information about using Event Viewer, click the Action menu in Event Viewer, and then click Help.

To open the even viewer in the Windows NT operating system,, click Start > Programs > Administrative Tools > Event Viewer. All ActiveX to EJB bridge events will have the text "WebSphere XJB" in the source column and will appear in the Application log. For information about using Event Viewer, click the Help menu in Event Viewer.

- Java Trace Log

The Java trace log displays information that you can use to debug method calls, class lookups, and argument coercion problems. Since the Java portion of the bridge mirrors the functionality of the COM IDispatch interface, the information

in the trace log is similar to what you have come to expect from an IDispatch interface. To understand the trace log, you need a fundamental understanding of IDispatch.

To enable user-logging, add the following parameters to the XJBInit() input string array:

```
"-DtraceString=com.ibm.ws.client.xjb.*=event=enabled"  
"-DtraceFile=C:\MyTrace.txt"
```

ActiveX client programming best practices

In general, the best way to access Java components is to use the Java language. It is recommended that you do as much programming as possible in the Java language and use a small simple interface between your COM Automation container (for example, Visual Basic) and the Java code. This interface avoids any overhead and performance problems that can occur when moving across the interface.

- Visual Basic guidelines
- Active Server Pages guidelines
- J2EE guidelines

bestvb">

The following guidelines are intended to help optimize your use of the ActiveX to EJB bridge with Visual Basic:

- Launch the Visual Basic replication through the `launchClientXJB.bat` file. If you want to run your Visual Basic application through the Visual Basic debugger, run the Visual Basic Integrated Development Environment (IDE) within the ActiveX to EJB bridge environment. After you create your Visual Basic project, you can launch it from a command line; for example, `launchClientXJB MyApplication.vbp`. You can also launch the Visual Basic application alone in the ActiveX to EJB environment, by changing the Visual Basic shortcut on the Windows Start menu so that the `launchClientXJB.bat` file precedes the call to the `VB6.EXE` file.

- Exit the Visual Basic IDE before debugging programs.

Because the Java virtual machine (JVM) code attaches to the running process, you must exit out of the Visual Basic editor before debugging your program. If you run then exit your program within the Visual Basic IDE, the JVM code continues to run and you reattach the same JVM code when `XJBInit()` is called by the debugger. This causes problems if you try to update `XJBInit()` arguments (for example, `classpath`) because the changes are not be applied until you restart Visual Basic.

- Store the `XJB.JClassFactory` object globally.

Because you cannot unload or reinitialize the JVM code, cache the resulting `XJB.JClassFactory` object as a global variable. The overhead of treating this object as a global variable or passing a single reference around is much less than recreating a new `XJB.JClassFactory` object and calling the `XJBInit()` argument more than once.

estwc">

The following guidelines intend to help optimize your use of the ActiveX to EJB bridge with CScript and Windows Scripting Host (WSH):

- Launch in ActiveX to EJB environment.

Launch the VBScriptfiles in the ActiveX to EJB bridge environment, to run VBScript files in `.vbs` files. Two common ways exist to launch your script:

- launchClientXJB MyScript.vbs
- launchClientXJB cscript MyScript.vbs

bestas">

The following guidelines intend to help optimize your use of the ActiveX to EJB bridge with Active Server Pages software:

- Use the ActiveX to EJB Helper functions from the Active Server Pages Application.

Because Active Server Pages (ASP) code typically use VBScript, you can use the included helper functions in any VBScript environment with minor changes. For more information about these helper functions, see *Helper functions for data type conversion*. To run outside of the ASP environment, remove or change all references to the Server, Request, Response, Application and Session objects; for example, change `Server.CreateObject` to `CreateObject`.

- Set JRE path globally in system.

The XJB.JClassFactory object must be able to find the Java run time Dynamic Link Library (DLL) when initializing. In Internet Information Server, you cannot specify a path for its processes independently; you must set the process paths in the system PATH variable. You can only have a single JVM version available on a machine using the ASP application. Also, remember that after you change the system PATH variable you must reboot the Internet Information Server machine so that Internet Information Server can see the change.

- Set the system TEMP environment variable.

If the system TEMP environment variable is not set, Internet Information Server stores all temporary files in the WINNT directory, which is usually not desired.

- Use high isolation or an isolated process.

When using the ActiveX to Java bridge with Active Server Pages software, creating your Web application in its own process is recommended. You can only load one JVM instruction in a single process and if you want to have more than one application running with different JVM environment options (for example, different classpaths), then you need to have separate processes.

- Use the Application Unload option.

When debugging your application, use **Unload** when viewing your ASP application properties in the Internet Information Server administration console to unload the process from memory and thereby unload the JVM code.

- Run one process per application.

Use only one ASP application per J2EE application or JVM environment, in your ASP environment. If you need separate classpaths or JVM settings you need separate ASP applications (virtual directories with high isolation or an isolated process).

- Store the XJB.JClassFactory object in application scope.

Because of the one-to-one relationship required between a JVM instruction and a process, and because the JVM code can never detach or shut down from a process independently, cache the XJB.JClassFactory object at application scope and call the XJBInit() method only once.

Because the ActiveX to EJB bridge employs a free-threaded marshaler, take advantage of the multi-threaded nature of Internet Information Server and the ASP environment. If you choose to reinitialize the XJB.JClassFactory object at Page scope (local variables), then the XJBInit() method can only initialize your local XJB.JClassFactory variable. It is more efficient to use the XJBInit() method once.

- Use VBScript conversion functions.
Because VBScript code only supports variant data types, use the CStr(), CByte(), CBool(), CCur(), CInt(), CInq(), CSng() and CDbl() functions to tell the activeX to EJB bridge which data type you are using; for example `oMyObject.Foo(CDbl(1.234))`.

The following guidelines are intended to help optimize your use of the ActiveX to EJB bridge with the J2EE environment;

- Store client container objects globally.

Because you can only have one JVM instruction per process, and a single J2EE client container (`com.ibm.websphere.client.applicationclient.launchClient`) per JVM instruction, initialize your J2EE client container only once and reuse it. For ASP applications, store the J2EE client container in an application level variable and initialize it only once (either on the `Application_OnStart()` event in the `global.asa` file or by checking to see if it `IsEmpty()`).

A side effect to storing the client container object globally is that you cannot change the client container parameters without destroying the object and creating a new one. These parameters include the EAR file, `BootstrapHost`, `classpath`, and so on. If you run a Visual Basic application and want to change the client container parameters, you must end the application and restart it. If you run an Active Server Pages application, you must first unload the application from Internet Information Server (see "Use the Application Unload Button" under Active Server Pages guidelines). Then load the Active Server Pages application with the different client container parameters. The parameters set the first time the Active Server Pages application loads. Since the client container is stored on the Internet Information Server, all the browser clients share the parameters using the Active Server Pages application. This behavior is normal for Active Server Pages code, but can be confusing when you try to run to different WebSphere Application Servers using the same Active Server Pages applicatio, which is unsupported.

- Reuse custom temp directory for EAR file extraction.

By default, the client container launches and extracts the application EAR file to your temp directory and then sets up the thread `ClassLoader` to use the extracted EAR file directory and JAR files included in the client JAR manifest. This process is time consuming and because of some limitations with JVM shutdown through Java Native Interface (JNI) and file locking, these files are never cleaned up.

Specifically, each time the client container `launch()` method is called, it extracts the EAR file to a random directory name in your temporary directory on your hard drive. The current Java thread class loader is then changed to point to this extracted directory which in turn locks the files within. In a normal J2EE Java client, these files automatically clean up after the application exits. This cleanup occurs when the client container shutdown hook is called (which never happens in the ActiveX to EJB bridge), which leaves the temporary directory there.

To avoid these problems, you can specify a directory to extract the EAR file by setting the `com.ibm.websphere.client.applicationclient.archivedir` Java system property before calling the client container `launch()` method. If the directory does not exist or is empty, you extract the EAR file normally. If the EAR file was previously extracted, the directory is reused. This feature is particularly important for server processes (for example, ASP), which can stop and restart, potentially calling `launchClient()` several times.

If you need to update your EAR file, delete the temporary directory first. The next time you create the client container object, it extracts the new EAR file to the temporary directory. If you do not delete the temporary directory or change

the system property value to point to a different temporary directory, the client container reuses the currently extracted EAR file, and does not use your changed EAR file.

Note: When specifying the `com.ibm.websphere.client.applicationclient.archivedir` property, make sure that the directory you specify *is unique* for each EAR file you use. For example, do not point `MyEar1.ear` and `MyEar2.ear` files to the same directory.

If you choose not to use this system property, go regularly to your Windows temp directory and delete the `WSTMP*` subdirectories. Over a relatively short period of time, these subdirectories can waste a very significant amount of space on the hard drive.

Developing applet client code

Before you begin

Applet clients have the following setup requirements:

- These clients are currently available on the Windows NT or Windows 2000 platforms. Check the prerequisites page for information on new platform support.
- They require one of these browsers:
 - Internet Explorer version 5.0+
 - Netscape Navigator 4.7+
- You must install the browser before installing the client code.

Unlike typical applets that reside on either Web servers or WebSphere Application Servers and can only communicate using the HTTP protocol, applet clients are capable of communicating over the HTTP protocol and the RMI-IIOP protocol. This additional capability gives the applet direct access to enterprise beans.

Steps for this task

1. Run the application server client installation.
2. Select the applet client option.
3. Install an applet client.
4. Install the WebSphere Application Server Plug-in for the browser.

From the WebSphere Application Server Java Plug-in Control panel, enter the following:

```
-Djava.security.policy=<product_installation_dir>\properties\client.policy
-Dwas.install.root=<product_installation_dir>
-Djava.ext.dirs=<product_installation_dir>\classes;
<product_installation_dir>\java\jre\lib\ext;
<product_installation_dir>\java\jre\lib;
<product_installation_dir>\lib;
<product_installation_dir>\properties
-Dcom.ibm.CORBA.securityEnabled=false
-Dcom.ibm.CORBA.ConfigURL=file:<product_installation_dir>\properties\
sas.client.props
-classpath <product_installation_dir>\properties
```

Note: The above entries are automatically placed into the WebSphere Application Server Java Plug-in control panel for the user who installed the WebSphere Application Server client. If this sample is being run by a user other than the person who installed the client, the user must enter the entries.

- The *Java Run Time Parameters* field is similar to the command prompt when using command line options. Therefore, you can enter most options available from the command prompt (for example, `-cp`, `classpath`, and others) in this field as well.
- Access the control panel from the **Start** menu. Click **start > Control panel > WebSphere Java Plug-in**.
- The applet container is the Web browser and the Java plug-in combination. You must first install the WebSphere Application Server Applet client so that the browser recognizes the IBM Java Plug-in.

View the Samples gallery for more information about application clients. Before you run the `basicCalculator` Sample, ensure the JMS Server is started.

Accessing secure resources using the TCP/IP protocol for applet clients

You cannot use the WebSphere Application Server JSSE from an applet, and applets cannot authenticate to WebSphere Application Server using Secure Sockets Layer (SSL). To access resources that are not secure, such as enterprise beans, applets are configured with security disabled. Applets can access secure resources using standard TCP/IP protocol. This will cause the `userid` and `password` to flow from the client machine to the application server in the clear. This may compromise your password should someone be monitoring the network. It is important to understand the security implications before you modify the client configuration to use the standard TCP/IP protocol to access secure resources on your application server.

Steps for this task

1. Make a copy of the following file so that you can use it for an applet:

```
<product_install_directory>/properties/sas.client.props
```

It is recommended to use a copy of the `sas.client.props` file for your applet because this file is used for all WebSphere Application Server client applications.

2. Edit the copy of `sas.client.props` file that you made with the following changes:
 - Change `com.ibm.CSI.protocol` to `com.ibm.CSI.protocol=csiv2`.
 - Change `com.ibm.CSI.performTransportAssocSSLTLSSupported` to `com.ibm.CSI.performTransportAssocSSLTLSSupported=false`.
 - Add `com.ibm.CSI.claimTransportAssocSSLTLSSupported=false`.
3. Click **Start > Control panel > WebSphere Java Plug-in** to open the WebSphere Application Server Java control panel.
 - Change `-Dcom.ibm.CORBA.securityEnabled=false` to `-Dcom.ibm.CORBA.securityEnabled=true`.
 - To use the file you created in step number 1, modify the following value:


```
-Dcom.ibm.CORBA.ConfigURL=file:<product_install_directory>\properties\sas.client.props
```

For more information on the `sas.client.props` file and WebSphere Application Server security, view **Security**.

Applet client security requirements

When code is loaded, it is assigned permissions based on the security policy in effect. This policy specifies the permissions that are available for code from various

locations. You can initialize this policy from an external policy file. By default, the client uses the `<product_installation_dir>/properties/client.policy` file. You must update this file with the following permissions:

- The `SocketPermission` grants permission to open a port and make a connection to a host machine, which is your WebSphere Application Server. In the following example, `yourserver.yourcompany.com` is the complete hostname of your WebSphere Application Server:

```
permission java.util.PropertyPermission "*", "read";
permission java.net.SocketPermission "yourserver.yourcompany.com", "connect";
```

Applet client tag requirements

Standard applets require the HTML `<APPLET>` tag to identify the applet to the browser. The `<APPLET>` tag invokes the Java Virtual Machine (JVM) of the browser.

- For applets to communicate with EJBs in the WebSphere Application Server environment, the `<APPLET>` tag must be replaced with the following tags:

```
<OBJECT>
<EMBED>
```

- The `classid` and `type` attributes cannot be modified, and must be entered as described in the applet client example. The `codebase` attribute on the `<OBJECT>` tag must be excluded. Do not confuse the `codebase` attribute on the `<OBJECT>` tag with the `codebase` attribute on the `<PARAM>` tag. Although both are called `codebase`, they are separate entities.
- The following code example illustrates the applet code. In this example, `MyApplet.class` is the applet code, `applet.jar` is the file that contains the applet code, and `EJB.jar` is the file that contains the enterprise bean code:

```
<OBJECT classid="clsid:8AE2D840-EC04-11D4-AC77-006094334AA9"
width="600" height="500">
<PARAM NAME=CODE VALUE=MyAppletClass.class>
<PARAM NAME="archive" VALUE='Applet.jar, EJB.jar'>
<PARAM TYPE="application/x-java-applet;version=1.3">
<PARAM NAME="scriptable" VALUE="false">
<PARAM NAME="cache-option" VALUE="Plugin">
<PARAM NAME="cache-archive" VALUE="Applet.jar, EJB.jar">
<COMMENT>
<EMBED type="application/x-websphere-client" CODE=MyAppletClass.class
ARCHIVE="Applet.jar, EJB.jar" WIDTH="600" HEIGHT="500"
scriptable="false">
<NOEMBED>
</COMMENT>
</NOEMBED>WebSphere Java Application/Applet Thin Client for
Windows is required.
</EMBED>
</OBJECT>
```

- The value of the `type` attribute on the `</EMBED>` tag can also be, for example:
`<EMBED type="application/x-websphere-client, version=4.0" ...`

Applet client code requirements

The code used by an applet to talk to an enterprise bean is the same as that used by a standalone Java program or a servlet, except for one additional property called `java.naming.applet`. This property informs the `InitialContext` and the `Object Request Broker (ORB)` that this client is an applet rather than a standalone Java application or servlet.

- When you initialize an instance of the `InitialContext` class, the first two lines in this code snippet illustrate what both a standalone Java program and a servlet issue to specify the computer name, domain, and port. In this example, `<yourserver.yourdomain.com>` is the computer name and domain where

WebSphere Application Server resides, and 900 is the configured port. After the bootstrap values (<yourserver.yourdomain.com>:900) are defined, the client to server communications occur within the underlying infrastructure. In addition to the first two lines, for applets, you must add the highlighted third line to your code. That line identifies this program as an applet, for example:

```
prop.put(Context.INITIAL_CONTEXT_FACTORY,
"com.ibm.websphere.naming.WsnInitialContextFactory");
prop.put(Context.PROVIDER_URL, "iiop://<yourserver.yourdomain.com>:900)
prop.put(Context.APPLET, this);
```

Developing J2EE application client code

A *J2EE application client* program operates similarly to a standard J2EE program in that it runs its own Java Virtual Machine code and is invoked at its main method. The Java Virtual Machine application client program differs from a standard Java program because it uses the Java Naming and Directory Interface (JNDI) name space to access resources. In a standard Java program, the resource information is coded in the program.

Steps for this task

1. Writing the client application program

Write the J2EE application client program on any development machine. At this stage, you do not require access to the WebSphere Application Server.

A *J2EE application client* program operates similarly to a standard J2EE program in that it runs its own Java Virtual Machine code and is invoked at its main method. The Java Virtual Machine application client program differs from a standard Java program because it uses the Java Naming and Directory Interface (JNDI) name space to access resources. In a standard Java program, the resource information is coded in the program.

Storing the resource information separately from the client application program makes the client application program portable and more flexible.

Using the `javax.naming.InitialContext` class, the client application program uses the lookup operation to access the Java Naming and Directory Interface (JNDI) name space. The `InitialContext` class provides the lookup method to locate resources.

The following example illustrates how a client application program uses the `InitialContext` class:

```
import javax.naming.*

public class myAppClient
{
    public static void main(String argv[])
    {
        InitialContext initCtx = new InitialContext();
        Object homeObject=initCtx.lookup("java:comp/env/ejb/BasicCalculator");
        BasicCalculatorHome bcHome = (BasicCalculatorHome)
        javax.rmi.PortableRemoteObject.narrow(homeObject, BasicCalculatorHome.class);
        BasicCalculatorHome bc = bcHome.create();
        ...
    }
}
```

In this example, the program looks up an enterprise bean called `BasicCalculator`. The `BasicCalculator` EJB reference is located in the client JNDI name space at `java:comp/env/ejb/BasicCalculator`. Since the actual enterprise bean runs on the server, the application client run time returns a reference to the `BasicCalculator` home interface.

If the client application program lookup was for a resource reference or an environment entry, then lookup returns an instance of the configured type as defined by the client application deployment descriptor. For example, if the program lookup was a JDBC datasource, the lookup would return an instance of `javax.sql.DataSource`.

2. Assemble the application client using the Application Assembly Tool

The JNDI name space knows what to return on a lookup because of the information assembled by the Application Assembly Tool (AAT).

Assemble the J2EE application client on any development machine with the AAT installed.

When you use the Application Assembly Tool to assemble your application client, you provide the *application client* run time with the required information to initialize the execution environment for your client application program. Refer to the Application Assembly Tool description for implementation details.

Following is a list of things to keep in mind when you configure resources used by your client application program:

- When configuring resource references, resource environment references, and EJB references in the Application Assembly Tool, the General tab contains a required Name field. This field specifies where the application client run time binds the reference to the real object in the `java:comp/env` portion of the JNDI name space. The application client run time always binds these references relative to `java:comp/env`. For the programming example above, specify `ejb/BasicCalculator` in the Name field on the General tab of the Application Assembly Tool, which requires the program to perform a lookup of `java:comp/env/ejb/BasicCalculator`. If the Name field is set to `myString`, the resulting lookup is `java:comp/env/myString`.
- When configuring Resource references in the Application Assembly Tool, the Name field on the General tab is used for:
 - Binding a reference of that object type into the JNDI name space.
 - Retrieving client specific resource configuration information that was configured using the Application Client Resource Configuration Tool.
- When configuring a resource reference in the Application Assembly Tool, the value in the Name field on the General tab must match the value in the JNDI Name field on the General tab for the resource in the Application Client Resource Configuration Tool.
- When configuring URL resources using the Client Resource Configuration Tool, the URL provider panel enables you to specify a protocol and a class to handle that protocol. If you want to use the default protocols, such as HTTP, you can leave those fields blank.
- When configuring resource references using the Application Assembly Tool, the General tab contains a field called Authorization. You can set this field to either Container or Application. If you set the field to Container, then the application client run time uses authorization information configured in the Application Client Resource Configuration tool for the resource. If the field is set to Application, then the application client run time expects the user application to provide authorization information for the resource. The application client run time ignores any authorization information configured with the Application Client Resource Configuration tool for that resource.
- When configuring resource environment references using the Application Assembly Tool, you must specify the location of the actual object in the server JNDI namespace using the Binding tab. A resource environment reference maps a logical name (the Name field on the general tab) used by the client application to the physical name of an object (the JNDI Name field on

the Bindings tab). Not all objects bound into the server JNDI namespace are intended for use by an application client. For example, the WebSphere Application Server client run-time does not support the use of Java 2 Connector (J2C) objects on the client. The object needs to be remotable, and the client-side implementations must be made available on the application client run-time classpath.

- Resource environment references are different than resource references. Resource environment references allow your application client to use a logical name to look-up a resource bound into the server JNDI namespace. A resource reference allows your application to use a logical name to look-up a local J2EE resource. The J2EE specification does not specify a particular implementation of a resource. The following is a table of the supported resource types and identifies the resources to which the WebSphere Application Server provides a client implementation.

Resource Type	Client Configuration Notes	Client implementation provided by WebSphere Application Server
javax.sql.DataSource	Supports specification of any Datasource implementation class	No
java.net.URL	Supports specification of custom protocol handlers	Provided by Java Runtime Environment files
javax.mail.Session	Supports custom protocol configuration	Yes - POP3, SMTP, IMAP
javax.jms.QueueConnectionFactory, javax.jms.TopicConnectionFactory, javax.jms.Queue, javax.jms.Topic	Supports configuration of WebSphere Embedded Messaging, IBM MQ Series and other JMS providers	Yes - WebSphere Embedded Messaging

3. Assembling the Enterprise Archive (EAR)

The application is contained in an enterprise archive or .ear file. The .ear file is composed of:

- Enterprise bean, application client, and user-defined modules or .jar files
- Web applications or .war files
- Metadata describing the applications or application .xml files

You must assemble the .ear file on the server machine.

4. Distribute the EAR file

The client machines configured to run this client must have access to the .ear file.

If all the machines in your environment share the same image and platform, run the Application Client Resource Configuration Tool (ACRCT) on one machine to configure the external resources, and then distribute the configured .ear file to the other machines.

If your environment is set up with a variety of client installations and platforms, run the ACRCT for each unique configuration.

You can either distribute the .ear files to the correct client machines, or make them available on a network drive.

Distributing the .ear files is the responsibility of the system and network administrator.

5. Deploy the application client.

6. Configure the application client resources

If the client application defines the local resources, run the ACRCT (clientConfig command) on the local machine to reconfigure the .ear file. Use the ACRCT to change the configuration. For example, the .ear file can contain a DB2 resource, configured as C:\DB2. If, however, you installed DB2 in the D:\Program Files\DB2 directory, use the ACRCT to create a local version of the .ear file.

What to do next

After developing the J2EE application client code, launch the application client.

J2EE application client class loading

When you run your J2EE application client using the WebSphere Application Server **launchClient** command, a hierarchy of class loaders is created to load classes used by your application.

The following list describes the hierarchy of class loaders:

- The topmost class loader, the bootstrap class loader, contains the JAR files that make up the Java Virtual Machine code, such as rt.jar, plus those JAR files defined by the -Xbootclasspath parameter on the Java command. The WebSphere Application client run time sets this value to the WAS_BOOTCLASSPATH environment variable.
- The *extensions class loader* class loader is a child to the bootstrap class loader. This class loader contains JAR files in the java/jre/lib/ext directory or those JAR files defined by the -Djava.ext.dirs parameter on the Java command. The WebSphere Application Client run time does not set -Djava.ext.dirs parameters, so it uses the JAR files in the java/jre/lib/ext directory.
- The *system class loader* class loader contains JAR files and classes that are defined by the -classpath parameter on the java command. The Application Client run time sets this parameter to the WAS_CLASSPATH environment variable.
- The *WebSphere class loader* class loader loads the WebSphere Application Client run time and any classes placed in the WebSphere Application Client user directories. The directories used by this class loader are defined by the WAS_EXT_DIRS environment variable. The WAS_BOOTCLASSPATH, WAS_CLASSPATH, and the WAS_EXT_DIRS environment variables are set in the installation_root/bin/setupCmdLine command shell for WebSphere Application Server server installations, or in the installation_root/bin/setupClient command shell for client installations.

As the J2EE application client run time initializes, additional class loaders are created as children of the WebSphere class loader. If your client application uses resources such as Java Database Connectivity (JDBC) API, Java Message Service (JMS) API, or Uniform Resource Locator (URL), a different class loader is created to load each of those resources. Finally, the application client run time sets the WebSphere class loader to load classes within the .ear file by processing the client JAR manifest repeatedly. The system classpath, defined by the CLASSPATH environment variable is never used and is not part of the hierarchy of class loaders.

To package your client application correctly, you must understand which class loader loads your classes. When Java loads a class, the class loader used to load that class is assigned to it. Any classes subsequently loaded by that class will use that class loader or any of its parents, but it will not use children class loaders.

In some cases the WebSphere Application Client run time can detect when your client application class is loaded by a different class loader from the one created for it by the WebSphere Application Client run time. When this detection occurs, you see the following message:

```
WSCL0205W: The incorrect class loader was used to load [0]
```

This message occurs when your client application class is loaded by one of the parent class loaders in the hierarchy. This situation is typically caused by having the same classes in the .ear file and on the hard drive. If one of the parent class loaders locates a class, that class loader loads it before the application client run time class loader. In some cases, your client application will still function correctly. In most cases, however, you receive "class not found" exceptions.

Configuring the classpath fields

When packaging your J2EE client application, you must configure various classpath fields. Ideally, you should package everything required by your application into your .ear file. This is the easiest way to distribute your J2EE client application to your clients. However, you should not package such resources as JDBC APIs, JMS APIs, or URLs. In the case of these resources, use classpath references to access those classes on the hard drive. You might also have other classes installed on your client machines that you do not need to redistribute. In this case, you also want to use classpath references to access the classes on the hard drive, as described below.

Referencing classes within the EAR file

WebSphere J2EE applications do not use the system class path. Use the MANIFEST Classpath entry to refer to other JARs within the .ear file. Configure these values using the module Classpath fields in the Application Assembly Tool. For example, if your client application needs to access the path of the enterprise bean JAR, add the deployed enterprise bean module name to your application client Classpath field in the Application Assembly Tool. The format of the Classpath field for each of the different modules (Application Client, enterprise bean, Web) is the same:

- The values must refer to .jar and .class files that are contained within the .ear file.
- The values must be relative to the root of the .ear file.
- The values cannot refer to absolute paths in the file systems.
- Multiple values must be separated by spaces, not colons or semi-colons.

Note: This is the Java method for allowing applications to function platform-independent.

Typically, you add modules (.jar files) to the root of the .ear file. In this case, you only need to specify the name of the module (.jar file) in the Classpath field. If you choose to add a module with a path, you need to specify the path relative to the root of the .ear file.

For referencing .class files, you must specify the directory relative to the root of the .ear file. With the Application Assembly Tool you can add individual class files to the .ear file. It is recommended that these additional class files are packaged in a .jar file. Add this .jar file to the module Classpath fields. If you add .class files to the root of the .ear file, add ./ to the module Classpath fields. Consider the following example directory structure in which the file myapp.ear contains an application client JAR file named client.jar and a mybeans.jar EJB

module. Additional classes reside in class1.jar and utility/class2.zip files. A class named xyz.class is not packaged in a JAR file but is in the root of the EAR file. Specify `./ mybeans.jar utility/class2.zip class1.jar` as the value of the Classpath property. The search order is: myapp.ear/client.jar myapp.ear/xyz.class myapp.ear/mybeans.jar myapp.ear/utility/class2.zip myapp.ear/class1.jar

Referencing classes that are not in the EAR file

Use the launchClient `-CCclasspath` parameter. This parameter is specified at run time and takes platform-specific classpath values, which means multiple values are separated by semi-colons or colons. There are many similarities between the client and the server in this respect.

Resource classpaths

When you configure resources used by your client application using the Application Client Resource Configuration Tool, you can specify classpaths that are required by the resource. For example, if your application is using a JDBC to a DB2 database, add db2java.zip to the classpath field of the database provider. These classpath values are platform-specific and require semi-colons or colons to separate multiple values.

Using the launchClient API

If you use the `launchClient` shell and bat command, the WebSphere class loader hierarchy is created for you. However, if you use the launchClient API, you must perform this setup yourself. You should mimic the launchClient shell command in defining the Java system properties.

Developing pluggable application client code

Steps for this task

1. Install the pluggable application client from the WebSphere Application Client CD by selecting option **Pluggable Application Client** from the **Custom client installation** panel.
2. Set the Java application pluggable client environment by using the `setupClient` shell, located in:
`install_root\AppClient\bin\setupClient.bat` (on Windows systems)
`install_root/AppClient/bin/setupClient.sh` (on UNIX platforms)
3. Add your specific Java client application JAR files to the CLASSPATH and start your Java client application from this environment, after setting the environment variables.

4. Run the following Java command to invoke your client application:

```
%JAVA_HOME%/bin/java -Xbootclasspath/p:%WAS_BOOTCLASSPATH% -classpath  
<list of your application jars and classes>  
-Djava.ext.dirs=%WAS_EXT_DIRS% -Djava.naming.provider.url=iiop://  
<your Websphere server machine name>  
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory  
%SERVER_ROOT% %CLIENTSAS% <fully qualified class name to run>  
  
$JAVA_HOME/bin/java -Xbootclasspath/p:$WAS_BOOTCLASSPATH -classpath  
<list of your application jars and classes>  
-Djava.ext.dirs=$WAS_EXT_DIRS -Djava.naming.provider.url=iiop://  
<your Websphere server machine name>  
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory  
$SERVER_ROOT $CLIENTSAS <fully qualified class name to run>
```


Usage scenario

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

Developing thin application client code

Steps for this task

1. Install the Java application thin client from the WebSphere Application Client CD by selecting option **J2EE/Thin application client for the WebSphere Application Client**.
2. Set the Java application thin client environment by using the **setupClient** shell, located in:

```
install_root\AppClient\bin\setupClient.bat (on Windows)
install_root/AppClient/bin/setupClient.sh (on UNIX platforms)
```

3. Add your specific Java client application JAR files to the CLASSPATH and start your Java client application from this environment, after setting the environment variables.
4. Run the following Java command to invoke your client application:

```
%JAVA_HOME%/bin/java -Xbootclasspath/p:%WAS_BOOTCLASSPATH% -classpath
<list of your application jars and classes>
-Djava.ext.dirs=%WAS_EXT_DIRS% -Djava.naming.provider.url=iiop://
<your Websphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
%SERVER_ROOT% %CLIENTSAS% <fully qualified class name to run>

$JAVA_HOME/bin/java -Xbootclasspath/p:$WAS_BOOTCLASSPATH -classpath
<list of your application jars and classes>
-Djava.ext.dirs=$WAS_EXT_DIRS -Djava.naming.provider.url=iiop://
<your Websphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
$SERVER_ROOT $CLIENTSAS <fully qualified class name to run>
```

Usage scenario

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

Assembling Application Client Modules

Before you begin

If you want to use existing J2EE 1.2 Web modules in your J2EE 1.3 application, migrate them to J2EE 1.3 first.

Note: This task only applies to J2EE application clients.

Assemble a client module to contain application client code. (Group enterprise beans, Web components, and resource adapter code in separate modules).

Steps for this task

1. ("Starting the Application Assembly Tool (AAT)").
2. Select **File>New>Application Client**.

The navigation tree now displays various sets of properties for configuring the new application client.

3. **(Optional)** Use the property dialog shown in the AAT workspace to change the default file name and location.
 - a. It is recommended that you change the display name so that it differs from the file name.
 - b. If you like, change the temporary location of the application client from the default location, `install_root/bin`.
4. Enter the main class filename and location.
 - a. Click **Browse** to locate the class file.
 - b. Select the archive containing the class files and click **Select**.
 - c. Select the files you need from the archive and click **OK**.
5. Define the assembly properties for the application client.
 - a. Right-click the assembly property in the navigation pane.
 - b. Select **New** in the right-click menu.
 - c. Enter values in the property dialog box.
 - d. Click **OK**
6. Add files for the application client.
 - a. Right-click **Files** in the navigation pane.
 - b. Select **Add Files** from the right-click menu.
 - c. Locate the directory where the files are located and click **Select**.
 - d. Select the files to add and click **Add**.
 - e. In the **Selected files** window, click **OK**.
7. ("Saving applications after assembly")

What to do next

Assemble zero or more other new modules of your choice:

- ("Assembling EJB modules")
- "Assembling Resource Adapter modules" (not in this document)

You can also (migrate existing modules).

Another option is to proceed directly to (assembling a new application module). While assembling an application module, you can create any new modules that you need.

Application client assembly settings

Use this page to specify assembly properties for J2EE application clients.

File name (Required, String)

Specifies the file name of the application client module, relative to the top level of the Enterprise Archive (EAR) file.

If this is a standalone module, the file name is the full path name of the archive.

Alternative DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module Java Archive (JAR) file.

This file is the postassembly version of the deployment descriptor file. (You can edit the original deployment descriptor file to resolve dependencies and security information. Directing the use of the alternative deployment descriptor allows you

to keep the original deployment descriptor file intact). The value of the Alternative DD property must be the full path name of the deployment descriptor file relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read directly from the module JAR file.

Classpath

Specifies the full classpath containing the dependent code that is not contained in the application client JAR file.

Specify the values relative to the root of the EAR file and separate the values with spaces. Absolute values that reference files or directories on the hard drive are ignored. To specify classes that are not in JAR files but are in the root of the EAR file, use a period and forward slash (.). Consider the following example directory structure in which the file myapp.ear contains an application client JAR file named client.jar. Additional classes reside in class1.jar and class2.zip files. A class named xyz.class is not packaged in a JAR file but is in the root of the EAR file myapp.ear/client.jar myapp.ear/class1.jar myapp.ear/class2.zip myapp.ear/xyz.class. Specify class1.jar class2.zip ./ as the value of the Classpath property. (Name only the directory for .class files.)

Display name (Required, String)

Specifies a short name that is intended for display by GUIs.

Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels).

The image is used as an icon to represent the application client in a GUI.

Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels).

The image is used as an icon to represent the application client in a GUI.

Description

Contains text describing the application client.

Main class (Required, String)

Specifies the full path name of the main class for this application client.

Deploying application clients

Before you begin

After developing an application client, deploy this application on client machines. *Deployment* consists of pulling together the various artifacts that the application client requires.

The *Application Client Resource Configuration Tool* (ACRCT) defines resources for the application client. These configurations are stored in the application client .ear file. The application client run time uses these configurations for resolving and creating an instance of the resources for the application client.

Note: This task only applies to J2EE application clients. Only perform this task if you configured your J2EE application client to use resource references.

Steps for this task

1. Start the ACRCT and open an EAR file.
2. Configure new data source providers.
3. Configure mail providers and sessions.
4. Configure URL providers and sessions.
5. Configure Java messaging client resources.
6. Configure new environment entries.
7. (Optional) Remove application client resources.
8. Save the EAR file.

Starting the Application Client Resource Configuration Tool and opening an EAR file

Before you begin

Note: This task only applies to J2EE application clients.

Steps for this task

1. Open a command prompt and change to the `install_root\bin` directory.
2. Run the `clientConfig.bat` file for a Windows system or the `clientConfig.sh` file for a UNIX system.
3. Open an EAR file within the Application Client Resource Configuration Tool (ACRCT):
 - Click **File > Open**.
 - Select the file and click **Open**.
4. Save your changes to the file and close the tool:
 - Click **File > Save**.
 - Click **File > Exit**.

JDBC providers for application clients

The J2EE application client does not support Java 2 Connection Factories. In addition, WebSphere Application Server and WebSphere Application Server clients do not provide client database drivers to be used directly from a J2EE application client. If your application client uses a database directly, you must provide the database drivers on the client machine. This can involve contacting your database vendor to acquire client database driver code and licenses. Instead of accessing the database directly, it is recommended that your client application use an enterprise bean. Accessing a database through an enterprise bean eliminates the need to have database drivers on the client machine, since the database access is handled by the enterprise bean running on the WebSphere Application Server. For a current list of providers that are supported on the WebSphere Application Server go the following site:

Supported hardware, software, and APIs

Data sources for application clients

A data source is used by an application to access the data from the database.

A data source is created under a JDBC provider which provides the specific JDBC driver implementation class. The data source represents the JCA connection factory for the Relational Resource Adapter.

Multiple data sources can be created under the same JDBC provider. Each JDBC provider supports the interfaces defined by Sun Microsystems listed below. These interfaces enable the application to run in a single-phase or two-phase transaction protocol.

- *ConnectionPoolDataSource* - this is a data source that enables applications to participate in all transactions, including a two-phase commit transaction. When this kind of data source is involved in a global transaction, there is no transaction recovery provided by the transaction manager. It is the application's responsibility to provide the backup recovery process if multiple resource managers are involved.
- *XADataSource* - this is a data source that allows applications to participate in a single-phase or a global (two-phase) transaction environment. When this data source is involved in a global transaction, transaction recovery will be provided by the transaction manager.

Before WebSphere Application Server Version 5.0, the function of data access was provided by a single connection manager (CM) architecture. In release 5.0, this connection manager architecture is still available to support WebSphere Application Server Version 4.0 applications, but a new connection manager architecture is provided based on the JCA architecture supporting the new J2EE 1.3 application style.

These two separate CM architectures are represented by two different types of data sources in the WebSphere Application Server. Administrators must understand the nature of their applications, EJB modules, and EJB beans in order to choose the right data source.

- Data source (Version 4.0) - this type of data source runs under the WebSphere connection manager architecture. Applications using this data source will maintain the same behavior as if they were running in WebSphere Application Server release 4.0.
- New data source - this type of data source utilizes the JCA standard architecture to provide J2EE 1.3 support. It runs under the JCA connection manager and the relational resource adapter. Applications using this type of data source may experience some behavioral differences because of the J2EE 1.3 architecture.

Choice of data source

- J2EE 1.2 application - all EJB beans, JDBC applications, or Servlets 2.2 must use the **4.0** data source.
- J2EE 1.3 application -
 - EJB 1.1 Module - all EJB 1.x beans must use the **4.0** data source.
 - EJB 2.0 Module - all EJB beans including CMP version 2.0 and 1.x must use the **new** data source.
 - JDBC applications and Servlet 2.3 - must use the **new** data source.

Configuring new data source providers (JDBC providers) for application clients

Before you begin

During this task, you create new data source providers, also known as JDBC providers, for your application client. In a separate administrative task, install the Java code for the required data source provider on the client machine on which the application client resides.

Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new data source provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new data source provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the **Data Source Providers** folder. Do one of the following:
 - Right-click the folder and click **New Provider**.
 - Click **Edit > New** on the menu bar.
5. Configure the data source provider properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

Configuring new data source providers

During this task, you will create new data source providers, also known as JDBC drivers, for your application client. In a separate administrative task, install the Java code for the required data source provider on the client machine where the application client resides.

Steps for this task

1. Start the ACRCT, click **File > Open**, and select the EAR file for which you want to configure the new data source provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new data source provider from the tree.
3. Expand the JAR file to view its contents.
4. Right click the **Data Source Providers** folder and select **New Provider**.
5. Configure the data source provider properties in the resulting property dialog.
6. Click **OK**.
7. Click **File > Save** to save your changes.

Example: Configuring data source provider and data source settings: The purpose of this article is to help you to configure data source provider and data source settings.

- Required fields:
 - Data Source Provider Properties page: name
 - Data Source Properties page: name, jndiName
- Special cases:
 - The user name and password fields have no equivalent xmi tags. You must specify these fields in the custom properties.
 - The password is encrypted when you use the Application Client Resource Configuration Tool (ACRCT). If you do not use the ACRCT, the field cannot be encrypted.
- Example:

```
<resources.jdbc:JDBCProvider xmi:id="JDBCProvider_1"
name="jdbcProvider:name" description="jdbcProvider:description"
implementationClassName="jdbcProvider:ImplementationClass">
<classpath>jdbcProvider:classpath</classpath>
<factories xmi:type="resources.jdbc:WAS40DataSource"
xmi:id="WAS40DataSource_1" name="jdbcFactory:name"
jndiName="jdbcFactory:jndiName" description="jdbcFactory:
description" databaseName="jdbcFactory:databasename">
```

```

<propertySet xmi:id="J2EEResourcePropertySet_13">
<resourceProperties xmi:id="J2EEResourceProperty_13"
name="jdbcFactory:customName" value="jdbcFactory:customValue"/>
<resourceProperties xmi:id="J2EEResourceProperty_14"
name="user" value="jdbcFactory:user"/>
<resourceProperties xmi:id="J2EEResourceProperty_15"
name="password" value="{xor}NTs9PBk+PCswLSZ1MT4y0g==" />
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_14">
<resourceProperties xmi:id="J2EEResourceProperty_16"
name="jdbcProvider:customName"
value="jdbcProvider:customeValue"/>
</propertySet>
</resources.jdbc:JDBCProvider>

```

Data source provider settings for application clients: Use this page to create a data source under a JDBC provider which provides the specific JDBC driver implementation class.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Data Source Providers >** and click **New**. The following fields appear on the **General** tab:

Name: Specifies the display name for the data source.

For example you can set this field to *Test Data Source*.

Data type String

Description: Specifies a text description for the resource.

Data type String

Class Path: A list of paths or jarfile names which together form the location for the resource provider classes.

Implementation class: Use this setting to perform database specific functions.

Data type String
 Default Dependent on JDBC driver implementation class

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Data source properties for application clients: Use this page to create or modify the V5.0 data sources.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected

JAR file > **Data Source Providers** > *Data source provider instance*. Right-click **Data Sources** and click **New**. The following fields appear on the **General** tab:

Name: Specifies the display name of this data source.

Data type String

Description: Specifies a text description of the data source.

Data type String

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Database Name: The name of the database to which you want to connect.

User: Use the user ID with the Password property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the User ID property, you must also specify a value for the Password property. The connection factory User ID and Password properties are used if the calling application does not provide a userid and password explicitly.

Password: Use the password with the User ID property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the User ID property, you must also specify a value for the Password property.

Re-Enter Password: Confirms the password.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Configuring new data sources for application clients

Before you begin

During this task, you create new data sources for your application client.

Steps for this task

1. Click the data source provider for which you want to create a data source in the tree. Do one of the following:
 - Configure a new data source provider.
 - Click an existing data source provider.
2. Expand the data source provider to view its **Data Sources** folder.
3. Click the folder. Do one of the following:

- Right-click the folder and click **New Factory**.
 - Click **Edit > New** on the menu bar.
4. Configure the data source properties in the resulting property dialog.
 5. Click **OK** when you finish.
 6. Click **File > Save** on the menu bar to save your changes.

Mail providers and mail sessions for the Application Client Assembly Tool

A JavaMail service provider is a driver that allows a JavaMail API to interact with mail servers running on a particular mail protocol. WebSphere Application Server includes service providers, also known as *protocol providers*, for mail protocols including Simple Mail Transfer Protocol (SMTP), Internet Message Access Protocol (IMAP), and Post Office Protocol 3 (POP3).

Mail provider encapsulates a collection of protocol providers. For example, WebSphere Application Server has a built-in mail provider that encompasses the three protocol providers: SMTP, IMAP and POP3. These protocol providers are installed as the default and are sufficient for most applications.

If you have a particular application that requires custom protocol providers, you must first follow the steps outlined in "JavaMail API Design Specification, V1.2, Chapter 5 - The Mail Session" to install your own protocol providers. See JavaMail: Resources for learning, for a link to this documentation.

Ensure every mail session is defined under a parent mail provider. Select a mail provider first and then create your new mail session.

Configuring mail providers and sessions for application clients

Before you begin

Use the Application Client Resource Configuration Tool (ACRCT) to edit the configurations of JavaMail sessions and providers for your application clients to use.

Steps for this task

1. Open the ACRCT.
2. Open an EAR file.
3. Locate the JavaMail objects in the tree that displays.
For example, if your file contains JavaMail sessions, expand **Resources > application.jar > JavaMail Providers > java_mail_provider_instance > JavaMail Sessions**.

In this example, *java_mail_provider_instance* is a particular JavaMail provider.

Results

The JavaMail session instances are located in the **JavaMail Sessions** folder.

Mail provider settings for application clients

Use this page to implement the JavaMail API and create mail sessions.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Mail Providers >** and click **New**. The following fields appear on the **General** tab:

Name: The name of the JavaMail resource provider.

Description: An optional description for the resource provider.

Class Path: Specifies a list of paths or JAR file names which together form the location for the resource provider classes.

Protocol: Specifies the name of the protocol.

Classname: Specifies the name of the class implementing the protocol. Leave this field blank if you want to use the default implementation.

Type: This menu contains the following two values: TRANSPORT or STORE.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Mail session settings for application clients

Use this page to configure mail session properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Mail Providers > mail provider instance**. Right-click **Mail Sessions** and click **New**. The following fields appear on the **General** tab:

Name: Represents the administrative name of the JavaMail session object.

Description: Provides an optional description for your administrative records.

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Mail Transport Host: Specifies the server to connect to when sending mail.

Mail Transport Protocol: Specifies the transport protocol to use when sending mail.

Mail Transport User: Specifies the user ID to use when the mail transport host requires authentication.

Mail Transport Password: Specifies the password to use when the mail transport host requires authentication.

Re-Enter Password: Confirms the password.

Mail From: Specifies the mail originator.

Mail Store Host: Mail account host (or "domain") name.

Mail Store User: The user ID of the mail account.

Mail Store Password: The password of the mail account.

Re-Enter Password: Confirms the password.

Mail Store Protocol: Specifies the protocol to be used when receiving mail.

Mail Debug: When true, JavaMail interaction with mail servers, along with these mail session properties will be printed to stdout.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Example: Configuring JavaMail provider and JavaMail session settings for application clients

The purpose of this article is to help you configure JavaMail provider and JavaMail session settings.

- Required fields:
 - JavaMail Provider Properties page: name, and at least one protocol provider
 - JavaMail Session Properties page: name, jndiName, mail transport protocol, mail store protocol
- Special cases:
 - The password is encrypted when using the ACRCT tool. Without the tool, you cannot encrypt this field.
- Example:

```
<resources.mail:MailProvider xmi:id="MailProvider_1"
name="Default Mail Provider"
description="IBM JavaMail Implementation">
<classpath>mailProvider:classpath</classpath>
<factories xmi:type="resources.mail:MailSession"
xmi:id="MailSession_1"
name="mailSession:name" jndiName="mailSession:jndiName"
description="mailSession:description"
mailTransportHost="mailSession:mailTransportHost"
mailTransportUser="mailSession:mailTransportUser"
mailTransportPassword="{xor}Mj42Mww6LCw2MDF1MT4yOg=="
mailFrom="mailSession:
mailFrom" mailStoreHost="mailSession:mailStoreHost"
mailStoreUser="mailSession:mailStoreUser"
mailStorePassword="{xor}Mj42Mww6LCw2MDF1MT4yOg=="
debug="true"
mailTransportProtocol="ProtocolProvider_1"
mailStoreProvider="ProtocolProvider_1">
<propertySet xmi:id="J2EEResourcePropertySet_1">
<resourceProperties xmi:id="J2EEResourceProperty_1"
name="mailSession:customName"
value="mailSession:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_2">
<resourceProperties xmi:id="J2EEResourceProperty_2"
```

```

name="mailProvider:customName"
value="mailProvider:customValue"/>
</propertySet>
<protocolProviders xmi:id="ProtocolProvider_1"
protocol="smtp" classname="smtp:className"/>
<protocolProviders xmi:id="ProtocolProvider_2"
protocol="pop3" classname="pop3:className"/>
<protocolProviders xmi:id="ProtocolProvider_3"
protocol="imap" classname="imap:className"/>
</resources.mail:MailProvider>

```

Configuring new mail sessions for application clients

Before you begin

During this task, you configure new mail sessions for your application client. The mail sessions are associated with the preconfigured default mail provider supplied by the product.

Steps for this task

1. Start the tool and open the EAR file for the new JavaMail session.
The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new JavaMail session.
3. Expand the JAR file to view its contents.
4. Click **JavaMail Providers > MailProvider > JavaMail Sessions**. Do one of the following:
 - Right-click the **JavaMail Sessions** folder and select **New Factory**.
 - Click **Edit > New** on the menu bar.
5. Configure the JavaMail session properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

URLs for application clients

A *Uniform Resource Locator* (URL) is an identifier that points to an electronically accessible resource, such as a directory file on a machine in a network, or a document stored in a database.

URLs appear in the format *scheme:scheme_information*.

You can represent a *scheme* as *http*, *ftp*, *file*, or another term that identifies the type of resource and the mechanism by which you can access the resource.

In a World Wide Web browser location or address box, a URL for a file available using HyperText Transfer Protocol (HTTP) starts with *http:*. An example is *http://www.ibm.com*. Files available using File Transfer Protocol (FTP) start with *ftp:*. Files available locally start with *file:*.

The *scheme_information* commonly identifies the Internet machine making a resource available, the path to that resource, and the resource name. The *scheme_information* for HTTP, FTP and File generally starts with two slashes (*//*), then provides the Internet address separated from the resource path name with one slash (*/*). For example,

```
http://www-4.ibm.com/software/webservers/appserv/library.html.
```

For HTTP and FTP, the path name ends in a slash when the URL points to a directory. In such cases, the server generally returns the default index for the directory.

URL providers for the Application Client Resource Configuration Tool

A URL provider implements the functionality for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP). This provider, comprised of a pair of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

Configuring new URL providers for application clients

Before you begin

During this task, you create URL providers and URLs for your client application. In a separate administrative task, you must install the Java code for the required URL provider on the client machine on which the client application resides.

Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new URL provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new URL provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the folder called **URL Providers**. Do one of the following:
 - Right-click the folder and click **New Provider**.
 - Click **Edit -> New** on the menu bar.
5. Configure the URL provider properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File -> Save** on the menu bar to save your changes.

Configuring URL providers and sessions using the Application Client Resource Configuration Tool

Before you begin

Use the Application Client Resource Configuration Tool (ACRCT) to edit the configurations of URL providers and URLs to be used by your application clients.

Steps for this task

1. Open the ACRCT.
2. Open an EAR file.
3. Locate the URL objects in the tree that displays.

For example, if your file contains URL providers and URLs, expand **Resources -> application.jar -> URL Providers -> url_provider_instance** where *url_provider_instance* is a particular URL provider.
4. If you expand the tree further, you will also see the **URLs** folders containing the URL instances for each URL provider instance.

URL settings for application clients: Use this page to implement the functionality for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP).

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **URL Providers** > *URL provider instance*. Right-click **URLs** and click **New**. The following fields appear on the **General** tab.

This provider, comprised of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

Name: Administrative name for the URL

Description: Optional description of the URL, for your administrative records

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

URL: A Uniform Resource Locator (URL) name that points to an internet or intranet resource. For example: `http://www.ibm.com`

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

URL provider settings for application clients: Use this page create new URLs..

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **URL Providers** > and click **New**. The following fields appear on the **General** tab.

A URL provider implements the functionality for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP). This provider, comprised of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

Name: Administrative name for the URL

Description: Optional description of the URL, for your administrative records

Class Path: A list of paths or jarfile names which together form the location for the resource provider classes.

Protocol: Protocol supported by this stream handler. For example, "nntp", "smtp", "ftp", etc.

To use the default protocol, leave this field blank.

Stream handler class: Fully qualified name of a User-defined Java class that extends `java.net.URLStreamHandler` for a particular URL protocol, such as ftp.

To use the default stream handler, leave this field blank.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Example: Configuring URL and URL provider settings for application clients

The purpose of this article is to help you to configure URL and URL provider settings.

- Required fields:
 - URL Properties page: name, jndiName, url
 - URL Provider Properties page: name

- Example:

```
<resources.url:URLProvider xmi:id="URLProvider_1" name="urlProvider:name"
description="urlProvider:description" streamHandlerClassName="urlProvider:
streamHandlerClass" protocol="urlProvider:protocol">
<classpath>urlProvider:classpath</classpath>
<factories xmi:type="resources.url:URL" xmi:id="URL_1" name="urlFactory:name"
jndiName="urlFactory:jndiName" description="urlFactory:description"
spec="urlFactory:url">
<propertySet xmi:id="J2EEResourcePropertySet_18">
<resourceProperties xmi:id="J2EEResourceProperty_20"
name="urlFactory:customName" value="urlFactory:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_19">
<resourceProperties xmi:id="J2EEResourceProperty_21"
name="urlProvider:customName" value="urlProvider:customValue"/>
</propertySet>
</resources.url:URLProvider>
```

Configuring new URLs with the Application Client Resource Configuration Tool

Before you begin

During this task, you create URLs for your client application.

Steps for this task

1. Click the URL provider for which you want to create a URL in the tree. Do one of the following:
 - Configure a new URL provider.
 - Click an existing URL provider.
2. Expand the URL provider to view the **URLs** folder.
3. Click the folder. Do one of the following:
 - Right-click the folder and click **New Factory**.
 - Click **Edit -> New** on the menu bar.
4. Configure the URL properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File -> Save** in the menu bar to save your changes.

WebSphere asynchronous messaging using the Java Message Service API for the Application Client Resource Configuration Tool

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. The JMS interface provides a common way for Java programs (clients and J2EE applications) to create, send, receive, and read asynchronous requests, as JMS messages.

This topic provides an overview of asynchronous messaging using JMS support provided by the WebSphere Application Server.

The base support for asynchronous messaging using the JMS API provides the common set of JMS interfaces and associated semantics that define how a JMS client can access the facilities of a JMS provider. This support enables WebSphere J2EE applications, as JMS clients, to exchange messages asynchronously with other JMS clients, by using JMS destinations (queues or topics). An J2EE application can use JMS queue destinations for point-to-point messaging and JMS topic destinations for Pub and Sub messaging. A J2EE application can explicitly poll for messages on a destination then retrieve messages for processing by business logic beans (enterprise beans).

With the base JMS/XA support, the J2EE application uses standard JMS calls to process messages, including any responses or outbound messaging. An enterprise bean can handle responses acting as a sender bean, or within the enterprise bean that receives the incoming messages. Optionally, this process can use two-phase commit within the scope of a transaction. This level of functionality for asynchronous messaging is called *bean-managed messaging*, and gives an enterprise bean complete control over the messaging infrastructure; for example, connection and session pool management. The common container has no role in *bean-managed messaging*.

WebSphere Application Server also supports automatic asynchronous messaging using message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) and JMS listeners (part of the JMS application server facilities). Messages are automatically retrieved from JMS destinations, optionally within a transaction, then sent to the message-driven bean in a J2EE application, without the application having to explicitly poll JMS destinations.

Configuring Java messaging client resources

Before you begin

In a separate administrative task, install the Java Message Service (JMS) client on the client machine where the application client resides. The messaging product vendor must provide an implementation of the JMS client. For more information, see your messaging product documentation.

During this task, you create new JMS provider configurations for your application client. The application client can use a messaging service through the Java Message Service APIs. A JMS provider provides two kinds of J2EE factories. One is a *JMS connection factory*, and the other is a *JMS destination factory*.

Steps for this task

1. Start the ACRCT, click **File > Open**, and select the EAR file for which you want to configure the new JMS provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new JMS provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the **JMS Providers** folder and click **New Provider**.
5. Configure the JMS provider properties in the resulting property dialog.
6. Click **OK**.
7. Click **File > Save**.

Configuring new JMS providers with the Application Client Resource Configuration Tool

Before you begin

During this task, you will create new JMS provider configurations for your application client. The application client can make use of a messaging service through the Java Message Service APIs. A JMS provider provides two kinds of J2EE factories. One is a JMS Connection factory, and the other is a JMS destination factory.

In a separate administrative task, you must install the JMS client on the client machine where the application client resides. The messaging product vendor must provide an implementation of the JMS client. For more information, see your messaging product documentation.

Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new JMS provider. The EAR file contents will be displayed in a tree view.
2. From the tree, select the JAR file in which you want to configure the new JMS provider.
3. Expand the JAR file to view its contents.
4. Click the folder called **JMS Providers**. Do one of the following:
 - Right-click the folder and select **New Provider**.
 - On the menu bar, click **Edit -> New**.
5. In the resulting property dialog, configure the JMS provider properties.
6. When finished, click **OK**.
7. On the menu bar, click **File -> Save** to save your changes.

JMS provider settings for application clients

Use this page to configure properties of the JMS provider, if you want to use a JMS provider other than the internal WebSphere JMS provider or the MQSeries JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **JMS Providers > click New**. The following fields appear on the **General** tab.

Name: The name by which the JMS provider is known for administrative purposes.

Data type String

Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Description: A description of the JMS provider, for administrative purposes

Data type	String
Units	En_US ASCII characters
Default	JMS provider
Range	1 through 30 ASCII characters

Class Path: A list of paths or jarfile names which together form the location for the resource provider classes.

Context factory class: The Java classname of the initial context factory for the JMS provider.

For example, for an LDAP service provider the value has the form:
com.sun.jndi.ldap.LdapCtxFactory.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Provider URL: The JMS provider URL for external JNDI lookups.

For example, an LDAP URL for a JMS provider has the form:
ldap://hostname.company.com/contextName.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere queue connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the internal WebSphere JMS provider that is installed with WebSphere Application Server. These configuration properties control how connections are created to the associated JMS queue destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected

JAR file > **JMS Providers** > *JMS provider instance*. Right-click **WAS Queue Connection Factories** and click **New**. The following fields appear on the **General** tab.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the internal WebSphere JMS provider. A queue connection factory for the internal WebSphere JMS provider has the following properties:

Name: The name by which this queue connection factory is known for administrative purposes. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Description: A description of this connection factory for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

User: The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Password: The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Units	En_US ASCII characters
Default	Null

Range 1 through 30 ASCII characters

Re-Enter Password: Confirms the password.

Node: The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

Data type String

Application Server: Enter the name of the application server. This name is not the host name of the machine, but the name of the configured application server.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere topic connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the internal WebSphere JMS provider. These configuration properties control how connections are created to the associated JMS topic destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **WAS Topic Connection Factories** and click **New**. The following fields appear on the **General** tab.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the associated JMS provider. A topic connection factory for the internal WebSphere JMS provider has the following properties.

Name: The name by which this queue connection factory is known for administrative purposes. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String
Units En_US ASCII characters.
Default Null
Range 1 through 30 ASCII characters

Description: A description of this topic connection factory for administrative purposes.

Data type String
Units En_US ASCII characters
Default Null

Range 1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

User: The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createTopicConnection()`. The JMS client flows the userid and password to the JMS server.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Password: The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Re-Enter Password: Confirms the password.

Node: The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Application Server: Enter the name of the application server. This name is not the host name of the machine, but the name of the configured application server.

Port: Which of the two ports that connections use to connect to the JMS Server. The QUEUED port is for full-function JMS publish/subscribe support, the DIRECT port is for non-persistent, non-transactional, non-durable subscriptions only.

Note: Message-driven beans cannot use the direct listener port for publish/subscribe support. Therefore, any topic connection factory configured with **Port** set to `Direct` cannot be used with message-driven beans.

Data type	Enum
Units	Not applicable
Default	QUEUED
Range	QUEUED The listener port used for full-function JMS-compliant, publish/subscribe support.
	DIRECT The listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for publish/subscribe support.
	The TCP/IP port numbers for these ports are defined on the WebSphere Internal JMS Server.

Client Id: The JMS client identifier used for connections to the MQSeries queue manager.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere queue destination settings for application clients

Use this panel to view or change the configuration properties of the selected queue destination for use with the WebSphere JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **WAS Queue Destinations** and click **New**. The following fields appear on the **General** tab.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the internal WebSphere JMS provider. A queue for use with the internal WebSphere JMS provider has the following properties.

Name: The name by which the queue is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

Description: A description of the queue, for administrative purposes

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Persistence: Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	

Application defined
Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined
[WebSphere MQ JMS queue destination only]
Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent
Messages on the destination are persistent.

Non persistent
Messages on the destination are not persistent.

Priority: Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	

Application defined
The priority of messages on this destination is defined by the application that put them onto the destination.

Specified
The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified Priority: If the **Priority** property is set to **Specified**, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to **Specified**, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry: Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	

Application defined

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

Specified

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

Unlimited

Messages on this queue have no expiry timeout, so those messages never expire.

Specified Expiry: If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	Null
Range	Greater than or equal to 0 <ul style="list-style-type: none"> • 0 indicates that messages never timeout • Other values are an integer number of milliseconds

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere topic destination settings for application clients

Use this panel to view or change the configuration properties of the selected topic destination for use with the internal WebSphere JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **WAS Topic Destinations** and click **New**. The following fields appear on the **General** tab.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the internal WebSphere JMS provider has the following properties.

Name: The name by which the topic is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Description: A description of the topic, for administrative purposes

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Topic Name: The name of the topic as defined to the JMS provider.

Data type	String
Units	ASCII characters
Default	Null
Range	1 to 30 ASCII characters

Persistence: Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	

Application defined

Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined

[WebSphere MQ JMS queue destination only]
Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent

Messages on the destination are persistent.

Non persistent

Messages on the destination are not persistent.

Priority: Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	<p>Application defined</p> <p>The priority of messages on this destination is defined by the application that put them onto the destination.</p> <p>Specified</p> <p>The priority of messages on this destination is defined by the Specified priority property. <i>If you select this option, you must define a priority on the Specified priority property.</i></p>

Specified Priority: If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry: Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	<p>Application defined</p> <p>The expiry timeout for messages on this queue is defined by the application that put them onto the queue.</p> <p>Specified</p> <p>The expiry timeout for messages on this queue is defined by the Specified expiry property. <i>If you select this option, you must define a timeout on the Specified expiry property.</i></p> <p>Unlimited</p> <p>Messages on this queue have no expiry timeout, so those messages never expire.</p>

Specified Expiry: If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	Null
Range	<p>Greater than or equal to 0</p> <ul style="list-style-type: none"> • 0 indicates that messages never timeout • Other values are an integer number of milliseconds

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

MQSeries queue connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the MQSeries JMS provider. These configuration properties control how connections are created to the associated JMS queue destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Queue Connection Factories** and click **New**. The following fields appear on the **General** tab.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the MQSeries JMS provider. A queue connection factory for the MQSeries JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries for JMS resources. For more information about configuring MQSeries JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

Name: The name by which this queue connection factory is known for administrative purposes. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Description: A description of this connection factory for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

User: The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Password: The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Re-Enter Password: Confirms the password.

Queue Manager: The name of the MQSeries queue manager for this connection factory.

Connections created by this factory connect to that queue manager.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Host: The name of the host on which the WebSphere MQ queue manager runs, for client connection only.

Data type	String
Units	A valid TCP/IP hostname
Default	Null
Range	1 through 30 ASCII characters

Port: The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

Data type	String
Units	En_US ASCII characters

Default	Null
Range	1 through 30 ASCII characters

Channel: The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Transport type: Whether WebSphere MQ client connection or JNDI bindings is used for connection to the WebSphere MQ queue manager.

Data type	Enum
Units	Not applicable
Default	BINDINGS
Range	<p>CLIENT WebSphere MQ client connection is used to connect to the queue manager.</p> <p>BINDINGS JNDI bindings are used to connect to the queue manager.</p> <p>DIRECT For WebSphere MQ Event Broker using DIRECT mode..</p>

Client ID: The JMS client identifier used for connections to the MQSeries queue manager.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

CCSID: The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* book, SC33-1873, which is available from the following WebSphere MQ messaging platform-specific books Web pages:

- <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>

- <http://www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

or from the WebSphere MQ collection kit, SK2T-0730.

Message Retention: Select this tick box to specify that unwanted messages are to be left on the queue. Otherwise, unwanted messages are dealt with according to their disposition options.

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Selected
	Unwanted messages are left on the queue.
	Cleared
	Unwanted messages are dealt with according to their disposition options.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

MQSeries topic connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the MQSeries JMS provider. These configuration properties control how connections are created to the associated JMS topic destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Topic Connection Factories** and click **New**. The following fields appear on the **General** tab.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the MQSeries JMS provider. A topic connection factory for the MQSeries JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries JMS resources. For more information about configuring MQSeries JMS resources, see the *MQSeries Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

Name: The name by which this topic connection factory is known for administrative purposes. The name must be unique within the JMS provider.

Data type	String
Units	En_US ASCII characters

Default	Null
Range	1 through 30 ASCII characters

Description: A description of this topic connection factory for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The JNDI name that is used to bind the topic connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

User: The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User** property, you must also specify a value for the **Password** property.

The connection factory **User** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createTopicConnection()`. The JMS client flows the userid and password to the JMS server.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Password: The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Re-Enter Password: Confirms the password.

Queue Manager: The name of the MQSeries queue manager for this connection factory. Connections created by this factory connect to that queue manager.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Host: The name of the host on which the WebSphere MQ queue manager runs, for client connection only.

Data type	String
Units	A valid TCP/IP hostname
Default	Null
Range	1 through 30 ASCII characters

Port: The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Channel: The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Transport Type: Whether MQSeries client connection or JNDI bindings is used for connection to the MQSeries queue manager.

Data type	Enum
Units	Not applicable
Default	BINDINGS
Range	CLIENT MQSeries client connection is used to connect to the MQSeries queue manager. BINDINGS JNDI bindings are used to connect to the MQSeries queue manager.

Client Id: The JMS client identifier used for connections to the MQSeries queue manager.

Data type	String
-----------	--------

Units	A valid JMS client ID, as En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

CCSID: The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* book, SC33-1873, which is available from the following WebSphere MQ messaging platform-specific books Web pages:

- <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>
- <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>

or from the WebSphere MQ collection kit, SK2T-0730.

Broker Control Queue: The name of the broker control queue, to which all command messages (except publications and requests to delete publications) are sent

The name of the broker control queue. Publisher and subscriber applications, and other brokers, send all command messages (except publications and requests to delete publications) to this queue.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Queue Manager: The name of the MQSeries queue manager that provides the Pub/Sub message broker.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Pub Queue: The name of the broker's input queue that receives all publication messages for the default stream

The name of the broker's input queue (stream queue) that receives all publication messages for the default stream. Applications can also send requests to delete publications on the default stream to this queue.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Sub Queue: The name of the broker queue from which non-durable subscription messages are retrieved

The name of the broker's queue from which non-durable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker CCSubQ: The name of the broker's queue from which non-durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

The name of the broker queue from which non-durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Version: Whether the message broker is provided by the MQSeries MA0C Supportpac or newer versions of WebSphere message broker products

Data type	Enum
Units	Not applicable
Default	Advanced
Range	<p>Advanced The message broker is provided by newer versions of WebSphere message broker products (MQ Integrator and MQ Publish and Subscribe)</p> <p>Basic The message broker is provided by the MQSeries MA0C SupportPac (MQSeries - Publish/Subscribe)</p>

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

MQSeries queue destination settings for application clients

Use this panel to view or change the configuration properties of the selected queue destination for use with the MQSeries JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers** > *JMS provider instance*. Right-click **MQ Queue Destinations** and click **New**. The following fields appear on the **General** tab.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the MQSeries JMS provider. A queue for use with the MQSeries JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries JMS resources. For more information about configuring MQSeries JMS resources, see the *MQSeries Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

Name: The name by which the queue is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

Description: A description of the queue, for administrative purposes

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Persistence: Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED

Range

Application defined

Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined

[WebSphere MQ JMS queue destination only]
Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent

Messages on the destination are persistent.

Non persistent

Messages on the destination are not persistent.

Priority: Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type

Enum

Units

Not applicable

Default

APPLICATION_DEFINED

Range

Application defined

The priority of messages on this destination is defined by the application that put them onto the destination.

Specified

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified Priority: If the **Priority** property is set to **Specified**, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to **Specified**, messages sent to this queue have the priority value specified by this property.

Data type

Integer

Units

Message priority level

Default

Null

Range

0 (lowest priority) through 9 (highest priority)

Expiry: Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type

Enum

Units

Not applicable

Default

APPLICATION_DEFINED

Range

Application defined

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

Specified

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

Unlimited

Messages on this queue have no expiry timeout, so those messages never expire.

Specified Expiry: If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	Null
Range	Greater than or equal to 0
	<ul style="list-style-type: none">• 0 indicates that messages never timeout• Other values are an integer number of milliseconds

Base Queue Name: The name of the queue to which messages are sent, on the queue manager specified by the **Base queue manager name** property

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Base Queue Manager Name: The name of the MQSeries queue manager to which messages are sent

This queue manager provides the queue specified by the **Base queue name** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	A valid MQSeries Queue Manager name, as 1 through 48 ASCII characters

CCSID: The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* book, SC33-1873, which is available from the following WebSphere MQ messaging platform-specific books Web pages:

- <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>
- <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>

or from the WebSphere MQ collection kit, SK2T-0730.

Integer encoding: If native encoding is not enabled, select whether integer encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL Normal integer encoding is used.

REVERSED
Reversed integer encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Decimal encoding: If native encoding is not enabled, select whether decimal encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL Normal decimal encoding is used.

REVERSED
Reversed decimal encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Floating point encoding: If native encoding is not enabled, select the type of floating point encoding.

Data type	Enum
Units	Not applicable
Default	IEEEENORMAL
Range	IEEEENORMAL IEEE normal floating point encoding is used.

IEEEEREVERSED
IEEE reversed floating point encoding is used.

S390 S390 floating point encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Native encoding: Select this checkbox to indicate that the queue destination should use native encoding (appropriate encoding values for the Java platform)..

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Cleared Native encoding is not used, so specify the properties below for integer, decimal, and floating point encoding.
	Selected Native encoding is used (to provide appropriate encoding values for the Java platform).

For more information about encoding properties, see the MQSeries *Using Java* document.

Target client: Whether the receiving application is JMS-compliant or is a traditional WebSphere MQ application

Data type	Enum
Units	Not applicable
Default	MQSeries
Range	MQSeries The target is a non-JMS, traditional WebSphere MQ application.
	JMS The target is a JMS-compliant application.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

MQSeries topic destination settings for application clients

Use this panel to view or change the configuration properties of the selected topic destination for use with the MQSeries JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Topic Destinations** and click **New**. The following fields appear on the **General** tab.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the MQSeries JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries JMS resources. For more information about configuring MQSeries JMS resources, see the *MQSeries Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

Name: The name by which the topic is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Description: A description of the topic, for administrative purposes

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Persistence: Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	

Application defined

Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined

[WebSphere MQ JMS queue destination only]
Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent

Messages on the destination are persistent.

Non persistent

Messages on the destination are not persistent.

Priority: Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED

Range

Application defined

The priority of messages on this destination is defined by the application that put them onto the destination.

Specified

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified Priority: If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry: Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	

Application defined

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

Specified

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

Unlimited

Messages on this queue have no expiry timeout, so those messages never expire.

Specified Expiry: If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	Null
Range	Greater than or equal to 0

- 0 indicates that messages never timeout
- Other values are an integer number of milliseconds

Base Topic Name: The name of the topic to which messages are sent

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

CCSID: The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* book, SC33-1873, which is available from the following WebSphere MQ messaging platform-specific books Web pages:

- <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>
- <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>

or from the WebSphere MQ collection kit, SK2T-0730.

Integer encoding: If native encoding is not enabled, select whether integer encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL

Normal integer encoding is used.

REVERSED

Reversed integer encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Decimal encoding: If native encoding is not enabled, select whether decimal encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL

Range

NORMAL
Normal decimal encoding is used.

REVERSED
Reversed decimal encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Floating point encoding: If native encoding is not enabled, select the type of floating point encoding.

Data type Enum
Units Not applicable
Default IEEEENORMAL
Range

IEEEENORMAL
IEEE normal floating point encoding is used.

IEEEEVERSED
IEEE reversed floating point encoding is used.

S390 S390 floating point encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Native encoding: Select this checkbox to indicate that the queue destination should use native encoding (appropriate encoding values for the Java platform)..

Data type Enum
Units Not applicable
Default Cleared
Range

Cleared
Native encoding is not used, so specify the properties above for integer, decimal, and floating point encoding.

Selected
Native encoding is used (to provide appropriate encoding values for the Java platform).

For more information about encoding properties, see the MQSeries *Using Java* document.

BrokerDurSubQueue: The name of the broker queue from which durable subscription messages are retrieved

The name of the broker queue from which durable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

Data type String
Units En_US ASCII characters
Default Null
Range 1 through 48 ASCII characters

BrokerCCDurSubQueue: The name of the broker queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

The name of the broker queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Target Client: Whether the receiving application is JMS-compliant or is a traditional MQSeries application

Data type	Enum
Units	Not applicable
Default	MQSeries
Range	MQSeries The target is a non-JMS, traditional MQSeries application. JMS The target is a JMS-compliant application.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Generic JMS connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected JMS connection factory for use with the associated JMS provider. These configuration properties control how connections are created to the associated JMS destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > new JMS Provider instance**. Right click **JMS Connection Factories > click New**. The following fields appear on the **General** tab.

A JMS connection factory is used to create connections to JMS destinations. The JMS connection factory is created by the associated JMS provider. A JMS connection factory for a generic JMS provider (other than the internal WebSphere JMS provider or the MQSeries JMS provider) has the following properties:

Name: The name by which this JMS connection factory is known for administrative purposes. The name must be unique within the associated JMS provider.

Data type	String
Units	En_US ASCII characters
Default	Null

Range 1 through 30 ASCII characters

Description: A description of this connection factory for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

JNDI Name: The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

User: The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Password: The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 30 ASCII characters

Re-Enter Password: Confirms the password entered in the Password field.

External JNDI Name: The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

Connection Type: Whether this JMS destination is a queue (for point-to-point) or topic (for pub/sub).

Select one of the following options:

Queue

A JMS queue destination for point-to-point messaging.

Topic A JMS topic destination for pub/sub messaging.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Generic JMS destination settings for application clients

Use this panel to view or change the configuration properties of the selected JMS destination for use with the associated JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > new JMS Provider instance**. Right click **JMS Destinations > click New**. The following fields appear on the **General** tab.

A JMS destination is used to configure the properties of a JMS destination for the associated generic JMS provider. Connections to the JMS destination are created by the associated JMS connection factory. A JMS destination for use with a generic JMS provider (not the internal WebSphere JMS provider or MQSeries JMS provider) has the following properties.

Name: The name by which the queue is known for administrative purposes.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

Description: A description of the queue, for administrative purposes

JNDI Name: The JNDI name of the actual (physical) name of the JMS destination bound into JNDI.

External JNDI Name: The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

Destination Type: Whether this JMS destination is a queue (for point-to-point) or topic (for pub/sub).

Select one of the following options:

Queue

A JMS queue destination for point-to-point messaging.

Topic A JMS topic destination for pub/sub messaging.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Example: Configuring JMS Provider, JMS Connection Factory and JMS Destination settings for application clients

The purpose of this article is to help you to configure JMS Provider, JMS Connection Factory and JMS Destination settings.

- Required fields:
 - JMS Provider Properties page: name, and at least one protocol provider
 - JMS Connection Factory Properties page: name, jndiName, destination type
 - JMS Destination Properties page: name, jndiName, destination type
- Special cases:
 - The destination type must be QUEUE, or TOPIC.
- Example:

```
<resources.jms:JMSProvider xmi:id="JMSProvider_3" name="genericJMSProvider:name"
description="genericJMSProvider:description"
externalInitialContextFactory="genericJMSProvider:contextFactoryClass"
externalProviderURL="genericJMSProvider:providerUrl">
<classpath>genericJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms:GenericJMSDestination"
xmi:id="GenericJMSDestination_1" name="jmsDestination:name"
jndiName="jmsDestination:jndiName" description="jmsDestination:description"
externalJNDIName="jmsDestination:externalJndiName" type="QUEUE">
<propertySet xmi:id="J2EEResourcePropertySet_15">
<resourceProperties xmi:id="J2EEResourceProperty_17"
name="jmsDestination:customName" value="jmsDestination:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms:GenericJMSConnectionFactory"
xmi:id="GenericJMSConnectionFactory_1" name="jmsCF:name"
jndiName="jmsCF:jndiName" description="jmsCF:description"
userID="jmsCF:user" password="{xor}NTiShB11MT4y0g=="
externalJNDIName="jmsCF:externalJndiName" type="QUEUE">
<propertySet xmi:id="J2EEResourcePropertySet_16">
```

```

<resourceProperties xmi:id="J2EEResourceProperty_18" name="jmsCF:customName"
value="jmsCF:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_17">
<resourceProperties xmi:id="J2EEResourceProperty_19"
name="genericJMSProvider:customName" value="genericJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

Configuring new connection factories for application clients

Before you begin

During this task, you create a new JMS connection factory configuration for your application client.

Steps for this task

1. Click the JMS provider for which you want to create a connection factory in the tree. Do one of the following:
 - Configure a new JMS provider.
 - Click an existing JMS provider.
2. Expand the JMS provider to view its **JMS Connection Factories** folder.
3. Click the folder. Do one of the following:
 - Right-click the folder and click **New Factory**.
 - Click **Edit > New** on the menu bar.
4. Configure the JMS connection factory properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

Configuring new Java Message Service destinations for application clients

Before you begin

During this task, you create new Java Message Service (JMS) destination configuration for your application client.

Steps for this task

1. Click the JMS provider in the tree for which you want to create a destination. Do one of the following:
 - Configure a new JMS provider.
 - Click an existing JMS provider.
2. Expand the JMS provider to view its **JMS Destinations** folder.
3. Click the folder. Do one of the following:
 - Right-click the folder and click **New Factory**.
 - Click **Edit > New** on the menu bar.
4. Configure the JMS destination properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

Example: Configuring MQ Queue and Topic connection factories and destination factories for application clients

The purpose of this article is to help you configure MQ Queue connection factory, MQ Topic connection factory, MQ Queue destination factory, and MQ Topic destination factory settings.

- Required fields:
 - MQ Queue Connection Factory Properties page: name, jndiName, transport type
 - MQ Topic Connection Factory Properties page: name, jndiName, broker version
 - MQ Queue Factory Properties page: name, jndiName, persistence, priority, expiry, baseQueueName, targetClient
 - MQ Topic Factory Properties page: name, jndiName, persistence, priority, expiry, baseQueueName, targetClient
- Special cases:
 - The transport type must be CLIENT, or BINDINGS.
 - The Broker Version must be MA0C, or MQSI.
 - The port must be a numerical value between -2417483648 and 2417483647.
 - The CCSID must be a numerical value between -2417483648 and 2417483647.
 - The persistence value must be APPLICATION_DEFINED, QUEUE_DEFINED, PERSISTENT or, NONPERSISTENT.
 - The priority must be APPLICATION_DEFINED, QUEUE_DEFINED, or SPECIFIED.
 - The expiry must be APPLICATION_DEFINED, UNLIMITED, or SPECIFIED.
 - The integer encoding must be Normal, or Reversed.
 - The decimal encoding must be Normal, or Reversed.
 - The floating encoding must be IEEENormal, IEEEReversed, S390.
 - The target client must be JMS or MQ.
 - On the MQ Queue Connection Factory Properties page, only set the queueManager, host, and portWhen (required) fields if the transport type is CLIENT.
 - On the MQ Topic Connection Factory Properties page, only set the queueManager, host, and port (required) fields if the transport type is CLIENT.
 - On the the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, only set the Integer encoding, decimal encoding, and floating point encoding (required) fields if you do not set nativeEncoding.
 - On the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, the specified priority entry field must be an integer between 0 and 9 if priority is set to SPECIFIED .
 - On the the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, the specified expiry entry field must be a value greater than 0 if expiry is set to SPECIFIED.
- Example:

```
<resources.jms:JMSProvider xmi:id="JMSProvider_1" name="MQ JMS Provider"
description="mqJMSProvider:description"
externalInitialContextFactory="mqJMSProvider:contextFactoryClass"
externalProviderURL="mqJMSProvider:providerUrl">
<classpath>mqJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms.mqseries:MQQueueConnectionFactory"
xmi:id="MQQueueConnectionFactory_1" name="mqQCF:name"
jndiName="mqQCF:jndiName" description="mqQCF:description"
userID="mqQCF:user" password="{xor}Mi40HB11MT4y0g=="
```

```

queueManager="mqQCF:queueManager" host="mqQCF:host" port="1"
channel="mqQCF:channel" transportType="CLIENT"
clientID="mqQCF:clientId" CCSID="2">
<propertySet xmi:id="J2EEResourcePropertySet_3">
<resourceProperties xmi:id="J2EEResourceProperty_3" name="mqQCF:customName"
value="mqQCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQTopicConnectionFactory"
xmi:id="MQTopicConnectionFactory_1" name="mqTCF:name"
jndiName="mqTCF:jndiName" description="mqTCF:description"
userID="mqTCF:user" password="{xor}Mi4LHB11NTE7NhE+Mjo=" host="mqTCF:host"
port="1" transportType="CLIENT" channel="mqTCF:channel"
queueManager="mqTCF:queueManager" brokerControlQueue="mqTCF:brokerControlQueue"
brokerQueueManager="mqTCF:brokerQueueManager"
brokerPubQueue="mqTCF:brokerPubQueue"
brokerSubQueue="mqTCF:brokerSubQueue" brokerCCSubQ="mqTCF:brokerCCSubQ"
brokerVersion="MA0C" clientID="mqTCF:clientId" CCSID="2">
<propertySet xmi:id="J2EEResourcePropertySet_4">
<resourceProperties xmi:id="J2EEResourceProperty_4" name="mqTCF:customName"
value="mqTCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQQueue" xmi:id="MQQueue_1"
name="mqQ:name" jndiName="mqQ:jndiName" description="mqQ:description"
persistence="APPLICATION_DEFINED" priority="SPECIFIED" specifiedPriority="1"
expiry="SPECIFIED" specifiedExpiry="1" baseQueueName="mqQ:baseQueueName"
baseQueueManagerName="mqQ:baseQueueManagerName" CCSID="1"
integerEncoding="Normal" decimalEncoding="Normal"
floatingPointEncoding="IEEENormal" targetClient="JMS">
<propertySet xmi:id="J2EEResourcePropertySet_5">
<resourceProperties xmi:id="J2EEResourceProperty_5"
name="mqQ:customName" value="mqQ:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQTopic"
xmi:id="MQTopic_1" name="mqT:name" jndiName="mqT:jndiName"
description="mqT:description" persistence="APPLICATION_DEFINED"
priority="SPECIFIED" specifiedPriority="1" expiry="SPECIFIED"
specifiedExpiry="2" baseTopicName="mqT:baseTopicName" CCSID="3"
integerEncoding="Normal" decimalEncoding="Normal"
floatingPointEncoding="IEEENormal" targetClient="JMS"
brokerDurSubQueue="mqT:brokerDurSubQueue"
brokerCCDurSubQueue="mqT:brokerCCDurSubQueue">
<propertySet xmi:id="J2EEResourcePropertySet_6">
<resourceProperties xmi:id="J2EEResourceProperty_6" name="mqT:customName"
value="mqT:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_7">
<resourceProperties xmi:id="J2EEResourceProperty_7"
name="mqJMSProvider:customName" value="mqJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

Example: Configuring WAS Queue and Topic connection factories and destination factories for application clients

The purpose of this article is to help you to configure WAS Queue connection factory, WAS Topic connection factory, WAS Queue destination factory, and WAS Topic destination factory settings.

- Required fields:
 - JMS Provider Properties page: name
 - WAS Queue Connection Factory Properties page: name, jndiName, node

- WAS Topic Connection Factory Properties page: name, jndiName, node, port
- WAS Queue Factory Properties page: name, jndiName, node, persistence, priority, expiry
- WAS Topic Factory Properties page: name, jndiName, topic name, persistence, priority, expiry
- Special cases:
 - The port must be QUEUED or DIRECT.
 - The CCSID must be a numerical value between -2417483648 and 2417483647.
 - The persistence value must be APPLICATION_DEFINED, PERSISTENT, or NONPERSISTENT.
 - The priority must be APPLICATION_DEFINED, or SPECIFIED.
 - The expiry must be APPLICATION_DEFINED, UNLIMITED, or SPECIFIED.
 - On the WAS Topic Factory Properties, and the WAS Queue Factory Properties pages, the specified priority entry field must be an integer between 0 and 9 if priority is set to SPECIFIED .
 - On the WAS Topic Factory Properties, and the WAS Queue Factory Properties pages, the specified expiry entry field must be an value greater than 0 if expiry is set to SPECIFIED.
- Example:

```

<resources.jms:JMSProvider xmi:id="JMSProvider_2" name="WebSphere JMS Provider"
description="wasJMSProvider:description"
externalInitialContextFactory="wasJMSProvider:contextfactoryclass"
externalProviderURL="wasJMSProvider:providerUrl">
<classpath>wasJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms.internalmessaging:WASQueueConnectionFactory"
xmi:id="WASQueueConnectionFactory_1" name="wasQCF:name"
jndiName="wasQCF:jndiName" description="wasQCF:description"
userID="wasQCF:user" password="{xor}KD4sDhwZZS0i0=" node="wasQCF:Node">
<propertySet xmi:id="J2EEResourcePropertySet_8">
<resourceProperties xmi:id="J2EEResourceProperty_8" name="wasQCF:customName"
value="wasQCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASTopicConnectionFactory"
xmi:id="WASTopicConnectionFactory_1" name="wasTCF:name"
jndiName="wasTCF:jndiName" description="wasTCF:description"
userID="wasTCF:user" password="{xor}KD4sCxwZZTE+Mjo="
node="wasTCF:node" port="QUEUED" clientID="wasTCF:clientId">
<propertySet xmi:id="J2EEResourcePropertySet_9">
<resourceProperties xmi:id="J2EEResourceProperty_9" name="wasTCF:customName"
value="wasTCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASQueue"
xmi:id="WASQueue_1" name="wasQ:name" jndiName="wasQ:jndiName"
description="wasQ:description" node="wasQ:node"
persistence="APPLICATION_DEFINED" priority="SPECIFIED"
specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="1">
<propertySet xmi:id="J2EEResourcePropertySet_10">
<resourceProperties xmi:id="J2EEResourceProperty_10" name="wasQ:customName"
value="wasQ:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASTopic"
xmi:id="WASTopic_1" name="wasT:name" jndiName="wasT:jndiName"
description="wasT:description" topic="wasT:topicName"
persistence="APPLICATION_DEFINED" priority="SPECIFIED"
specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="1">
<propertySet xmi:id="J2EEResourcePropertySet_11">
<resourceProperties xmi:id="J2EEResourceProperty_11" name="wasT:customName"

```

```

value="wasT:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_12">
<resourceProperties xmi:id="J2EEResourceProperty_12"
name="wasJMSProvider:customName" value="wasJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

Configuring new resource environment providers for application clients

Before you begin

During this task, you create new resource environment provider configurations for your application client.

To configure a new resource environment provider, perform the following steps:

Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new JMS provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file in which you want to configure the new JMS provider.
3. Expand the JAR file to view its contents.
4. Click the folder called **Resource Environment Providers**. Do one of the following:
 - Right-click the folder and click **New Provider**.
 - Click **Edit > New** on the menu bar.
5. Configure the JMS provider properties in the resulting property dialog.
6. Click **OK** when finished.
7. Click **File > Save** on the menu bar to save your changes.

Resource environment provider settings for application clients

Use this page to specify resource environment entry properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Resource Environment Providers >** and click **New**. The following fields appear on the **General** tab:

Name: Specifies the administrative name for the resource environment provider.

Description: Specifies a description of the resource environment provider for your administrative records.

Class Path: Specifies the path to the JAR file that contains the implementation classes for the resource environment provider.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Configuring new resource environment entries for application clients

Before you begin

During this task, you create new resource environment entries for your client application.

Steps for this task

1. Click the resource environment provider in the tree, for which you want to create a resource environment entry. Do one of the following:
 - Configure a new resource environment provider.
 - Click an existing resource environment provider.
2. Expand the resource environment provider to view the **resource environment entries** folder.
3. Click the folder. Do one of the following:
 - Right-click the folder and select **New Factory**.
 - Click **Edit > New** on the menu bar.
4. Configure the data source properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

Resource environment entry settings for application clients

Use this page to specify resource environment entry properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Resource Environment Providers** > *resource environment instance*. Right-click **Resource environment entry** > and click **New**. The following fields appear on the **General** tab:

Name: Specifies the administrative name for the resource environment entry.

Description: Specifies a description of the URL for your administrative records.

JNDI Name: Specifies the Java Naming and Directory Interface (JNDI) name for the resource, including any naming subcontexts.

Use this name to link to the binding information of the platform. The binding associates the resources defined in the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Custom Properties: Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Managing application clients

Before you begin

Perform the following tasks after deploying application clients.

Note: This task only applies to J2EE application clients.

Steps for this task

1. Update data source and data source provider configurations.
2. Update URLs and URL provider configurations.
3. Update mail session configurations.
4. Update JMS provider, connection factories, and destination configurations.
5. Update MQ JMS provider, MQ connection factories, and MQ destination configurations.
6. Update Resource Environment Entry and Resource Environment Provider configurations.
7. (Optional) Remove application client resources.

Updating data source and data source provider configurations with the Application Client Resource Configuration Tool

Before you begin

During this task, you update the configuration of an existing data source or data source provider.

Steps for this task

1. Start the tool and open the EAR file containing the data source or data source provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the data source or data source provider to update.
3. Expand the JAR file to view its contents until you locate the particular data source or data source provider to update. Do one of the following:
 - Right-click the object and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
4. Update the properties in the resulting property dialog. For detailed field help, go to:
 - Data source provider properties
 - Data source properties
5. Click **OK** when finished.
6. Click **File > Save** on the menu bar to save your changes.

Updating URLs and URL provider configurations for application clients

Steps for this task

1. Start the tool and open the EAR file containing the URL or URL provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the URL or URL provider to update.
3. Expand the JAR file to view its contents.

4. Keep expanding the JAR file contents until you locate the particular URL or URL provider to update. Do one of the following:
 - a. Right-click the object and click **Properties**
 - b. Click **Edit > Properties** on the menu bar.
5. Update the properties in the resulting property dialog.
6. Click **OK** when finished.
7. Click **File > Save** to save your changes on the menu bar.

Updating mail session configurations for application clients

Before you begin

During this task, you update the configuration of an existing JavaMail session.

Note:

You cannot update the name of the default JavaMail provider. Also, you cannot delete the default JavaMail provider from the tree.

Steps for this task

1. Start the tool and open the EAR file containing the JavaMail session.
The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the JavaMail session to update.
3. Expand the JAR file to view its contents.
4. Keep expanding the JAR file contents until you locate the particular JavaMail session to update. Do one of the following:
 - a. Right-click the object and click **Properties**
 - b. Click **Edit > Properties** from the menu bar.
5. Update the properties in the resulting property dialog.
6. Click **OK** when finished.
7. Select **File > Save** from the menu bar to save your changes.

Updating Java Message Service provider, connection factories, and destination configurations for application clients

Before you begin

During this task, you update the configuration of an existing Java Message Service (JMS) provider, connection factory, or destination.

Steps for this task

1. Start the tool and open the EAR file containing the JMS provider, connection factory, or destination. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the JMS provider, connection factory, or destination to update.
3. Expand the JAR file to view its contents until you locate the particular JMS provider, connection factory, or destination to update. When you find it, do one of the following:
 - Right-click the object and click **Properties**.
 - Click **Edit > Properties** on the menu bar.

4. Update the properties in the resulting property dialog. For detailed field help, see:
 - JMS provider properties
 - WAS Queue connection factory properties
 - WAS Topic connection factory properties
 - WAS Queue destination properties
 - WAS Topic destination properties
5. Click **OK**.
6. Click **File > Save** to save your changes.

Updating MQ Java Message Service provider, MQ connection factories, and MQ destination configurations for application clients

Before you begin

During this task, you will update the configuration of an existing MQ JMS provider, MQ connection factory, or MQ destination.

Steps for this task

1. Start the tool and open the EAR file containing the MQ JMS provider, MQ connection factory, or MQ destination. The EAR file contents will be displayed in a tree view.
2. Select from the tree the JAR file containing the MQ JMS provider, MQ connection factory, or MQ destination to update.
3. Expand the JAR file to view its contents until you locate the particular MQ JMS provider, MQ connection factory, or MQ destination that you want to update. Do one of the following:
 - Right-click the object and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
4. Update the properties in the resulting property dialog. For detailed field help, see:
 - JMS provider properties
 - MQ Queue connection factory properties
 - MQ Topic connection factory properties
 - MQ Queue destination properties
 - MQ Topic destination properties
5. Click **OK**.
6. Click **File > Save** to save your changes.

Updating Resource Environment Entry and Resource Environment Provider configurations for application clients

Before you begin

During this task, you update the configuration of an existing Resource Environment Entry or Resource Environment Provider.

Steps for this task

1. Start the tool and open the EAR file containing the Resource Environment Entry or Resource Environment Provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the Resource Environment Entry or Resource Environment provider to update.
3. Expand the JAR file to view its contents until you locate the Resource Environment Entry or Resource Environment Provider to update. Do one of the following:
 - Right-click the object and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
4. Update the properties in the resulting property dialog. For detailed field help, see:
 - Resource environment provider properties
 - Resource environment entry properties
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

Example: Configuring Resource Environment settings

The purpose of this article is to help you configure Resource Environment settings.

- Required fields:
 - Resource Environment Provider page: name
 - Resource Environment Entry page: name, jndiName
- Example:

```
<resources.env:ResourceEnvironmentProvider xmi:id="ResourceEnvironmentProvider_1"
name="resourceEnvProvider:name" description="resourceEnvProvider:description">
<classpath>resourceEnvProvider:classpath</classpath>
<factories xmi:type="resources.env:ResourceEnvEntry"
xmi:id="ResourceEnvEntry_1" name="resourceEnvEntry:name"
jndiName="resourceEnvEntry:jndiName" description="resourceEnvEntry:description">
<propertySet xmi:id="J2EEResourcePropertySet_20">
<resourceProperties xmi:id="J2EEResourceProperty_22"
name="resourceEnvEntry:customName" value="resourceEnvEntry:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_21">
<resourceProperties xmi:id="J2EEResourceProperty_23"
name="resourceEnvProvider:customName" value="resourceEnvProvider:customValue"/>
</propertySet>
</resources.env:ResourceEnvironmentProvider>
```

Example: Configuring Resource Environment custom settings for application clients

The purpose of this article is to help you configure Resource Environment custom settings.

- The custom page applies to every resource type. You can specify as many custom names and values as you need.
- Example:

```
<propertySet xmi:id="J2EEResourcePropertySet_20">
<resourceProperties xmi:id="J2EEResourceProperty_22"
name="resourceEnvEntry:customName" value="resourceEnvEntry:customValue"/>
</propertySet>
```

Removing application client resources

Before you begin

Note: This task only applies to J2EE application clients.

Steps for this task

1. Start the Application Client Resource Configuration Tool (ACRCT) and open the EAR file from which you want to remove an object. The EAR file contents display in a tree view.
If you already have an EAR file open, and have made some changes, click **File > Save** to save your work before proceeding to delete an object.
2. Locate the object that you want to remove in the tree.
3. Right-click the object, then click **Delete**.
4. Click **File > Save**.

What to do next

The option to delete an item does not offer a confirmation dialog. As a safeguard, consider saving your work right before you begin this task. If you change your mind after removing an item, you can close the EAR file without saving your changes, canceling your deletion. Remember to close the EAR file immediately after the deletion, or you also lose any unsaved work that you performed since the deletion.

Running application clients

Before you begin

The J2EE specification requires support for a client container that runs standalone Java applications (known as J2EE application clients) and provides J2EE services to the applications. J2EE services include naming, security, and resource connections.

You are ready to run your application client using this tool after you have:

1. Written the application client program.
2. Assembled and installed an application module (.ear file) in the application server run time.

Note: This task only applies to J2EE application clients.

Steps for this task

1. Open a command window and invoke the following script to launch J2EE application clients using the launchClient shell:

```
install_root/bin/launchClient.bat
```

The launchClient batch command starts the application client run time, which:

- Initializes the client run time.
- Loads the class that you designated as the main class with the Application Assembly Tool (AAT).
- Runs the main method of the application client program.

When your program terminates, the application client run time cleans up the environment and the Java Virtual Machine code ends.

2. Pass parameters to the **launchClient** command. You can pass parameters to your application client program as well. The `launchClient` command allows you to do both.

The `launchClient` command requires that the first parameter is either:

- An EAR file specifying the application client to launch.
- A request for `launchClient` usage information.

All other parameters intended for the **launchClient** command must begin with the `-CC` prefix.

Parameters that are not EAR files, or usage requests, or that do not begin with the `-CC` prefix, are ignored by the application client run time, and are passed directly to the application client program.

The **launchClient** command retrieves parameters from three places:

- a. The command line
- b. A properties file
- c. System properties

The parameters are resolved in the order listed above, with command line values having the highest priority and system properties the lowest. This prioritization allows you to set and override default values.

3. Specify the server name.

By default, the **launchClient** command uses the environment variable `COMPUTERNAME` for the `BootstrapHost` property value. This setting is effective for testing your application client when it is installed on the same computer as the server. However, in other cases override this value with the name of your server.

You can override the `BootstrapHost` value by invoking `launchClient` with the following parameters:

```
launchClient myapp.ear-CCBootstrapHost=abc.midwest.mycompany.com
```

You can also override the default by specifying the value in a properties file and passing the file name to the `launchClient` shell.

Note: Security is controlled by the server. You do not need to configure security on the client because the client assumes that security is enabled. If security is not enabled, the server ignores the security request, and the application client works as expected.

Usage scenario

You can store `launchClient` values in a properties file, a good method for distributing default values. You can then override one or more values on the command line. The format of the file is one `launchClient -CC` parameter per line without the `-CC` prefix. For example:

```
verbose=true  
classpath=c:\mydir\util.jar;c:\mydir\harness.jar;c:\production\G19\global.jar  
BootstrapHost=abc.westcoast.mycompany.com tracefile=c:\WebSphere\mylog.txt
```

launchClient tool

Syntax

This section describes the command line syntax for the Java TM2 Platform, Enterprise Edition (J2EE) launchClient tool.

The command line invocation syntax for the launchClient tool follows:

```
launchClient [<userapp.ear> [-help|-?] [-CCname=value] [app args]
```

where *userapp.ear* is the path and the name of the EAR file that contains the application client, *name* is the name of the parameter, *value* is the value to which the parameter id set, and *app args* are arguments that pass to the application client.

To print the usage information, the first parameter must be a path and a name to an EAR file, `-help`, or `-?`. All the other parameters are optional and can appear in any order. The application client run time ignores any optional parameters that do not begin with a `-CC` prefix, and passes them to the application client.

Parameters

Supported arguments include:

-CCsoapConnectorPort

The soap connector port. If you do not specify this argument, the WebSphere Application Server default value is used.

-CCverbose

This option displays additional information messages. The default is false.

-CCclasspath

A class path value. When you launch an application, the system class path is not used. If you want to access classes that are not in the EAR file or part of the resource class paths, specify the appropriate class path here. Multiple paths can be concatenated.

-CCjar

The name of the client JAR file that resides within the EAR file for the application you wish to launch. Use this argument when you have multiple client JAR files in the EAR file.

-CCaltDD

The name of an alternate deployment descriptor. This parameter is used with the `-CCjar` parameter to specify the deployment descriptor to use. Use this argument when a client jar file is configured with more than one deployment descriptor. Set the value to null to use the client JAR file standard deployment descriptor

-CCBootstrapHost

The name of the host server you want to connect to initially. The format is: `<your.server.ofchoice.com<`

-CCBootstrapPort

The server port number. If you do not specify this argument, the WebSphere Application Server default value is used.

-CCproviderURL

Provides bootstrap server information that the initial context factory can use to obtain an initial context. WebSphere Application Server initial context factory can use either a CORBA object URL or an IIOP URL. CORBA object URLs are more flexible than IIOP URLs and are the recommended URL format to use. This value can contain more than one bootstrap server address. This feature can be used when attempting to obtain an initial context from a server cluster. You can specify bootstrap

server addresses, for all servers in the cluster, in the URL. The operation will succeed if at least one of the servers is running, eliminating a single point of failure. The address list does not process in a particular order. For naming operations, this value overrides the `-CCBootstrapHost` and `-CCBootstrapPort` parameters. An example of a CORBA object URL specifying multiple systems follows:

```
-CCproviderURL=corbaloc:iiop:myserver.mycompany.com:9810,:mybackupserver.  
mycompany.com:2809
```

This value is mapped to the `java.naming.provider.url` system property.

-CCinitonly

Use this option to initialize application client run time for ActiveX application clients without launching the client application. The default is `false`.

-CCtrace

Use this option to obtain debug trace information. You might need this information when reporting a problem to IBM Service. The default is `false`.

-CCtracefile

The name of the file to write trace information. The default is to output to the console.

-CCpropfile

Name of a properties file that contains `launchClient` properties. Specify the properties without the `-CC` prefix in the file. For example: `verbose=true`.

-CCsecurityManager

Enables and runs the WebSphere Application Server with a security manager. The default is `disable`.

-CCsecurityMgrClass

The fully qualified name of a class that implements a security manager. Only use this argument if the `-CCsecurityManager` parameter is set to `enable`. The default is `java.lang.SecurityManager`.

-CCsecurityMgrPolicy

The name of a security manager policy file. Only use this argument if the `-CCsecurityManager` parameter is set to `enable`. When you enable this parameter, the `java.security.policy` system property is set. The default is `install_root>/properties/client.policy`.

-CCD Use this option to have the WebSphere Application Server set the specified system property during initialization. Do not use the `=` character after the `-CCD`. For example: `-CCDcom.ibm.test.property=testvalue`. You can specify multiple `-CCD` parameters. The general format of this parameter is `-CCD<property key>=<property value>`.

-CCexitVM

Use this option to have the WebSphere Application Server call `System.exit()` after the client application completes. The default is `false`.

-CCdumpJavaNameSpace

Prints out the Java portion of the WebSphere Application Server Java Naming and Directory Interface (JNDI) name space. The `true` value uses the short format which prints out the binding name and the type of the object bound at that location. The `long` value uses the long format which

prints out the binding name, bound object type, local object, type, and string representation of the local object, for example: IORs, and string values. The default value is false.

Examples

The following examples demonstrate correct syntax.

On the Windows operating system:

```
launchClient c:\earfiles\myapp.ear -CCBootstrapHost=myWASServer  
-CCverbose=true app_parm1 app_parm2
```

On the UNIX operating system:

```
./launchClient.sh /usr/earfiles/myapp.ear  
-CCBootstrapHost=myWASServer-CCverbose=true app_parm1 app_parm2
```

Application client troubleshooting tips

This section provides some debugging tips for resolving common J2EE application client problems. To use this troubleshooting guide, review the trace entries for one of the J2EE application client exceptions, and then locate the exception in the guide. Some of the errors in the guide are samples; the actual error you receive can be slightly different than what is shown here. Also, it can be useful to rerun the **launchClient** command specifying the **-CCverbose=true** option. This option provides additional information when the J2EE application client run time is initializing

Error: `java.lang.NoClassDefFoundError`

Explanation	This exception is thrown when Java code cannot load the specified class.
Possible causes	<ul style="list-style-type: none">• Invalid or non-existent class• Classpath problem• Manifest problem

Recommended response

Check to determine if the specified class exists in a JAR file within your EAR file. If it does, make sure the path for the class is correct. For example, if you get the exception:

```
java.lang.NoClassDefFoundError:  
WebSphereSamples.Hello  
EJB.HelloHome
```

ensure the class HelloHome exists in one of the JAR files in your EAR file. If it exists, ensure the path for the class is WebSphereSamples.HelloEJB.

If both the class and path are correct, then it is a classpath issue. Most likely, you do not have the failing class JAR file specified in the client JAR file manifest. To verify this situation, perform the following steps:

1. Open your EAR file with the Application Assembly Tool and click on the Application Client.
2. Add the names of the other JAR files in the EAR file to the Classpath field.

This exception is generally caused by a missing EJB module name from the Classpath field.

If you have multiple JAR files to enter in the Classpath field, be sure to separate the JAR names with spaces.

If you still have the problem, you have a situation where a class is loaded from the file system instead of the EAR file. This is a very difficult situation to debug because the offending class is not the one specified in the exception. Instead, another class is loaded from the file system before the one specified in the exception. To correct this problem, review the classpaths specified with the `-CCclasspath` option and the classpaths configured with the Application Client Resource Configuration Tool. Look for classes that also exist in the EAR file. You must resolve the situation where one of the classes is found on the file system instead of in the `.ear` file. Remove entries from the classpaths, or include the `.jar` files and classes in the `.ear` file instead of referencing them from the file system.

If you use the `-CCclasspath` parameter or resource classpaths in the Application Client Resource Configuration Tool, and you have configured multiple JAR files or classes, verify they are separated with the correct character for your operating system. Unlike the classpath field in the Application Assembly Tool, these classpath fields use platform-specific separator characters, usually a colon (on UNIX platforms) or a semi-colon (on Windows systems).

Note: The system classpath is not used by the Application Client run time if you use the `launchClient` batch or shell files. In this case, the system classpath would not cause this problem. However, if you load the `launchClient` class directly, you do have to search through the system classpath as well.

Error: com.ibm.websphere.naming.CannotInstantiateObjectException: Exception occurred while attempting to get an instance of the object for the specified reference object. [Root exception is javax.naming.NameNotFoundException: xxxxxxxxxx]

Explanation	This exception occurs when you perform a lookup on an object that is not installed on the host server. Your program can look up the name in the local client Java Naming and Directory Interface (JNDI) name space, but received a <code>NameNotFoundException</code> exception because it is not located on the host server. One typical example is looking up an enterprise bean that is not installed on the host server that you access. This exception might also occur if the JNDI name you configured in your Application Client module does not match the actual JNDI name of the resource on the host server.
Possible causes	<ul style="list-style-type: none"> • Incorrect host server invoked • Resource is not defined • Resource is not installed • Application server is not started • Invalid JNDI configuration
Recommended response	If you are accessing the wrong host server, run the launchClient command again with the -CCBootstrapHost parameter specifying the correct host server name. If you are accessing the correct host server, use the WebSphere dumpnamespace command line tool to see a listing of the host server JNDI name space. If you do not see the failing object name, the resource is either not installed on the host server or the appropriate application server is not started. If you determine the resource is already installed and started, your JNDI name in your client application does not match the global JNDI name on the host server. Use the Application Assembly Tool to compare the JNDI bindings value of the failing object name in the client application to the JNDI bindings value of the object in the host server application. They must match.

Error: javax.naming.ServiceUnavailableException: A communication failure occurred while attempting to obtain an initial context using the provider url: "iiop://[invalidhostname]". Make sure that the host and port information is correct and that the server identified by the provider URL is a running name server. If no port number is specified, the default port number 2809 is used. Other possible causes include the network environment or workstation network configuration. Root exception is org.omg.CORBA.INTERNAL: JORB0050E: In Profile.getIPAddress(), InetAddress.getByName[invalidhostname] threw an UnknownHostException. minor code: 4942F5B6 completed: Maybe

Explanation	This exception occurs when you specify an invalid host server name.
Possible causes	<ul style="list-style-type: none"> • Incorrect host server invoked • Invalid host server name

Recommended response

Run the **launchClient** command again and specify the correct name of your host server with the **-CBootstrapHost** parameter.

Error: javax.naming.CommunicationException: Could not obtain an initial context due to a communication failure. Since no provider URL was specified, either the bootstrap host and port of an existing ORB was used, or a new ORB instance was created and initialized with the default bootstrap host of "localhost" and the default bootstrap port of 2809. Make sure the ORB bootstrap host and port resolve to a running name server. Root exception is org.omg.CORBA.COMM_FAILURE: WRITE_ERROR_SEND_1 minor code: 49421050 completed: No

Explanation

This exception occurs when you run the **launchClient** command to a host server that does not have the Application Server started. You also receive this exception when you specify an invalid host server name. This situation might occur if you do not specify a host server name when you run **launchClient**. The default behavior is for **launchClient** to run to localhost, because WebSphere Application Server does not know the name of your host server. This default behavior only works when you are running the client on the same computer with WebSphere Application Server is installed.

Possible causes

- Incorrect host server invoked
- Invalid host server name
- Invalid reference to localhost
- Application server is not started
- Invalid bootstrap port

Recommended response

If you are not running to the correct host server, run the **launchClient** command again and specify the name of your host server with the **-CBootstrapHost** parameter. Otherwise, start the Application Server on the host server and run the **launchClient** command again.

Error: javax.naming.NameNotFoundException: Name comp/env/ejb not found in context "java:"

Explanation

This exception is thrown when the Java code cannot locate the specified name in the local JNDI name space.

Possible causes

- No binding information for the specified name
- Binding information for the specified name is incorrect
- Wrong class loader was used to load one of the program classes
- A resource reference does not include any client configuration information

Recommended response

Open the EAR file with the Application Assembly Tool and check the bindings for the failing name. Ensure this information is correct. If you are using Resource References, open the EAR file with the Application Client Resource Configuration Tool, and make sure the Resource Reference has client configuration information and the name of the Resource Reference exactly matches the JNDI name of the client configuration. If it is correct, you might have a class loader issue.

Error: java.lang.ClassCastException: Unable to load class: org.omg.stub.WebSphereSamples.HelloEJB._HelloHome_Stub at com.ibm.rmi.javax.rmi.PortableRemoteObject.narrow (portableRemoteObject.java:269)

Explanation

This exception occurs when the application program attempts to narrow to the EJB home class and the class loaders cannot find the EJB client side bindings.

Possible causes

- The files, *_Stub.class and _Tie.class, are not in the EJB .jar file

Recommended response

- Class loader could not find the classes
Look at the EJB .jar file located in the .ear file and verify the class contains the EJB client side bindings. These are class files whose names end in _Stub and _Tie. If these files are not present, then use the Application Assembly Tool to generate the binding classes. For more information, see article [Generating deployment code for modules](#). If the binding classes are in the EJB .jar file, then you might have a class loader issue.

Error: WSCL0210E: The Enterprise archive file [EAR file name] could not be found. com.ibm.websphere.client.applicationclient.ClientContainerException: com.ibm.etools.archive.exception.OpenFailureException

Explanation

This error occurs when the application client run time cannot read the Enterprise Archive (EAR) file.

Possible causes

The most likely cause of this error is that the system cannot find the EAR file cannot be found in the path specified on the **launchClient** command.

Recommended response

Verify that the path and file name specified on the **launchclient** command are correct. If you are running on the Windows NT operating system and the path and file name are correct, use a short version of the path and file name (8 character file name and 3 character extension).

The launchClient command appears to hang and does not return to the command line when the client application has finished.

Explanation	<p>When running your application client using the launchClient command the WebSphere Application Server run time might need to display the security login dialog. To display this dialog the WebSphere Application Server run time creates an Abstract Window Toolkit (AWT) thread. When your application returns from its main method to the application client run time, the application client run time attempts to return to the operating system and end the Java Virtual Machine code. However, since there is an AWT thread, the Java Virtual Machine code will not end until <code>System.exit</code> is called.</p>
Possible causes	<p>The Java Virtual Machine code does not end because there is an AWT thread. Java code requires that <code>System.exit()</code> be called to end AWT threads.</p>
Recommended response	<ul style="list-style-type: none">• Modify your application to call <code>System.exit(0)</code> as the last statement.• Use the <code>-CCexitVM=true</code> parameter when you call the launchClient command.

Chapter 7. Developing and managing Web services

Developing and managing Apache SOAP 2.3-based Web services explains how to work with Web services that are written directly to Apache SOAP, Version 2.3 API. This activity is described in the following tasks:

1. Developing a Simple Object Access Protocol client.
2. Assembling Web services.
3. Deploying Web services applications.
4. Administering deployed Web services (SOAP-XML administrative tool).
5. Securing SOAP services.

An additional Web services component, *IBM WebSphere Web Services for J2EE Technology Preview*, is available for use with Version 5.0. It is a separate download available at <http://www7b.boulder.ibm.com/wsdd/downloads/techreviews.html>, and includes additional documentation.

The Web services technology preview supports emerging Java Web services standards like JAX-RPC and Web services for J2EE. It is recommended that new development efforts use the Web Services Technology Preview and follow these standards.

You can find more information about working with Web services by visiting the Internet sites listed in the InfoCenter article *Web services: Resources for Learning*.

Web services

Web services are self-contained, modular applications that you can describe, publish, locate, and invoke over a network.

Web services could be weather reports or stock quotes. Transaction Web services, supporting business-to-business (B2B) or business-to-client (B2C) operations, could be airline reservations or purchase orders.

Web services reflect a new, service-oriented approach to programming, based on the idea of building applications by discovering and implementing network-available services, or by invoking available applications to accomplish some task. This service-oriented approach is independent of specific programming languages or operating systems. Instead, Web services rely on pre-existing transport technologies, such as HTTP, and standard data encoding techniques, such as XML, for invoking the implementation.

The key components of a Web service are:

- Simple Object Access Protocol (SOAP)
- Web Services Description Language (WSDL)
- Universal Discovery, Description and Integration Protocol (UDDI)

For more information about Web services, see *Web services: Resources for learning*.

Developing a Simple Object Access Protocol client

Before you begin

Develop a client that has a Web Services Description Language file or service implementation.

Creating clients to access the Simple Object Access Protocol (SOAP) services published in WebSphere Application Server is a straightforward process. The Apache SOAP implementation, integrated with WebSphere Application Server, contains a client API to assist in SOAP client application development.

The SOAP API documentation is available in WebSphere Application Server Java documentation.

To create a client that interacts with a SOAP Remote Procedure Call (RPC) service:

Steps for this task

1. Obtain the interface description of the SOAP service.
This description provides the signatures of the methods that you want to invoke. You can either look at a Web Services Description Language (WSDL) file for the service, or view the service itself to see its implementation.
2. Create the call object.
The SOAP call object is the main interface to the underlying SOAP RPC code.
3. Set the target Uniform Resource Identifier (URI) in the call object using the `setTargetObjectURI()` method.
Pass the Uniform Resource Name (URN), a type of URI, that the service uses for its identifier, in the deployment descriptor.
4. Set the method name that you want to invoke in the call object using the `setMethodName()` method.
This method must be one of the methods exposed by the service located at the URN from the previous step.
5. Create the necessary parameter objects for the RPC call and then set them in the call object using the `setParams()` method.
Ensure you have the same number and same type of parameters as those required by the service.
6. Execute the call object `invoke()` method and retrieve the response object.
Remember the RPC call is synchronous, so it can take some time to complete.
7. Check the response for a fault using the `getFault()` method, and then extract any results or returned parameters.
While most of the providers only return a result, the DB2 stored procedure provider can also return output parameters.

What to do next

Interacting with a document-oriented SOAP service requires you to use lower-level Apache SOAP API calls. You must first construct an envelope object containing the contents of the message, including the body and any headers, that you wish to send. Then create a message object where you invoke the `send()` method to perform the actual transmission.

To create a secure SOAP service, do the following:

1. Create a simple object.
2. Define an envelope editor.
3. Specify a pluggable envelope editor.

4. Define the transports.

Your code can look like the following example:

```
EnvelopeEditor editor=  
new PluggableEnvelopeEditor(new InputSource(conf), home);  
SOAPTransport transport =  
new FilterTransport(editor, new SOAPHTTPConnection());  
call.setSOAPTransport(transport);
```

The characteristics of the secure session are specified by the conf configuration file.

Administering enterprise beans through Simple Object Access Protocol

Calling enterprise beans through Simple Object Access Protocol (SOAP) is handled in the same manner as calling Java bean methods through SOAP. The SOAP run-time handles the bean cases for you, such as calling an enterprise bean create method if the create was not called previously.

A Web service can be a simple stateless session bean that performs number processing and returns a data value. When the client code makes a call to the data processing method of this service and an instance of the stateless session is not available, the SOAP run time does the following:

- Calls the enterprise bean create method to obtain a stateless session
- Calls the requested method

Assembling Web services

Use the Simple Object Access Protocol (SOAP) EAR Enabler tool to assemble Web services. The following tasks comprise assembling Web services:

Steps for this task

1. Create or locate the software resource to be exposed as a service.

To expose the service, create a programming artifact, one of the supported types, or locate an existing piece of code of the supported type.

2. (Assemble an Enterprise Archive (EAR) file).

Package the code artifact into an EAR file. This step is a deployment packaging requirement of WebSphere Application Server. Use the Application Assembly Tool (AAT) to package the artifact.

3. Create the Apache SOAP deployment descriptor for the desired service.

In order to deploy an artifact as a SOAP service, create a Apache SOAP deployment descriptor that describes the service you are creating. This step exposes the programming artifact as a service. The descriptor describes and defines the parts of the code that will be invoked with the SOAP calls.

The information contained in the deployment descriptor varies, depending on the type of artifact you are exposing. For more information about various deployment descriptor types see the article Simple Object Access Protocol deployment descriptor. For example, the following deployment descriptor might be used with the *StockQuoteSample*:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"  
id="urn:service-urn" [type="message"]>  
<isd:provider type="java"  
scope="Request | Session | Application"  
methods="exposed-methods">
```

```

<isd:java class="implementing-class" [static="true|false"]/>
</isd:provider>
<isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>

```

4. Execute the SoapEarEnabler tool to enable your Web service.

Your code artifact must first be packaged into an EAR file. Next, using the deployment descriptor as input, add the necessary pieces to the EAR file to enable the artifact as a Web service. To facilitate this process, use the Java based tool called SoapEarEnabler. Depending on whether you secure the Web service, this tool will add two Web modules: soap.war and soap-sec.war to the EAR file. These Web modules include the SOAP deployment descriptors, plus the necessary parts to deploy the service into the WebSphere Application Server run time.

The service does not become available until the soap-enabled EAR file is installed, and the server is restarted.

Simple Object Access Protocol deployment descriptor

Apache Simple Object Access Protocol (SOAP) utilizes XML documents called deployment descriptors to provide the SOAP run-time with information on client services.

Deployment descriptors provide a variety of information including:

- Service Uniform Resource Name (URN) used to route the request when it arrives
- Method and class details, if the service is being provided by a Java class
- User ID and password information, if the service provider must connect to a database

There are four types of deployment descriptors that you can use in WebSphere Application Server.

- Standard Java class
- Enterprise bean
- Bean Scripting Framework (BSF)
- DB2 stored

The contents of the deployment descriptor depend on the type of artifact that is being exposed using SOAP. Samples of these deployment descriptors are available in the SOAP 2.3 sample. You can link to the sample from the article [Web services: Resources for Learning](#).

Standard Java class deployment descriptor

A deployment descriptor that exposes a service implemented with a standard Java class, including a normal Java bean. An example of the standard Java class deployment descriptor follows:

```

<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn" [type="message"]>
  <isd:provider type="java"
    scope="Request | Session | Application"
    methods="exposed-methods">
    <isd:java class="implementing-class" [static="true|false"]/>
  </isd:provider>
<isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>

```

where:

- *service-urn* is the uniform resource name (URN) given to a service. All services deployed within a single EAR file must have URNs that are unique within that EAR file.
- *exposed-methods* is a list of methods, separated by spaces, which are being exposed.
- *implementing-class* is a fully-qualified class name (that is, a *packagename.classname*) that provides the methods that are being exposed.
- *type* is an optional attribute on the service <element>. Type is set to the value "message" if the service is document-oriented instead of RPC-invoked.
- *static* is an optional attribute called on the java> element, which may be set to either "true" or "false", depending on whether the methods are exposed or not exposed. If exposed, this attribute indicates whether the method is static or not static.
- *scope* is an attribute on the <provider> element, which indicates the lifetime of the instantiation of the implementing class.
- "Request" indicates the object is removed after the request completes.
- "Session" indicates the object lasts for the current lifetime of the HTTP session.
- "Application" indicates the object lasts until the servlet, which is servicing the requests, is terminated.

EJB deployment descriptor

A deployment descriptor that exposes a service implemented with an Enterprise Java Bean. An example of an EJB deployment descriptor is as follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn">
  <isd:provider type="provider-class"
    scope="Application"
    methods="exposed-methods">
    <isd:option key="JNDIName" value="jndi-name" />
    <isd:option key="FullHomeInterfaceName" value="home-name" />
  </isd:provider>
  <isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

where:

- *service-urn* and *exposed-methods* have the same meaning as in the standard Java class deployment descriptor
- *provider-class* with a stateless session bean follows:
com.ibm.soap.providers.WASStatelessEJBProvider
- *jndi-name* is the registered JNDI name of the enterprise bean
- *home-name* is the fully-qualified class name of the enterprise bean's home.

The default values for the iiop URL and context provider keys are:

```
<isd:option key="ContextProviderURL" value="iiop://localhost:2809" />
<isd:option key="FullContextFactoryName" value="com.ibm.websphere.naming.WsnInitialContextFactory" />
```

To use your own values, you must specify:

```
<isd:option key="ContextProviderURL" value="<URL to the JNDI provider>" />
<isd:option key="FullContextFactoryName" value="<Context factory full class name>" />
```

Bean Scripting Framework (BSF) script deployment descriptor

A deployment descriptor that exposes a service implemented with a BSF script. An example follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn">
  <isd:provider type="script"
    scope="Request | Session | Application"
    methods="exposed-methods">
    <isd:script language="language-name" [source="source-filename"]>[script-body]
    </isd:script>
  </isd:provider>
  <isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

where:

- *service-urn*, *exposed-methods*, and *scope* have the same meaning as in the standard Java class deployment descriptor.
- *language-name* is the name of the BSF-supported language that is used to write the script.
- *source* or *script-body* is an attribute that is mandatory for the `<script>` element. The *script-body* attribute contains the actual script that is used to provide the service.
- *source-filename* refers to the file which contains the service implementation if the deployment descriptor has the *source* attribute.

DB2 stored procedure deployment descriptor

A deployment descriptor that exposes one or more DB2 stored procedures as a service. An example follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn">
  <isd:provider type="com.ibm.soap.providers.WASDB2SPPProvider"
    scope="Application"
    methods="* | exposed-methods">
    <isd:option key="userID" value="db-userid" />
    <isd:option key="password" value="db-password" />
    [<isd:option key="fullContextFactoryName" value="context-factory" />]
    <isd:option key="datasourceJNDI" value="jndi-name" />]
    [<isd:option key="dbDriver" value="db-driver" />]
    <isd:option key="dbURL" value="db-url" />]
  </isd:provider>
  <isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

where:

- *service-urn* and *exposed-methods* have the same meaning as in the standard Java class deployment descriptor.
- *db-userid* is a valid user ID used to access the database where the stored procedures reside.
- *db-password* is a valid password for the specified user ID
The *db-userid* and *db-password* entries are optional. These entries can be set in the datasource. In WebSphere Application Server, the preferred way for administering the *db-userid* and *db-password* entries is with a datasource. Changing the user ID and password is easier when the information is located in a datasource rather than in a separate deployment descriptor file.
- *context-factory* is the name of the context factory used to access the database
- *jndi-name* is the datasource used to access the database
- *db-driver* is the database driver used to access the database.

A *db-driver* is not required if a data source JNDI name is specified.

- *db-url* is the URL that specifies the database to access.

The *methods* attribute on the <provider> element can contain a list of space-separated procedure names to expose, or an * (asterisk). An asterisk indicates all available stored procedures should be exposed.

Simple Object Access Protocol-enabling the services in an Enterprise Application Archive file

The Simple Object Access Protocol (SOAP) EAR enabler tool is a Java application that enables a set of SOAP services within an Enterprise Application Archive (EAR) file.

The SoapEAREnabler (SEE) tool guides you through the required steps to enable one or more services within an application. The tool makes a backup copy of your original EAR file in the event you need to remove or add services at a later time.

Before invoking the SEE tool, create an SOAP deployment descriptor for each service to be enabled.

Follow these steps to enable a Web service:

Steps for this task

1. Invoke the SEE tool from `was_install/bin` directory using one of the following commands:
 - Use the command **SoapEarEnabler** on the Windows platform
 - Use the command **SoapEarEnabler.sh** on UNIX platforms.
2. Start with an existing EAR file, either one created with the Application Assembly Tool (AAT), or a previously-created, valid, J2EE-compliant EAR file. The SEE tool will **not** accept a SOAP-enabled EAR file as input.
3. After you enable Web services, install the EAR file in WebSphere Application Server.
4. You will be prompted as to whether you wish to add the administration client to the EAR. This is a Web-based client that will allow you to list all active services for a specific context from a browser window. With this interface, you can stop and start existing services. You might choose to not add this interface for security reasons, or you might want to secure the interface before making a service available.

What to do next

You have invoked the SEE from the WebSphere Application Server bin directory and now you should know how the tool operates and its prompts and commands.

The tool operates in two modes:

1. Interactive
2. Silent

Specify all required command line arguments to use the tool in **silent** mode.

In **interactive** mode, the SEE tool prompts all required information. The following dialog is an example of using the tool in interactive mode.

Note: In this dialog, user input is in *italics*, and tool output is in **bold**.

SoapEarEnabler (On Windows NT)
SoapEarEnabler.sh (On UNIX platforms)

Please enter the name of your ear file:

..\work\stockquote.ear

How many services would you like your application to contain (1...n)?

1

Now prompting for info for service #1:

Please enter the file name of the SOAP deployment descriptor xml file:

..\work\StockQuoteDD.xml

Is this service an EJB (y=yes / n=no)?

n

If there are additional jar files that need to be added to the classpath, specify the number to add (answer 0 for none) (0...n)?

1

Classpath addition 1: Here are all the archive files found, please choose one to add to the classpath: ([1] samples.jar, [2] stockquote.war):

1

Should this service be secured ("y/n" versus "y= yes, n= no)?

n

Please enter a context root for your non-secured services (e.g. /soap):

/soapsamples

Do you wish to install the administration client?

Warning! You should not install this client in a production ear unless you intend to secure the URI to it.

Install the administration client ("y/n" versus "y= yes, n= no)?

y

In **Silent mode**, supply the arguments in the same order as for the interactive prompts.

Note: In silent mode, the SOAP administrative GUI will not install. Also, you will not be prompted for the SOAP administrative GUI.

The following example describes how to use the tool in silent mode:

soapearenabler [args]

where the arguments must be specified in the following order:

<ear-file-name>

<number-of-services>

The following block is repeated based on the number of services specified.

<deployment-descriptor-file-name>

<service-is-an-ejb-(y/n)>

The following argument should be supplied only if service-is-an-ejb-(y/n) is y.

<ejb-jar-file-uri-(already-in-ear)>

<number-of-additional-jar-files>

(0, 1, 2...)>

The following argument is repeated number-of-additional-jar-files times.

<classpath-entry-uri-(already-in-ear)>

<secure-this-service-(y/n)>

This following argument is supplied only if secure-this-service-(y/n) is n for any service.

<context-root-for-non-secured-services,

ex: /soap>

This following argument is supplied only if secure-this-service-(y/n) is y for any service.

<context-root-for-secured-services, ex:

/soapsec>

Examples

Silent mode example

The following is an example of deploying one enterprise bean as a non-secured service:

```
soapearenabler soap.ear 1
  d:\xml-soap\java\samples\ejbadder\deploymentdescriptor.xml
y adderservice-ejb.jar 1 samples.jar n /soap
```

The following is an example of deploying one enterprise bean as a non-secured service, and one Java class as a secured service:

```
soapearenabler soap.ear 2
  d:\xml-soap\java\samples\ejbadder\deploymentdescriptor.xml
y adderservice-ejb.jar 1 samples.jar n
  d:\xml-soap\java\samples\stockquote\deploymentdescriptor.xml
n 1 samples.jar y /soap /soap-sec
```

Deploying two Java classes as non-secured services example

```
soapearenabler soap.ear 2
  d:\xml-soap\java\samples\stockquote\deploymentdescriptor.xml
n 1 samples.jar n
  d:\xml-soap\java\samples\addressbook\deploymentdescriptor.xml
n 1 samples.jar n /soap
```

The line breaks in the above examples have been modified for this documentation. Typically, commands are issued on a single line.

Deploying Web services applications

To deploy a programming artifact as a Simple Object Access Protocol (SOAP)-accessible Web service in WebSphere Application Server:

Steps for this task

1. Install the service-enabled EAR file.

Use the SOAP-enabled EAR file that you created and install it in WebSphere Application Server. See the Infocenter article Simple Object Access Protocol-enabling the services in an Enterprise Application Archive file for more information about the SOAP-enabled EAR file.

2. Update the Web server plug-in configuration.

Run the GenPluginCfg.bat file on Windows NT or the GenPluginCfg.sh script on UNIX platforms to regenerate the plug-in configuration.

3. Restart the application server.

Administering deployed Web services (XML-SOAP administrative tool)

Administer the deployed Web service with the XML-Simple Object Access Protocol (SOAP) administrative tool.

Use the SOAPEarEnabler tool to add administrative interfaces to your EAR files. Then, use the XML-SOAP administrative tool with these EAR files to do the following tasks for each context root:

- List configured services, showing active and stopped services
- Stop a service
- Start a service
- View the Apache SOAP deployment descriptor for a service

Access the XML-SOAP administrative tool through a Web browser by specifying:
`http://localhost/<contextroot>/admin/index.html`

where *context root* is the context specified when installing the SOAP-enabled .ear file.

What to do next

You cannot use the XML-SOAP administrative tool to add or remove a service. Use the SoapEarEnabler tool to add or remove services. A stopped service is persisted across starts and stops of the application server. Therefore, if you stop a service, it remains stopped until the next time you use the XML-SOAP administrative tool to start it again.

You can add the XML-SOAP administrative tool interface to an enterprise application when you SOAP-enable the EAR file. In interactive mode, you are asked whether you want to add the XML-SOAP administrative tool interface. Replying *yes* will add the necessary Java Server Pages (JSP) files and bindings allowing access to the XML-SOAP administrative tool interface for the application. The interface is an optional addition because you might not want to expose it in a production environment. Optionally, secure the XML-SOAP administrative tool using the Application Assembly Tool (AAT) to assign roles.

Securing Simple Object Access Protocol services

The SOAP security extension, included with WebSphere Application Server is a security architecture based on the SOAP security specification, and widely-accepted security technologies such as Secure Sockets Layer (SSL).

There are three options for security when using HTTP as the transport protocol:

- HTTP basic authentication
- SSL (HTTPS)
- SOAP signature

You can combine these security options according to your specific security requirements.

Because the Simple Object Access Protocol (SOAP) specification does not address all security issues, several proposals evolved to bridge the security gaps. See *Web services: Resources for Learning* for more information about additional documentation.

Migrating SOAP security

You can use the following steps when migrating a WebSphere Application Server, Version 4.0 application that uses SOAP 2.2 security, to a WebSphere Application Server, Version 5.0 application that uses SOAP 2.3 security.

Steps for this task

1. If you have any scripts or .bat files that set up the CLASSPATH before running your application, make sure that the `j2ee.jar` file, located in the Version 5.0 `<WebSphere/AppServer>/install root` directory is in the CLASSPATH.
2. (Optional) Replace security key files, if needed.

You might need new security key files if you get the following error:

```
"conf/cl-sig-config.xml"  
javax.security.cert.CertPathBuilderException:  
No end-entity certificate matching the selection criteria  
could be found
```

The three security key files are named `SOAPclient`, `SOAPserver`, and `sslserver.p12`. To locate the key files:

- a. If the application is already installed these files are found in the directory:
`<WebSphere/AppServer>/installedApps/<yourAppName>.ear/soapsec.war/key/`
 - b. If the application is not yet installed, these files are found by expanding:
`<yourAppName.ear>`. Then, expand the `soapsec.war`. The files are in the `key/` directory.
3. Stop and restart the server if you made any updates or replaced the key files.

Securing Simple Object Access Protocol services with HTTP basic authentication

Many applications require users to provide identifying information. You cannot provide access control for individual services. You can only provide access control for the router servlets, for example, the `rpcrouter` servlet URI. If you can get to a servlet, you can access any of the Web services served through the servlet. Therefore, if you have a set of secure services, you have to partition them differently so that they are accessed through a URI that is secured, for example, `/secureRPCRouter`. An example of a service that is not secure or accessible to everyone is `/unprotectedRCRouter`.

Using the (Application Assembly Tool (AAT)), you can set authorization levels by assigning roles to HTTP methods and by assigning users to roles. You can then *authenticate* users, verifying they are authorized to view specific information. There are many ways to prompt users for authentication data.

Securing Simple Object Access Protocol services on Secured Socket Layer

with HTTP basic authentication

To make a request over HTTPS using the Secured Socket Layer support of Apache Simple Object Access Protocol (SOAP), you need a separate Java Secure Socket Extension (JSSE) provider.

WebSphere Application Server includes the `ibmjsse.jar` in the IBM Developer Kit for Windows, Java Technology Edition.

The *SOAP on SSL* scenario is useful for many business-to-business (B2B) applications because:

- The data in transit is protected from eavesdropping or forgery by SSL.
- The client identity is authenticated through user ID and password, which are encrypted by the SSL transport.

For example, if an inventory application is configured as a Web service, the service provider has the following two SOAP service entries:

- `https://foo.com/inventory/inquiry`
- `https://foo.com/inventory/update`

Each SOAP service entry should be deployed as a separate enterprise application (EAR) because each service has a different access control policy, which is: anyone can inquire about the inventory but only the inventory clerks can update the contents.

The SOAP enablement model limits you to one context root for the unsecured services and another for the secured services. In this example, you want to make the inquiry service unsecured and the update service secured. If you want different levels of security for a secured service, then you must deploy the entries in the secured service as separate EAR files.

To enable the *SOAP on SSL* scenario:

Steps for this task

1. Configure the Web server (`httpd.conf`) so that it only allows SSL access to these servlets.
2. Configure the security role for the `RPCRouterServlet` in the inquiry services EAR file.

By doing this step, the `RPCRouterServlet` for the *inquiry* service is accessible by everyone, while the `RPCRouterServlet` for the *update* service requires authentication based on the HTTP basic authentication (userID and password).

In this case, the *update* application does not know the identity of the requester; it only knows that access is granted. In other words, the *update* application is not concerned with the identity of the user because it knows WebSphere Application Server is ensuring that only authenticated users have access.

Securing SOAP services on SSL with SOAP Signature

Applications might need non-repudiable proof of exchanged messages. One example is a Web service that accepts part orders. The business partners establish a form of trust relationship based on public keys. This can be done using the public key infrastructure (PKI) through a third party certificate authority (CA), or by exchanging public keys with a secure channel. The following service is deployed with a signature verification function:

`https://foo.com/partorder`

Configure signature verification with the following information:

- Scope of signature (indicates the portion of the SOAP envelope that must be authenticated. The default is the content of *SOAP-ENV:Body*).
- Trusted keys or trusted root keys.
- Default key to verify signature if no KeyInfo is specified.
- Other policies regarding signature validation.
- Behavior when signature verification fails.
- Additional requirements on signature (as for example, specific requirements on hash/C14N algorithms to be used, timestamp validity, and so forth).

If the signature is missing or if signature verification fails, the signature verification function can be configured so that the servlet returns a SOAP fault.

To send part orders to the `https://foo.com/partorder` service, the service requester should sign his SOAP messages with a signature component. The signature component is initialized using two templates:

1. `<ds:SignedInfo>` template
2. `<ds:KeyInfo>` template

The `<ds:SignedInfo>` template controls the following:

- What parts of the SOAP envelope must be signed
- What algorithms (canonicalization, transformation, digest, sign) should be used

The `<ds:KeyInfo>` template controls the following:

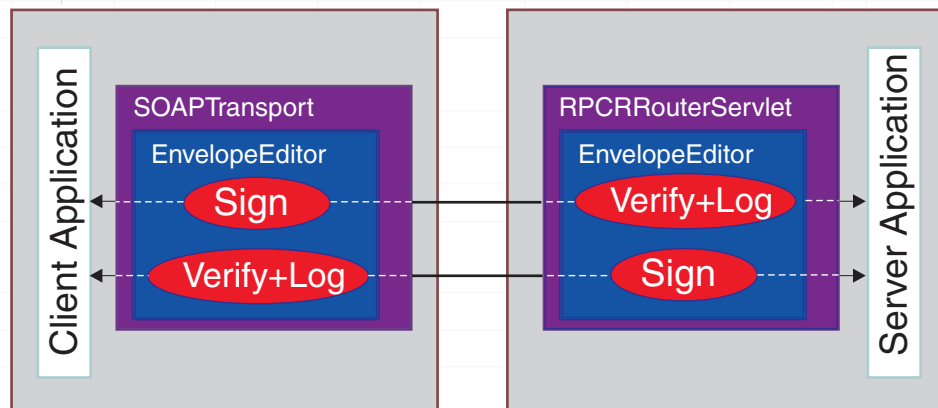
- Whether or not to include the entire certificate chain in `<ds:KeyInfo>`
- Decision to include only certificate and serial number
- Public key value
- Decision to provide no key information (so that the default key must be used for verification).

You can combine the service request with HTTP basic authentication, if necessary.

Simple Object Access Protocol signature architecture

An overview of the Simple Object Access Protocol (SOAP) signature architecture is illustrated in the figure below.

Transport Hook and Security Components



Using the SOAP transport hook, you can plug-in the following security components:

- Signer
- Verifier, with logging capability

The transport hook is called the *EnvelopeEditor*. A *PluggableEnvelopeEditor* is also provided, which allows you to plug in your security components. As illustrated, the *EnvelopeEditor* is encapsulated in the *SOAPTransport* on the client side. On the server side, *EnvelopeEditor* is encapsulated in *RPC/MessageRouterServlet*. The same components can be used on either side.

When a client application sends a request, the request is signed and transmitted to the server. At the server side, the request is verified and delivered to a server application or, in the case of a Remote Procedure Call (RPC), to a Java object. The response is processed in the same manner. The verifier component also has a logging function to log the verified messages in a file. Signatures and verifier components are configurable. You can specify encryption, digest message algorithm, certificate path policy, and other security technologies.

For more information about *EnvelopeEditor*, Signature header handler and Verification header handler, see [Web services: Resources for Learning](#).

UDDI4J specifications

The following considerations are specific to the support for UDDI4J provided by WebSphere Application Server, Version 5:

- **UDDI4J class libraries provided.**

WebSphere Application Server provides two UDDI4J class libraries:

uddi4jv2.jar

This class library contains classes which support Version 2 of the UDDI specification.

uddi4j.jar

This class library is provided for compatibility with WebSphere Application Server, Version 5, and supports Version 1 of the UDDI specification. The classes in this library are deprecated.

- **UDDI4J error handling**

When invoking UDDIProxy inquiry methods, UDDIException is thrown if errors are received from the UDDI proxy. UDDIException can contain a DispositionReport with information about the error. APIs that do not return a data object, provide the disposition report. SOAPException is thrown if a communication error occurs or if the resulting data cannot be parsed as a valid SOAP message.

See the article *Web services: Resources for Learning*, to learn more about the use of Simple Object Access Protocol (SOAP), Universal Description, Discovery and Integration (UDDI) Registry, UDDI4J, programming specifications and examples.

Web services: Resources for learning

Use the following links to find relevant supplemental information about getting started with Web services. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas. The following sections are covered in this reference:

View links to additional information about:

- Web services overview: Purpose, planning and designing to use Web services
- Developing and using Web services applications and clients
- Programming instructions and examples
- Administration
- Samples
- Other references
- ✓ **IBM Web Services architecture debuts**

<http://www.ibm.com/developerworks/webservices/library/w-int.html?dwzone=webservices>

Introducing IBM Web services, a distributed software architecture of service components. This brief overview and in-depth interview on IBM DeveloperWorks cover the fundamental concepts of Web services architecture

and what they mean for developers. The interview with IBM professional Rod Smith explores which types of developers Web services targets, how Web services reduces development time, what developers could be doing with Web services now, and takes a glance at the economics of dynamically discoverable services.

- ✓ **Web services (r)evolution, Part 1**

<http://www-106.ibm.com/developerworks/library/ws-peer1.html>

This article focuses on the benefits and challenges of building Web services applications. Web services might be an evolutionary step in designing distributed applications, however, they are not without their problems. Outlined are the difficulties developers face in creating a truly workable distributed system of Web services. This article also outlines author Grahm Glass' plan for building peer-to-peer Web applications.

- 📄 **SOAP**

<http://www.w3.org/TR/SOAP>

This article is a detailed overview of SOAP, which includes programming specifications.

- ✓ **Building a SOAP client**

<http://www-106.ibm.com/developerworks/library/x-soapcl/index>

This article describes a simple, general purpose SOAP client in Java that uses no specialized SOAP libraries. Instead of creating the SOAP request XML document for you under the hood, this client lets you create your own request with any XML or text editor. Instead of merely giving you the remote method return values, the client shows you the actual SOAP response XML document. The short Java program shows exactly what SOAP is all about: opening up an HTTP connection, sending the appropriate XML to invoke a remote method, and then reading the XML response returned by the server.

- 📄 **Web Services Description Language**

<http://www.w3.org/TR/wsdl>

This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications.

- 📄 **Universal Description, Discovery and Integration**

<http://www.uddi.org/about.html>

This article is a detailed overview of Universal Description, Discovery and Integration (UDDI).

- 📄 **UDDI4J: Matchmaking for Web services**

<http://www-106.ibm.com/developerworks/library/ws-uddi4j>

Reviewed in this article are the basics of UDDI, the Java API to UDDI, and how you can use this technology to start building, testing, and deploying your own Web services.

- 📄 **Web Services Description Language**

<http://www.w3.org/TR/wsdl>

This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications.

- 📄 **Universal Description, Discovery and Integration**

<http://www.uddi.org/about.html>

This article is a detailed overview of Universal Description, Discovery and Integration (UDDI).

- 📄 **UDDI4J: Matchmaking for Web services**

<http://www-106.ibm.com/developerworks/library/ws-uddi4j>

Reviewed in this article are the basics of UDDI, the Java API to UDDI, and how you can use this technology to start building, testing, and deploying your own Web services.

- ✓ **SOAP Security Extensions: Digital Signature**

<http://www.w3.org/TR/SOAP-dsig>

This document specifies the syntax and processing rules of a SOAP header entry to carry digital signature information within a SOAP 1.1 Envelope.

- 📄 **Simple Object Access Protocol (SOAP) 1.1**

<http://www.w3.org/TR/SOAP-dsig>

This site offers detailed information about SOAP Signature, including Envelope Editor, Signature Header Handler and Verification Header Handler.

- 📄 **Apache Software Foundation**

<http://www.apache.org>

- Two SOAP samples are available. The samples include information about implementing SOAP services and SOAP security. Refer to the 📄 **Samples Gallery**

<http://www7b.software.ibm.com/wsdd/library/samples/AppServer.html> for these samples.

- 📄 **Web services insider, Part 1: Reflections on SOAP**

<http://www-106.ibm.com/developerworks/webservices/library/ws-ref1>

What is the current state of the *Web services revolution*? Find out at this Web site that features the column *Web services insider, Part 1*. The author answers this question by reviewing the tools and technologies that have emerged over the past year, highlighting their differences and similarities.

- 📄 **The Web services insider, Part 2: A summary of the W3C Web Services Workshop**

<http://www-106.ibm.com/developerworks/webservices/library/ws-ref2>

This is a brief summary of a W3C Web services Workshop.

Chapter 8. Enabling Web services to use the Web Services Invocation Framework

Use this topic to learn how to invoke Web services to use the Web Services Invocation Framework

The Web Services Invocation Framework (WSIF) is a WSDL-oriented Java API that allows you to invoke Web services dynamically, regardless of what format (for example EJB) the service is implemented in, or what mechanism (for example JMS) is used to access it.

WSIF enables you, as a Web services developer, to move away from the usual Web services programming model of working directly with the SOAP APIs, towards a model where you interact with representations of the services. So you can work with the same programming model regardless of how the service is implemented and accessed.

If you want to know more about the issues that WSIF addresses, see [Goals of WSIF](#).

If you want to know how WSIF addresses these issues, see [An overview of WSIF](#).

To use WSIF, see the following topics:

- [Using WSIF to invoke Web services](#).
- [WSIF system management and administration](#).
- [WSIF API](#).

For more information about working with WSIF, visit the Internet sites listed in [WSIF: Resources for Learning](#).

Goals of WSIF

SOAP bindings for Web services are part of the WSDL specification. So when most developers think of using a Web service, they immediately think of assembling a SOAP message and sending it across the network to the service endpoint, using some SOAP client API. For example: with Apache SOAP the client creates and populates a Call object which encapsulates the service endpoint, the identification of the SOAP operation to be invoked, the parameters that have to be sent, and so on.

While this works for SOAP, it is limited in its use as a general model for invoking Web services for the following reasons:

- Web services are not just SOAP services.
- Tying client code to a particular protocol implementation is restricting.
- Incorporating new bindings into client code is hard.
- Multiple bindings can be used in flexible ways.
- A freer Web services environment enables intermediaries.

The goals of WSIF are therefore:

- To give a binding-independent mechanism for Web service invocation.

- To free client code from the complexities of any particular protocol used to access a Web service.
- To enable dynamic selection between multiple bindings to a Web service.
- To help the development of Web service intermediaries.

WSIF - Web services are not just SOAP services

You can deploy as a Web service any program with a WSDL description of its functional aspects and access protocols; and in the J2EE environment, the same component is available over multiple transports and protocols.

For example, you can have a database stored procedure, which is then exposed as a stateless session bean, and then deployed into a SOAP router to become a SOAP service. At each stage, the fundamental service is the same. All that changes is the access mechanism: from JDBC to RMI-IIOP and then to SOAP.

The WSDL specification defines a SOAP binding for Web services, but you can add binding extensions to the WSDL so that, for example, you can offer an enterprise bean as a Web service using RMI/IIOP as the access protocol. You can even treat a single Java class as a Web service, with in-thread Java method invocations as the access protocol. With this broader definition of a Web service, you need a binding-independent mechanism for service invocation.

WSIF - tying client code to a particular protocol implementation is restricting

If your client code is tightly bound to a client library for a particular protocol implementation, it can become hard to maintain. For example if you move from Apache SOAP to a different SOAP implementation, the process can take a lot of time and effort. To avoid these problems, you need a protocol implementation-independent mechanism for service invocation.

WSIF - incorporating new bindings into client code is hard

As is explained in Web services are not just SOAP services, if you want to make an application that uses a custom protocol work as a Web service, you can add extensibility elements to WSDL to define the new bindings. But in practice, achieving this is hard. For example you have to design the client APIs for using this protocol; and if your application uses just the abstract interface of the Web service, you have to write tools to generate the stubs that enable an abstraction layer. These are tasks that can take a lot of time and effort. What you need is a service invocation mechanism that allows bindings to be updated or new bindings to be plugged in easily.

WSIF - multiple bindings can be used in flexible ways

Imagine that you have successfully deployed an application that uses a Web service offering multiple bindings. For example, imagine that you have a SOAP binding for the service and a local Java binding that lets you to treat the local service implementation (a Java class) as a Web service.

The local Java binding for the service can only be used if the client is deployed in the same environment as the service itself, and if this is the case it is far more efficient to communicate with the service by making direct Java calls than using the SOAP binding.

If your clients could switch the actual binding used based on run-time information, they could choose the most efficient available binding for each situation. In order to take advantage of Web services that offer multiple bindings, you need a service invocation mechanism that allows you to switch between the available service bindings at runtime, without having to generate or recompile a stub.

WSIF - a freer Web services environment enables intermediaries

Web services offer application integrators a loosely-coupled paradigm. In such environments, intermediaries can be very powerful. Intermediaries can add value to the service invocation without specific programming. Facilities such as logging, high-availability and transformation can be provided by an intermediary. WSIF is designed to make building intermediaries both possible and simple.

An overview of WSIF

WSIF provides a Java API for invoking Web services, independent of the format of the service or the transport protocol through which it is invoked. It addresses all of the issues identified in the goals of WSIF.

WSIF provides the following features:

- It has an API that provides binding-independent access to any Web service.
- It is closely based on WSDL, so it can invoke any service that can be described in WSDL.
- It allows stubless (completely dynamic) invocation of a Web service.
- You can plug a new or updated implementation of a binding into WSIF at runtime.
- You can defer the choice of a binding until runtime.

WSIF is designed to work both in an unmanaged environment (standalone) and inside a managed container. You can use JNDI to find the WSIF service, or else read in the WSDL definition.

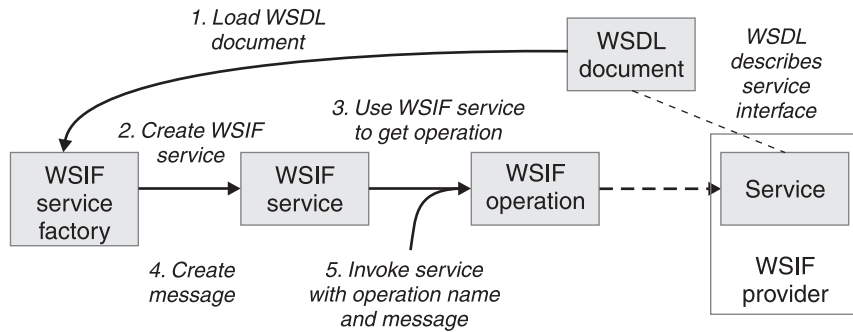
For more conceptual information about WSIF and WSDL, see the following topics:

- WSIF and WSDL
- WSIF architecture
- Using WSIF with Web services that offer multiple bindings
- WSIF usage scenarios
- Dynamic invocation

WSIF architecture

The WSIF architecture is shown in the following figure. The components of this architecture are described after the figure.

WSIF architecture. The WSIF architecture, showing a Web service being invoked by loading a WSDL document, creating a WSIF service, using the service to get a WSIF operation, then invoking the target Web service by providing the WSIF operation with the target service operation's name and the message that it needs.



WSIF provider

A WSIF provider is an implementation of a WSDL binding that can run a WSDL operation through a binding-specific protocol. WebSphere Application Server includes WSIF providers for SOAP over HTTP, SOAP over JMS, Java, enterprise beans, and Native JMS. For more information, see *Using the WSIF providers*.

WSIFOperation

The runtime representation of an operation, called *WSIFOperation* is responsible for invoking a service based on a particular binding. For more information, see *WSIF API reference: Using ports*.

WSIFService

The *WSIFService* is responsible for generating an instance of *WSIFOperation* to be used for a particular invocation of a service operation. For more information, see *Finding a port factory or service*.

WSDL documents

The Web service WSDL document contains the location of the Web service. The binding document defines the protocol and format for operations and messages defined by a particular portType.

Using WSIF with Web services that offer multiple bindings

You can use WSIF to enable client applications to switch between service bindings at runtime, to enable them to use the optimum binding to invoke operations on a Web service provider.

For example, a Web service provider could offer a SOAP binding for the service and a local Java binding that allows you to treat the local service implementation (a Java class) as a Web service. If the client is deployed in the same environment as the service, the local Java binding for the service can be used and provides more efficient communication with the service by making direct Java calls rather than using the SOAP binding.

WSIF and WSDL

WSDL is the acronym for <http://www.w3.org/TR/wsdl>.

In **WSDL** a service is defined in three distinct parts:

- The **PortType**. The PortType defines the abstract interface offered by the service. A PortType defines a set of *Operations*. Each operation can be In-Out (request-response), In-Only, Out-Only and Out-In (Solicit-Response). Each operation defines the input and/or output *Messages*. A message is defined as a set of *Parts* and each part has a schema-defined type.

- The **Binding**. A binding defines how to map between the abstract PortType and a real service format and protocol. For example the SOAP binding defines the encoding style, the SOAPAction header, the namespace of the body (the targetURI), and so on.
- The **Port**. This defines the actual location (endpoint) of the available service - for example, the HTTP URL on which a SOAP service is available.

Currently in WSDL, each Port has one and only one binding, and each binding has a single PortType. But (more importantly) each Service (PortType) can have multiple Ports, each of which represents an alternative location and binding for accessing that service.

WSIF follows the semantics of WSDL as much as possible:

- The WSIF dynamic invocation API directly exposes runtime equivalents of the model from WSDL. For example, invocation of an operation involves executing an "Operation" with an Input Message.
- WSDL has extension points that allow new ports and bindings to be added so that WSDL can describe new systems. The equivalent concept in WSIF is a provider, that allows WSIF to understand a class of extensions and therefore support new service implementation types.

As a metadata-based invocation framework, WSIF follows the design of the metadata. As WSDL is extended, WSIF is updated to follow.

Note: The **implicit and primary** type system of WSIF is XML Schema, not Java. WSIF supports invocation using dynamic proxies, which support Java type systems, but when you use the WSIFMessage interface it is your responsibility to populate WSIFMessage objects with data based on the XML Schema types as defined in the WSDL document. So you should define types of objects by a canonical and fixed mapping from Schema types into the runtime.

For more information on WSDL, see (WSIF: Resources for learning).

WSIF usage scenarios

This topic describes two brief scenarios that illustrate the role WSIF plays in the emerging Web services environment.

Scenario: Redevelopment and redeployment

If you are implementing Web services today you are probably working with simple prototypes. As your Web services move into production, you will need to reimplement and redeploy them. WSIF uses the same API calls with different underlying technologies. So if you use WSIF you can reimplement and redeploy your services without changing the client code, and you can use existing highly reliable and high-performance infrastructures like RMI-IIOP and JMS without sacrificing the location-independence that the Web service model offers.

Scenario: Service Flow composition

A service flow typically invokes a Web service, then passes the response from one Web service into the next Web service, perhaps performing some transformation in the middle.

There are two key aspects to this that WSIF provides:

- A representation of the service invocation based on the metadata in WSDL.
- The ability to build invocations based on the portType only, which can be used on any implementation.

For example, imagine that you build a "meta-service" that uses a number of services to build a process. Initially several of those services are simple JavaBean prototypes that are written and exposed through SOAP, but you plan to reimplement some of them as EJB components, and to out-source others.

If you use SOAP, it ties up multiple threads for every onward invocation, as they pass through the webserver and servlet engine into the SOAP router. If you use WSIF to call the beans directly, you get much better performance compared to SOAP and you don't lose access or location transparency. Using WSIF, you can move the JavaBean implementations to EJB implementations without changing the client code, and to move some of the Web services from local implementations to external SOAP services you just update the WSDL.

Dynamic invocation

In WSIF, dynamic invocation means providing the following levels of support when invoking Web services:

1. Support for WSDL extensions and bindings that were not known at build time.
2. Support for Web services that were not known at build time.

WSIF supports (1) above through the use of providers.

The providers support (2) above by using the WSDL description to access the target service.

Using WSIF to invoke Web services

You invoke a Web service dynamically by using the WSIF API directly. You only specify the location of the WSDL file for the service, the name of the operation to be invoked, and any operation arguments needed. All the information needed to access the Web service is available through WSDL; the abstract interface, the binding, and the service endpoint.

This kind of invocation does not generate stub classes and does not need a separate compilation cycle.

More information on using WSIF to invoke Web services is given in the following topics:

- Using the WSIF providers.
- Developing a WSIF service.
- Using complex types.
- Using JNDI.
- Interacting with the WebSphere J2EE container.
- Running WSIF as a client.

Using the WSIF providers

A WSIF provider is an implementation of a WSDL binding that can run a WSDL operation through a binding-specific protocol.

Providers implement the interface between the WSIF API and the actual implementation of a service. Providers are pluggable within the WSIF framework, and are registered based upon the namespace of the WSDL extension that they implement.

WebSphere Application Server includes the following WSIF providers:

- SOAP (over HTTP) provider.
- JMS providers (SOAP over JMS, and Native JMS).
- Java provider.
- EJB provider.

Note:

- Some providers use the J2EE programming model to utilize J2EE services.
- If a provider is available, but its required class libraries are not, the provider is disabled.

Using the SOAP provider

The SOAP provider allows WSIF stubs and dynamic clients to invoke SOAP services. The provider supports SOAP 1.1 over HTTP. The WSIF SOAP Provider utilizes ApacheSOAP 2.3 for parsing and creating SOAP messages, but is not limited to invoking services served by ApacheSOAP.

The WSIF SOAP provider supports:

- SOAP-ENC encoding
- RPC style

The SOAP provider is not transactional.

Note: Before you deploy to WebSphere Application Server a Web service that you expect to be used by multiple clients connecting over SOAP, you must set up your application's deployment descriptor file (dds.xml) to handle multiple connections correctly. For more information see WSIF troubleshooting tips.

What to do next

For an example of the sort of code changes that need to be made in the WSDL file for a SOAP provider, see The SOAP over JMS provider - writing the WSDL extension.

Using the JMS providers

JMS is an API for transport technology. The mapping to a JMS destination is defined during deployment and maintained by the container.

The JMS Destination Web service endpoint may be realized in any of the following ways:

- The JMS destination for the queue can be the Web service implementation.
- The JMS destination can be (but is not required to be) associated with a Message Driven Bean by the EJB container, thereby allowing the Message Driven Bean to be the Web service implementation.
- (For SOAP over JMS) The JMS Destination can unwrap the JMS message and route the SOAP message to a stateless session bean Web service implementation.

The JMS Destination must respect the interaction model expected by the client and defined by the WSDL. It must return a response if one is required.

When the JMS Destination creates the JMS response message the following rules must be adhered to:

- The response message must be sent to the **JMSReplyTo** from the incoming request.
- The **JMSCorrelationID** of the response message must be set to the value of the **JMSMessageID** from the request message.
- The response must be sent with a **deliveryMode** equal to the **JMSDeliveryMode** of the request message.
- The response must be sent with a **priority** equal to the **JMSPriority** of the request message.
- The **timetolive/JMSExpiration** must be set a value equal to the **JMSExpiration** of the request message.

The client does not see any of these headers. The container receives the JMS message and (for SOAP over JMS) removes the SOAP message to send to the client.

What to do next

See also the following topics:

- Using the SOAP over JMS provider
- Using the native JMS provider
- The JMS providers - configuring the client and server

Using the SOAP over JMS provider: [Before you begin](#)

For information on working with the JMS API, see [Using the JMS providers](#).

The SOAP message, including the SOAP envelope, is wrapped with a JMS message and put on the appropriate Queue. The container receives the JMS message and removes the SOAP message to send to the client.

What to do next

For detailed implementation information, see the following topics:

- The SOAP over JMS provider - writing the WSDL extension
- The JMS providers - configuring the client and server

The SOAP over JMS provider - writing the WSDL extension: SOAP protocol messages are carried on the JMS transport with the JMS body type **TextMessage** if the message is strictly XML.

Usage scenario

WSDL Extensions

WSDL Binding

The WSDL Binding for SOAP over JMS varies only slightly from the SOAP over HTTP binding. The transport element under the soap:binding indicates that JMS is being used.

RPC style: If the style is set to "rpc", then it is assumed that an operation is being invoked on the Web service endpoint. The Java parameters and response holders are encoded in the same way as for the standard WSIF SOAP binding:

```
<soap:binding style="rpc"
              transport="http://schemas.xmlsoap.org/soap/jms" />
```

When a SOAP/JMS binding is being used, the <wsdl:port> must contain a <jms:address> element to identify the JMS Queue to be used.

Address:

The address within the port within the service within the WSDL provides the information required for a client to correctly connect to the Web service using the JMS programming model. Typically, it will be the stubs generated to support the SOAP/JMS binding that will act as the JMS client. This does not preclude the Web service client from using the JMS programming model directly. The address element under the service and port must take this form:

```
<jms:address
  destinationStyle= "queue"
  jmsVendorURI = "http://ibm.com/ns/mqseries"?
  initialContextFactory= "com.ibm.NamingFactory"?
  jndiProviderURL= "iiop://something:900/wherever"?
  jndiConnectionFactoryName= "orange"
  jndiDestinationName= "fred"
/>
```

where "?" means "optional".

The vendor URI is a string which uniquely identifies the JMS implementation. WSIF ignores this URI, which is used by the client developer and perhaps the client implementation to determine if it has access to the correct JMS provider in the client runtime.

The connectionFactory attribute gives the name of a JMS ConnectionFactory object, which can be looked up within the JNDI context given by the jndiContext attribute. This ConnectionFactory can be used to create a JMS connection to the JMS provider instance that owns the Queue. In a simple configuration this ConnectionFactory is the actual ConnectionFactory used by the server message listener, and by the clients. However both server and clients can use different ConnectionFactories, provided that they all create Connections to the same JMS provider instance.

The JNDI usage pattern is this:

```
<jms:address destinationStyle= "queue"
  jndiConnectionFactoryName= "orange"
  jndiDestinationName= "fred"
>
```

This usage pattern requires the runtime to have a default JNDI provider configured. The provider URL and context factory can also be added to this:

```

<jms:address destinationStyle="queue"
  initialContextFactory = "com.ibm.Naming"
  jndiProviderURL= "iiop://server:900/"
  jndiConnectionFactoryName= "orange"
  jndiDestinationName= "fred"
>

```

Headers and Properties

JMS headers and properties can be set using the <jms:property> extension within the binding. This maps a part of a message into a JMS property. For example:

```

<jms:property name="Priority"
  {part="requestPriority" | value="fixedValue"} />

```

The header can also be set in the jms:address component with a literal value:

```

<jms:property name="Priority" value="fixedValue" />

```

This binding extension is shared with the Native JMS binding.

WSDL Binding example

Here is an example of a WSDL that defines a SOAP over JMS binding:

```

<!-- Example: SOAP over JMS Text Message -->

<?xml version="1.0" encoding="UTF-8"?>
<wsdl:definitions
  name="StockQuoteInterfaceDefinitions"
  targetNamespace="urn:StockQuoteInterface"
  xmlns:tns="urn:StockQuoteInterface"
  xmlns:xsd="http://www.w3.org/2000/10/XMLSchema"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
  <wsdl:message name="GetQuoteInput">
    <part name="symbol" type="xsd:string" />
  </wsdl:message>
  <wsdl:message name="GetQuoteOutput">
    <part name="value" type="xsd:float" />
  </wsdl:message>

  <wsdl:portType name="StockQuoteInterface">
    <wsdl:operation name="GetQuote">
      <wsdl:input message="tns:GetQuoteInput" />
      <wsdl:output message="tns:GetQuoteOutput" />
    </wsdl:operation>
  </wsdl:portType>

  <wsdl:binding name="StockQuoteSoapJMSBinding" type="tns:StockQuoteInterface">
    <soap:binding style="rpc"
      transport="http://schemas.xmlsoap.org/soap/jms" />
    <wsdl:operation name="GetQuote">
      <soap:operation soapAction="urn:StockQuoteInterface#GetQuote" />
      <wsdl:input>
        <soap:body use="encoded" namespace="urn:StockQuoteService"
          encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
      </wsdl:input>
      <wsdl:output>
        <soap:body use="encoded" namespace="urn:StockQuoteService"
          encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
      </wsdl:output>
    </wsdl:operation>
  </wsdl:binding>
</wsdl:definitions>
<wsdl:service name="StockQuoteService">

```

```

        <wsdl:port name="StockQuoteServicePort"
            binding="sqi:StockQuoteSoapJMSBinding">
            <jms:address destinationStyle="queue"
                jndiConnectionFactoryName="myQCF"
                jndiDestinationName="myQ"
                initialContextFactory= "com.ibm.NamingFactory"
                jndiProviderURL= "iiop://something:9000/" />
        </wsdl:port>
    </wsdl:service>
</wsdl:definitions>

```

Using the native JMS provider: Before you begin

For information on working with the JMS API, see [Using the JMS providers](#).

The native JMS provider allows WSIF clients to treat a JMS destination as a Web service.

What to do next

For detailed implementation information, see the following topics:

- The native JMS provider - writing the WSDL extension
- The JMS providers - configuring the client and server

The native JMS provider - writing the WSDL extension: The WSDL extensions for JMS are identified with the namespace prefix "jms". For example <jms:binding>.

Operations

The supported operations are either One-way operations (send for JMS point-to-point messaging or publish for JMS publish/subscribe messaging) or Request-Response operations (send/receive for JMS point-to-point messaging).

The WSDL operations therefore specify an input message only, or a WSDL input and output message.

Message Header Data

JMS does not make assumptions about message headers. If the JMS provider is MQSeries, then each JMS message carries an RFH2 header, however, data in this message header is only accessed indirectly by getting and setting JMS message properties.

Quality of Service Attributes

The Native JMS binding of WSDL operations provides the capability to set message properties. See "JMS Message Header Fields and Properties" below.

JMS Message Types

JMS message body types are specified in the WSDL binding using the jms:binding WSDL extension:

```

<wsdl:binding ... >
    <jms:binding type="messageBodyType" />
    ...
</wsdl:binding>

```

where `messageBodyType` is either `ObjectMessage` or `TextMessage`.

Input and Output Messages

Input and output messages are used in the WSDL input and output sections of the WSDL binding, respectively:

```
<wsdl:input ... >
  <jms:input parts="part1 part2 ..." />
</wsdl:input>

<wsdl:output ... >
  <jms:output parts="part1 part2 ..." />
</wsdl:output>
```

For JMS Text messages and JMS Object messages created from one or more WSDL message parts, the `jms:input` and `jms:output` WSDL extension specifies the message parts to be used for the JMS message. If no parts are defined, then all message parts will be used. This is used in the WSDL input and output sections of the WSDL binding.

The WSDL message has one part that contains the complete message body. This message body might be the result of a mapping from some other representation (see "Data Mapping" below).

```
<wsdl:input ... >
  <jms:input parts="part1" />
</wsdl:input>
```

Fault Messages

Operations that describe message interfaces with a native JMS binding do not have fault messages. No assumptions are made about the message schema or the semantics of message properties, so no distinction can be made between output and fault messages.

Data Mapping

Mapping of data types is specified using the WSDL format binding extension:

```
<wsdl:binding ... >
  <jms:binding type="..." />

  <format:typeMapping encoding="Java" style="Java">
    <format:typeMap typeName="..." formatType="targetType"/>
  </format:typemapping>
  ...
</wsdl:binding>
```

The value of *targetType* is dependent on the JMS message type (discussed above) as follows:

- For JMS Object messages, the target data type implements `java.io.Serializable`.
- For JMS Text messages, the target data type is always `java.lang.String`.

Note that the format binding is also used in other bindings that deal with Java interfaces.

JMS Message Header Fields and Properties

You can set JMS Properties in the WSDL. When you want your application to pass a property into WSIF as a part on the WSIF message, use a `<jms:property>`. When you want to hard code an actual property value into the WSDL, use a `<jms:propertyValue>`. You can put `<jms:property>` and `<jms:propertyValue>` in the `<input>` and `<output>` in the binding operation, and in the `<jms:address>`. `<jms:property>` takes precedence over `<jms:propertyValue>` and properties in the binding operation take precedence over properties in the `<jms:address>`.

For the association with WSDL parts, the `<jms:property>` and `<jms:propertyValue>` WSDL extensions are used as shown below:

```
<wsdl:input ... >
    <jms:property name="propertyName" part="partName" />
    <jms:propertyValue name="propertyName"
        type="xsdType" value="actualValue" />
</wsdl:input>
<wsdl:output ... >
    <jms:property name="propertyName" part="partName" />
</wsdl:output>
```

where *propertyName* identifies the JMS property that is associated with the header field, and *partName* identifies the message part that is associated with the property.

The *propertyName* is the name of one of the predefined JMS message header fields or a user-defined property. Possible values and Java types for the predefined message header fields are shown in the following table:

Value	Java type
JMSMessageId	java.lang.String
JMSTimeStamp	long
JMSCorrelationId	byte [] or java.lang.String
JMSReplyTo	javax.jms.Destination
JMSDestination	javax.jms.Destination
JMSDeliveryMode	int
JMSRedelivered	boolean
JMSType	java.lang.String
JMSExpiration	long

See the JMS specification for restrictions that apply for setting JMS header field values. Attempts to set such restricted values are ignored.

For application-defined JMS message properties, the Java types used in the Native JMS binding implementation (used for calls to the corresponding JMS methods) is derived from the XML schema type in the abstract interface (`wsdl:part`) and the type mapping information in the format binding (`format:typemap`).

The `jms:propertyValue` contains a specification of a literal value and its associated XML schema type.

For `wsdl:output`, only `jms:property` can be specified.

Transactions

Independent of other quality of service attributes, the asynchronous processing of request-response operations has implications for callers running in a transaction scope. The send request part and the receive response part must be separated into **two transactions** as the send needs to be committed in order for the request message to become visible.

So implementations that process WSDL for asynchronous request-response operations (such as WSIF) must take the following additional actions:

- They must return a correlation id to the user, and provide a **callback** that allows users to pass in the response message in order to process the "second half" of the operation.
- (Optionally) They might implement their own response message "listener" in order to recognize the arrival of response messages, and to manage the correlation to the request message.

Usage scenario

WSDL Examples

Example 1: JMS Text Message (Request-Response)

The JMS text message contains a **java.lang.String**. In this example, the WSDL message contains only one part that represents the whole message body.

```
<!-- Example 1: JMS Text Message -->
<wsdl:definitions ... >

  <!-- simple or complex types for input and output message -->
  <wsdl:types> ... </wsdl:types>

  <wsdl:message name="JmsOperationRequest"> ... </wsdl:message>
  <wsdl:message name="JmsOperationResponse"> ... </wsdl:message>

  <wsdl:portType name="JmsPortType">
    <wsdl:operation name="JmsOperation">
      <wsdl:input name="Request"
        message="tns:JmsOperationRequest"/>
      <wsdl:output name="Response"
        message="tns:JmsOperationResponse"/>
    </wsdl:operation>
  </wsdl:portType>

  <wsdl:binding name="JmsBinding" type="JmsPortType">
    <jms:binding type="TextMessage" />

    <format:typemapping style="Java" encoding="Java">
      <format:typemap name="xsd:String" formatType="String" />
    </format:typemapping>

    <wsdl:operation name="JmsOperation">
      <wsdl:input message="JmsOperationRequest">
        <jms:input parts="requestMessageBody" />
      </wsdl:input>
      <wsdl:output message="JmsOperationResponse">
        <jms:output parts="responseMessageBody" />
      </wsdl:output>
    </wsdl:operation>
  </wsdl:binding>

  <wsdl:service name="JmsService">
    <wsdl:port name="JmsPort" binding="JmsBinding">
      <jms:address destinationStyle="queue"

```

```

        jndiConnectionFactoryName="myQCF"
        jndiDestinationName="myDestination"/>
    </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

Example 2: JMS Message with Accessing Application-Defined JMS Properties (Request-Response)

As an extension of the previous JMS message example, the following WSDL describes a request-response operation where specific JMS property values of the request and response message are set for the request message and retrieved from the response message.

The JMS properties in the request message are set according to the values in the input message. Likewise, selected JMS properties of the response message are copied to the corresponding values of the output message. The direction of the mapping is determined by the appearance of the **jms:property** tag in the input or output section, respectively.

```

<!-- Example 2: JMS Message with JMS Properties -->
<wsdl:definitions ... >

    <!-- simple or complex types for input and output message -->
    <wsdl:types> ... </wsdl:types>

    <wsdl:message name="JmsOperationRequest">
        <wsdl:part name="myInt" type="xsd:int"/>
        ...
    </wsdl:message>

    <wsdl:message name="JmsOperationResponse">
        <wsdl:part name="myString" type="xsd:String"/>
        ...
    </wsdl:message>

    <wsdl:portType name="JmsPortType">
        <wsdl:operation name="JmsOperation">
            <wsdl:input name="Request"
                message="tns:JmsOperationRequest"/>
            <wsdl:output name="Response"
                message="tns:JmsOperationResponse"/>
        </wsdl:operation>
    </wsdl:portType>

    <wsdl:binding name="JmsBinding" type="JmsPortType">
        <!-- the JMS message type may be any of the above -->
        <jms:binding type="..." />

        <format:typemapping style="Java" encoding="Java">
            <format:typemap name="xsd:int" formatType="int" />
            ...
        </format:typemapping>

    <wsdl:operation name="JmsOperation">
        <wsdl:input message="JmsOperationRequest">
            <jms:property message="tns:JmsOperationRequest" parts="myInt" />
            <jms:propertyValue name="myLiteralString"
                type="xsd:string" value="Hello World" />
            ...
        </wsdl:input>
        <wsdl:output message="JmsOperationResponse">
            <jms:property message="tns:JmsOperationResponse" parts="myString" />

```



```

        ...
        </wsdl:output>
    </wsdl:operation>
</wsdl:binding>

<wsdl:service name="JmsService">
    <wsdl:port name="JmsPort" binding="JmsBinding">
        <jms:address destinationStyle="queue"
            jndiConnectionFactoryName="myQCF"
            jndiDestinationName="myDestination"/>
    </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

The JMS providers - configuring the client and server: To enable a service to be invoked through JMS by a WSIF client application, complete the following steps:

1. Create a JMS queue for WSIF to send to the request message to.

Note: This can be done using MQSeries Explorer, or an equivalent tool from your JMS implementation package.
2. Create a JNDI definition for that JMS queue, and a JNDI queue connection factory.

Note: If you are using MQSeries, this can be done using the jmsadmin tool.
3. Put the JNDI names of the queue and queue connection factory, as well as your JNDI configuration, in the WSDL file.

You should also be aware of the following specific ways in which WSIF interacts with JMS:

- Only input JMS properties are supported.
- WSIF needs two queues when invoking an operation: one for the request message and one for the reply. The replyTo queue is by default a temporary queue which WSIF creates on behalf of the application. You can specify a permanent queue by setting the JMSReplyTo property to the JNDI name of a queue.
- WSIF uses the default values for properties set by the JMS implementation. However in MQSeries and some other JMS implementations, messages are by default persistent, and the default temporary queue is temporary dynamic and so cannot have persistent messages written to it. So your JMS listener may fail to write a persistent response message to the temporary replyTo queue.

Note: If you are using MQSeries, you need to create a temporary model queue which is permanent dynamic, then pass this model as the tempmodel of your queue connection factory. This will ensure that persistent messages can be written to the permanent dynamic temporary replyTo queue.

Using the Java provider

The WSIF Java Provider allows WSIF to invoke Java classes and JavaBeans. This means that, in a "thin-client" environment such as a JVM or Tomcat test runtime, you can define "shortcuts" to local Java code.

The WSIF Java Provider is not intended to be used in a J2EE environment. There is a difference between a client using the WSIF Java Provider to invoke a Java component, and implementing a Web service as a Java component on the server side.

The Java binding exploits the format binding for type mapping. The format binding allows WSDL to define the mapping between XML Schema types and Java types.

The Java provider requires the targeted Java classes to be in the class path of the client. The Java method is invoked synchronously, in-process, in-thread, with the current thread and ORB contexts.

The Java provider is not transactional.

What to do next

For examples of the code changes that need to be made in the WSDL file, see The Java provider - writing the WSDL extension.

The Java provider - writing the WSDL extension: The Java provider allows the invocation of a method on a local Java object. In order to use the Java provider you require the following binding specified in the WSDL:

Usage scenario

```
<!-- Java binding -->
<binding .... >
  <java:binding />
  <format:typeMapping style="Java" encoding="Java"/>?
    <format:typeMap name="qname" formatType="nmtoken"/>*
  </format:typeMapping>
  <operation>*
    <java:operation
      methodName="nmtoken"
      parameterOrder="nmtoken"
      returnPart="nmtoken"?
      methodType="instance|constructor" />
    <input name="nmtoken"? />?
    <output name="nmtoken"? />?
    <fault name="nmtoken"? />?
  </operation>
</binding>
```

where "?" means "optional" and "*" means "0 or more".

Note:

- The format:typeMap **name** attribute is a qualified name of a simple or complex type used by one of the Java operations.
- The format:typeMap **formatType** attribute is the fully qualified Class name for the Java Class that the element specified by **name** maps to.
- The java:operation **methodName** attribute is the name of the method on the Java object that is called by the operation.
- The java:operation **parameterOrder** attribute contains a whitespace-separated list of part names that define the order in which they are passed to the Java Object's method.
- The java:operation **methodType** attribute must be set to either "instance" or "constructor". The value specifies whether the method being invoked on the object is an instance method or a constructor for the object.

```

<service ... >
  <port>*
    <java:address
      className="nmtoken"/>
    </port>
</service>

```

Note: The java:address **className** attribute specifies the fully qualified class name of the object containing the method to invoke.

Using the EJB provider

The EJB Provider allows WSIF clients to invoke enterprise beans. The EJB client JAR must be available in the client runtime with the current provider. The enterprise bean is invoked using normal EJB invocation methods, using RMI-IIOP, with the current security and transaction contexts. If the EJB provider is invoked within a transaction then the transaction is passed to the onward service and the standard EJB transaction attribute applies.

If there are multiple implementations of the service, it is up to the provider of the service to make sure that they offer the same semantics. For example, in the case of transactionality the bean deployer should specify TX_REQUIRES_NEW to force a new transaction.

What to do next

For examples of the sort of code changes that need to be made in the WSDL file, see The EJB provider - writing the WSDL.

The EJB provider - writing the WSDL extension: The EJB provider allows the invocation of an enterprise bean through RMI/IIOP. In order to use the EJB provider you require the following binding specified in the WSDL:

Usage scenario

```

<!-- EJB binding -->
<binding .... >
  <ejb:binding />
  <format:typeMapping style="Java" encoding="Java"/>?
    <format:typeMap name="qname" formatType="nmtoken"/>*
  </format:typeMapping>
  <operation>*
    <ejb:operation
      methodName="nmtoken"
      parameterOrder="nmtoken"
      returnPart="nmtoken"?
      interface="remote|home" />
    <input name="nmtoken"? />?
    <output name="nmtoken"? />?
    <fault name="nmtoken"? />?
  </operation>
</binding>

```

where "?" means "optional" and "*" means "0 or more".

Note:

- The format:typeMap **name** attribute is a qualified name of a simpleType or complexType used by one of the EJB operations.
- The format:typeMap **formatType** attribute is the fully qualified Class name for the Java Class that the element specified by **name** maps to.

- The `ejb:operation` **methodName** attribute is the name of the method on the enterprise bean that will be called by the operation.
- The `ejb:operation` **parameterOrder** attribute contains a whitespace-separated list of part names which define the order in which they are passed to the EJB method.
- The `ejb:operation` **interface** attribute must be set to either "remote" or "home". The value specifies the interface of the enterprise bean on which the method named by the **method** attribute is accessible.


```
<service ... >
  <port>*
    <ejb:address
      class="nmtoken"
      jndiName="nmtoken"
      initialContextFactory="nmtoken" ?
      jndiProviderURL="nmtoken" ? />
    </port>
  </service>
```

Note:

- The `ejb:address` **class** attribute specifies the fully qualified class name home interface class of the enterprise bean.
- The `ejb:address` **jndiName** attribute specifies the full JNDI name which is used to look up the enterprise bean.
- The `ejb:address` **initialContextFactory** attribute is optional and specifies the initial context factory class.
- The `ejb:address` **jndiProviderURL** attribute is optional and specifies the jndi provider URL

Developing a WSIF service

A WSIF service is a Web service that uses WSIF. To develop a WSIF service, you first develop the Web service (or use an existing Web service), then develop the WSIF client based on the WSDL document for that Web service.

There are also two pre-built WSIF samples available for download from the  <http://www.ibm.com/websphere/developer/library/samples/AppServer.html> page of the IBM WebSphere Developer Domain Web site. These are as follows:

- The Address Book sample.
- The Stock Quote sample.

For more information on using the pre-built samples, see the documentation that is included in the download package.

To develop a WSIF service, complete the following steps:

Steps for this task

1. Develop the Web service.

Use Web services tools to discover, create, and publish the Web service. You can develop Java bean, enterprise bean, and URL Web services. You can use Web service tools to create a skeleton Java bean and a sample application from a WSDL document. For example, an enterprise bean can be offered as a Web service, using RMI/IIOP as the access protocol. Or you can use a Java class as a Web service, with native Java invocations as the access protocol.

You can use the WebSphere Studio Application Developer to create a Web service from a Java bean, as described in its StockQuote service tutorial. The bean that you use in this scenario returns the last trading price from the Internet Web site www.xmltoday.com given a stock symbol. Using the Web Service wizard, you generate a binding WSDL document named `StockQuoteService-binding.wsdl` and a service WSDL document named `StockQuoteService-service.wsdl` from the bean `StockQuoteService.java`. You then deploy the Web service to a Web server, generate a client proxy to the Web service, and generate a sample application that accesses the `StockQuoteService` through the client proxy. You test the StockQuote Web service, publish it using the IBM UDDI Explorer, and then discover the StockQuote Web service in the IBM UDDI Test Registry.

2. Develop the WSIF client.

The information you need to develop a WSIF client is given in the following topics:

- Developing the WSIF client - the Address Book sample gives example code to show how you define a Web service in WSDL.
- Using the WSIF providers describes the available providers, and gives example code of how their WSDL extensions are coded.
- WSIF API defines the main interfaces that your client uses to support the invocation of Web services defined in WSDL.

Note: The Address Book sample is written for synchronous interaction. If you are using a JMS provider, your WSIF client might need to act asynchronously. WSIF provides two main features that meet this requirement:

- A **correlation service** that assigns identifiers to messages so that the request can be matched up with the (eventual) response.
- A **response handler** that picks up the response from the Web service at a later time.

For more information see the WSIF API topic `WSIFOperation - Asynchronous interactions` reference.

Developing the WSIF client - the Address Book sample

The following code fragments show you how to use the WSIF API to invoke the AddressBook sample Web service dynamically.

Usage scenario

This is example code for dynamic invocation of the AddressBook sample Web service using WSIF:

```
try {
    String wsdlLocation="clients/addressbook/AddressBookSample.wsdl";

    // The starting point for any dynamic invocation using wsif is a
    // WSIFServiceFactory. We create ourselves one via the newInstance
    // method.
    WSIFServiceFactory factory = WSIFServiceFactory.newInstance();

    // Once we have a factory, we can use it to create a WSIFService object
    // corresponding to the AddressBookService service in the wsdl file.
    // Note: since we only have one service defined in the wsdl file, we
    // do not need to use the namespace and name of the service and can pass
    // null instead. This also applies to the port type, although values have
    // been used below for illustrative purposes.
    WSIFService service = factory.getService(
        wsdlLocation, // location of the wsdl file
```

```

    null,          // service namespace
    null,          // service name
    "http://www.ibm.com/namespace/wsif/samples/ab", // port type namespace
    "AddressBookPT" // port type name
);

// The AddressBook.wsdl file contains the definitions for two complexType
// elements within the schema element. We will now map these complexTypes
// to Java classes. These mappings are used by the Apache SOAP provider
service.mapType(
    new javax.xml.namespace.QName(
        "http://www.ibm.com/namespace/wsif/samples/ab/types",
        "address"),
    Class.forName("com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress"));

service.mapType(
    new javax.xml.namespace.QName(
        "http://www.ibm.com/namespace/wsif/samples/ab/types",
        "phone"),
    Class.forName("com.ibm.www.namespace.wsif.samples.ab.types.WSIFPhone"));

    // We now have a WSIFService object. The next step is to create a WSIFPort
    // object for the port we wish to use. The getPort(String portName) method
    // allows us to generate a WSIFPort from the port name.
WSIFPort port = null;

if (portName != null) {
    port = service.getPort(portName);
}
if (port == null) {
    // If no port name was specified, attempt to create a WSIFPort from
    // the available ports for the port type specified on the service
    port = getPortFromAvailablePortNames(service);
}

// Once we have a WSIFPort, we can create an operation. We are going to execute
// the addEntry operation and therefore we attempt to create a WSIFOperation
// corresponding to it. The addEntry operation is overloaded in the wsdl ie. there
// are two versions of it, each taking different parameters (parts). This overloading
// requires that we specify the input and output message names for the operation
// in the createOperation method so that the correct operation can be resolved.
// Since the addEntry operation has no output message, we use null for its name.
WSIFOperation operation =
    port.createOperation("addEntry", "AddEntryWholeNameRequest", null);

// Create messages to use in the execution of the operation. This should
// be done by invoking the createXXXXXMessage methods on the WSIFOperation.
WSIFMessage inputMessage = operation.createInputMessage();
WSIFMessage outputMessage = operation.createOutputMessage();
WSIFMessage faultMessage = operation.createFaultMessage();

// Create a name and address to add to the addressbook
String nameToAdd = "Chris P. Bacon";
WSIFAddress addressToAdd =
    new WSIFAddress (1,
        "The Waterfront",
        "Some City",
        "NY",
        47907,
        new WSIFPhone (765, "494", "4900"));

// Add the name and address to the input message
inputMessage.setObjectPart("name", nameToAdd);
inputMessage.setObjectPart("address", addressToAdd);

// Execute the operation, obtaining a flag to indicate its success
boolean operationSucceeded =

```

```

        operation.executeRequestResponseOperation(
            inputMessage,
            outputMessage,
            faultMessage);

    if (operationSucceeded) {
        System.out.println("Successfully added name and address to addressbook\n");
    } else {
        System.out.println("Failed to add name and address to addressbook");
    }

    // Start from fresh
    operation = null;
    inputMessage = null;
    outputMessage = null;
    faultMessage = null;

    // This time we will lookup an address from the addressbook.
    // The getAddressFromName operation is not overloaded in the
    // wsdl and therefore we can simply specify the operation name
    // without any input or output message names.
    operation = port.createOperation("getAddressFromName");

    // Create the messages
    inputMessage = operation.createInputMessage();
    outputMessage = operation.createOutputMessage();
    faultMessage = operation.createFaultMessage();

    // Set the name to find in the addressbook
    String nameToLookup = "Chris P. Bacon";
    inputMessage.setObjectPart("name", nameToLookup);

    // Execute the operation
    operationSucceeded =
        operation.executeRequestResponseOperation(
            inputMessage,
            outputMessage,
            faultMessage);

    if (operationSucceeded) {
        System.out.println("Successful lookup of name '" + nameToLookup + "' in addressbook");

        // We can obtain the address that was found by querying the output message
        WSIFAddress addressFound = (WSIFAddress) outputMessage.getObjectPart("address");
        System.out.println("The address found was:");
        System.out.println(addressFound);
    } else {
        System.out.println("Failed to lookup name in addressbook");
    }
} catch (Exception e) {
    System.out.println("An exception occurred when running the sample:");
    e.printStackTrace();
}
}

```

The code above refers to the following sample method:

```

WSIFPort getPortFromAvailablePortNames(WSIFService service)
    throws WSIFException {
    String portChosen = null;

    // Obtain a list of the available port names for the service
    Iterator it = service.getAvailablePortNames();
    {
        System.out.println("Available ports for the service are: ");
        while (it.hasNext()) {

```

```

        String nextPort = (String) it.next();
        if (portChosen == null)
            portChosen = nextPort;
        System.out.println(" - " + nextPort);
    }
}
if (portChosen == null) {
    throw new WSIFException("No ports found for the service!");
}
System.out.println("Using port " + portChosen + "\n");

// An alternative way of specifying the port to use on the service
// is to use the setPreferredPort method. Once a preferred port has
// been set on the service, a WSIFPort can be obtained via getPort
// (no arguments). If a preferred port has not been set and more than
// one port is available for the port type specified in the WSIFService,
// an exception is thrown.
service.setPreferredPort(portChosen);
WSIFPort port = service.getPort();
return port;
}

```

The web service itself uses the following classes:

WSIFAddress:

```

public class WSIFAddress implements Serializable {

    //instance variables
    private int streetNum;
    private java.lang.String streetName;
    private java.lang.String city;
    private java.lang.String state;
    private int zip;
    private WSIFPhone phoneNumber;

    //constructors
    public WSIFAddress () { }

    public WSIFAddress (int streetNum,
                        java.lang.String streetName,
                        java.lang.String city,
                        java.lang.String state,
                        int zip,
                        WSIFPhone phoneNumber) {
        this.streetNum = streetNum;
        this.streetName = streetName;
        this.city = city;
        this.state = state;
        this.zip = zip;
        this.phoneNumber = phoneNumber;
    }

    public int getStreetNum() {
        return streetNum;
    }

    public void setStreetNum(int streetNum) {
        this.streetNum = streetNum;
    }

    public java.lang.String getStreetName() {
        return streetName;
    }

    public void setStreetName(java.lang.String streetName) {
        this.streetName = streetName;
    }
}

```



```

}

public java.lang.String getCity() {
    return city;
}

public void setCity(java.lang.String city) {
    this.city = city;
}

public java.lang.String getState() {
    return state;
}

public void setState(java.lang.String state) {
    this.state = state;
}

public int getZip() {
    return zip;
}

public void setZip(int zip) {
    this.zip = zip;
}

public WSIFPhone getPhoneNumber() {
    return phoneNumber;
}

public void setPhoneNumber(WSIFPhone phoneNumber) {
    this.phoneNumber = phoneNumber;
}
}

```

WSIFPhone:

```

public class WSIFPhone implements Serializable {

    //instance variables
    private int areaCode;
    private java.lang.String exchange;
    private java.lang.String number;

    //constructors
    public WSIFPhone () { }

    public WSIFPhone (int areaCode,
                      java.lang.String exchange,
                      java.lang.String number) {
        this.areaCode = areaCode;
        this.exchange = exchange;
        this.number = number;
    }

    public int getAreaCode() {
        return areaCode;
    }

    public void setAreaCode(int areaCode) {
        this.areaCode = areaCode;
    }

    public java.lang.String getExchange() {
        return exchange;
    }
}

```

```

public void setExchange(java.lang.String exchange) {
    this.exchange = exchange;
}

public java.lang.String getNumber() {
    return number;
}

public void setNumber(java.lang.String number) {
    this.number = number;
}
}

```

WSIFAddressBook:

```

public class WSIFAddressBook {
    private Hashtable name2AddressTable = new Hashtable();

    public WSIFAddressBook() {
    }

    public void addEntry(String name, WSIFAddress address)
    {
        name2AddressTable.put(name, address);
    }

    public void addEntry(String firstName, String lastName, WSIFAddress address)
    {
        name2AddressTable.put(firstName+" "+lastName, address);
    }

    public WSIFAddress getAddressFromName(String name)
        throws IllegalArgumentException
    {
        if (name == null)
        {
            throw new IllegalArgumentException("The name argument must not be " +
                "null.");
        }
        return (WSIFAddress)name2AddressTable.get(name);
    }
}

```

And here's the corresponding WSDL file for the Web service:

```

<?xml version="1.0" ?>

<definitions targetNamespace="http://www.ibm.com/namespace/wsif/samples/ab"
    xmlns:tns="http://www.ibm.com/namespace/wsif/samples/ab"
    xmlns:typens="http://www.ibm.com/namespace/wsif/samples/ab/types"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema"
    xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
    xmlns:format="http://schemas.xmlsoap.org/wsdl/formatbinding/"
    xmlns:java="http://schemas.xmlsoap.org/wsdl/java/"
    xmlns:ejb="http://schemas.xmlsoap.org/wsdl/ejb/"
    xmlns="http://schemas.xmlsoap.org/wsdl/">

<types>
<xsd:schema
    targetNamespace="http://www.ibm.com/namespace/wsif/samples/ab/types"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema">

<xsd:complexType name="phone">
<xsd:element name="areaCode" type="xsd:int"/>
<xsd:element name="exchange" type="xsd:string"/>
<xsd:element name="number" type="xsd:string"/>

```

```

</xsd:complexType>

<xsd:complexType name="address">
  <xsd:element name="streetNum" type="xsd:int"/>
  <xsd:element name="streetName" type="xsd:string"/>
  <xsd:element name="city" type="xsd:string"/>
  <xsd:element name="state" type="xsd:string"/>
  <xsd:element name="zip" type="xsd:int"/>
  <xsd:element name="phoneNumber" type="typens:phone"/>
</xsd:complexType>

</xsd:schema>
</types>

<message name="AddEntryWholeNameRequestMessage">
  <part name="name" type="xsd:string"/>
  <part name="address" type="typens:address"/>
</message>

<message name="AddEntryFirstAndLastNamesRequestMessage">
  <part name="firstName" type="xsd:string"/>
  <part name="lastName" type="xsd:string"/>
  <part name="address" type="typens:address"/>
</message>

<message name="GetAddressFromNameRequestMessage">
  <part name="name" type="xsd:string"/>
</message>

<message name="GetAddressFromNameResponseMessage">
  <part name="address" type="typens:address"/>
</message>

<portType name="AddressBookPT">
  <operation name="addEntry">
    <input name="AddEntryWholeNameRequest"
      message="tns:AddEntryWholeNameRequestMessage"/>
  </operation>
  <operation name="addEntry">
    <input name="AddEntryFirstAndLastNamesRequest"
      message="tns:AddEntryFirstAndLastNamesRequestMessage"/>
  </operation>
  <operation name="getAddressFromName">
    <input name="GetAddressFromNameRequest" message="tns:GetAddressFromNameRequestMessage"/>
    <output name="GetAddressFromNameResponse" message="tns:GetAddressFromNameResponseMessage"/>
  </operation>
</portType>

<binding name="SOAPHttpBinding" type="tns:AddressBookPT">
  <soap:binding style="rpc"
    transport="http://schemas.xmlsoap.org/soap/http"/>
  <operation name="addEntry">
    <soap:operation soapAction=""/>
    <input name="AddEntryWholeNameRequest">
      <soap:body use="encoded"
        namespace="http://www.ibm.com/namespace/wsif/samples/ab"
        encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
    </input>
  </operation>
  <operation name="addEntry">
    <soap:operation soapAction=""/>
    <input name="AddEntryFirstAndLastNamesRequest">
      <soap:body use="encoded"
        namespace="http://www.ibm.com/namespace/wsif/samples/ab"
        encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
    </input>
  </operation>

```

```

<operation name="getAddressFromName">
  <soap:operation soapAction=""/>
  <input name="GetAddressFromNameRequest">
    <soap:body use="encoded"
      namespace="http://www.ibm.com/namespace/wsif/samples/ab"
      encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"/>
  </input>
  <output name="GetAddressFromNameResponse">
    <soap:body use="encoded"
      namespace="http://www.ibm.com/namespace/wsif/samples/ab"
      encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"/>
  </output>
</operation>
</binding>

<binding name="JavaBinding" type="tns:AddressBookPT">
  <java:binding/>
  <format:typeMapping encoding="Java" style="Java">
    <format:typeMap typeName="typens:address"
      formatType="com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress" />
    <format:typeMap typeName="xsd:string" formatType="java.lang.String" />
  </format:typeMapping>
  <operation name="addEntry">
    <java:operation
      methodName="addEntry"
      parameterOrder="name address"
      methodType="instance" />
    <input name="AddEntryWholeNameRequest"/>
  </operation>
  <operation name="addEntry">
    <java:operation
      methodName="addEntry"
      parameterOrder="firstName lastName address"
      methodType="instance" />
    <input name="AddEntryFirstAndLastNamesRequest"/>
  </operation>
  <operation name="getAddressFromName">
    <java:operation
      methodName="getAddressFromName"
      parameterOrder="name"
      methodType="instance"
      returnPart="address" />
    <input name="GetAddressFromNameRequest"/>
    <output name="GetAddressFromNameResponse"/>
  </operation>
</binding>

<binding name="EJBBinding" type="tns:AddressBookPT">
  <ejb:binding/>
  <format:typeMapping encoding="Java" style="Java">
    <format:typeMap typeName="typens:address"
      formatType="com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress" />
    <format:typeMap typeName="xsd:string" formatType="java.lang.String" />
  </format:typeMapping>
  <operation name="addEntry">
    <ejb:operation
      methodName="addEntry"
      parameterOrder="name address"
      interface="remote" />
    <input name="AddEntryWholeNameRequest"/>
  </operation>
  <operation name="addEntry">
    <ejb:operation
      methodName="addEntry"
      parameterOrder="firstName lastName address"
      interface="remote" />
    <input name="AddEntryFirstAndLastNamesRequest"/>
  </operation>

```

```

</operation>
<operation name="getAddressFromName">
  <ejb:operation
    methodName="getAddressFromName"
    parameterOrder="name"
    interface="remote"
    returnPart="address" />
  <input name="GetAddressFromNameRequest"/>
  <output name="GetAddressFromNameResponse"/>
</operation>
</binding>
<service name="AddressBookService">
  <port name="SOAPPort" binding="tns:SOAPHttpBinding">
    <soap:address
      location="http://localhost/wsif/samples/addressbook/soap/servlet/rpcrouter"/>
  </port>
  <port name="JavaPort" binding="tns:JavaBinding">
    <java:address className="services.addressbook.WSIFAddressBook"/>
  </port>
  <port name="EJBPort" binding="tns:EJBBinding">
    <ejb:address class="services.addressbook.ejb.AddressBook"
      jndiName="ejb/samples/wsif/AddressBook"
      classLoader="services.addressbook.ejb.AddressBook.ClassLoader"/>
  </port>
</service>
</definitions>

```

Using complex types

WSIF supports the use of user defined complex types through the mapping of complex types to Java classes. This mapping must be specified by the user. The method to use to create these mappings depends on the provider being used. For the Java and EJB providers, the mappings are specified in the wsdl file in the binding element. The syntax for specifying the mapping is as follows:

```

<binding .... >
  <ejb:binding|java:binding/>
  <format:typeMapping style="Java" encoding="Java"/>?
  <format:typeMap name="qname" formatType="nmtoken"/> *
</format:typeMapping>
  ...
</binding>

```

where "?" means "optional" and "*" means "0 or more".

The `format:typeMap name` attribute is a qualified name of a complex type or simple type used by one of the operations.

The `format:typeMap formatType` attribute is the fully qualified Class name for the Java Class that the element specified by `name` maps to.

If using the Apache SOAP provider then the mapping of a complex type to a Java Class is specified in the client code through two methods on the `org.apache.wsif.WSIFService` interface:

```
public void mapType(QName elementType, Class javaType)
```

and

```
public void mapPackage(String namespaceURI, String packageName)
```

The **mapType** allows the user to specify a mapping between a WSDL element and method takes a QName representing the complex type or simple type and the corresponding Java Class it maps to.

The **mapPackage** method allows the user to specify a more general mapping between a namespace and a Java package. Any custom complex or simple types whose namespace matches that of the mapping will be mapped to a Java Class in the corresponding package. The name of the actual class is derived from the name of the complex type using standard xml to Java naming conventions.

Using JNDI

This example task shows you how to use WSIF to bind a reference to a Web service, then look up the reference using JNDI.

You access a Web service through information given in the WSDL document for the service. If you don't know where to find the WSDL document for the service, but you know that it has been registered in a UDDI registry, you look it up in the registry. Java programs access java objects and resources in a similar manner, but using a JNDI interface.

The following example shows how, using WSIF, you can bind a reference to a Web service then look up the reference using JNDI.

Usage scenario

Specifying the argument values for the Web Service

The Web service is represented in WSIF by an instance of the `org.apache.wsif.naming.WSIFServiceRef` class. This simple Referencable object has the following constructor:

```
public WSIFServiceRef(
    String WSDL,
    String sNS,
    String sName,
    String ptNS,
    String ptName)
{
    [...]
}
```

where

- *WSDL* is the location of the WSDL file containing the definition of the service.
- *sNS* is the full namespace for the service definition (null can be specified if only one service is defined in the WSDL file).
- *sName* is the local name for the service definition (null can be specified if only one service is defined in the WSDL file).
- *ptNS* is the full namespace for the port type within the service that you want to use (null can be specified if only one port type is available for the service).
- *ptName* is the local name for the port type (null can be specified if only one port type is available for the service).

For example, if the WSDL file for the Web service is available from the URL `http://localhost/WSDL/Example.WSDL` and contains the following service and port type definitions -

```

<definitions targetNamespace="http://hostname/namespace/example"
              xmlns:abc="http://hostname/namespace/abc"
[...]
```

```

  <portType name="ExamplePT">
    <operation name="exampleOp">
      <input name="exampleInput" message="tns:ExampleInputMsg"/>
    </operation>
  </portType>
[...]
```

```

  <service name="abc:ExampleService">
[...]
```

```

  </service>
[...]
```

```

</definitions>
```

- then you specify the following argument values for `WSIFServiceRef`:

- *WSDL* is `http://localhost/WSDL/Example.WSDL`
- *sNS* is `http://hostname/namespace/abc`
- *sName* is `ExampleService`
- *ptNS* is `http://hostname/namespace/example`
- *ptName* is `ExamplePT`

Binding the service using JNDI

To bind the service reference in the naming directory using JNDI, you can use the WebSphere Application Server `JndiHelper` `com.ibm.websphere.naming.JndiHelper` class as follows:

```

[...]
```

```

import com.ibm.websphere.naming.JndiHelper;
import org.apache.wsif.naming.*;
[...]
```

```

try {
  Context startingContext = new InitialContext();
  WSIFServiceRef ref = new WSIFServiceRef("http://localhost/WSDL/Example.WSDL",
                                         "http://localhost/WSDL/Example.WSDL",
                                         "http://hostname/namespace/abc",
                                         "ExampleService",
                                         "http://hostname/namespace/example",
                                         "ExamplePT");
  JndiHelper.recursiveRebind(startingContext, "myContext/mySubContext/myServiceRef", ref);
}
catch (NamingException e) {
  // Handle error.
}
[...]
```

Looking up the service using JNDI

The following code fragment shows the lookup of a service using JNDI:

```

[...]
```

```

try {
[...]
```

```

  InitialContext ic = new InitialContext();
  WSIFService myService = (WSIFService) ic.lookup("myContext/mySubContext/myServiceRef");
[...]
```

```

}
catch (NamingException e) {
  // Handle error.
}
[...]
```

Interacting with the WebSphere J2EE container

Interaction with a container is limited to the following aspects:

Steps for this task

1. The WebSphere Application Server administrative console and WCCM allow users to define Web services to WebSphere. This is described in Using JNDI and WSIF system management and administration. As part of the definition of a service, the administrator may define a "Preferred Port".
2. WSIF makes log and trace calls to the WebSphere Server J2EE services, as described in Trace and logging for WSIF.
3. Some providers use the J2EE programming model to utilize J2EE services. The EJB provider uses JNDI and calls to remote EJBs.
4. WSIF wraps the use of container services so that when WSIF is run in an unmanaged (thin) environment, the operation can succeed.

Running WSIF as a client

WSIF runs in the WebSphere Application Server application client container, and in similar clients from other suppliers.

To simplify the process of launching client applications in the WebSphere Application Server application client, use the `launchClient` tool as described in (Running application clients).

WSIF system management and administration

WSIF is provided as a standalone JAR file called `wsif.jar`. The JAR file contains the core WSIF classes, and the Java, EJB, SOAP over HTTP and SOAP over JMS providers. Additional providers are packaged as separate JAR files.

When you install WebSphere Application Server, `wsif.jar` is put on the WebSphere or JVM class path.

WSIF requires no further configuration. WSIF is a thin abstraction layer between application code and the relevant invocation infrastructure.

For specific information on other aspects of managing your WSIF system, see the following topics:

- Maintaining the WSIF properties file.
- Enabling security for WSIF.
- Trace and logging for WSIF.
- WSIF troubleshooting tips.
- WSIF (Web Services Invocation Framework) messages.

Maintaining the WSIF properties file

WSIF properties are stored in a properties file (in `wsif.jar`) called `wsif.properties`. This file is kept on the class path, so that WSIF can find it and the client administrator can use it to configure WSIF.

The initial contents of `wsif.properties` are printed out below. All the possible properties are listed and described.

Usage scenario


```

# Two properties are used to override which WSIFProvider is selected when there
# exists multiple providers supporting the same namespace URI. These properties are:
#
#   wsif.provider.default.CLASSNAME=N
#   wsif.provider.uri.M.CLASSNAME=URI
#
# CLASSNAME is the WSIFProvider class name
# N is the number of following default wsif.provider.uri.M.CLASSNAME properties
# M is a number from 1 to N to uniquely identify each wsif.provider.uri.M.CLASSNAME
#   property key.
# For example the following two properties would override the default SOAP provider
# to be the Apache SOAP provider:
#
# wsif.provider.default.org.apache.wsif.providers.soap.apacheaxis.WSIFDynamicProvider_ApacheAxis=1
# wsif.provider.uri.1.org.apache.wsif.providers.soap.apacheaxis.WSIFDynamicProvider_ApacheAxis=\
# http://schemas.xmlsoap.org/wsdl/soap/
#

# maximum number of milliseconds to wait for a response to a synchronous request.
# Default value if not defined is to wait forever.
wsif.syncrequest.timeout=10000

# maximum number of seconds to wait for a response to an async request.
# if not defined on invalid defaults to no timeout
wsif.asyncrequest.timeout=60

```

Enabling security for WSIF

WSIF interacts with a security manager as follows:

- WSIF runs in the current J2EE security context without modifying it.
- When WSIF is run under a J2EE container, Port implementations can utilize security context to pass on security tokens or credentials as necessary.
- WSIF implementations can automatically convert J2EE security context into appropriate context for onward services.

For WSIF to interact effectively with the WebSphere Application Server's security manager, the following permissions must be set in the `server.policy` file:

- FilePermission to load the WSDL (this is only required when a WSDL file is referred to using the `file:///` protocol)
- RuntimePermission "getClassLoader" for the current thread's context class loader (this is not required by the EJB portion).
- RuntimePermission "accessDeclaredMembers" (this is required by both portions for handling enterprise beans)
- PropertyPermission for system properties (this is required by SOAP and many others; read and write access is required for the SOAP and Java portion, only write access is required for the EJB portion)
- NetPermission "specifyStreamHandler" (this must be in either the SOAP and Java portion, or the EJB portion, but it need not be in both).
- SocketPermission "*host_name*", "resolve" (this is not required by the SOAP and Java portion)
- SocketPermission "*host_name:port_no*", "connect" (this is required by both portions)

where *host_name* is your host name (for example `localhost`), and *port_no* is your port number (for example `9080`).

WSIF troubleshooting tips

For information on resolving WebSphere-level problems, see [Diagnosing and fixing problems](#).

To identify and resolve WSIF-related problems, you can use the standard WebSphere Application Server trace and logging facilities. If you encounter a problem that you think might be related to WSIF, you can check for error messages in the WebSphere Application Server administrative console, and in the application server's `stdout.log` file. You can also enable the application server debug trace to provide a detailed exception dump.

A list of the WSIF runtime system messages, with details of what each message means, is given in [Message reference for WSIF](#).

Here is a checklist of major WSIF activities, with advice on common problems associated with each activity:

Create service

Handcrafted WSDL can cause numerous problems. To help ensure that your WSDL is valid, use a tool such as WebSphere Studio Application Developer (WSAD) to create your service.

Define transport mechanism

For JMS, check that you have set up JNDI correctly, and created all the necessary queues.

For SOAP, make sure that the deployment descriptor (`dds.xml`) is correct - preferably by creating it using WSAD or similar tooling.

Create client - Java code

Follow the correct format for creating a WSIF service, port, operation and message. For examples of correct code, see the [Address Book sample](#).

Compile code (client and service)

Check that the build path against code is correct, and that it contains the correct levels of JAR files.

Create a valid EAR file for your service in preparation for deployment to a web server.

Deploy service

When you install and deploy the service EAR file, check carefully any messages given when the service is deployed.

Server setup and start

Make sure that the WebSphere Application Server file `server.policy` (in the `/properties` directory) has the correct security settings within it. For more information see [Enabling security for WSIF](#).

WSIF setup

Check that file `wsif.properties` is correctly set up. For more information see [Maintaining the WSIF properties file](#).

Run client

Either check that you have defined the class path correctly to include references to your client classes, WSIF JAR files and any other necessary JARs, or (preferably) run your client using the WebSphere Application Server (`launchClient`) tool.

Here is a list of common errors, and information on their probable causes:

"No class definition" errors received when running client code

This is likely to be a problem with the class path setup. Check that the relevant JAR files are included.

"Can't find WSDL" error

Some likely causes are:

- The application server is not running.
- The server location and port number in the WSDL are not correct.
- The WSDL is badly formed (check the error messages in the application server's `stdout.log` file).
- The application server has not been restarted since the service was installed.

You might also try the following checks:

- Can you load the WSDL into your Web browser from the location specified in the error message?
- Can you load the corresponding WSDL binding files into your Web browser?

Your Web service's EAR file does not install correctly onto the application server.

It is likely that the EAR file is badly formed. Verify the installation by completing the following steps:

- For an EJB binding, run the WebSphere Application Server tool `\bin\dumpnamespace`. This tool lists the current contents of the JNDI directory.
- For a SOAP over HTTP binding, open `http://pathToServer/WebServiceName/admin/list.jsp` (if you have the SOAP administration pages installed). This page lists all currently installed Web services.
- For a SOAP over JMS binding, complete the following checks:
 - Check that the Queue Manager is running.
 - Check that the necessary queues are defined.
 - Check the JNDI setup.
 - Use the `jmsadmin` tool's "display context" option to list the current JNDI definitions.
 - Check that the RPCrouter is running.

There is a permissions problem or security error.

Check that the WebSphere Application Server file `server.policy` (in the `/properties` directory) has the correct security settings within it. For more information see *Enabling security for WSIF*.

Using WSIF with multiple clients causes a SOAP parsing error.

Before you deploy a Web service to WebSphere Application Server, you must decide on the Web service's scope. The application's deployment descriptor file (`dds.xml`) includes the following line:

```
<isd:provider type="java" scope="Application" .....
```

You can set Scope to "Application" or "Session". The default setting is "Application", and this is correct if each request to the Web service does not require objects to be maintained for longer than a single instance. If Scope is set to "Application" the objects are not available to another request during the execution of the single instance, and they are released on completion. If your Web service needs objects to be maintained for

multiple requests, and to be unique within each request, you must set the scope to "Session". If Scope is set to "Session", the objects are not available to another request during the life of the session, and they are released on completion of the session. If scope is set to "Application" but it should be set to "Session", you might get the following SOAP error:

```
SOAPException: SOAP-ENV:ClientParsing error, response was:  
FWK005 parse may not be called while parsing.;  
nested exception is:
```

```
[SOAPException: faultCode=SOAP-ENV:Client; msg=Parsing error, response was:
```

```
FWK005 parse may not be called while parsing.; targetException=org.xml.sax.SAXException:  
FWK005 parse may not be called while parsing.]
```

Trace and logging for WSIF

If you want to enable trace for the WSIF API within WebSphere Application Server, and have trace, stdout and stderr for the application server written to a well-known location, see Enabling trace.

WSIF offers tracepoints at opening and closing of Ports, invocation of services, and responses from services.

To trace the WSIF API, you need to specify the following trace string:

```
wsif=all=enabled
```

WSIF also includes a SimpleLog utility through which you can run trace when using WSIF outside of WebSphere Application Server. To enable this, complete the following steps:

Steps for this task

1. create file commons-logging.properties with the following contents:

```
org.apache.commons.logging.LogFactory=org.apache.commons.logging.impl.LogFactoryImpl  
org.apache.commons.logging.Log=org.apache.commons.logging.impl.SimpleLog
```
2. create file simplelog.properties with the following contents:

```
org.apache.commons.logging.simplelog.defaultlog=trace  
org.apache.commons.logging.simplelog.showShortLogname=true  
org.apache.commons.logging.simplelog.showdatetime=true
```
3. Put both these files, and commons-logging.jar, on the class path.

Results

The SimpleLog mechanism writes trace to file System.err.

WSIF (Web Services Invocation Framework) messages

WebSphere system messages are logged from a variety of sources, including application server components and applications. Messages logged by application server components and associated IBM products start with a unique message identifier that indicates the component or application that issued the message. For more information about the message identifier format, see the topic Message Format.

The rest of this topic contains a list of the WSIF runtime system messages, with details of what each message means.

WSIF0001E: An extension registry was not found for the element type "{0}"

Explanation: Parameters: {0} element type. No extension registry was found for the element type specified.

User Response: Add the appropriate extension registry to the port factory in your code.

WSIF0002E: A failure occurred in loading WSDL from "{0}"

Explanation: Parameters: {0} location of the WSDL file. The WSDL file could not be found at the location specified or did not parse correctly

User Response: Check that the location of the WSDL file is correct. Check that any network connections required are available. Check that the WSDL file contains valid WSDL.

WSIF0003W: An error occurred finding pluggable providers: {0}

Explanation: Parameters: {0} specific details about the error. There was a problem locating a WSIF pluggable provider using the J2SE 1.3 JAR file extensions to support service providers architecture. The WSIF trace file will contain the full exception details.

User Response: Verify that a META-INF/services/org.apache.wsif.spi.WSIFProvider file exists in a provider jar, that each class referenced in the META-INF file exists in the class path, and that each class implements org.apache.wsif.spi.WSIFProvider. The class in error will be ignored and WSIF will continue locating other pluggable providers.

WSIF0004E: WSDL contains an operation type "{0}" which is not supported for "{1}"

Explanation: Parameters: {0} name of the operation type specified. {1} name of the portType for the operation. An operation type which is not supported has been specified in the WSDL.

User Response: Remove any operations of the unsupported type from the WSDL. If the operation is required then make sure all messages have been correctly specified for the operation.

WSIF0005E: An error occurred when invoking the method "{1}" . (" {0} ")

Explanation: Parameters: {0} name of communication type. For example EJB or ApacheSOAP. {1} name of the method that failed. An error was encountered when invoking a method on the web service using the communication shown in brackets.

User Response: Check that the method exists on the web service and that the correct parts have been added to the operation as described in the WSDL. Network problems might be a cause if the method is remote and so check any required connections.

WSIF0006W: Multiple WSIFProvider found supporting the same namespace URI "{0}" . Found (" {1} ")

Explanation: Parameters: {0} the namespace URI. {1} a list of the WSIFProvider found.. There are multiple org.apache.wsif.spi.WSIFProvider classes in the service provider path that support the same namespace URI.

User Response: A following WSIF0007I message will be issued notifying which WSIFProvider will be used. Which WSIFProvider is chosen is based on settings in the wsif.properties file, or if not defined in the properties, the last WSIFProvider found will be used. See the wsif.properties file for more details on how to define which provider should be used to support a namespace URI.

WSIF0007I: Using WSIFProvider "{0}" for namespaceURI "{1}"

Explanation: Parameters: {0} the classname of the WSIFProvider being used. {1} the namespaceURI the provider will be used to support.. Either a

previous WSIF0006W message has been issued or the SetDynamicWSIFProvider method has been used to override the provider used to support a namespaceURI.

User Response: None. See also WSIF0006W.

WSIF0008W: WSIFDefaultCorrelationService removing correlator due to timeout. ID:"{0}"

Explanation: Parameters: {0} the id of the correlator being removed from the correlation service. A stored correlator is being removed from the correlation service due to its timeout expiring.

User Response: Determine why no response has been received for the asynchronous request within the timeout period. The wsif.asyncrequest.timeout property of the wsif.properties file defines the length of the timeout period.

WSIF0009I: Using correlation service - "{0}"

Explanation: Parameters: {0} the name of the correlation service being used. This identifies the name of the correlation service that will be used to process asynchronous requests.

User Response: None. If a correlation service other than the default WSIF supplied one is required, ensure that it is correctly registered in the JNDI java:comp/wsif/WSIFCorrelationService namespace.

WSIF0010E: Exception thrown while processing asynchronous response - "{0}"

Explanation: Parameters: {0} the error message string of the exception. While processing the response from an executeRequestResponseAsync call an exception was thrown.

User Response: Use the exception error message string to determine the cause of the error. The WSIF trace will have more details on the error including the exception stack trace.

WSIF0011I: Preferred port "{0}" was not available

Explanation: Parameters: {0} the user's preferred port. The preferred port set by the user on org.apache.wsif.WSIFService is not available

User Response: None unless this message appears for long periods of time in which case the user might want to pick a different port as their preferred port.

WSIF API

The WSIF API supports the invocation of Services defined in WSDL. WSIF is intended to be used in both WSIF clients and also in Web service intermediaries.

The WSIF API is driven by the abstract service description in WSDL; it is completely independent of the actual binding used. This makes the API more natural to work with, because it uses WSDL terms to refer to message parts, operations, and so on.

The WSIF API was designed for the WSDL usage model; to pick a port that supports the port type needed, then invoke the operation by providing the necessary abstract input message consisting of the required parts, without worrying about how the message is mapped to a specific binding protocol.

Other Web service APIs, for example SOAP APIs, are not designed on WSDL, but for a specific binding protocol with its associated syntax; for example, target URIs and encoding styles.

The WSIF API's main interfaces are described in the following topics:

- Creating a message for sending to a port (the `WSIFMessage` interface).
- WSIF API reference: Finding a port factory or service (the `WSIFService` interface and the `WSIFServiceFactory` class).
- WSIF API reference: Using ports (the `WSIFPort` interface and the `WSIFOperation` interface).

Note: You must ensure that your application uses only one thread to call WSIF.

For additional technical details of the WSIF API, see the WSIF javadoc (`../javadoc/wsi/index.html`).

WSIF API reference: Creating a message for sending to a port

For message management (that is, message construction and parsing) the underlying API is modeled on WSDL semantics. There is a simple and direct mapping from the WSDL model to WSIF classes.

In WSDL, a Message describes the abstract type of the input or output to an operation. The corresponding WSIF class is **WSIFMessage**, which represents in memory the actual input or output of an operation. A `WSIFMessage` is a container for a set of named parts. The `WSIFMessage` interface separates the actual representation of the data from the abstract type defined by WSDL. WSDL defines Messages as XML Schema types. There are three natural ways to represent a WSDL message in a runtime environment:

- The generated Java class based on a WSDL to Java mapping such as that provided by JAX-RPC.
- The XML representation of the data, for example using SOAP Encoding.

Each option offers benefits in different scenarios. The Java class is the natural approach when WSIF is being used in a standard Java client. However, in other scenarios where WSIF is being used in an intermediary, it may be more efficient to keep a WSDL Message in the SOAP encoded format.

The style used to define messages must be consistent within the message, so all the parts in one message must be consistent. A string - `getRepresentationStyle()` - always returns `null`. This indicates that parts on this `WSIFMessage` are represented as Java objects.

Parts are added to a `WSIFMessage` with **`setObjectPart`** or **`setTypePart`**. Each part is named. Part names within a message are unique. If a part is **`set`** multiple times, the last time is the one that stands.

Parts are retrieved from a `WSIFMessage` by name with **`getObjectPart`** or **`getTypePart`**. If the named part does not exist, the `getXXXXPart` method returns a `WSIFException`.

Parts can be retrieved from the Message by use of Iterators through the **`getParts()`** and **`getPartNames()`** methods.

The order in which parts are set is not currently important, but the Message implementation may be more efficient if the parts are set in the parameter order specified by WSDL.

WSIFMessages are cloneable and serializable. If the parts set are not cloneable, the implementation should try to clone them using serialization. If the parts are not serializable either, then a CloneNotSupportedException will be thrown if cloning is attempted.

WSIFMessages can be sent between JVMs.

In addition to the containing parts, a WSIFMessage also has a message name. This is required for operation overloading - which is supported by WSDL and WSIF.

Here is the Javadoc for the WSIFMessage interface.

WSIF API reference: Finding a port factory or service

To find a port you use WSIFService, which is a factory for ports. The port factory models and supports the WSDL approach in which a service is available on one or more ports. The factory hides the implementation of the port from the user. WSIF supports "dynamic" ports that are based on a particular protocol/transport and configured using the WSDL at runtime. For example, the dynamic SOAP port can invoke any SOAP service based on the WSDL description of that service. The Service allows you to hide and modify this at runtime.

Here is the WSIFService interface.

To find a service from a WSDL document at a URL, or from a code-generated codebase, you can use the WSIFServiceFactory class.

WSIFService interface

The WSIFService is responsible for generating an instance of WSIFOperation to be used for a particular invocation of a service operation.

The WSIF service stores a list of providers that can each generate a WSIF operation for a particular WSDL binding. This service looks up providers by the provider type, so, for example, it knows about one provider that handles SOAP ports, and other providers that handle Java ports that you define. In a managed environment, the container can configure WSIFService.

Here is the Javadoc for the WSIFService interface.

A WSIFService implementation can choose a preferred port based on a number of criteria. The WSIFService implementation can set the preferred port, or it can be set by calling **setPreferredPort**.

The **getPort** method returns an instance of WSIFPort that is used to invoke a service on the port. Variants of the **getPort** method are used to define the characteristics of the port to be created. **getPort** with no arguments returns the "preferred" port. **getPort** with a string argument returns the port named by the string containing the WSDL identifier for the selected port. The return value is null if the port name is not valid.

If a port is chosen (either by WSIFService default behavior, or by setPreferredPort), then the WSIFService implementation validates that the relevant provider exists and is configured. If the provider fails this validation check, WSIFService chooses

any other port whose provider is defined. For example, if the preferred port is SOAP over JMS, but the JMS libraries are not available, WSIF chooses another port. If no preferred port is set, or the preferred port is not available, the WSIF implementation chooses the first available port listed in the WSDL.

The method `getAvailablePortNames()` returns an Iterator of Strings which is the list names of WSDL ports, filtered by the set of available providers.

The `getDefinition()` method returns the WSDL definition that was used to provide the Service. If the WSDL definition is not available, this returns a null.

WSIFServiceFactory class

To find a service from a WSDL document at a URL, or from a code-generated codebase, you can use the WSIFServiceFactory class.

Note: When creating a WSIFService from a WSIFServiceFactory you can specify a ClassLoader object to use in locating the WSDL file. You need to do this when the WSDL file is in a JAR file. In such a case, the location of the WSDL file is specified relative to the root of the JAR file, using forward slashes with the preceding slash removed.

For example:

```
com/myCompany/wsd1/MyWSDLFile.wsd1
```

rather than

```
/com/myCompany/wsd1/MyWSDLFile.wsd1
```

Here is the Javadoc for the WSIFServiceFactory class.

The WSIFServiceFactory returns a null if no service is found with that identifier.

WSIF API reference: Using ports

A WSIFPort handles the details of invoking an operation. The port provides access to the actual implementation of the service. A WSDL can provide many different WSDL Bindings, and these bindings can drive multiple Ports. The client can choose a Port, the service stub can choose a Port, or WSIF can choose a default Port.

The Port offers an interface to retrieve an Operation object. A WSIFOperation offers the ability to execute the given operation.

Note:

If the Port is serialized and deserialized at a later time, then WSIF ensures that the client provides the correct information to the server to identify the instance. If the server instance is no longer available, then it is up to the server to decide whether to throw a fault or provide a new instance. That behavior can depend on the type of service.

For example, for an enterprise bean the WSIFPort stores the EJB Home, and uses it to "select" the bean before each invocation. It is the responsibility of the client to serialize or maintain the Port instance if it wants instance support. The client must create a new operation and messages for each invocation.

Here is the WSIFPort interface.

Here is the WSIFOperation interface.

WSIFPort interface

The port implements a factory method for WSIFOperations.

Here is the Javadoc for the WSIFPort interface.

The `createOperation(String)` method returns a new instance of a `WSIFOperation` object. If the `operationName` is not valid, it will throw an exception. If the operation is overloaded, it will throw an exception. The `createOperation(String, String, String)` allows WSIF to support overloaded WSDL operations.

Note: You can overload based on the input parameters, but not on the output parameters.

It is the duty of the client to call the **close** method when a Port is no longer in use. In many cases, where the transport is sessionless, like HTTP, this will have no effect. However, if the Port is using a session-based protocol such as MQSeries, JMS, or ECI, this allows the Port to cache an open connection to the server and then close it as required. Responsible stubs will call this if appropriate.

WSIFOperation interface

`WSIFOperation` is the runtime representation of an operation. It is responsible for invoking a service based on a particular binding. It provides methods to create input, output and fault messages, and to invoke the operation.

Here is the Javadoc for the WSIFOperation interface.

The **createInputMessage**, **createOutputMessage**, and **createFaultMessage** methods are factory methods to create the messages required by the invocation methods. All invocation methods require an input message.

The **executeRequestResponseOperation** invokes "In Out" operations.

The **executeInputOnlyOperation** invokes "In only" operations.

If the invocation method is **executeRequestResponseOperation** then an output and a fault message are instantiated and passed on the call to the **executeRequestResponseOperation** method. The output message contains the response message when the **executeRequestResponseOperation** returns true. If **executeRequestResponseOperation** returns false, then a fault occurred and is returned in the fault message.

All of these **executeNmmn** methods fail with an exception if there is an error in processing the request in the WSIF provider.

The **executeRequestResponseAsync** is a feature that allows "In Out" operations to be invoked with the reply handled using an alternate thread.

Use of the **setContext** and **getContext** methods is discussed in `WSIFOperation - Context`.

Use of the **executeRequestResponseAsync** feature is discussed further in `WSIFOperation - Asynchronous interactions`.

WSIFOperation - Context: Although WSDL does not define context, a number of uses of WSIF require the ability to pass context to the Port that is invoking the

service. For example, a SOAP/HTTP port may require an HTTP username and password. This information is specific to the invocation, but not a parameter of the service. In general, context is defined as a set of name-value pairs. However, because Web services tend to define the types of data using XML Schema types, WSIF represents the name-value pairs of the context using the same representation that WSIFMessages use; that is a set of named Parts, each of which equates to an instance of an XML Schema type.

The WSIFOperation methods **setContext** and **getContext** allow you to pass context information to the binding. The Port implementation can use this context - for example to update a SOAP header. There is no definition of how a Port can utilize the context.

The parameter of the **setContext** and **getContext** methods is a WSIFMessage, and this has named parts defining the context information. The WSIFConstants class defines constants for the part names that can be set in a context WSIFMessage.

The following code shows how to set the user name and password to be used for HTTP basic authentication:

```
// set a basic authentication header
WSIFMessage headers = new WSIFDefaultMessage();
headers.setObjectPart( WSIFConstants.CONTEXT_HTTP_USER, "username" );
headers.setObjectPart( WSIFConstants.CONTEXT_HTTP_PSWD, "password" );
operation.setContext( headers );
```

The WSIFOperation will ignore context parts that it does not support, for example, the above code would be ignored by the WSIF Java provider.

The WSIFConstants class includes the following constants which can be used for context part names:

- CONTEXT_HTTP_USER
- CONTEXT_HTTP_PSWD
- CONTEXT_SOAP_HEADERS

The HTTP header values are expected to have a type **String**, and the SOAP header value is expected to have a type of **java.util.List** which should contain entries of type **org.w3c.dom.Element**.

WSIFOperation - Asynchronous interactions reference: WSIF allows asynchronous operation. In this mode of operation, the client puts the request message as part of one transaction, and carries on with the thread of execution. The response message is then handled by a different thread, with a separate transaction. The SOAP/JMS and Native JMS providers are the only WSIF providers that currently support asynchronous operation.

The WSIFPort has method **supportsAsync** to test if asynchronous operation is supported.

An asynchronous operation is initiated with the WSIFOperation method **executeRequestResponseAsync**. This method is designed to let an RPC method be invoked asynchronously. The method returns before the operation has completed, and the thread of execution continues.

The response to the asynchronous request is processed by one of the WSIFOperation methods **fireAsyncResponse** or **processAsyncResponse**.

To initiate the request, there are two forms of the `executeRequestResponseAsync` method:

```
public WSIFCorrelationId executeRequestResponseAsync (WSIFMessage input, WSIFResponseHandler handler)

and

public WSIFCorrelationId executeRequestResponseAsync (WSIFMessage input)
```

The first of these takes an input message and a **WSIFResponseHandler**. The `WSIFResponseHandler` is invoked on another thread when the operation completes. When using this method the client listener would call the **fireAsyncResponse** method which will call the `WSIFResponseHandler` **executeAsyncResponse** method.

Here is the Javadoc for the `WSIFResponseHandler` interface.

The other form of **executeRequestResponseAsync** that only takes a `WSIFMessage` input message does not use a response handler, and the client listener should process the response by calling the `WSIFOperation` method **processAsyncResponse**. This will update the `WSIFMessage` output and fault messages with the result of the request.

WSIF supports correlation between the request and response in the asynchronous request-response case. When the request is sent, the `WSIFOperation` is serialized into the **WSIFCorrelationService**. The **executeRequestResponseAsync** methods returns a **WSIFCorrelationId** object which identifies the serialized `WSIFOperation`. The client listener can use this to match a response to a particular request.

The correlation service is located with the **getCorrelationService()** method of the **WSIFCorrelationServiceLocator** class in the `org.apache.wsif.utils` package.

In a managed container a default correlation service is defined in the default JNDI namespace using the name: `java:comp/wsif/WSIFCorrelationService`. If this is not available, then WSIF will use the `WSIFDefaultCorrelationService`.

Here is the Javadoc for the `WSIFCorrelationService` interface.

and this is the correlator id:

```
public interface WSIFCorrelator extends Serializable {
    public String getCorrelationId();
}
```





The client must implement their own response message "listener" or MDB in order to recognize the arrival of response messages themselves. This would manage the correlation of the response message to the request and call of one of the asynchronous response processing methods. As an example of the type of thing required of a client listener, the following code fragment shows what could be in the **onMessage** method of a JMS listener:

```
public void onMessage(Message msg) {
    WSIFCorrelationService cs = WSIFCorrelationServiceLocator.getCorrelationService();
    WSIFCorrelationId cid = new JmsCorrelationId( msg.getJMSCorrelationID() );
    WSIFOperation op = cs.get( cid );
    op.fireAsyncResponse( msg );
}
```

WSIF: Resources for learning

Use the following links to find supplementary information about getting started with WSIF. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

-  **SOAP** <http://www.w3.org/TR/SOAP>
This article is a detailed overview of SOAP, which includes programming specifications.
-  **Web Services Description Language** <http://www.w3.org/TR/wsdl>
This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications
-  **The Apache Software Foundation** <http://www.apache.org>
The Apache Software Foundation provides support for the Apache community of open-source software projects. Of particular interest are the **Apache SOAP** project and its follow-on project **Apache AXIS**. The WSIF source code has been donated by IBM to the Apache Software Foundation, and is maintained as a sub-project of Apache AXIS.
-  **JSR109** <http://jcp.org/jsr/detail/109.jsp> .
This is the Java Community Process specification for implementing Web services in Java.

See also the set of Web services links given in (Web services: Resources for learning)

Chapter 9. Classloading

Classloaders affect the (packaging of applications) and the run-time behavior of packaged applications deployed on application servers.

Steps for this task

1. Read about classloaders. The article "Classloading: Resources for learning" points to additional sources.
2. If necessary, migrate classloader Module Visibility Mode settings for Version 4.0.x applications to Version 5.0 application or WAR classloader policies.
3. When assembling an enterprise application resource (EAR) file that has EJB modules, (set the classpath for the classloader to use during packaging).
4. **(Optional)** If an application module uses a resource, create a resource provider that specifies the directory name of the resource drivers. Do not specify the resource JAR file names. All JAR files in the specified directory will be added into the classpath of the WebSphere Application Server extensions classloader.
5. Configure classloaders of an application server for the run-time environment.
 - a. Click **Servers > Application Servers > *server_name*** and, on the settings page for an application server, set the application classloader policy and application classloader mode.

The application classloader policy controls the isolation of applications running in the system. When set to SINGLE, applications are not isolated; a single application classloader is used to contain all EJB modules, dependency JAR files, and shared libraries in the system. When set to MULTIPLE, applications are isolated from each other; each application receives its own classloader to load that application's EJB modules, dependency JAR files, and shared libraries.

The application classloader mode specifies the classloader mode when the application classloader policy is SINGLE. PARENT_FIRST causes the classloader to first delegate the loading of classes to its parent classloader before attempting to load the class from its local classpath. PARENT_LAST causes the classloader to first attempt to load classes from its local classpath before delegating the classloading to its parent. This allows an application classloader to override and provide its own version of a class that exists in the parent classloader.

- b. On the settings page for an application server, click **Classloader**. On the Classloader page, click **New**.
- c. On the settings page for a classloader, specify the classloader mode. PARENT_FIRST causes the classloader to delegate the loading of classes to its parent classloader before attempting to load the class from its local classpath. PARENT_LAST causes the classloader to attempt to load classes from its local classpath before delegating the classloading to its parent. Then, click **OK**.
- d. **(Optional)** On the settings page for a classloader, click **Libraries**. From the Library Ref page, click **Add**. On the settings page for a library reference, specify variables for the library reference as needed and click **OK**. Repeat the previous step until you define a library reference instance for each library file that your application needs. To define a library reference, you must first define one or more shared libraries.

6. When configuring an installed enterprise application for deployment in the run-time environment, (set the classloader mode and the WAR classloader policy).
7. When configuring an installed Web module for deployment in the run-time environment, (set the classloader mode).

Classloaders

Classloaders are part of the Java virtual machine (JVM) code and are responsible for finding and loading class files. Classloaders affect the packaging of applications and the run-time behavior of packaged applications deployed on application servers.

The run-time environment of WebSphere Application Server uses the following classloaders to find and load new classes for an application in the following order:

1. The **bootstrap, extensions, and CLASSPATH classloaders** created by the JVM.
The bootstrap classloader uses the boot classpath (typically classes in `jre/lib`) to find and load classes. The extensions classloader uses the system property `java.ext.dirs` (typically `jre/lib/ext`) to find and load classes. The CLASSPATH classloader uses the CLASSPATH environment variable to find and load classes.
The CLASSPATH classloader contains the J2EE APIs of the WebSphere Application Server product (inside `j2ee.jar`). Because the J2EE APIs are in this classloader, you can add libraries that depend on J2EE APIs to the classpath system property to extend a server's classpath. However, a preferred method of extending a server's classpath is to add a shared library.
2. A **WebSphere-specific extensions classloader**.
The WebSphere extensions classloader loads the WebSphere run time and J2EE classes that are required at run time. The extensions classloader uses a `ws.ext.dirs` system property to determine the path used to load classes. Each directory in the `ws.ext.dirs` classpath and every JAR file or ZIP file in these directories is added to the classpath used by this classloader.
The WebSphere extensions classloader also loads resource provider classes into a server if an application module installed on the server refers to a resource that is associated with the provider and if the provider specifies the directory name of the resource drivers.
3. One or more **application module classloaders** that load elements of enterprise applications running in the server.
The application elements can be Web modules, EJB modules, resource adapters, and dependency JAR files. Application classloaders follow J2EE class loading rules to load classes and JAR files from an enterprise application. The WebSphere run-time enables you to associate a shared library classpath with an application.

Each classloader is a child of the classloader above it. That is, the application module classloaders are children of the WebSphere-specific extensions classloader, which is a child of the CLASSPATH Java classloader. Whenever a class needs to be loaded, the classloader usually delegates the request to its parent classloader. If none of the parent classloaders can find the class, the original classloader attempts to load the class. Requests can only go to a parent classloader; they cannot go to a child classloader. If the WebSphere classloader is requested to find a class in a J2EE module, it cannot go to the application module classloader to find that class and a `ClassNotFoundException` occurs. Once a class is loaded by a classloader, any new classes that it tries to load reuse the same classloader or go up the precedence list until the class is found.

Classloader isolation policies

The number and function of the application module classloaders depends on the classloader policies specified in the server configuration. Classloaders provide multiple options for isolating applications and modules to enable different application packaging schemes to run on an application server.

Two classloader policies control the isolation of applications and modules:

Application classloader policy

Application classloaders consist of EJB modules, dependency JAR files, resource adapters, and shared libraries. Depending on the application classloader policy, an application classloader can be shared by multiple applications (SINGLE) or unique for each application (MULTIPLE). The application classloader policy controls the isolation of applications running in the system. When set to SINGLE, applications are not isolated. When set to MULTIPLE, applications are isolated from each other.

WAR classloader policy

By default, Web module classloaders load the contents of the WEB-INF/classes and WEB-INF/lib directories. The application classloader is the parent of the Web module classloader. You can change the default behavior by changing the application's WAR classloader policy.

The WAR classloader policy controls the isolation of Web modules. If this policy is set to APPLICATION, then the Web module contents also are loaded by the application classloader (in addition to the EJB files, RAR files, dependency JAR files, and shared libraries). If the policy is set to MODULE, then each web module receives its own classloader whose parent is the application classloader.

Note: WebSphere server classloaders never load application client modules.

For each application server in the system, you can set the application classloader policy to SINGLE or MULTIPLE. When the application classloader policy is set to SINGLE, then a single application classloader loads all EJB modules, dependency JAR files, and shared libraries in the system. When the application classloader policy is set to MULTIPLE, then each application receives its own classloader used for loading that application's EJB modules, dependency JAR files, and shared libraries.

This application classloader can load each application's Web modules if that WAR module's classloader policy is also set to APPLICATION. If the WAR module's classloader policy is set to APPLICATION, then the application's loader loads the WAR module's classes. If the WAR classloader policy is set to MODULE, then each WAR module receives its own classloader.

The following example shows that when the application classloader policy is set to SINGLE, a single application classloader loads all EJB modules, dependency JAR files, and shared libraries of all applications on the server. The single application classloader can also load Web modules if an application has its WAR classloader policy set to APPLICATION. Applications having a WAR classloader policy set to MODULE use a separate classloader for Web modules.

Application classloader policy: SINGLE

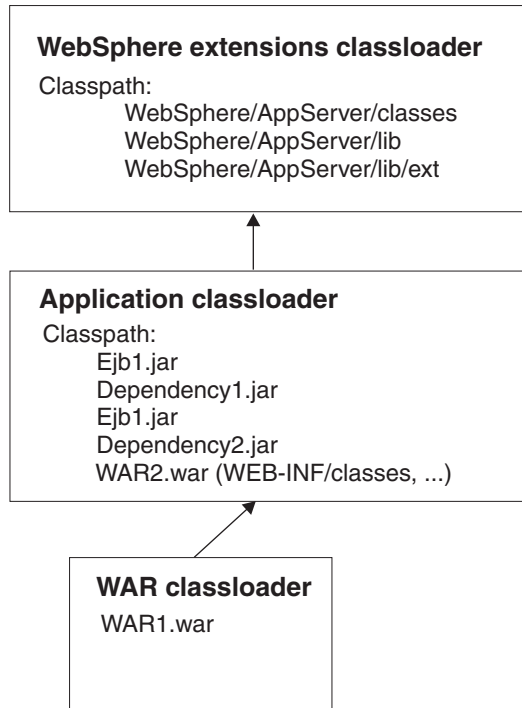
```
Application 1
Module: EJB1.jar
Module: WAR1.war
```



```

MANIFEST Class-Path: Dependency1.jar
WAR Classloader Policy = MODULE
Application 2
Module: EJB2.jar
MANIFEST Class-Path: Dependency2.jar
Module: WAR2.war
WAR Classloader Policy = APPLICATION

```




The following example shows that when the application classloader policy of an application server is set to `MULTIPLE`, each application on the server has its own classloader. An application classloader also loads its Web modules if the application's WAR classloader policy is set to `APPLICATION`. If the policy is set to `MODULE`, then a Web module uses its own classloader.

Application classloader policy: `MULTIPLE`

```

Application 1
Module: EJB1.jar
Module: WAR1.war
MANIFEST Class-Path: Dependency1.jar
WAR Classloader Policy = MODULE
Application 2
Module: EJB2.jar
MANIFEST Class-Path: Dependency2.jar
Module: WAR2.war
WAR Classloader Policy = APPLICATION

```



Classloader modes

There are two possible values for a classloader mode:

PARENT_FIRST

The PARENT_FIRST classloader mode causes the classloader to first delegate the loading of classes to its parent classloader before attempting to load the class from its local classpath. This is the default for classloader policy and for standard JVM classloaders.

PARENT_LAST

The PARENT_LAST classloader mode causes the classloader to first attempt to load classes from its local classpath before delegating the classloading to its parent. This policy allows an application classloader to override and provide its own version of a class that exists in the parent classloader.

The following settings determine a classloader's mode:

- If the application classloader policy of an application server is SINGLE, the application classloader policy of an application server defines the mode for an application classloader.
- If the application classloader policy of an application server is MULTIPLE, the classloader mode of an application defines the mode for an application classloader.
- If the WAR classloader policy of an application is MODULE, the WAR classloader policy of a Web module defines the mode for a WAR classloader.

Classloader collection

Use this page to manage classloader instances on an application server. A classloader determines whether an application classloader or a parent classloader finds and loads Java class files for an application.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Classloader**.

Classloader ID

States a string unique to the server identifying the classloader instance. The product assigns the identifier.

Classloader Mode

Specifies the classloader mode when the application classloader policy is SINGLE. PARENT_FIRST causes the classloader to delegate the loading of classes to its parent classloader before attempting to load the class from its local classpath. PARENT_LAST causes the classloader to attempt to load classes from its local classpath before delegating the classloading to its parent; this allows an application classloader to override and provide its own version of a class that exists in the parent classloader.

Classloader settings

Use this page to configure a classloader for applications that reside on an application server.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Classloader > *classloader_ID***.

Classloader ID

States a string unique to the server identifying the classloader instance. The product assigns the identifier.

Data type	String
-----------	--------

Classloader Mode

Specifies the classloader mode when the application classloader policy is SINGLE. PARENT_FIRST causes the classloader to delegate the loading of classes to its parent classloader before attempting to load the class from its local classpath. PARENT_LAST causes the classloader to attempt to load classes from its local classpath before delegating the classloading to its parent; this allows an application classloader to override and provide its own version of a class that exists in the parent classloader.

Data type	String
Default	PARENT_FIRST

Migrating the classloader Module Visibility Mode setting

WebSphere Application Server Version 4.0.x had a server-wide configuration setting called **Module Visibility Mode**. For Version 5.0, you use application or WAR classloader policies instead of module visibility modes. The Version 5.0 policies provide additional flexibility because you can configure applications running in a server for an application classloader policy of SINGLE or MULTIPLE and for a WAR classloader policy of APPLICATION or MODULE.

To migrate module visibility modes in your Version 4.0.x applications to their equivalents in Version 5.0, change the settings for your Version 4.0.x applications and modules to the Version 5.0 values shown in the table below.

Version 4.0.x module visibility mode	Version 5.0 application classloader policy	Version 5.0 WAR classloader policy
Server	SINGLE	APPLICATION
Compatibility	SINGLE	MODULE
Application	MULTIPLE	APPLICATION
Module*	MULTIPLE	MODULE
J2EE	MULTIPLE	MODULE

*There is no exact equivalent for the Version 4.0.x Module mode because it isolated EJB modules within an application.

Classloading: Resources for learning


Use the following links to find relevant supplemental information about classloaders. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.



View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming model and decisions
- Programming instructions and examples
- Programming specifications



Planning, business scenarios, and IT architecture

-  **WebSphere Application Server V5.0: Architecture and Overview**
http://developerworks.cybercentral.com/ibm0502/amt/ibmpresentations/683_1.pdf

Programming model and decisions

-  **J2EE Class Loading Demystified**
http://www7b.boulder.ibm.com/wsdd/library/techarticles/0112_deboer/deboer.html
-  **Understanding J2EE Application Server Class Loading Architectures**
<http://www.theserverside.com/resources/article.jsp?l=ClassLoading>
-  **Exploiting the Java Virtual Machine**
<http://www.develop.com/downloads/DevWPJav.pdf>



Programming instructions and examples

-  **Developing and Deploying Modular J2EE Applications with WebSphere Studio Application Developer and WebSphere Application Server**
http://www7b.boulder.ibm.com/wsdd/library/techarticles/0206_robinson/robinson.html
-  **Understanding the Java ClassLoader**

<http://www-105.ibm.com/developerworks/education.nsf/java-onlinecourse-bytitle/06B49359139A1AD186256A310049AC9B?OpenDocument>

- ✓ **IBM WebSphere Application Server Programming**
<http://www.mcgraw-hill.co.uk/html/0072224592.html>

Programming specifications

-  **Sun's J2EETM Platform Specification**
<http://java.sun.com/j2ee/download.html#platformspec>
-  **Sun's J2EETM Extension Mechanism Architecture**
<http://java.sun.com/j2se/1.4/docs/guide/extensions/spec.html>

Chapter 10. Using EJB query

The EJB query language is used to specify a query over container-managed entity beans. The language is similar to SQL. An EJB query is independent of the bean's mapping to a persistent store.

An EJB query can be used in three situations:

- To define a finder method of an EJB entity bean.
- To define a select method of an EJB entity bean.
- To dynamically specify a query using the `executeQuery()` dynamic API.

Finder and select queries are specified in the bean's deployment descriptor using the `<ejb-ql>` tag. Queries specified in the deployment descriptor are compiled into SQL during deployment. Dynamic queries require the interface provided by WebSphere Application Server Enterprise.

WebSphere's EJB query language is compliant with the EJB QL defined in Sun's EJB 2.0 specification and has additional capabilities as listed in the topic [Comparison of EJB 2.0 specification and WebSphere Query Language](#).

In your WebSphere application, you can define an EJB query in the following ways:

- **Application Assembly Tool.** When assembling an EJB 2.0 entity bean, specify the `<ejb-ql>` tag for the `finder()` or `select()` method.
- **WebSphere Studio Application Developer.** When defining an entity bean, specify the `<ejb-ql>` tag for the `finder` or `select` method.
- **Dynamic query service.** Add the `executeQuery()` method to your application. The dynamic query API is provided as an Enterprise Extension to WebSphere Application Server.

Before using EJB query, familiarize yourself with query language concepts, starting with the topic, [EJB Query Language](#).

Usage scenario

See the topic ["Example: EJB queries"](#).

EJB query language

An EJB query is a string that contains the following elements:

- a `SELECT` clause that specifies the EJBs or values to return;
- a `FROM` clause that names the bean collections;
- an optional `WHERE` clause that contains search predicates over the collections;
- an optional `GROUP BY` and `HAVING` clause (see [Aggregation functions](#));
- an optional `ORDER BY` clause that specifies the ordering of the result collection.

The `SELECT` clause is optional in order to maintain compatibility with WebSphere Application Server Version 4.

Collections of entity beans are identified in EJB queries through the use of their abstract schema name in the query FROM clause.

The elements of EJB query language are discussed in more detail in the following related topics.

Example: EJB queries

Here is an example EJB schema, followed by a set of example queries:

DeptBean schema

Entity bean name (EJB name)	DeptEJB (not used in query)
Abstract schema name	DeptBean
Implementation class	com.acme.hr.deptBean (not used in query)
Persistent attributes (cmp fields)	<ul style="list-style-type: none"> deptno - Integer (key) name - String budget - BigDecimal
Relationships	<ul style="list-style-type: none"> emps - 1:Many with EmpEJB mgr - Many:1 with EmpEJB

EmpBean schema

Entity bean name (EJB name)	EmpEJB (not used in query)
Abstract schema name	EmpBean
Implementation class	com.acme.hr.empBean (not used in query)
Persistent attributes (cmp fields)	<ul style="list-style-type: none"> empid - Integer (key) name - String salary - BigDecimal bonus - BigDecimal hireDate - java.sql.Date birthDate - java.util.Calendar address - com.acme.hr.Address
Relationships	<ul style="list-style-type: none"> dept - Many:1 with DeptEJB manages - 1:Many with DeptEJB

Address is a serializable object used as cmp field in EmpBean. The definition of address is as follows:

```
public class com.acme.hr.Address extends Object implements Serializable {
    public String street;
    public String state;
    public String city;
    public integer zip;
    public double distance (String start_location) { ... };
    public String format ( ) { ... };
}
```

The following query returns all departments:

```
SELECT OBJECT(d) FROM DeptBean d
```

The following query returns departments whose name begins with the letters "Web". Sort the result by name:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.name LIKE 'Web%' ORDER BY d.name
```

The keywords SELECT and FROM are shown in uppercase in the examples but are case insensitive. If a name used in a query is a reserved word, the name must be enclosed in double quotes to be used in the query. There is a list of reserved words later in this document. Identifiers enclosed in double quotes are case sensitive. This example shows how to use a cmp field that is a reserved word:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d."select" > 5
```

The following query returns all employees who are managed by Bob. This example shows how to navigate relationships using a path expression:

```
SELECT OBJECT (e) FROM EmpBean e WHERE e.dept.mgr.name='Bob'
```

A query can contain a parameter which refers to the corresponding value of the finder or select method. Query parameters are numbered starting with 1:

```
SELECT OBJECT (e) FROM EmpBean e WHERE e.dept.mgr.name= ?1
```

This query shows navigation of a multivalued relationship and returns all departments that have an employee that earns at least 50000 but not more than 90000:

```
SELECT OBJECT(d) FROM DeptBean d, IN (d.emps) AS e  
WHERE e.salary BETWEEN 50000 and 90000
```

There is a join operation implied in this query between each department object and its related collection of employees. If a department has no employees, the department does not appear in the result. If a department has more than one employee that earns more than 50000, that department appears multiple times in the result.

The following query eliminates the duplicate departments:

```
SELECT DISTINCT OBJECT(d) from DeptBean d, IN (d.emps) AS e WHERE e.salary > 50000
```

Find employees whose bonus is more than 40% of their salary:

```
SELECT OBJECT(e) FROM EmpBean e where e.bonus > 0.40 * e.salary
```

Find departments where the sum of salary and bonus of employees in the department exceeds the department budget:

```
SELECT OBJECT(d) FROM DeptBean d where d.budget <  
( SELECT SUM(e.salary+e.bonus) FROM IN(d.emps) AS e )
```

A query can contain DB2 style date-time arithmetic expressions if you use java.sql.* datatypes as CMP fields and your datastore is DB2. Find all employees who have worked at least 20 years as of January 1st, 2000:

```
SELECT OBJECT(e) FROM EmpBean e where year( '2000-01-01' - e.hireDate ) >= 20
```

If the datastore is not DB2 or if you prefer to use java.util.Calendar as the CMP field, then you can use the java millisecond value in queries. The following query finds all employees born before Jan 1, 1990:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.birthDate < 631180800232
```

Find departments with no employees:

```
SELECT OBJECT(d) from DeptBean d where d.emps IS EMPTY
```


Find all employees whose earn more than Bob:

```
SELECT OBJECT(e) FROM EmpBean e, EmpBean b
WHERE b.name = 'Bob' AND e.salary + e.bonus > b.salary + b.bonus
```

Find the employee with the largest bonus:

```
SELECT OBJECT(e) from EmpBean e WHERE e.bonus =
(SELECT MAX(e1.bonus) from EmpBean e1)
```

The above queries all return EJB objects. A finder method query must always return an EJB Object for the home. A select method query can in addition return CMP fields or other EJB Objects not belonging to the home.

The following would be valid select method queries for EmpBean. Return the manager for each department:

```
SELECT d.mgr FROM DeptBean d
```

Return department 42 manager's name:

```
SELECT d.mgr.name FROM DeptBean d WHERE d.deptno = 42
```

Return the names of employees in department 42:

```
SELECT e.name FROM EmpBean e WHERE e.dept.deptno=42
```

Another way to write the same query is:

```
SELECT e.name from DeptBean d, IN (d.emps) AS e WHERE d.deptno=42
```

Finder and select queries allow only a single CMP field or EJBObject in the SELECT clause.

FROM clause

The FROM clause specifies the collections of objects to which the query is to be applied. Each collection is identified either by an abstract schema name and an identification variable, called a range variable, or by a collection member declaration that identifies a multivalued relationship and an identification variable.

Conceptually, the semantics of the query is to first form a temporary collection of tuples R. Tuples are composed of elements from the collections identified in the FROM clause. Each tuple contains one element from each of the collections in the FROM clause. All possible combinations are formed subject to the constraints imposed by the collection member declarations. If any schema name identifies a collection for which there are no records in the persistent store, then the temporary collection R will be empty.

Example: FROM clause

DeptBean contains records 10, 20 and 30 in the persistent store. EmpBean contains records 1, 2 and 3 that are related to department 10 and records 4, 5 that are related to department 20. Department 30 has no related employees.

```
FROM DeptBean d, EmpBean e
```

This forms a temporary collection R that contains 15 tuples.

```
FROM DeptBean d, DeptBean d1
```

This forms a temporary collection R that contains 9 tuples.

```
FROM DeptBean d, IN (d.emps) AS e
```

This forms a temporary collection R that contains 5 tuples. Department 30 because it contains no employees will not be in R. Department 10 will be contained in R three times and department 20 will be contained in R twice.

After forming the temporary collection the search conditions of the WHERE clause will be applied to R and this will yield a new temporary collection R1. The ORDER BY and SELECT clauses are applied to R1 to yield the final result set.

An identification variable is a variable declared in the FROM clause using the operator IN or the optional AS.

```
FROM DeptBean AS d, IN (d.emps) AS e
```

is equivalent to:

```
FROM DeptBean d, IN (d.emps) e
```

An identification variable that is declared to be an abstract schema name is called a range variable. In the query above "d" is a range variable. An identification variable that is declared to be a multivalued path expression is called a collection member declaration. "d" and "e" in the example above are collection member declarations.

Note that the following path expression is illegal as a collection member declaration because it is not multivalued:

```
e.dept.mgr
```

Inheritance in EJB query

If an EJB inheritance hierarchy has been defined for an abstract schema, using the abstract schema name in a query statement implies the collection of objects for that abstract schema as well as all subtypes.

Example: Inheritance

Suppose that bean ManagerBean is defined as a subtype of EmpBean and ExecutiveBean is a subtype of ManagerBean in an EJB inheritance hierarchy. The following query returns employees as well as managers and executives:

```
SELECT OBJECT(e) FROM EmpBean e
```

Path expressions

An identification variable followed by the navigation operator (.) and a cmp or relationship name is a path expression.

A path expression that leads to a cmr field can be further navigated if the cmr field is single-valued. If the path expression leads to a multi-valued relationship, then the path expression is terminal and cannot be further navigated. If the path expression leads to a cmp field whose type is a value object, it is possible to navigate to attributes of the value object.

If the path expression leads to a cmp field whose type is a value object it is possible to further navigate to attributes and methods of the value object.

Example: Value object

Assume that address is a cmp field for EmpBean, which is a value object.

```
SELECT object(e) FROM EmpBean e
WHERE e.address.distance('San Jose') < 10 and e.address.zip = 95037
```

It is best to use the composer pattern to map value object attributes to relational columns if you intend to search on value attributes. If you store value objects in serialized format, then each value object must be retrieved from the database and deserialized. Value object methods can only be done in dynamic queries.

A path expression can also navigate to a bean method. The method must be defined on either the remote or local bean interface. Methods can only be used in dynamic queries. You cannot mix both remote and local methods in a single query statement.

If the query contains remote methods, the dynamic query must be executed using the query remote interface. Using the query remote interface causes the query service to activate beans and create instances of the remote bean interface

Likewise, a query statement with local bean methods must be executed with the query local interface. This causes the query service to activate beans and local interface instances.

Do not use get methods to access cmp and cmr fields of a bean.

If a method has overloaded definitions, the overloaded methods must have different number of parameters.

Methods must have non-void return types and method arguments and return types must be either primitive types byte, short, int, long, float, double, boolean, char or wrapper types from the following list:

Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Character, java.util.Calendar, java.sql.Date, java.sql.Time, java.sql.Timestamp, java.util.Date

If any input argument to a method is NULL, it is assumed the method returns a NULL value and the method is not invoked.

A collection valued path expression can be used in the FROM clause as a collection member declaration, and with the IS EMPTY, MEMBER OF, and EXISTS predicates in the WHERE clause.

FROM EmpBean e WHERE e.dept.mgr.name='Bob'	OK
FROM EmpBean e WHERE e.dept.emps.name='BOB'	INVALID — cannot navigate through emps because it is multivalued
FROM EmpBean e, IN (e.dept.emps) e1 WHERE e1.name='BOB'	OK
FROM EmpBean e WHERE e.dept.emps IS EMPTY	OK

WHERE clause

The WHERE clause contains search conditions composed of the following:

- literal values
- input parameters
- expressions

- basic predicates
- quantified predicates
- BETWEEN predicate
- IN predicate
- LIKE predicate
- NULL predicate
- EMPTY collection predicate
- MEMBER OF predicate
- EXISTS predicate
- IS OF TYPE predicate

If the search condition evaluates to TRUE, the tuple is added to the result set.

Literals

A string literal is enclosed in single quotes. A single quote that occurs within a string literal is represented by two single quotes; For example: 'Tom''s'. A string literal cannot exceed the maximum length that is supported by the underlying persistent datastore.

A numeric literal can be any of the following:

- an exact value such as 57, -957, +66
- any value supported by Java long
- a decimal literal such as 57.5, -47.02
- an approximate numeric value such as 7E3, -57.4E-2

A decimal or approximate numeric value must be in the range supported by Java double.

A boolean literal can be the keyword TRUE or FALSE and is case insensitive.

Input parameters

Input parameters are designated by the question mark followed by a number; For example: ?2

Input parameters are numbered starting at 1 and correspond to the arguments of the finder or select method; therefore, a query must not contain an input parameter that exceeds the number of input arguments.

An input parameter can be a primitive type of byte, short, int, long, float, double, boolean, char or wrapper types of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Char, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp or an EJBObject.

An input parameter must not have a NULL value. To search for the occurrence of a NULL value the NULL predicate should be used.

Expressions

Conditional expressions can consist of comparison operators and logical operators (AND, OR, NOT).

Arithmetic expressions can be used in comparison expressions and can be composed of arithmetic operations and functions, path expressions that evaluate to a numeric value and numeric literals and numeric input parameters.

String expressions can be used in comparison expressions and can be composed of string functions, path expressions that evaluate to a string value and string literals and string input parameters. A cmp field of type char is handled as if it were a string of length 1.

Boolean expressions can be used with = and <> comparison and can be composed of path expressions that evaluate to a boolean value and TRUE and FALSE keywords and boolean input parameters.

Reference expressions can be used with = and <> comparison and can be composed of path expressions that evaluate to a cmr field, an identification variable and an input parameter whose type is an EJB reference

Four different expression types are supported for working with date-time types. For portability the java.util.Calendar type should be used. DB2 style date, time and timestamp expressions are supported if the datastore is DB2 and the CMP field is of type java.util.Date, java.sql.Date, java.sql.Time or java.sql.Timestamp.

A Calendar type can be compared to another Calendar type, an exact numeric literal or input parameter of type long whose value is the standard Java long millisecond value.

The following query finds all employees born before Jan 1, 1990:
SELECT OBJECT(e) FROM EmpBean e WHERE e.birthDate < 631180800232

Date expressions can be used in comparison expressions and can be composed of operators + - , date duration expressions and date functions, path expressions that evaluate to a date value, string representation of a date and date input parameters.

Time expressions can be used in comparison expressions and can be composed of operators + - , time duration expressions and time functions, path expressions that evaluate to a time value, string representation of time and time input parameters.

Timestamp expressions can be used in comparison expressions and can be composed of operators + - , timestamp duration expressions and timestamp functions, path expressions that evaluate to a timestamp value, string representation of a timestamp and timestamp input parameters.

Standard bracketing () for ordering expression evaluation is supported.

The operators and their precedence order from highest to lowest are:

- Navigation operator (.)
- Arithmetic operators in precedence order:
 - + - unary
 - * / multiply, divide
 - + - add, subtract
- Comparison operators: =, >, <, >=, <=, <>(not equal)
- Logical operator NOT
- Logical operator AND
- Logical operator OR

Null value semantics: The following describe the semantics of NULL values:

- Comparison or arithmetic operations with an unknown (NULL) value yield an unknown value
- Path expressions that contain NULL evaluate to NULL
- The IS NULL and IS NOT NULL operators can be applied to path expressions and return TRUE or FALSE. Boolean operators AND, OR and NOT use three valued logic.

AND	True	False	Unknown
True	True	False	Unknown
False	False	False	False
Unknown	Unknown	False	Unknown

OR	True	False	Unknown
True	True	True	True
False	True	False	Unknown
Unknown	True	Unknown	Unknown

	NOT
True	False
False	True
Unknown	Unknown

Example: Null value semantics

```
select object(e) from EmpBean where e.salary > 10 and e.dept.budget > 100
```

If salary is NULL the evaluation of `e.salary > 10` returns unknown and the employee object is not returned. If the cmr field dept or budget is NULL evaluation of `e.dept.budget > 100` returns unknown and the employee object is not returned.

```
select object(e) from EmpBean where e.dept.budget is null
```

If dept or budget is NULL evaluation of `e.dept.budget is null` returns TRUE and the employee object is returned.

```
select object(e) from EmpBean e , in (e.dept.emps) e1 where e1.salary > 10
```

If dept is NULL, then the multivalued path expression `e.dept.emps` results in an empty collection (not a collection that contains a NULL value). An employee with a null dept value will not be returned.

```
select object(e) from EmpBean e where e.dept.emps is empty
```

If dept is NULL the evaluation of the predicate in unknown and the employee object is not returned.

```
select object(e) from EmpBean e , EmpBean e1 where e member of e1.dept.emps
```

If dept is NULL evaluation of the member of predicate returns unknown and the employee is not returned.

Date time arithmetic and comparisons: DATE, TIME and TIMESTAMP values may be compared with another value of the same type. Comparisons are chronological. Date time values can also be incremented, decremented, and subtracted.

If the datastore is DB2, then DB2 string representation of DATE, TIME and TIMESTAMP types can also be used. A string representation of a date or time can use ISO, USA, EUR or JIS format. A string representation of a timestamp uses ISO format.

Format	Date format	Date examples	Time format	Time examples
ISO	yyyy-mm-dd	1987-02-24 1987-2-24	hh.mm.ss	13.50.00 13.50
USA	mm/dd/yyyy	2/24/1987	hh:mm AM or PM	1:50 pm 02:10 AM
EUR	dd.mm.yyyy	24.02.1987 24.2.1987	hh.mm.ss	13.50.00 13.55
JIS	yyyy-mm-dd	1987-02-24	hh:mm:ss	13:50 13:50:05

Example 1: Date time arithmetic comparisons

```
e.hiredate > '1990-02-24'
```

The timestamp of February 24th, 1990 1:50 pm can be represented as follows:

```
'1990-02-24-13.50.00.000000' or  
'1990-02-24-13.50.00'
```

If the datastore is DB2, DB2 decimal durations can be used in expressions and comparisons. A date duration is a decimal(8,0) number that represents the difference between two dates in the format YYYYMMDD. A time duration is a decimal(6,0) number that represents the difference between two time values as HHMMSS. A timestamp duration is a decimal(20,6) number representing the differences between two timestamp values as YYYYMMDDHHMMSS.ZZZZZZ (ZZZZZZ is the number of microseconds and is to the right of the decimal point) .

Two date values (or time values or timestamp values) can be subtracted to yield a duration. If the second operand is greater than the first the duration is a negative decimal number. A duration can be added or subtracted from a datetime value to yield a new datetime value.

Example 2: Date time arithmetic comparisons

DATE('3/15/2000') - '12/31/1999' results in a decimal number 215 which is a duration of 0 years, 2 months and 15 days.

Durations are really decimal numbers and can be used in arithmetic expressions and comparisons.

```
( DATE('3/15/2000') - '12/31/1999' ) + 14 > 215 evaluates to TRUE.
```

```
DATE('12/31/1999') + DECIMAL(215,8,0) results in a date value 3/15/2000.
```

TIME('11:02:26') - '00:32:56' results in a decimal number 102930 which is a time duration of 10 hours, 29 minutes and 30 seconds.

```
TIME('00:32:56') + DECIMAL(102930,6,0) results in a time value of 11:02:26.
```

```
TIME('00:00:59') + DECIMAL(240000,6,0) results in a time value of 00:00:59.
```

```
e.hiredate + DECIMAL(500,8,0) > '2000-10-01' means compare the hiredate plus 5 months to the date 10/01/2000.
```

Basic predicates

Basic predicates can be of two forms

```
expression-1 comparison-operator expression-2  
expression-3 comparison-operator ( subselect )
```

The subselect must not return more than one value and the subselect can not return a type of an EJB reference. Boolean types and reference types only support = and <> comparisons.

Example: Basic predicates

```
d.name='Java Development'  
e.salary > 20000  
e.salary > ( select avg(e.salary) from EmpBean e)
```

Quantified predicates

A quantified predicate compares a value with a set of values produced by a subselect.

```
expression comparison-operator SOME | ANY | ALL ( subselect )
```

The expression must not evaluate to a reference type.

When SOME or ANY is specified the result of the predicate is as follows:

- TRUE if the comparison is true for at least one value returned by the subselect.
- FALSE if the subselect is empty or if the comparison is false for every value returned by the subselect.
- UNKNOWN if the comparison is not true for all of the values returned by the subselect and at least one of the comparisons is unknown because of a null value.

When ALL is specified the result of the predicate is as follows:

- TRUE if the subselect returns empty or if the comparison is true to every value returned by the subselect.
- FALSE if the comparison is false for at least one value returned by the subselect.
- UNKNOWN if the comparison is not false for all values returned by the subselect and at least one comparison is unknown because of a null value.

BETWEEN predicate

The BETWEEN predicate determines whether a given value lies between two other given values.

```
expression [NOT] BETWEEN expression-2 AND expression-3
```

Example: BETWEEN predicate

```
e.salary BETWEEN 50000 AND 60000
```

is equivalent to:

```
e.salary >= 50000 AND e.salary <= 60000  
e.name NOT BETWEEN 'A' AND 'B'
```

is equivalent to:

```
e.name < 'A' OR e.name > 'B'
```

IN predicate

The IN predicate compares a value to a set of values and can have one of two forms:


```
expression [NOT] IN ( subselect )
expression [NOT] IN ( value1, value2, .... )
```

ValueN can either be a literal value or an input parameter. The expression can not evaluate to a reference type.

Example: IN predicate

```
e.salary IN ( 10000, 15000 )
```

is equivalent to

```
( e.salary = 10000 OR e.salary = 15000 )
e.salary IN ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary = ANY ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
e.salary NOT IN ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary <> ALL ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

LIKE predicate

The LIKE predicate searches a string value for a certain pattern.

```
string-expression [NOT] LIKE pattern [ ESCAPE escape-character ]
```

The pattern value is a string literal or parameter marker of type string in which the underscore (`_`) stands for any single character and percent (`%`) stands for any sequence of characters (including empty sequence). Any other character stands for itself. The escape character can be used to search for character `_` and `%`. The escape character can be specified as a string literal or an input parameter.

If the string-expression is null, then the result is unknown.

If both string-expression and pattern are empty, then the result is true.

Example: LIKE predicate

- `'' LIKE ''` is true
- `'' LIKE '%'` is true
- `e.name LIKE '12%3'` is true for `'123'` `'12993'` and false for `'1234'`
- `e.name LIKE 's_me'` is true for `'some'` and `'same'`, false for `'soome'`
- `e.name LIKE '/_foo'` escape `'/'` is true for `'_foo'`, false for `'afoo'`
- `e.name LIKE '//_foo'` escape `'/'` is true for `'/afoo'` and for `'/bfoo'`
- `e.anme LIKE '///_foo'` escape `'/'` is true for `'/_foo'` but false for `'/afoo'`

NULL predicate

The NULL predicate tests for null values.

```
single-valued-path-expression IS [NOT] NULL
```

Example: NULL predicate

```
e.name IS NULL
e.dept.name IS NOT NULL
e.dept IS NOT NULL
```

EMPTY collection predicate

To test if a multivalued relationship is empty, use the following syntax:

```
collection-valued-path-expression IS [NOT] EMPTY
```

Example: Empty collection predicate

To find all departments with no employees:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.emps IS EMPTY
```

MEMBER OF predicate

This expression tests whether the object reference specified by the single valued path expression or input parameter is a member of the designated collection. If the collection valued path expression designates an empty collection the value of the MEMBER OF expression is FALSE.

```
{ single-valued-path-expression | input_parameter } [ NOT ] MEMBER [ OF ]  
collection-valued-path-expression
```

Example: MEMBER OF predicate

Find employees that are not members of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e , DeptBean d  
WHERE e NOT MEMBER OF d.emps AND d.deptno = ?1
```

Find employees whose manager is a member of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e, DeptBean d  
WHERE e.dept.mgr MEMBER OF d.emps and d.deptno=?1
```

EXISTS predicate

The exists predicate tests for the presence or absence of a condition specified by a subselect.

```
EXISTS ( subselect )  
EXISTS collection-valued-path-expression
```

The result of EXISTS is true if the subselect returns at least one value or the path expression evaluates to a nonempty collection, otherwise the result is false.

To negate an EXISTS predicate, precede it with the logical operator NOT.

Example: EXISTS predicate

Return departments that have at least one employee earning more than 1000000:

```
SELECT OBJECT(d) FROM DeptBean d  
WHERE EXISTS ( SELECT 1 FROM IN (d.emps) e WHERE e.salary > 1000000 )
```

Return departments that have no employees:

```
SELECT OBJECT(d) FROM DeptBean d  
WHERE NOT EXISTS ( SELECT 1 FROM IN (d.emps) e)
```

The above query can also be written as follows:

```
SELECT OBJECT(d) FROM DeptBean d WHERE NOT EXISTS d.emps
```

IS OF TYPE predicate

The IS OF TYPE predicate is used to test the type of an EJB reference. It is similar in function to the Java instance of operator. IS OF TYPE is used when several abstract beans have been grouped into an EJB inheritance hierarchy. The type

names specified in the predicate are the bean abstract names. The ONLY option can be used to specify that the reference must be exactly this type and not a subtype.

```
identification-variable IS OF TYPE ( [ONLY] type-1, [ONLY] type-2, ..... )
```

Example: IS OF TYPE predicate

Suppose that bean ManagerBean is defined as a subtype of EmpBean and ExecutiveBean is a subtype of ManagerBean in an EJB inheritance hierarchy.

The following query returns employees as well as managers and executives:

```
SELECT OBJECT(e) FROM EmpBean e
```

If you are interested in objects which are employees and not managers and not executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE( ONLY EmpBean )
```

If you are interested in object which are managers or executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE( ManagerBean)
```

The above query is equivalent to the following query:

```
SELECT OBJECT(e) FROM ManagerBean e
```

If you are interested in managers only and not executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE( ONLY ManagerBean)
```

or:

```
SELECT OBJECT(e) FROM ManagerBean e  
WHERE e IS OF TYPE (ONLY ManagerBean)
```

Scalar functions

EJB query contains scalar built-in functions for doing type conversions, string manipulation, and for manipulating date-time values. The list of scalar functions is documented in the topic EJB query: Scalar functions.

Example: Scalar functions

Find employees hired in 1999:

```
SELECT OBJECT(e) FROM EmpBean e where YEAR(e.hireDate) = 1999
```

The only scalar functions that are guaranteed to be portable across backend datastore vendors are the following:

- ABS
- SQRT
- CONCAT
- LENGTH
- LOCATE
- SUBSTRING

The other scalar functions should be used only when DB2 is the backend datastore.

EJB query: Scalar functions

EJB query contains scalar built-in functions, as listed below, for doing type conversions, string manipulation, and for manipulating date-time values.

Numeric functions

ABS (< any numeric datatype >) -> < any numeric datatype >

SQRT (< any numeric datatype >) -> Double

Type conversion functions

CHAR (< any numeric datatype >) -> string

CHAR (< string >) -> string

CHAR (< any datetime datatype > [, Keyword k]) -> string

Datetime datatype is converted to its string representation in a format specified by the keyword k. The valid keywords values are ISO, USA, EUR or JIS. If k is not specified the default is ISO.

BIGINT (< any numeric datatype >) -> Long

BIGINT (< string >) -> Long

The following function converts the argument to an integer n by truncation and returns the date that is n-1 days after January 1, 0001:

DATE (< date string >) -> Date

DATE (< any numeric datatype>) -> Date

The following function returns date portion of a timestamp:

DATE(timestamp) -> Date

DATE (< timestamp-string >) -> Date

The following function converts number to decimal with optional precision p and scale s.

DECIMAL (< any numeric datatype > [, p [,s]]) -> Decimal

The following function converts string to decimal with optional precision p and scale s.

DECIMAL (< string > [, p [, s]]) -> Decimal

DOUBLE (< any numeric datatype >) -> Double

DOUBLE (< string >) -> Double

FLOAT (< any numeric datatype >) -> Double

FLOAT (< string >) -> Double

Float is a synonym for DOUBLE.

INTEGER (< any numeric datatype >) -> Integer

INTEGER (< string >) -> Integer

REAL (< any numeric datatype >) -> Float

SMALLINT (< any numeric datatype >) -> Short

SMALLINT (< string >) -> Short

TIME (< time >) -> Time

TIME (< time-string >) -> Time

TIME (< timestamp >) -> Time

TIME (< timestamp-string >) -> Time

TIMESTAMP (< timestamp >) -> Timestamp

TIMESTAMP (< timestamp-string >) -> Timestamp

String functions

CONCAT (<string>, <string>) -> String

The following function returns a character string representing absolute value of the argument not including its sign or decimal point. For example, `digits(-42.35)` is "4235".

```
DIGITS ( Decimal d ) -> String
```

The following function returns the length of the argument in bytes. If the argument is a numeric or datetime type, it returns the length of internal representation.

```
LENGTH ( < string > ) -> Integer
```

The following function returns a copy of the argument string where all upper case characters have been converted to lower case.

```
LCASE ( < string > ) -> String
```

The following function returns the starting position of the first occurrence of argument 1 inside argument 2 with optional start position. If not found, it returns 0.

```
LOCATE ( String s1 , String s2 [ , Integer start ] ) -> Integer
```

The following function returns a substring of `s` beginning at character `m` and containing `n` characters. If `n` is omitted, the substring contains the remainder of string `s`. The result string is padded with blanks if needed to make a string of length `n`.

```
SUBSTRING ( String s , Integer m [ , Integer n ] ) -> String
```

The following function returns a copy of the argument string where all lower case characters have been converted to upper case.

```
UCASE ( < string > ) -> String
```

Date - time functions

The following function returns the day portion of its argument. For a duration, the return value can be -99 to 99.

```
DAY ( Date ) -> Integer  
DAY ( < date-string > ) -> Integer  
DAY ( < date-duration > ) -> Integer  
DAY ( Timestamp ) -> Integer  
DAY ( < timestamp-string > ) -> Integer  
DAY ( < timestamp-duration > ) -> Integer
```

The following function returns one more than number of days from January 1, 0001 to its argument.

```
DAYS ( Date ) -> Integer  
DAYS ( < Date-string > ) -> Integer  
DAYS ( Timestamp ) -> Integer  
DAYS ( < timestamp-string > ) -> Integer
```

The following function returns the hour part of its argument. For a duration, the return value can be -99 to 99.

```
HOUR ( Time ) -> Integer  
HOUR ( < time-string > ) -> Integer  
HOUR ( < time-duration > ) -> Integer  
HOUR ( Timestamp ) -> Integer  
HOUR ( < timestamp-string > ) -> Integer  
HOUR ( < timestamp-duration > ) -> Integer
```

The following function returns the microsecond part of its argument.

```
MICROSECOND ( Timestamp ) -> Integer
MICROSECOND ( < timestamp-string > ) -> Integer
MICROSECOND ( < timestamp-duration > ) -> Integer
```

The following function returns the minute part of its argument. For a duration, the return value can be -99 to 99.

```
MINUTE ( Time ) -> Integer
MINUTE ( < time-string > ) -> Integer
MINUTE ( < time-duration > ) -> Integer
MINUTE ( Timestamp ) -> Integer
MINUTE ( < timestamp-string > ) -> Integer
MINUTE ( < timestamp-duration > ) -> Integer
```

The following function returns the month portion of its argument. For a duration, the return value can be -99 to 99.

```
MONTH ( Date ) -> Integer
MONTH ( < date-string > ) -> Integer
MONTH ( < date-duration > ) -> Integer
MONTH ( Timestamp ) -> Integer
MONTH ( < timestamp-string > ) -> Integer
MONTH ( < timestamp-duration > ) -> Integer
```

The following function returns the second part of its argument. For a duration, the return value can be -99 to 99.

```
SECOND ( Time ) -> Integer
SECOND ( < time-string > ) -> Integer
SECOND ( < time-duration > ) -> Integer
SECOND ( Timestamp ) -> Integer
SECOND ( < timestamp-string > ) -> Integer
SECOND ( < timestamp-duration > ) -> Integer
```

The following function returns the year portion of its argument. For a duration, the return value can be -9999 to 9999.

```
YEAR ( Date ) -> Integer
YEAR ( < date-string > ) -> Integer
YEAR ( < date-duration > ) -> Integer
YEAR ( Timestamp ) -> Integer
YEAR ( < timestamp-string > ) -> Integer
YEAR ( < timestamp-duration > ) -> Integer
```

Aggregation functions

Queries that return aggregate values can only be used with the dynamic query interface available in WebSphere Application Server Enterprise. However, aggregation functions can be used in non-dynamic queries if the aggregation function is used in a subselect or HAVING clause.

Aggregation functions operate on a set of values to return a single scalar value. The following is an example of an aggregation:

```
SELECT SUM (e.salary + e.bonus) FROM EmpBean e WHERE e.dept.deptno =20
```

This computes the total salary and bonus for department 20.

The aggregation functions are avg, count, max, min and sum. The syntax of an aggregation function is as follows:

```
aggregation-function ( [ ALL | DISTINCT ] expression )
```

or:

```
COUNT( * )
```

The DISTINCT option eliminates duplicate values before applying the function. ALL is the default and does not eliminate duplicates. Null values are ignored in computing the aggregate function except for COUNT(*) which returns a count of all elements in the set.

MAX and MIN can apply to any numeric, string or datetime datatype and returns the same datatype. SUM and AVG take a numeric type as input. SUM and AVG return numeric type. The actual numeric type returned by SUM and AVG depends on the underlying datastore. COUNT can take any datatype and returns an integer.

The set of values that is used for the aggregate function is determined by the collection that results from the FROM and WHERE clause of the subquery. This set can be divided into groups and the aggregation function applied to each group. This is done by using a GROUP BY clause in the subquery. The GROUP BY clause defines grouping members which is a list of path expressions. Each path expression specifies a field that is a primitive type of byte, short, int, long, float, double, boolean, char, or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Character, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time or java.sql.Timestamp.

Finder or select queries can not return aggregation function values. In other words, aggregation functions can not appear in the top level SELECT of a finder or select query but can be used in subqueries.

Example: Aggregation functions

```
SELECT e.dept.deptno, AVG ( e.salary) FROM EmpBean e GROUP BY e.dept.deptno
```

The above query computes the average salary for each department.

In dividing a set into groups, a NULL value is considered equal to another NULL value.

Just as the WHERE clause filters tuples from the FROM clause, the groups can be filtered using a HAVING clause that tests group properties involving aggregate functions or grouping members:

```
SELECT e.dept.deptno, AVG ( e.salary) FROM EmpBean e
GROUP BY e.dept.deptno
HAVING COUNT(*) > 3 AND e.dept.deptno > 5
```

This query returns average salary for departments that have more than 3 employees and the department number is greater than 5.

It is possible to have a HAVING clause without a GROUP BY clause in which case the entire set is treated as a single group to which the HAVING clause is applied.

SELECT clause

For finder and select queries, the syntax of the SELECT clause is as follows:

```
SELECT [ ALL | DISTINCT ]
{ single-valued-path-expression | OBJECT ( identification-variable ) }
```

The SELECT clause consists of either a single identification variable that is defined in the FROM clause or a single valued path expression that evaluates to a object reference or CMP value. The keyword DISTINCT can be used to eliminate duplicate references.

For a query that defines a finder method the query must return an object type consistent with the home for which the finder method associated with the query. In other words, a finder method for a department home can not return employee objects.

A scalar-subselect is a subselect that returns a single value.

Example: SELECT clause

Find all employees that earn more than John:

```
SELECT OBJECT(e) FROM EmpBean ej, EmpBean e
WHERE ej.name = 'John' and e.salary > ej.salary
```

Find all departments that have one or more employees who earn less than 20000:

```
SELECT DISTINCT e.dept FROM EmpBean e where e.salary < 20000
```

A select method query can have a path expression that evaluates to an arbitrary value:

```
SELECT e.dept.name FROM EmpBean e where e.salary < 2000
```

The above query returns a collection of name values for those departments having employees earning less than 20000.

ORDER BY clause

The ORDER BY clause specifies an ordering of the objects in the result collection:

```
ORDER BY [ order_element ,]* order_element
order_element ::= { path-expression | integer } [ ASC | DESC ]
```

The path expression must specify a single valued field that is a primitive type of byte, short, int, long, float, double, char or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Character, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp.

ASC specifies ascending order and is the default. DESC specifies descending order.

Integer refers to a selection expression in the SELECT clause.

Example: ORDER BY clause

Return department objects in decreasing deptno order:

```
SELECT OBJECT(d) FROM DeptBean d ORDER BY d.deptno DESC
```

Return employee objects sorted by department number and name:

```
SELECT OBJECT(e) FROM EmpBean e ORDER BY e.dept.deptno ASC, e.name DESC
```

Subqueries

A subquery can be used in quantified predicates, EXISTS predicate or IN predicate. A subquery should only specify a single element in the SELECT clause. When a path expression appears in a subquery, the identification variable of the path expression must be defined either in the subquery, in one of the containing subqueries, or in the outer query. A scalar subquery is a subquery that returns one value. A scalar subquery can be used in a basic predicate and in the SELECT clause of a dynamic query.

Example: Subqueries

```
SELECT OBJECT(e) FROM EmpBean e
WHERE e.salary > ( SELECT AVG(e1.salary) FROM EmpBean e1)
```

The above query returns employees who earn more than average salary of all employees.

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.salary >
( SELECT AVG(e1.salary) FROM IN (e.dept.emps) e1 )
```

The above query returns employees who earn more than average salary of their department.

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.salary =
( SELECT MAX(e1.salary) FROM IN (e.dept.emps) e1 )
```

The above query returns employees who earn the most in their department.

```
SELECT OBJECT(e) FROM EmpBean e
WHERE e.salary > ( SELECT AVG(e.salary) FROM EmpBean e1
WHERE YEAR(e1.hireDate) = YEAR(e.hireDate) )
```

The above query returns employees who earn more than the average of employees hired in same year.

EJB query restrictions

An EJB query is compiled into an SQL query and executed against the underlying datastore based on schema mapping of the abstract bean to the datastore schema. The semantics of comparison and arithmetic operations are that of the underlying datastore. In the case of SQL, note that two strings are equal if the shorter string padded with blanks equals the longer string. For example, 'A' is equal to 'A '. This differs from the equality of strings in the Java language. Arithmetic overflow operations are an error in SQL.

A cmp field can not be used in comparison operations or predicates (except for LIKE) if that cmp field is mapped to a long varchar or LOB column or any other column type for which the database server does not support predicates or comparison operations.

A cmp field of any type can be used in a SELECT clause. Fields that can be used in predicates, grouping, or ordering operations must be of the types listed below:

- Primitive types : byte, short, int, long, float, double, boolean, char
- Object types: Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Char, java.util.Calendar , java.util.Date
- JDBC types: java.sql.Date, java.sql.Time, java.sql.Timestamp

The field must be mapped to a table column that is compatible in type either by using a "top-down" default mapping generated by the WebSphere deploy tool, or using a "meet-in-the-middle" mapping between compatible types.

In order to search on attributes of a cmp field that is a user-defined value object, you should use a "meet-in-the-middle" mapping and use a composer to map each attribute to a compatible column. The default "top-down" mapping stores the object as a serialized object in a column of type blob, which does not allow searching.

If a cmp field is mapped to a column using a "meet-in-the-middle" mapping with a converter, that field can only be used with the NULL predicate or with basic predicates of the following form:

```
path-expression <comparison> literal_value
path-expression <comparison> input_parameter
```

In this situation, the converter method `toData()` is called to convert the righthand side of the predicate to an SQL value.

Example of allowable predicate on a cmp field with user defined converter:

```
e.name = 'Chris'
e.name > ?1
e.name IS NULL
```

Examples of unallowable predicates:

```
substring( e.name, 1, 3 ) = 'ABC'
e.salary > d.budget
```

A converter should preserve equality, collating sequence and null values when doing a conversion. Otherwise cmp fields created by the converter should not be used in WHERE, GROUP, HAVING or ORDER clauses of a query.

EJB Query: Reserved words

The following words are reserved in WebSphere EJB query:

all, as, distinct, empty, false, from, group, having, in, is, like, select, true, union, where

Avoid using identifiers that start with underscore (for example, `_integer`) as these are also reserved.

EJB query: BNF syntax

```
EJB QL ::= [select_clause] from_clause [where_clause]
          [group_by_clause] [having_clause] [order_by_clause]
from_clause ::= FROM identification_variable_declaration
              [,identification_variable_declaration]*
identification_variable_declaration ::= collection_member_declaration |
                                       range_variable_declaration
collection_member_declaration ::=
    IN ( collection_valued_path_expression ) [AS] identifier
range_variable_declaration ::= abstract_schema_name [AS] identifier
single_valued_path_expression ::=
    {single_valued_navigation | identification_variable}. ( cmp_field |
    method | cmp_field.value_object_attribute |
    cmp_field.value_object_method )
    | single_valued_navigation
single_valued_navigation ::=
    identification_variable.[ single_valued_cmr_field. ]*
    single_valued_cmr_field
collection_valued_path_expression ::=
    identification_variable.[ single_valued_cmr_field. ]*
    collection_valued_cmr_field
select_clause ::=
    SELECT { ALL | DISTINCT } {single_valued_path_expression |
    identification_variable |
    OBJECT ( identification_variable ) }
```

```

select_clause_eex ::=
    SELECT { ALL | DISTINCT } [ selection , ]* selection
selection ::=
    { expression [[AS] id ] | subselect }
order_by_clause ::=
    ORDER BY [ {single_valued_path_expression | integer}
              [ASC|DESC],]*
              {single_valued_path_expression | integer}
              [ASC|DESC]
where_clause ::= WHERE conditional_expression
conditional_expression ::= conditional_term |
    conditional_expression OR conditional_term
conditional_term ::= conditional_factor |
    conditional_term AND conditional_factor
conditional_factor ::= [NOT] conditional_primary
conditional_primary ::= simple_cond_expression |
    (conditional_expression)
simple_cond_expression ::= comparison_expression |
    between_expression | like_expression |
    in_expression | null_comparison_expression |
    empty_collection_comparison_expression |
    quantified_expression | exists_expression |
    is_of_type_expression | collection_member_expression
between_expression ::= expression [NOT] BETWEEN
    expression AND expression
in_expression ::= single_valued_path_expression [NOT] IN
    { (subselect) | ( atom ,)* atom ) }
atom = { string-literal | numeric-constant | input-parameter }
like_expression ::= expression [NOT] LIKE
    {string_literal | input_parameter}
    [ESCAPE {string_literal | input_parameter}]
null_comparison_expression ::=
    single_valued_path_expression IS [ NOT ] NULL
empty_collection_comparison_expression ::=
    collection_valued_path_expression IS [NOT] EMPTY
collection_member_expression ::=
    { single_valued_path_expression | input_paramter }
    [ NOT ] MEMBER [ OF ]
    collection_valued_path_expression
quantified_expression ::=
    expression comparison_operator {SOME | ANY | ALL}
    (subselect)
exists_expression ::=
    EXISTS {collection_valued_path_expression | (subselect)}
subselect ::=
    SELECT [{ ALL | DISTINCT }] expression
    from_clause [where_clause]
    [group_by_clause] [having_clause]
group_by_clause ::=
    GROUP BY [single_valued_path_expression,]*
    single_valued_path_expression
having_clause ::= HAVING conditional_expression
is_of_type_expression ::= identifier IS OF TYPE
    ([[ONLY] abstract_schema_name,]* [ONLY]
    abstract_schema_name)
comparison_expression ::= expression
    comparison_operator { expression | ( subquery ) }

```

```

comparison_operator ::= = | > | >= |
                    < | <= | <>
method ::= method_name( [[expression , ]* expression ] )
expression ::= term | expression {+|-} term
term ::= factor | term {*/} factor
factor ::= {+|-} primary
primary ::= single_valued_path_expression | literal |
          ( expression ) | input_parameter | functions
functions ::=
    ABS(expression) |
    AVG([ALL|DISTINCT] expression) |
    BIGINT(expression) |
    CHAR({expression [, {ISO|USA|EUR|JIS}] } ) |
    CONCAT (expression , expression ) |
    COUNT({[ALL|DISTINCT] expression | *}) |
    DATE(expression) |
    DAY({expression } |
    DAYS( expression ) |
    DECIMAL( expression [,integer[,integer]]) |
    DIGITS( expression ) |
    DOUBLE( expression ) |
    FLOAT( expression ) |
    HOUR ( expression ) |
    INTEGER( expression ) |
    LCASE ( expression ) |
    LENGTH(expression) |
    LOCATE( expression, expression [, expression] ) |
    MAX([ALL|DISTINCT] expression) |
    MICROSECOND( expression ) |
    MIN([ALL|DISTINCT] expression) |
    MINUTE ( expression ) |
    MONTH( expression ) |
    REAL( expression ) |
    SECOND( expression ) |
    SMALLINT( expression ) |
    SQRT ( expression ) |
    SUBSTRING( expression, expression[, expression]) |
    SUM([ALL|DISTINCT] expression) |
    TIME( expression ) |
    TIMESTAMP( expression ) |
    UCASE ( expression ) |
    YEAR( expression )

```

Comparison of EJB 2.0 specification and WebSphere query language

Item	EJB 2.0 specification	WebSphere Query	WebSphere Enterprise (Dynamic) Query
Bean methods	no	no	yes
Calendar comparisons	yes	yes	yes
Delimited identifiers	no	yes	yes
Dependent Value attributes	no	yes	yes
Dependent Value methods	no	no	yes
Dynamic Query APIs	no	no	yes
EXISTS predicate	no	yes	yes
Inheritance	no	yes	yes

Item	EJB 2.0 specification	WebSphere Query	WebSphere Enterprise (Dynamic) Query
Multiple element select clauses	no	no	yes
Order by	no	yes	yes
Scalar functions	yes *	yes	yes
Select clause	required	optional	optional
SQL Date/time expressions	no	yes	yes
String comparisons	= and <> only	= <> > <	= <> > <
Subqueries, aggregations, group by, and having clauses	no	yes	yes

* EJB 2.0 defines the following scalar functions: abs, sqrt, concat, length, locate and substring. WebSphere query and dynamic query support additional scalar functions as listed in the topic, EJB query: Scalar functions.

Chapter 11. Internationalizing applications

For your application to be used in multiple regions around the world, its user interfaces will need to support multiple locales and time zones. IBM WebSphere Application Server supports the maintenance and deployment of centralized message catalogs for the output of properly formatted, language-specific (*localized*) interface strings.

- If you are new to internationalization, read “Internationalization” before you continue.
- For general information about internationalization, see “Internationalization: Resources for learning” on page 387.

Steps for this task

1. Identify localizable text in your application.
2. Create the message catalogs necessary for the locales to be supported by your application.
3. In your application code, compose the language-specific strings for output.
4. Assemble your application into one or more application components.
5. Prepare the localizable-text package for deployment with your localized application.
In this step, you create a deployment JAR.
6. Assemble application modules and the deployment JAR into a J2EE application.
7. Deploy and manage the application.

Internationalization

An application that can present information to users according to regional cultural conventions is said to be *internationalized*: The application can be configured to interact with users from different localities in culturally appropriate ways. In an internationalized application, a user in one region sees error messages, output, and interface elements in the requested language. Date and time formats, as well as currencies, are presented appropriately for users in the specified region. A user in another region sees output in the conventional language or format for that region.

Historically, the creation of internationalized applications has been restricted to large corporations writing complex systems. Internationalization techniques have traditionally been expensive and difficult to implement, so they have been applied only to major development efforts. However, given the rise in distributed computing and in the use of the World Wide Web, application developers have been pressured to internationalize a much wider variety of applications. This requires making internationalization techniques much more accessible to application developers.

In an application that is not internationalized, the interface that the user sees is unalterably written into the application code. On the other hand, localizing the displayed strings adds a layer of abstraction into the design of the application. Instead of simply printing an error message, an internationalized application represents the error message with some language-neutral information; in the simplest case, each error condition corresponds to a key. To print a usable error message, then, the application looks up the key in the configured message catalog.

Each message catalog is a list of keys with associated strings. Different message catalogs provide strings for the different languages supported. The application looks up the key in the appropriate catalog, retrieves the corresponding error message in the requested language, and prints this string for the user.

Localization of text can be used for far more than translating error messages. For example, by using keys to represent each element in a graphical user interface (GUI) and by providing the appropriate message catalogs, the GUI itself (buttons, menus, and so on) can support multiple languages. Extending support to additional languages requires that you provide message catalogs for those languages; in many cases, the application itself needs no further modification.

Internationalization of an application is driven by two variables, the time zone and the locale. The time zone indicates how to compute the local time as an offset from a standard time like Greenwich Mean Time. The locale is a collection of information about language, currency, and the conventions for presenting information like dates. In a localized application, the locale also indicates the message catalog from which an application is to retrieve message strings. A time zone can cover many locales, and a single locale can span time zones. With both time zone and locale, the date, time, currency, and language for users in a specific region can be determined.

The localizable-text package is a set of Java classes and interfaces that can be used to localize the strings in distributed applications easily. Language-specific string catalogs can be stored centrally so that they can be maintained efficiently.

Identifying localizable text

Steps for this task

1. Determine which elements of the application need to be translated.

Good candidates for localization include the following:

- Graphical user interfaces: window titles, menus and menu items, buttons, on-screen instructions
- Prompts in command-line interfaces
- Application output: messages and logs

2. Assign a unique key to each element for use in message catalogs for the application.

The key provides a language-neutral link between the application and language-specific strings in the message catalogs. Establishing a naming convention for keys before creating the catalogs can make writing code with these keys much more intuitive for interface programmers.

Usage scenario

Suppose you are localizing the GUI for a banking system, and the first window contains a pull-down list to be used for selecting a type of account. The labels for the list and the account types in the list are good choices for localization. Three elements require keys: the list itself and two items in the list.

What to do next

Create message catalogs for the language-specific strings.

Creating message catalogs

Before you begin

Identify strings that need to be localized.

You can create a catalog as either a subclass of `java.util.ResourceBundle` or a Java properties file. The properties-file approach is more common, because properties files can be prepared by people without programming experience and swapped in without modifying the application code.

Steps for this task

1. For each string identified for localization, add a line to the message catalog that lists the string's key and value in the current language.

In a properties file, each line has the following structure:

```
key = string associated with the key
```

2. Save the catalog, giving it a locale-specific name.

To enable resolution to a specific properties file, the Java API specifies naming conventions for the properties files in a resource bundle as `bundleName_localeID.properties`. Give the set of message catalogs a collective name, for example, `BankingResources`. For information about locale IDs that are recognized by the Java APIs, see “Internationalization: Resources for learning” on page 387.

Usage scenario

The following English catalog (`BankingResources_en.properties`) supports the labels for the list and its two list items:

```
accountString = Accounts  
savingsString = Savings  
checkingString = Checking
```

The corresponding German catalog (`BankingResources_de.properties`) supports the labels as follows:

```
accountString = Konten  
savingsString = Sparkonto  
checkingString = Girokonto
```

What to do next

Write code to compose the language-specific strings.

Composing language-specific strings

Before you begin

Create message catalogs for the language-specific strings.

Steps for this task

1. In application code, create a `LocalizableTextFormatter` instance, passing in required localization values.
2. **(Optional)** Set other localization values as needed for more complex situations.
3. Generate a properly formatted, language-specific string.

What to do next

When the application is finished, deploy the formatter bean with your application.

Localization API support

The package *com.ibm.websphere.i18n.localizabletext* contains classes and interfaces for localizing text. This package makes extensive use of the internationalization features of the standard Java APIs from Sun Microsystems, including the following:

- `java.util.Locale`
- `java.util.TimeZone`
- `java.util.ResourceBundle`
- `java.text.MessageFormat`

For more information about the standard Java APIs, see “Internationalization: Resources for learning” on page 387.

The WebSphere localizable-text package wraps the Java support and extends it for efficient and simple use in a distributed environment. The primary class used by application programmers is `LocalizableTextFormatter`. Instances of this class are usually created in server programs, but client programs can also create them. Formatter instances are created for specific resource-bundle names and keys. Client programs that receive a `LocalizableTextFormatter` instance call its `format` method. This method uses the locale of the client application to retrieve the appropriate resource bundle and compose a locale-specific message based on the key.

For example, suppose that a distributed application supports both French and English locales; the server is using an English locale and the client, a French locale. The server creates two resource bundles, one each for English and French. When the client makes a request that triggers a message, the server creates a `LocalizableTextFormatter` instance that contains the name of the resource bundle and the key for the message and passes the instance back to the client.

When the client receives the `LocalizableTextFormatter` instance, it calls the object’s `format` method. By using the locale and name of the resource bundle, the `format` method determines the name of the resource bundle that supports the French locale and retrieves the message that corresponds to the key from the French resource bundle. Formatting of the message is transparent to the client.

In this simple example, the resource bundles reside centrally with the server. They do not have to exist with the client. Part of what the localizable-text package provides is the infrastructure to support centralized catalogs. This implementation uses an enterprise bean (a stateless session bean provided with the localizable-text package) to access the message catalogs. When the client calls the `format` method on the `LocalizableTextFormatter` instance, the following events occur:

1. The client application sets the time-zone and locale values in the `LocalizableTextFormatter` instance, either by passing them explicitly or through default values.
2. A call, `LocalizableTextFormatterEJBFinder`, is made to retrieve a reference to the formatter bean.
3. Information from the `LocalizableTextFormatter` instance, including the client’s time zone and locale, is sent to the formatting bean.
4. The formatting bean uses the name of the resource bundle, the message key, the time zone, and the locale to compose a language-specific message.

5. The formatter bean returns the formatted message to the client.
6. The formatted message is inserted into the `LocalizableTextFormatter` instance and returned by the `format` method.

A call to the `format` method requires at most one remote call, to contact the formatter bean. As an alternative, the `LocalizableTextFormatter` instance can cache formatted messages, eliminating the remote call for subsequent uses. In addition, you can set a fallback string so that the application can return a readable string even if it cannot access the appropriate message catalog.

The resource bundles can be stored locally. The `localizable-text` package provides a static variable that indicates whether the bundles are stored locally (`LocalizableConfiguration.LOCAL`) or remotely (`LocalizableConfiguration.REMOTE`). However, the setting of this variable applies to all applications running within the same Java virtual machine.

LocalizableTextFormatter class

The `LocalizableTextFormatter` class, found in the package `com.ibm.websphere.i18n.localizabletext`, is the primary programming interface for using the `localizable-text` package. Instances of this class contain the information needed to create language-specific strings from keys and resource bundles.

`LocalizableTextFormatter` extends `java.lang.Object` and implements the following interfaces:

- `java.io.Serializable`
- `com.ibm.websphere.i18n.localizabletext.LocalizableText`
- `com.ibm.websphere.i18n.localizabletext.LocalizableTextL`
- `com.ibm.websphere.i18n.localizabletext.LocalizableTextTZ`
- `com.ibm.websphere.i18n.localizabletext.LocalizableTextLTZ`

Creation and initialization of class instances

`LocalizableTextFormatter` supports the following constructors:

- `LocalizableTextFormatter()`
- `LocalizableTextFormatter(String resourceBundleName, String patternKey, String appName)`
- `LocalizableTextFormatter(String resourceBundleName, String patternKey, String appName, Object[] args)`

The `LocalizableTextFormatter` instance must have certain values, such as resource-bundle name, key, and the name of the formatting application. If you do not pass these values in by using the second constructor listed previously, you can set them separately by making the following calls:

- `setResourceBundleName(String resourceBundleName)`
- `setPatternKey(String patternKey)`
- `setApplicationName(String appName)`

You can use a fourth method, `setArguments(Object[] args)`, to set optional localization values after construction. See *Processing of application-specific values* at the end of this article. For a usage example, see “*Composing complex strings*” on page 383.

API for formatting text

The formatting methods in `LocalizableTextFormatter` generate a string from a set of message keys and resource bundles, based on some combination of locale and time-zone values. Each method corresponds to one of the four localizable-text interfaces implemented. The following list indicates the interface in which each formatting method is defined:

- `LocalizableText.format()`
- `LocalizableTextL.format(java.util.Locale locale)`
- `LocalizableTextTZ.format(java.util.TimeZone timeZone)`
- `LocalizableTextLTZ.format(java.util.Locale locale, java.util.TimeZone timeZone)`

The `format` method with no arguments uses the locale and time-zone values set as defaults for the Java virtual machine. All four methods throw `LocalizableException` objects.

Location of message catalogs and the `appName` value

Applications written with the localizable-text package can access message catalogs locally or remotely. In a distributed environment, use of remote, centrally located message catalogs is appropriate. All clients can use the same catalogs, and maintenance of the catalogs is simplified. Local formatting is useful in test situations and appropriate under some circumstances. In order to support either local or remote formatting, a `LocalizableTextFormatter` instance must indicate the name of the formatting application. For example, when an application formats a message by using remote catalogs, the message is actually formatted by an enterprise bean on the server. Although the localizable-text package contains the code to automate the lookup of the formatter bean and the issue of a call to it, the application needs to know the name of the formatter bean. Several methods in the `LocalizableTextFormatter` class use a value described as *appName*; this refers to the name of the formatting application, which is not necessarily the name of the application in which the value is set.

Caching of messages

`LocalizableTextFormatter` can optionally cache formatted messages so that they do not have to be reformatted when needed again. By default, caching is not enabled, but use `LocalizableTextFormatter.setCacheSetting(true)` to enable caching. When caching is enabled and the `format` method is called, the method determines whether the message has already been formatted. If so, the cached message is returned. If the message is not found in the cache, the message is formatted and returned to the caller, and a copy of the message is cached for future use.

If caching is disabled after messages have been cached, those messages remain in the cache until the cache is cleared by a call to `LocalizableTextformatter.clearCache()`. You can clear the cache at any time; the cache is automatically cleared when any of the following methods is called:

- `setResourceBundleName(String resourceBundleName)`
- `setPatternKey(String patternKey)`
- `setApplicationName(String appName)`
- `setArguments(Object[] args)`

API for providing fallback information

Under some circumstances, it can be impossible to format a message. The localizable-text package implements a fallback strategy, making it possible to get some information even if a message cannot be formatted correctly into the requested language. The LocalizableTextFormatter instance can optionally store fallback values for a message string, the time zone, and the locale. These can be ignored unless the LocalizableTextFormatter instance throws an exception. To set fallback values, call the following methods as appropriate:

- setFallbackString(String message)
- setFallbackLocale(Locale locale)
- setFallbackTimeZone(TimeZone timeZone)

For a usage example, see “Generating localized text” on page 385.

Processing of application-specific values

The localizable-text package provides native support for localization based on time zone and locale, but one can construct messages on the basis of other values as well. If you need to consider variables other than locale and time zone in formatting localized text, write your own formatter class.

Your formatter class can extend LocalizableTextFormatter or independently implement some or all of the same localizable-text interfaces. As a minimum, your class must implement java.io.Serializable and at least one of the localizable-text interfaces and its corresponding format method. If your class implements more than one localizable-text interface and format method, the order of evaluation of the interfaces is as follows:

1. LocalizableTextLTZ
2. LocalizableTextL
3. LocalizableTextTZ
4. LocalizableText

As an example, the localizable-text package provides a class that reports the time and date (LocalizableTextDateTimeArgument). In that class, date and time formatting is localized in accordance with three values: locale, time zone, and style.

Creating a formatter instance

Server programs typically create LocalizableTextFormatter instances that are sent to clients as the result of some operation; clients format the objects at the appropriate time. Less typically, client programs create LocalizableTextFormatter objects locally.

Steps for this task

1. **(Optional)** If needed for your application, write your own formatter class.

For more information about implementation, see “LocalizableTextFormatter class” on page 379.

2. In application code, call the appropriate constructor for the formatter class and set required localization values.

Some localization values, such as resource bundle name, key and formatting application, must be set, either through a constructor or soon after construction. Other localization values can be set only as needed.

Usage scenario

The following code creates a `LocalizableTextFormatter` instance by using the default constructor and then sets the required localization values:

```
import com.ibm.websphere.i18n.localizabletext.LocalizableException;
import com.ibm.websphere.i18n.localizabletext.LocalizableTextFormatter;
import java.util.Locale;

public void drawAccountNumberGUI(String accountType) {
    ...
    LocalizableTextFormatter ltf = new LocalizableTextFormatter();
    ltf.setPatternKey("accountNumber");
    ltf.setResourceBundleName("BankingSample.BankingResources");
    ltf.setApplicationName("BankingSample");
    ...
}
```

The application that is requesting a localized message can specify the locale and time zone for which the message is to be formatted, or the application can use the default values set for the Java virtual machine.

For example, a GUI can enable users to select the language in which to display the interface. A default value must be set initially so that the GUI can be created properly when the application first starts, but users can then change the locale for the GUI to suit their needs. The following code shows how to change the locale used by an application based on the selection of a menu item:

```
import java.awt.event.ActionListener;
import java.awt.event.ActionEvent;
...
import java.util.Locale;

public void actionPerformed(ActionEvent event) {
    String action = event.getActionCommand();
    ...
    if (action.equals("en_us")) {
        applicationLocale = new Locale("en", "US");
        ...
    }
    if (action.equals("de_de")) {
        applicationLocale = new Locale("de", "DE");
        ...
    }
    if (action.equals("fr_fr")) {
        applicationLocale = new Locale("fr", "FR");
        ...
    }
    ...
}
```

For more information, see “Generating localized text” on page 385.

What to do next

Set optional localization values.

Setting optional localization values

In addition to setting localization values that are required by `LocalizableTextFormatter`, you can set a number of optional values in application code, either through the constructor or by calling any of several methods for that purpose. With optional values, you can do the following:

- Compose complex strings from variable substrings

- Customize the formatting of strings, taking into account variables other than time zone and locale

Steps for this task

1. In application code, add the optional values into an array of type Object.

```
Object[] arg = {new String(getAccountNumber())};
```

2. Pass the array into a `LocalizableTextFormatter` instance.

You can pass the array through the appropriate constructor or by calling the `setArguments(Object[])` method. For a usage example, see “Composing complex strings”.

Note: Because the array is passed by value rather than by reference, any updates to the array variable after this point are not reflected in the `LocalizableTextFormatter` instance unless it is reset by calling the `setArguments(Object[])` method.

What to do next

Write code to generate the localized text.

Composing complex strings

Before you begin

Identify strings that need to be localized.

The localized-text package supports the substitution of variable substrings into a localized string that is retrieved from the message catalog by key.

Steps for this task

1. In the message catalog, specify the location of the substitution in the string to be retrieved by key.

Variable components are designated by curly braces (for example, `{0}`).

2. In application code, create a `LocalizableTextFormatter` instance, passing in an array that contains the variable value.

If the variable substring must itself be localized, you can create a nested `LocalizableTextFormatter` instance for it and pass the instance in instead of a value.

3. Generate a localized string.

When a format method is called on a formatter instance, the formatter takes each element of the array passed in the previous step and substitutes it for the placeholder with the matching index in the string retrieved from the message catalog. For example, the value at index 0 in the array replaces the `{0}` variable in the retrieved string.

Usage scenario

The following line from an English message catalog shows a string with a single substitution:

```
successfulTransaction = The operation on account {0} was successful.
```

The same key in message catalogs for other languages has a translation of this string with the variable at the appropriate location for each language.

The following code shows the creation of a single-element argument array and the creation and use of a `LocalizableTextFormatter` instance:

```
public void updateAccount(String transactionType) {
    ...
    Object[] arg = {new String(this.accountNumber)};
    ...
    LocalizableTextFormatter successLTF =
        new LocalizableTextFormatter ("BankingResources",
            "successfulTransaction",
            "BankingSample",
            arg);
    ...
    successLTF.format(this.applicationLocale);
    ...
}
```

Nesting formatter instances for localized substrings

Before you begin

Identify strings that need to be localized.

The ability to substitute variable substrings into the strings retrieved from message catalogs adds a level of flexibility to the localizable-text package, but this capability is of limited use unless the variable value itself can be localized. You can do this by nesting `LocalizableTextFormatter` instances.

Steps for this task

1. In the message catalog, add entries that correspond to potential values for the variable substring.
2. In application code, create a `LocalizableTextFormatter` instance for the variable substring, setting required localization values.
3. Create a `LocalizableTextFormatter` instance for the primary string, passing in an array that contains the formatter instance for the variable substring.

Usage scenario

The following line from an English message catalog shows a string entry with two substitutions and entries to support the localizable variable at index 0 (the second variable in the string, the account number, does not need to be localized):

```
successfulTransaction = The {0} operation on account {1} was successful.
depositOpString = deposit
withdrawOpString = withdrawal
```

The following code shows the creation of the nested formatter instance and its insertion (with the account number variable) into the primary formatter instance:

```
public void updateAccount(String transactionType) {
    ...
    // Successful deposit
    LocalizableTextFormatter opLTF =
        new LocalizableTextFormatter("BankingResources",
            "depositOpString",
            "BankingSample");
    Object[] args = {opLTF, new String(this.accountNumber)};
    ...
    LocalizableTextFormatter successLTF =
        new LocalizableTextFormatter ("BankingResources",
            "successfulTransaction",
            "BankingSample",
            args);
    ...
}
```

```

    ...
    successLTF.format(this.applicationLocale);
    ...
}

```

Generating localized text

Before you begin

Create a formatter instance and set localization values as needed.

Steps for this task

1. **(Optional)** If needed, customize the formatting behavior.
2. In application code, call the appropriate format method.

Usage scenario

You can provide fallback behavior for use if the appropriate message catalog is not available at formatting time.

The following code generates a localized string. If the formatting fails, the application retrieves and uses a fallback string instead of the localized string:

```

import com.ibm.websphere.i18n.localizabletext.LocalizableException;
import com.ibm.websphere.i18n.localizabletext.LocalizableTextFormatter;
import java.util.Locale;

public void drawAccountNumberGUI(String accountType){
    ...
    LocalizableTextFormatter ltf = new LocalizableTextFormatter();
    ...
    ltf.setFallbackString("Enter account number: ");
    try {
        msg = new Label(ltf.format(this.applicationLocale), Label.CENTER);
    }
    catch (LocalizableException le) {
        msg = new Label(ltf.getFallbackString(), Label.CENTER);
    }
    ...
}

```

What to do next

When the application is finished, deploy the formatter bean with your application.

Customizing the behavior of a formatting method

You can customize formatting behavior by passing your own formatter classes into a `LocalizableTextFormatter` instance through an array of optional values. This enables you to take variables other than locale and time zone into account when formatting localized text.

Steps for this task

1. Write your own formatter class.
For more information about implementation, see “`LocalizableTextFormatter` class” on page 379.
2. In application code, create an instance of your formatter class as appropriate and pass it with any other optional localization values into an instance of `LocalizableTextFormatter`.

When the `LocalizableTextFormatter` instance reads the instance that has been passed in, it attempts to call `format()` on the passed-in instance. The string returned is then processed with any other elements in the array.

Usage scenario

The `localizable-text` package provides an example of a user-defined class, called `LocalizableTextDateTimeArgument`. This class enables date and time information to be selectively formatted according to the style values defined in `java.text.DateFormat` as well as constants defined within `LocalizableTextDateTimeArgument` itself.

Preparing the localizable-text package for deployment

Before you begin

Write code to compose the language-specific strings.

The `LocalizableTextEJBDeploy` tool is used to create a deployment JAR for the Localizable Text service. You must deploy the enterprise bean in each enterprise application that requires support for localized text.

Steps for this task

1. Make sure the **LocalizableTextEJBDeploy** tool (`ltext.jar`) exists in the `lib` directory under the product's installation root directory.
2. Set up a working directory for the **LocalizableTextEJBDeploy** tool to use.
You will need to pass this location to the tool through a command-line interface.
3. Run the **LocalizableTextEJBDeploy** tool.
You might be asked if you want to regenerate deployment code for the `LocalizableText` bean. Do not redeploy the bean; if you do, the generated JNDI name will be incorrect.
To deploy the bean on multiple hosts and servers, run the tool for each host/server combination. This generates a unique JNDI name for each deployment.
After the tool is run, a deployment JAR is located in the working directory you specified.

What to do next

Assemble the deployment JAR in an enterprise application with other application components. As part of preparing for deployment, verify the following:

- Add the resource bundles for your application to the EAR as files.
- Add the location of the EAR to the server's class path. This is so the resource bundles can be located on the virtual host and server.

The same deployment JAR can be included in several enterprise applications.

LocalizableTextEJBDeploy command

Syntax

This topic describes the command-line syntax for the **LocalizableTextEJBDeploy** tool. The file that contains this tool (`ltext.jar`) must be located in the `lib` directory of the product installation root.

```
LocalizableTextEJBDeploy
-a applicationName
-h virtualHostName
-i installationDirectory
-s serverName
-w workingDirectory
```

Parameters

The required parameters, which can be specified in any order, follow:

applicationName

The name of the formatting session bean. This name is used in `LocalizableTextFormatter` instances to specify where the actual formatting takes place. If the name cannot be resolved at run time, the format method throws an exception.

virtualHostName

The name of the virtual host on which the formatting session bean is deployed. This value is case-sensitive on all operating platforms.

installationDirectory

The location at which the application server product is installed.

serverName

The name of the application server. If this argument is not specified, the default server name for the product is used.

workingDirectory

A location for the tool to use temporarily.

Internationalization: Resources for learning

Use the following links to find relevant supplemental information about internationalization. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.




View links to additional information about:

- Programming instructions and examples
- Programming specifications

Programming instructions and examples

- ✓ **Java internationalization tutorial**
(<http://java.sun.com/docs/books/tutorial/i18n/index.html>)
An online tutorial that explains how to use the Java 2 SDK Internationalization API.

Programming specifications

-  **Java 2 SDK, Standard Edition Documentation: Internationalization** (<http://java.sun.com/j2se/1.3/docs/guide/intl/>)
The Java internationalization documentation from Sun Microsystems, including a list of supported locales and encodings.
-  **Making the WWW truly World Wide** (<http://www.w3.org/International/>)
The W3C's effort to make World Wide Web technology work with the many writing systems, languages, and cultural conventions of the global community:
-  **developerWorks - Unicode** (<http://www.ibm.com/developerworks/unicode/>)
Articles on various subjects relating to Unicode, from IBM's developerWorks.

Chapter 12. Using the transaction service

These topics provide information about using transactions with WebSphere applications

WebSphere applications can use transactions to coordinate multiple updates to resources as atomic units (as indivisible units of work) such that all or none of the updates are made permanent.

In WebSphere Application Server, transactions are handled by three main components:

- A transaction manager that supports the enlistment of recoverable XAResources and ensures that each such resource is driven to a consistent outcome either at the end of a transaction or after a failure and restart of the application server.
- A container in which the J2EE application runs. The container manages the enlistment of XAResources on behalf of the application when the application performs updates to transactional resource managers (for example, databases). Optionally, the container can control the demarcation of transactions for enterprise beans configured for container-managed transactions.
- An application programming interface (UserTransaction) that is available to bean-managed enterprise beans and servlets. This allows such application components to control the demarcation of their own transactions.

For more information about using transactions with WebSphere applications, see the following topics:

- "Transaction support in WebSphere Application Server"
- "Developing components to use transactions"
- "Configuring transaction properties for an application server"
- "Using local transactions"
- "Setting transactional attributes in the deployment descriptor"
- "Using bean-managed transactions"
- "Managing active transactions"
- "Managing transaction logging for optimum server availability"
- "Troubleshooting transactions"
- "Transaction service exceptions"
- "UserTransaction interface - methods available"

Transaction support in WebSphere Application Server

A transaction is unit of activity within which multiple updates to resources can be made atomic (as an indivisible unit of work) such that all or none of the updates are made permanent. For example, multiple SQL statements to a relational database are committed atomically by the database during the processing of an SQL COMMIT statement. In this case, the transaction is contained entirely within the database manager and can be thought of as a *resource manager local transaction (RMLT)*. In some contexts, a transaction is referred to as a logical unit of work (LUW).

The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself).
- Entity beans use container-managed transactions.
- Web components (servlets) use bean-managed transactions.

WebSphere Application Server is a transaction manager that supports the coordination of resource managers through their XAResource interface and participates in distributed global transactions with other OTS 1.2 compliant transaction managers (for example J2EE 1.3 application servers). WebSphere applications can also be configured to interact with databases, JMS queues, and JCA connectors through their *local transaction* support when distributed transaction coordination is not required.

Resource managers that offer transaction support can be categorized into those that support two-phase coordination (by offering an XAResource interface) and those that support only one-phase coordination (for example through a LocalTransaction interface). The WebSphere Application Server transaction support provides coordination, within a transaction, for any number of two-phase capable resource managers. It also enables a single one-phase capable resource manager to be used within a transaction in the absence of any other resource managers, although a WebSphere transaction is not necessary in this case.

With the Last Participant Support of WebSphere Application Server Enterprise, you can coordinate the use of a single one-phase commit (1PC) capable resource with any number of two-phase commit (2PC) capable resources in the same global transaction. At transaction commit, the two-phase commit resources are prepared first using the two-phase commit protocol, and if this is successful the one-phase commit-resource is then called to `commit(one_phase)`. The two-phase commit resources are then committed or rolled back depending on the response of the one-phase commit resource.

The ActivitySession service of WebSphere Application Server Enterprise provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. It is a distributed context that can be used to coordinate multiple one-phase resource managers. The WebSphere EJB container and deployment tooling support ActivitySessions as an extension to the J2EE programming model. EJBs can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An application can then interact with a resource manager through its LocalTransaction interface for the period of a client-scoped ActivitySession rather than just the duration of an EJB method.

Resource manager local transaction (RMLT)

A resource manager local transaction (RMLT) is a resource manager's view of a local transaction; that is, it represents a unit of recovery on a single connection that is managed by the resource manager.

Resource managers include:

- Enterprise Information Systems that are accessed through a resource adapter, as

described in the J2EE Connector Architecture 1.0 at 
<http://java.sun.com/j2ee/connector/index.html>.

- Relational databases that are accessed through a JDBC datasource.
- JMS queue and topic destinations.

Resource managers offer specific interfaces to enable control of their RMLTs. J2EE connector resource adapters that include support for local transactions provide a `LocalTransaction` interface to enable applications to request that the resource adapter commit or rollback RMLTs. JDBC datasources provide a `Connection` interface for the same purpose.

The boundary at which all RMLTs must be complete is defined in WebSphere Application Server by a "Local transaction containment (LTC)".

Global transactions

If an application uses two or more resources, then an external transaction manager is needed to coordinate the updates to both resource managers in a *global transaction*.

Global transaction support is available to web and enterprise bean J2EE components. Enterprise bean components can be subdivided into beans that exploit container-managed transactions (CMT) or bean-managed transactions (BMT).

BMT enterprise beans and web components can use the Java Transaction API (JTA) `UserTransaction` interface to define the demarcation of a global transaction. The `UserTransaction` interface is obtained by a JNDI lookup of `java:comp/UserTransaction`. The `UserTransaction` is not available to the following components:

- CMT enterprise beans. Any attempt by such beans to obtain the interface results in an exception in accordance with the EJB specification.
- Client applications running outside the Web and EJB containers.

Ensure that programs that perform a JNDI lookup of the `UserTransaction` interface, use an `InitialContext` that resolves to a local implementation of the interface. Also ensure that such programs use a JNDI location appropriate for the EJB version.

Before the EJB 1.1 specification, the JNDI location of the `UserTransaction` interface was not specified. Each EJB container implementor defined it in an implementation-specific manner. Earlier versions of WebSphere Application Server, up to and including Version 3.5.x (without EJB 1.1), bind the `UserTransaction` interface to a JNDI location of `jta/usertransaction`. WebSphere Application Server Version 4, and later releases, bind the `UserTransaction` interface at the location defined by EJB 1.1, which is `java:comp/UserTransaction`. WebSphere Application Server, Version 5 no longer provides the `jta/usertransaction` binding within Web and EJB containers to applications at a J2EE level of 1.3 or later. For example, EJB 2.0 applications can use only the `java:comp/UserTransaction` location.

Local transaction containment (LTC)

A *local transaction containment (LTC)* is used to define the application server behavior in an unspecified transaction context.

(Unspecified transaction context is defined in the Enterprise JavaBeans 2.0

Specification at  <http://java.sun.com/products/ejb/2.0.html>.)

A LTC is a bounded unit-of-work scope within which zero, one, or more resource manager local transactions (RMLTs) can be accessed. The LTC defines the boundary at which all RMLTs must be complete; any incomplete RMLTs are resolved, according to policy, by the container. An LTC is local to a bean instance; it is not shared across beans even if those beans are managed by the same container. LTCs are started by the container before dispatching a method on a J2EE component (such as an enterprise bean or servlet) whenever the dispatch occurs in the absence of a global transaction context. LTCs are completed by the container depending on the application-configured LTC boundary; for example at the end of the method dispatch. There is no programmatic interface to the LTC support; rather LTCs are managed exclusively by the container and configured by the application deployer through transaction attributes in the application deployment descriptor.

A local transaction containment cannot exist concurrently with a global transaction. If application component dispatch occurs in the absence of a global transaction, the container always establishes an LTC. The only exceptions to this are as follows:

- Where application component dispatch occurs without container interposition; for example, for a stateless session bean create.
- J2EE 1.2 web components.
- J2EE 1.2 BMT enterprise beans.

Using local transactions

Local transaction containment (LTC) support, and its configuration through local transaction extended deployment descriptors, gives IBM WebSphere Application Server application programmers a number of advantages. This topic describes those advantages and how they relate to the settings of the local transaction extended deployment descriptors. This topic also describes points to consider to help you best configure transaction support for some example scenarios that use local transactions.

Develop an enterprise bean or servlet that accesses one or more databases that are independent and require no coordination.

If an enterprise bean does not need to use global transactions, it is often more efficient to deploy the bean with the Container Transaction deployment descriptor **Transaction** attribute set to Not supported instead of Required.

With the extended local transaction support of IBM WebSphere Application Server, applications can perform the same business logic in an unspecified transaction context as they can under a global transaction. An enterprise bean, for example, runs under an unspecified transaction context if it is deployed with a **Transaction** attribute of Not supported or Never.

The extended local transaction support provides a container-managed, implicit local transaction boundary within which application updates can be committed and their connections cleaned up by the container. Applications can then be designed with a greater degree of independence from deployment concerns. This makes using a **Transaction** attribute of Supports much simpler, for example, when the business logic may be called either with or without a global transaction context.

An application can follow a get-use-close pattern of connection usage regardless of whether or not the application runs under a transaction. The application can depend on the close behaving in the same way and not causing a rollback to occur on the connection if there is no global transaction.

There are many scenarios where ACID coordination of multiple resource managers is not needed. In such scenarios running business logic under a **Transaction** policy of Not supported performs better than if it had been run under a Required policy. This benefit is exploited through the **Local Transactions - Resolution-control** extended deployment setting of ContainerAtBoundary. With this setting, application interactions with resource providers (such as databases) are managed within implicit RMLTs that are both started and ended by the container. The RMLTs are committed by the container at the configured **Local Transactions - Boundary**; for example at the end of a method. If the application returns control to the container by an exception, the container rolls back any RMLTs that it has started.

This usage applies to both servlets and enterprise beans.

Use local transactions in a managed environment that guarantees clean-up.

Applications that want to control RMLTs, by starting and ending them explicitly, can use the default **Local Transactions - Resolution-control** extended deployment setting of Application. In this case, the container ensures connection cleanup at the boundary of the local transaction context.

J2EE specifications that describe application use of local transactions do so in the manner provided by the default setting of **Local Transactions - Resolution-control=Application** and **Local Transactions - Unresolved-action=Rollback**. By configuring the **Local Transactions - Unresolved-action** extended deployment setting to Commit, then any RMLTs started by the application but not completed when the local transaction containment ends (for example, when the method ends) are committed by the container. This usage applies to both servlets and enterprise beans.

Extend the duration of a local transaction beyond the duration of an EJB component method.

The J2EE specifications restrict the use of RMLTs to single EJB methods. This restriction is because the specifications have no scoping device, beyond a container-imposed method boundary, to which an RMLT can be extended. In WebSphere Application Server Enterprise, you can exploit the **Local Transactions - Boundary** extended deployment setting to give the following advantages:

- Significantly extend the use-cases of RMLTs
- Make conversational interactions with one-phase resource managers possible through ActivitySession support.

An ActivitySession is a WebSphere Application Server Enterprise programming model extension that provides a distributed context with a boundary that is longer than a single method. You can extend the use of RMLTs over the longer ActivitySession boundary, which can be controlled by a client. The ActivitySession boundary reduces the need to use distributed transactions where ACID operations on multiple resources are not needed. This benefit is exploited through the **Local Transactions - Boundary** extended deployment setting of ActivitySession. Such extended RMLTs can remain under the control of the application or be managed by the container depending on the use of the **Local Transactions - Resolution-control** deployment descriptor setting.

Coordinate multiple one-phase resource managers.

For resource managers that do not support XA transaction coordination, a

client can exploit ActivitySession-bounded local transaction contexts. Such contexts give a client the same ability to control the completion direction of the resource updates by the resource managers as the client has for transactional resource managers. A client can start an ActivitySession and call its entity beans under that context. Those beans can perform their RMLTs within the scope of that ActivitySession and return without completing the RMLTs. The client can later complete the ActivitySession in a commit or rollback direction and cause the container to drive the ActivitySession-bounded RMLTs in that coordinated direction.

To determine how best to configure the transaction support for an application, depending on what you want to do with transactions, consider the following points.

General points

- You want to start and end global transactions explicitly in the application (BMT session beans and servlets only).
For a session bean, set the **Transaction type** to Bean (to use bean-managed transactions) in the component's deployment descriptor. (You do not need to do this for servlets.)
- You want to access only one XA or non-XA resource in a method.
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to ContainerAtBoundary. In the Container transaction deployment descriptor, set **Transaction** to Supports.
- You want to access several XA resources atomically across one or more bean methods.
In the Container transaction deployment descriptor, set **Transaction** to Required, Requires new, or Mandatory.
- You want to access several non-XA resource in a method without having to worry about managing your own local transactions.
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to ContainerAtBoundary. In the Container transaction deployment descriptor, set **Transaction** to Not supported.
- You want to access several non-XA resource in a method and want to manage them independently.
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to Application and set **Local Transactions - Unresolved-action** to Rollback. In the Container transaction deployment descriptor, set **Transaction** to Not supported.

Points specific to WebSphere Application Server Enterprise

- You want to access one of more non-XA resources across multiple EJB method calls without having to worry about managing your own local transactions.
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to ContainerAtBoundary, **Local Transactions - Boundary** to ActivitySession, and **Bean Cache - Activate at** to ActivitySession. In the Container transaction deployment descriptor, set **Transaction** to Not supported and set **ActivitySession** attribute to Required, Requires new, or Mandatory.
- You want to access several non-XA resources across multiple EJB method calls and want to manage them independently.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to Application, **Local Transactions - Boundary** to ActivitySession, and **Bean Cache - Activate at** to ActivitySession. In the Container Transaction deployment descriptor, set **Transaction** to Not supported and set **ActivitySession** attribute to Required, Requires new, or Mandatory.

- You want to use a single non-XA resource and one or more XAResources.

Use the Last Participant Support of WebSphere Application Server Enterprise.

Local and global transaction considerations

Applications use resources, such as JDBC data sources or connection factories, that are configured through the Resources view of the WebSphere Application Server Administrative Console. How these resources participate in a global transaction depends on the underlying transaction support of the resource provider. For example, a JDBC provider can provide either XA or non-XA versions of a data source. A non-XA data source can support only resource manager local transactions (RMLTs), but an XA data source can support two-phase commit coordination.

If an application uses two or more resource providers that support only RMLTs, then atomicity cannot be assured because of the one-phase nature of these resources. To ensure atomic behavior, the application should use resources that support XA coordination and should access them within a global transaction.

If an application uses only one RMLT, the atomic behavior can be guaranteed by the resource manager, which can be accessed under a local transaction containment context.

An application can also access a single resource manager under a global transaction context, even if that resource manager does not support the XA coordination. An application can do this, because WebSphere Application Server performs an "only resource optimization" and interacts with the resource manager under a RMLT. Within a global transaction context, any attempt to use more than one resource provider that supports only RMLTs causes the global transaction to be rolled back.

At any moment, an instance of an enterprise bean can have work outstanding in either a global transaction context or a local transaction containment context, but never both. An instance of an enterprise bean can change from running under one type of context to the other (in either direction), if all outstanding work in the original context is complete. Any violation of this principle causes an exception to be thrown when the enterprise bean tries to start the new context.

Developing components to use transactions

These topics provide information about developing WebSphere application components to use transactions

The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself).

- Entity beans use container-managed transactions.
- Web components (servlets) use bean-managed transactions.

You configure whether a component uses container- or bean-managed transactions by setting an appropriate value on the Transaction type deployment attribute, as described in "Setting transactional attributes in the deployment descriptor". You can also configure other transactional deployment descriptor attributes.

If you want a session bean to manage its own transactions, you must write the code that explicitly demarcates the boundaries of a transaction as described in "Using bean-managed transactions".

Similarly, if you want a Web component to use transactions, you must write the code that explicitly demarcates the boundaries of a transaction as described in "Using bean-managed transactions".

Setting transactional attributes in the deployment descriptor

Use this task to configure the transactional deployment descriptor attributes associated with an EJB or Web module, to enable a J2EE application to use transactions.

To set transactional attributes in the deployment descriptor for an application component (enterprise bean or servlet), complete the following steps:

Steps for this task

1. (Start the Application Assembly Tool).
2. Create or edit the application EAR file.
For example, to change attributes of an existing application, click **File-> Open** then select the EAR file.
3. In the navigation pane, select the component instance; for example:
 - For a session bean, expand *ejb_module_instance*-> **Session beans** then select the bean instance.
 - For a servlet, expand *web_application*-> **Web Components** then select the servlet instance.

A property dialog notebook for the component is displayed in the property pane.

4. In the property pane, select the Advanced tab.
5. Set the **Transaction type** attribute, which defines the transactional manner in which the container invokes a method.
You can set this attribute to Container or Bean, as follows:
 - For a session bean to use container-managed transactions, set Container
 - For a session bean to use bean-managed transactions, set Bean
 - For an entity bean, set Container
 - For a Web component (servlet), set Bean
6. In the property pane, select the IBM Extensions tab.
7. Configure J2EE component extensions attributes for extended local transaction containment.

To enable management of local transaction containments, configure the following EJB extensions attributes. These attributes configure, for the

component, the behaviour of the container's local transaction containment (LTC) environment that the container establishes whenever a global transaction is not present.

Boundary

Specifies the duration of a local transaction context. You can set this attribute to either **Bean method** or **ActivitySession**, as described in ("Entity bean assembly settings").

Note: The ActivitySession option is not supported in the web container.

This property can be changed on WAS Enterprise only.

Resolution control

Specifies how the local transaction is to be resolved before the local transaction context ends: by the application through user code or by the EJB container. You can set this attribute to either **Application** or **ContainerAtBoundary**, as described in ("Entity bean assembly settings").

Unresolved action

Specifies the action that the container must take when the local transaction context scope ends, if resources are uncommitted by an application in a local transaction and the **Resolution control** is set to Application. You can set this attribute to either **Commit** or **Rollback**, as described in ("Entity bean assembly settings").

8. [For EJB components only] For container-managed transactions, configure how the container must manage the transaction boundaries when delegating a method invocation to an enterprise bean's business method:
 - a. In the navigation pane, select **Container Transactions**.
This displays a table of the methods for enterprise beans.
 - b. For each method of the enterprise bean set the **Transaction attribute** attribute to an appropriate value, as defined in ("Container transaction assembly settings").

Using bean-managed transactions

This topic describes how to enable a session bean or servlet to use bean-managed transactions, to manage its own transactions directly instead of letting the container manage the transactions.

Note: Entity beans cannot manage transactions (so cannot use bean-managed transactions).

To enable a session bean or servlet to use bean-managed transactions, complete the following steps:

Steps for this task

1. Set the **Transaction type** attribute in the component's deployment descriptor to Bean, as described in "Setting transactional attributes in the deployment descriptor" (not in this document).

2. Write the component code to actively manage transactions

When writing the code required by a component to manage its own transactions, remember the following basic rules:

- An instance of a stateless session bean cannot reuse the same transaction context across multiple methods called by an EJB client.

- An instance of a stateful session bean can reuse the same transaction context across multiple methods called by an EJB client.

The following code extract shows the standard code required to obtain an object encapsulating the transaction context. There are three basic steps involved:

- The component class must set the value of the `javax.ejb.SessionContext` object reference in the `setSessionContext` method.
- A `javax.transaction.UserTransaction` object is created by calling a lookup on `"java:comp/UserTransaction"`.
- The `UserTransaction` object is used to participate in the transaction by calling transaction methods such as `begin` and `commit` as needed. If an enterprise bean begins a transaction, it must also complete that transaction either by invoking the `commit` method or the `rollback` method.

```

...
import javax.transaction.*;
...
public class MyStatelessSessionBean implements SessionBean {
private SessionContext mySessionCtx =null;
...
public void setSessionContext (SessionContext ctx)throws EJBException {
mySessionCtx =ctx;
}
...
    public float doSomething(long arg1)throws FinderException,EJBException {
        UserTransaction userTran = (UserTransaction)initCtx.lookup(
            "java:comp/UserTransaction");
        ...
        //User userTran object to call transaction methods
        userTran.begin ();
        //Do transactional work
        ...
        userTran.commit ();
        ...
    }
    ...
}
}

```

Configuring transaction properties for an application server

Use this task to configure the transaction properties for an application server; for example, to define the location of the directory that contains the transaction log or to change default timeouts associated with transactions.

To configure the transaction properties for an application server, complete the following steps:

Steps for this task

1. Start the Administrative console
2. In the navigation pane, select **Servers-> Manage Local Server**
This displays the properties of the application server in the content pane.
3. Select the Transaction Service tab, to display the properties page for the transaction service, as two notebook pages:

Configuration

The values of properties defined in the configuration file. If you change these properties, the new values are applied when the application server next starts.

Runtime

The runtime values of properties. If you change these properties, the new values are applied immediately, but are overwritten with the Configuration values when the application server next starts.

4. Select the Configuration tab, to display the transaction-related configuration properties.
5. **(Optional)** If you want to change the directory in which transaction logs are written, type the full pathname of the directory in the **Transaction log directory** field.

You can check the current runtime value of **Transaction log directory**, by clicking the Runtime tab.

You can also specify a size for the transaction logs, as described in the following step.

Note: If you change the transaction log directory, you should apply the change and restart the application server as soon as possible, to minimize the risk of problems caused that might occur before the application server is restarted. For example, if a problem causes the server to fail (with in-flight transactions), the server next starts with the new log directory and is unable to automatically resolve in-flight transactions that were recorded in the old log directory.

6. **(Optional)** If you want to change the default file size of transaction log files, modify the **Transaction log directory** field to include a file size setting, in the following format:

directory_name;file_size

Where

- *directory_name* is the name of the transaction log directory
- *file_size* is the new default size specified in bytes. The *nK* or *nM* suffix can be used to indicate kilobytes or megabytes. If you do not specify a file size value, the default value of 1M is used.

For example, `c:\tranlogs;2M` indicates the files are to be created with 2M bytes size and stored in the directory `c:\tranlogs`.

In a non-production environment, you can use the transaction log directory value of `;0` to disable transaction logging. (There must be no directory name element before the size element of 0.) You should not disable transaction logging in a production environment, because this prevents recovery after a system failure and, therefore, data integrity cannot be guaranteed.

7. In the **Total transaction lifetime timeout** field, type the number of milliseconds a transaction can remain inactive before it is ended by the transaction service. A value of 0 (zero) indicates that there is no timeout limit.
8. In the **Client inactivity timeout** field, type the number of seconds after which a client is considered inactive and the transaction service ends any transactions associated with that client. A value of 0 (zero) indicates that there is no timeout limit.
9. Click **OK**.
10. Stop then restart the application server.

If you change the transaction log directory configuration property to an incorrect directory name, the application server will restart but be unable to open the transaction logs. You should change the configuration property to a valid directory name, then restart the application server.

Transaction service settings

Use this page to modify transaction service settings.

To view this administrative console page, click **Servers > Application Servers > server > Transaction Service**.

Transaction log directory

Specifies the name of a directory for this server where the transaction service stores log files for recovery.

A blank value in the server configuration is expanded by the transaction log at startup as the directory (install_root)/tranlog/(server_name).

Data type	String
-----------	--------

Total transaction lifetime timeout

Specifies the maximum duration, in seconds, for transactions on this application server.

Any transaction that is not requested to complete before this timeout is rolled back. If set to 0, there is no timeout limit.

Data type	Integer
Units	Seconds
Default	120
Range	0 to 2 147 483 647

Client inactivity timeout

Specifies the maximum duration, in seconds, between transactional requests from a remote client.

Any period of client inactivity that exceeds this timeout results in the transaction rolling back in this application server. If set to 0, there is no timeout limit.

Data type	Integer
Units	Seconds
Default	60
Range	0 to 2 147 483 647

Managing active transactions

Use this task to manage transactions that are active on an application server.

You can use this task to display a snapshot of all the transactions currently running on an application server. For each transaction, the following properties are shown: its local ID, global ID, and current status. You can also choose to finish transactions manually.

Under normal circumstances, transactions should run and complete (commit or rollback) automatically, without the need for intervention. However, in some circumstances, you may need to finish a transaction manually. For example, you may want to finish a transaction that has become stuck polling a resource manager that you know will not become available again within the desired timeframe.

Note: If you choose to finish a transaction on an application server, it is recorded as having completed in the transaction service logs for that server, so will not be eligible for recovery during server start up. If you finish a transaction, you are responsible for cleaning up any in-doubt transactions on the resource managers affected.

To manage the active transactions for an application server, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Manage Application Servers**
This displays a list of application servers in the content pane.
2. In the content pane, click *your_app_server*
This displays the properties of the application server, *your_app_server*.
3. In the content pane, click the **Runtime** tab.
This displays the runtime properties of the application server.
4. In the Additional Properties table, select **Transaction Service**
This displays the runtime properties of the Transaction Service.
5. Click **Manage Transactions**.
This displays a snapshot of all the transactions currently running on the server. For each transaction, the following properties are shown: its local ID, global ID, and current status.
6. **(Optional)** If you want to finish one or more transactions, select the checkbox provided on the entry for the transaction, then click **Finish**. Alternatively, to finish all transactions, select the checkbox in the header of the transactions table, then click **Finish**.

Managing transaction logging for optimum server availability

This topic describes some considerations and actions that you can use to manage transaction logging to help ensure that the availability of your application servers is optimized.

The transaction service writes information to the transaction log for every global transaction which involves two or more resources or is distributed across multiple servers. The transaction log is stored on disk and is used by the transaction service for recovery after a system or server crash. The transaction log for each application server consists of multiple files held in a single directory. You can change the directory that an application server uses to store the transaction log, as described in "Configuring transaction properties for an application server".

When a global transaction is completed, the information in the transaction log is not needed anymore so is marked for deletion. Periodically, this redundant information is garbage collected and the space reused by new transactions. The log files are created of fixed size at server startup, thus no further disk space allocation is required during the lifetime of the server. The default allocation is suitable for around 500 concurrent transactions.

If all the log space is in use when a transaction needs to save information, the transaction is rolled back and the message "WTRN0075W: The transaction log file is full. Transaction rolled back." is reported to the system error log. No more transactions can commit until more log space is made available when existing active transactions complete.

You can monitor the number of concurrent global transactions by using the performance monitoring counters for transactions. The "Global transaction commit time" counter is a measure of how long a transaction takes to complete and, therefore, how long the log is in use by a transaction. If this value is high, then transactions are taking a long time to complete, which can be due to resource manager or network failures. If you ensure this value is low, the log is more efficiently used and unlikely to become full.

You can change the default size of log files by updating the transaction log settings as described in "Configuring transaction properties for an application server". Take care if you increase the size above the default 1Mbyte setting, because this extends the time of the log file garbage collection process, and can lead to undesirable periodic "transaction stall" effects.

Configuring transaction aspects of servers for optimum availability

This topic describes some considerations and actions that you can take to configure transaction-related aspects of application servers for optimum availability.

To configure transaction-related aspects of application servers for optimum availability, complete the following steps:

Steps for this task

1. Store the transaction log files on a fast disk in a highly-available file system, such as a RAID device.

The transaction log may need to be accessed by every global transaction and be used for transaction recovery after a crash. Therefore, the disk the log files are being written to should be on a highly-available file system, such as a RAID device.

The performance of the disk also directly affects the transaction performance. In general, a global transaction makes two disk writes, one after the prepare phase when the outcome of the transaction is known (this information is forced to disk) and a further disk write at transaction completion. Therefore, the transaction logs should be placed on the fastest disks available and not make use of network mounted devices.

2. Mirror the transaction log files by using hardware disk mirroring or dual-ported disks.

If log files have been mirrored or can be recovered, they can be used when restarting a failed server or moved to another machine and another server started there to perform recovery.

Hardware disk mirroring or dual-ported disks can be used by specifying the appropriate file system directory for the transaction logs using the WebSphere Administrative Console.

3. Specify the optimum location of the transaction log directory for application servers.

The default transaction log directory for an application server configuration is unset. By default the application server places transaction log files in a subdirectory of the installed WebSphere tranlog directory (as defined by the WebSphere variable `TRANLOG_ROOT`), and the subdirectory name is the same as the server name. For example, the default directory for a server named `server1` on Windows 2000 is: `c:\WebSphere\AppServer\tranlog\server1`.

You can specify an optimum location of the transaction log directory for all application servers, either on a node or cell scope, by setting the WebSphere

variable `TRANLOG_ROOT`. You can specify a different location separately for each application server by setting the **Transaction Log Directory** property for the server.

4. Never allow more than one application server to concurrently use the same set of log files.

Because the transaction logs record the state of global transactions within a server, if the logs become lost or corrupt, then transactions that are in the prepared state before failure can leave resources in an in-doubt state and prevent further updates or access to the resources by other users or servers. These transactions may need to be manually resolved by either committing or rolling back the transactions at the affected resource managers. The failed server can then be cold-started, which creates new empty transaction logs.

If log files have been mirrored or can be recovered, they can be used when restarting the failed server or moved to an alternate server or machine and another server restarted to perform recovery, as described in the related tasks.

Never allow more than one application server to concurrently use the same set of log files, because each server will destroy the information recorded by the other, resulting in corrupt log files that are unusable for future recovery purposes.

5. Configure application servers to always use the same listening port address at each startup.

If you are running distributed transactions between multiple application servers, the remote object references saved in the transaction log need to be redirected to the originating server on recovery.

On Application Server Network Deployment, the node agents automatically redirect such remote object references to the appropriate application servers on recovery. However, if the distributed transaction is between application servers that are not on Application Server Network Deployment, then you must handle the redirection of remote object references for transaction recovery to complete. For example, you must do this if an application server is deployed on WebSphere Application Server (not the Network Deployment edition) and runs distributed transactions with non-WebSphere EJB or Corba servers.

In particular, the default restart action of an application server not on Application Server Network Deployment is to use a different listening port address to the port when the server shut down. This prevents transaction recovery completing. To overcome this, you should always configure application servers to always use the same listening port address at each startup (see the ORB property `com.ibm.CORBA.ListenerPort` in "Warning: no string named [rejb_setg] found." (not in this document)). You may need to make similar configuration changes to other application servers involved in transactions, to be able to access those servers during recovery.

Moving a transaction log from one server to another

This topic describes some considerations and actions that you can take to move the transaction logs for an application server to another server.

To move transaction logs from one application server to another, consider the following steps:

Steps for this task

1. Move all the transaction log files for the application server.

The transaction log directory for each server contains four log files; named `tranlog1`, `tranlog2`, `XAresource1`, and `XAresource2`. When moving transaction

logs from one server to another you must move all four files; otherwise recovery may not complete resulting in data inconsistency.

2. For a single server configuration, move the transaction logs to any server that has access to the same resource managers.

For a single server configuration (where there are no distributed transactions), the transaction logs can be moved to any server (on any node) that has access to the same resource managers as the original server. For example, the server needs communication and valid security access to databases or message queues.

All the transaction log files for the original server need to be moved to a directory accessible by the new server. This can be accomplished by either renaming the transaction log directory or copying all the log files to the new server's transaction log directory before starting the new server.

Note: To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

3. For a network-deployed server configuration, move the transaction logs to a server that has the same name and host IP address, and access to the same resource managers.

For a network-deployed server configuration, where there can be distributed transactions present in the logs, there are more restrictions. Distributed transactions that access multiple servers log information about each server involved in the transaction. This information includes the server name and the IP address of the machine on which the server is running. When recovery is taking place on server restart, the server uses this information to contact the distributed servers and similarly, the distributed servers try to contact the server with the same original name. So, if a server fails and the logs need to be recovered on an alternative server, that alternative server needs to have the same name and host IP address as the original server. The alternative server also needs to have access to the same resource managers as the original server. For example, the server needs communication and valid security access to databases or message queues.

Note: All servers within a cell must have unique names.

Note: To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

Restarting an application server on a different host

This topic describes some considerations and actions that you can take with transaction logs to restart an application server on a different host.

Moving transactions logs to a different host is similar to moving logs from one server to another, as described in "Moving a transaction log from one server to another".

This involves moving an original application server on one host to an alternative server, which has access to the same resource managers, on another host. For a network-deployed server configuration, the alternative server must have the same name and host IP address as the original server.

Note: To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

To restart an application server on a different host, complete the following steps:

Steps for this task

1. Ensure that the alternative application server is stopped.
2. Move all the transaction logs for the original server to the alternative application server, according to the considerations described in "Moving a transaction log from one server to another".
3. Restart the alternative application server.

Transactional interoperation with non-WebSphere application servers

To interoperate transactionally with a non-WebSphere application server, WebSphere Application Server switches dynamically between native transaction contexts and interoperable OTS contexts depending on the capability of the partner with which it is interoperating. The following system properties (that were needed to be set in WebSphere Application Server before version 5.0 to enable transactional interoperation), and the use of native contexts, are deprecated:

```
com.ibm.ejs.jts.jts.ControlSet.nativeOnly=false  
com.ibm.ejs.jts.jts.ControlSet.interoperabilityOnly=true
```

In a future release of WebSphere Application Server only interoperable OTS contexts will be supported.

Troubleshooting transactions

Use this overview task to help resolve a problem that you think is related to the Transaction service.

To identify and resolve transaction-related problems, you can use the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to transactions, complete the following stages:

Steps for this task

1. Check for transaction messages in the admin console.
The Transaction service produces diagnostic messages prefixed by "WTRN". The error message indicates the nature of the problem and provides some detail. The associated message information provides an explanation and any user actions to resolve the problem.
2. Check for Transaction messages in the activity log.
Activity log messages produced by the Transaction service are accompanied by log analyzer descriptions.
3. Check for more messages in the application server's stdout.log.
For more information about a problem, check the stdout.log file for the application server, which should contain more error messages and extra details about the problem.
4. Check for messages in the application server's transaction log directory for information about the transactions in-flight when the problem occurred.


Note: If you changed the transaction log directory and a problem caused the application server to fail (with in-flight transactions) before the server was restarted properly, the server will next start with the new log directory and be unable to automatically resolve in-flight transactions that were recorded in the old log directory. To resolve this, you can copy the transaction logs to the new directory then stop and restart the application server.

Transaction service exceptions

This topic lists the exceptions that can be thrown by the WebSphere Application Server transaction service. The exceptions are listed in the following groups:


- Standard exceptions
- Heuristic exceptions

If the EJB container catches a system exception from the business method of an enterprise bean, and the method is running within a container-managed transaction, the container rolls back the transaction before passing the exception on to the client. For more information about how the container handles the exceptions thrown by the business methods for beans with container-managed transaction demarcation, see the section *Exception handling* in the Enterprise JavaBeans 2.0

specification at  <http://java.sun.com/products/ejb/docs.html>. That section specifies the container's action as a function of the condition under which the business method executes and the exception thrown by the business method. It also illustrates the exception that the client receives and how the client can recover from the exception.

Standard exceptions

The standard exceptions such as `TransactionRequiredException`, `TransactionRolledbackException`, and `InvalidTransactionException` are defined in

the Java Transaction API (JTA) 1.0.1 Specification at  <http://java.sun.com/products/jta/>.

InvalidTransactionException

This exception indicates that the request carried an invalid transaction context.

TransactionRequiredException exception

This exception indicates that a request carried a null transaction context, but the target object requires an active transaction.

TransactionRolledbackException exception


This exception indicates that the transaction associated with processing of the request has been rolled back, or marked for roll back. Thus the requested operation either could not be performed or was not performed because further computation on behalf of the transaction would be fruitless.

Heuristic exceptions

A heuristic decision is a unilateral decision made by one or more participants in a transaction to commit or rollback updates without first obtaining the consensus outcome determined by the Transaction Service. Heuristic decisions are an issue only after the participant has been prepared and the second phase of commit processing is underway. Heuristic decisions are normally made only in unusual circumstances, such as repeated failures by the transaction manager to

communicate with a resource manager during two-phase commit. If a heuristic decision is taken, there is a risk that the decision differs from the consensus outcome, resulting in a loss of data integrity.

The following list provides a summary of the heuristic exceptions. For more detail,

see the Java Transaction API (JTA) 1.0.1 Specification at 
<http://java.sun.com/products/jta/>.

HeuristicRollback exception


This exception is raised on the commit operation to report that a heuristic decision was made and that all relevant updates have been rolled back.

HeuristicMixed exception

This exception is raised on the commit operation to report that a heuristic decision was made and that some relevant updates have been committed and others have been rolled back.

UserTransaction interface - methods available

For details about the methods available with the UserTransaction interface, see the WebSphere Application Server application programming interface reference

information (Javadoc) or the Java Transaction API (JTA) 1.0.1 Specification at 
<http://java.sun.com/products/jta/>.

Chapter 13. Using naming

Naming is used by clients of WebSphere Application Server applications most commonly to obtain references to objects related to those applications, such as Enterprise JavaBeans (EJB) homes. The following steps outline the context of Naming in the overall application development and deployment process. Steps for this task follow:

Steps for this task

1. Develop your application using either JNDI or CosNaming (CORBA) interfaces.

Use these interfaces to look up server application objects that are bound into the name space and obtain references to them. Most Java developers use the JNDI interface. However, the CORBA CosNaming interface is also available for performing Naming operations on WebSphere Application Server name servers or other CosNaming name servers.

2. (Assemble your application)

Application assembly is a packaging and configuration step that is a prerequisite to application deployment. If the application you are assembling is a client to an application running in another process, you should qualify the `jndiName` values in the deployment descriptors for the objects related to the other application. Otherwise, you may need to override the names with qualified names during application deployment. If the objects have fixed qualified names configured for them, you should use them so that the `jndiName` values do not depend on the other application's location within the topology of the cell.

3. (Deploy your application)

Put your assembled application onto the application server. If the application you are assembling is a client to an application running in another server process, be sure to qualify the `jndiName` values for the other application's server objects if they are not already qualified.

For more information on qualified names, see Lookup names support in deployment descriptors and thin clients.

4. Configure name space bindings.

This step is necessary in these cases:

- Your deployed application is to be accessed by legacy client applications running on previous versions of WebSphere Application Server. In this case, you must configure additional name bindings for application objects relative to the default initial context for legacy clients. (Version 5 clients have a different initial context from legacy clients.)
- The application requires qualified name bindings for such reasons as:
 - It will be accessed by J2EE client applications or server applications running in another server process
 - It will be accessed by thin client applications

In this case, you can configure name bindings as additional bindings for application objects. The qualified names for the configured bindings are *fixed*, meaning they do not contain elements of the cell topology that can change if the application is moved to another server. Objects as bound into the name

space by the system can always be qualified with a topology-based name. You must explicitly configure a name binding to use as a fixed qualified name.

For more information on qualified names, see Lookup names support in deployment descriptors and thin clients. For more information on configured name bindings, see Configured name bindings.

5. Troubleshoot any problems that develop.

If a Naming operation is failing and you need to verify whether certain name bindings exist, use the `dumpNameSpace` tool to generate a dump of the name space.

Naming

Naming is used by clients of WebSphere Application Server applications to obtain references to objects related to those applications, such as Enterprise JavaBeans (EJB) homes.

These objects are bound into a mostly hierarchical structure, referred to as a *name space*. In this structure, all non-leaf objects are called *contexts*. Leaf objects can be contexts and other types of objects. Naming operations, such as lookups and binds, are performed on contexts. All naming operations begin with obtaining an *initial context*. You can view the initial context as a starting point in the name space.

The name space structure consists of a set of *name bindings*, each consisting of a name relative to a specific context and the object bound with that name. For example, the name `myApp/myEJB` consists of one non-leaf binding with the name `myApp`, which is a context. The name also includes one leaf binding with the name `myEJB`, relative to `myApp`. The object bound with the name `myEJB` in this example happens to be an EJB home reference. The whole name `myApp/myEJB` is relative to the initial context, which you can view as a starting place when performing naming operations.

You can access and manipulate the name space through a *name server*. Users of a name server are referred to as *naming clients*. Naming clients typically use the Java Naming and Directory Interface (JNDI) to perform naming operations. Naming clients can also use the Common Object Request Broker Architecture (CORBA) `CosNaming` interface.

Typically, objects bound to the name space are resources and objects associated with installed applications. These objects are bound by the system, and client applications perform lookup operations to obtain references to them. Occasionally, server and client applications bind objects to the name space. An application can bind objects to transient or persistent partitions, depending on requirements.

In J2EE environments, some JNDI operations are performed with `java:` URL names. Names bound under these names are bound to a completely different name space which is local to the calling process. However, some lookups on the `java:` name space may trigger indirect lookups to the name server.

New features for name space support

The following are new features of the WebSphere Application Server v5 naming implementation:

- **Name space is distributed.**

For additional scalability, the name space for a cell is distributed among various servers. Every server has a name server. In previous releases, there was only one name server for an entire administrative domain.

In WebSphere Application Server versions prior to v5, all servers shared the same default initial context, and everything was bound relative to that same initial context. In WebSphere Application Server v5, the default initial context for a server is its server root. System artifacts, such as EJB homes and resources, are bound to the server root of the server with which they are associated.

- **Transient and persistent partitions.**

The name space is partitioned into transient areas and persistent areas. Server roots are transient. System-bound artifacts such as EJB homes and resources are bound under server roots. There is a cell persistent root, which you can use for cell-scoped persistent bindings, and a node persistent root, which you can use to bind objects with a node scope.

- **System name space structure.**

The name space for the entire cell is federated among all servers in the cell. Every server process contains a name server. All name servers provide the same logical view of the cell name space. The various server roots and persistent partitions of the name space are interconnected by means of a system name space. You can use the system name space structure to traverse to any context in the cell name space.

- **Configured bindings.**

You can use the configuration graphical interface and script interfaces to configure bindings in various root contexts within the name space. These bindings are read-only and are bound by the system at server startup.

- **Support for CORBA Interoperable Naming Service (INS) object URLs.**


WebSphere Application Server v5 contains support for Common Object Request Broker Architecture (CORBA) object URLs (corbaloc and corbname) as Java Naming and Directory Interface (JNDI) provider URLs and lookup names.

Name space logical view

The name space for the entire cell is federated among all servers in the cell. Every server process contains a name server. All name servers provide the same logical view of the cell name space. The various server roots and persistent partitions of the name space are interconnected by a system name space. You can use the system name space structure to traverse to any context in a the cell's name space. A logical view of the name space is shown in the following diagram.

Name Space Logical View

Logical View of a Cell's Name Space



The bindings in the preceding diagram appear with solid arrows, labeled in bold, and dashed arrows, labeled in gray. Solid arrows represent *primary bindings*. A primary binding is formed when the associated subcontext is created. Dashed arrows show *linked bindings*. A linked binding is formed when an existing context is bound under an additional name. Linked bindings are added for convenience or interoperability with previous WebSphere Application Server versions.

A cell name space is composed of contexts which reside in servers throughout the cell. All name servers in the cell provide the same logical view of the cell name space. A name server constructs this view at startup by reading configuration information. Each name server has its own local in-memory copy of the name space and does not require another running server to function. There are, however, a few exceptions. Server roots for other servers are not replicated among all the servers. The respective server for a server root must be running to access that server root context.

Name space partitions

There are four major partitions in a cell name space:

- System name space partition
- Server roots partition
- Cell persistent partition
- Node persistent partition

Each partition is discussed in more detail below.

System name space partition

The system name space contains a structure of contexts based on the cell topology. The system structure supports traversal to all parts of a cell name space and to the

cell root of other cells, which are configured as foreign cells. The root of this structure is the cell root. In addition to the cell root, the system structure contains a node root for each node in the cell. You can access other contexts of interest specific to a node from the node root, such as the node persistent root and server roots for servers configured in that node.

All contexts in the system name space are read-only. You cannot add, update, or remove any bindings.

Server roots partition

Each server in a cell has a server root context. A server root is specific to a particular server. You can view the server roots for all servers in a cell as being in a transient read/write partition of the cell name space. System artifacts, such as EJB homes for server applications and resources, are bound under the server root context of the associated server. A server application can also add bindings under its server root. These bindings are transient. Therefore, the server application creates all required bindings at application startup, so they exist anytime the application is running.

A server cluster is composed of many servers that are logically equivalent. Each member of the cluster has its own server root. These server roots are not replicated across the cluster. In other words, adding a binding to the server root of one member does not propagate it to the server roots of the other cluster members. To maintain the same view across the cluster, you should create all user bindings under the server root by the server application at application startup so that the bindings are present under the server root of each cluster member. Because of Workload Management (WLM) behavior, a JNDI client outside a cluster has no control over which cluster member's server root context becomes the target of the JNDI operation. Therefore, you should execute bind operations to the server root of a cluster member should from within that cluster member process only.

Distributing application objects among many server roots is a departure from previous WebSphere Application Server releases, where all system artifacts were bound under a single root. This change can affect the names that clients use to look up these objects.

Server-scoped bindings are relative to a server's server root.

Cell persistent partition

The root context of the cell persistent partition is the cell persistent root. A binding created under the cell persistent root is saved as part of the cell configuration and continues to exist until it is explicitly removed. Applications that need to create additional persistent bindings of objects generally associated with the cell can bind these objects under the cell persistent root.

It is important to note that the cell persistent area is not designed for transient, rapidly changing bindings. The bindings are more static in nature, such as part of an application setup or configuration, and are not created at run time.

Note: In WebSphere Application Server Network Deployment cells, to bind objects to the cell persistent root, ensure that the deployment manager and all node agents in the cell are running.

An important role of the cell persistent root is as the initial context for clients running in previous WebSphere Application Server versions. If you want to access an enterprise bean by WebSphere Application Server v4.0.x and 3.5.x clients, you must ensure that a binding for it has been added to the cell persistent root. You can configure these additional bindings as cell-scoped bindings.

Node persistent partition

The node persistent partition is similar to the cell partition except that each node has its own node persistent root. A binding created under a node persistent root is saved as part of that node configuration and continues to exist until it is explicitly removed.

Applications that need to create additional persistent bindings of objects associated with a specific node can bind those objects under that particular node's node persistent root. As with the cell persistent area, it is important to note that the node persistent area is not designed for transient, rapidly changing bindings. These bindings are more static in nature, such as part of an application setup or configuration, and are not created at run time.

Note: In WebSphere Application Server Network Deployment cells, to bind objects to a node persistent root, ensure the node agent for the node is running.

Unlike the cell persistent root, the node persistent root plays no special role in interoperability with WebSphere Application Server clients of previous releases. Node-scoped bindings are relative to a node's node persistent root.

Note: In the system name space, there is no persistent node root for the deployment manager node because no node agent or application servers run in that node.

Initial context support

All naming operations begin with obtaining an initial context. You can view the initial context as a starting point in the name space. Use the initial context to perform naming operations, such as looking up and binding objects in the name space.

Initial contexts registered with the ORB as initial references

The server root, cell persistent root, cell root, and node root are registered with the name server's ORB and can be used as an initial context. An initial context is used by CORBA and enterprise bean applications as a starting point for name space lookups. The keys for these roots as recognized by the ORB are shown in the following table:

Root Context	Initial Reference Key
Server Root	NameServiceServerRoot
Cell Persistent Root	NameServiceCellPersistentRoot
Cell Root	NameServiceCellRoot NameService
Node Root	NameServiceNodeRoot

A server root initial context is the server root context for the specific server you are accessing. Similarly, a node root initial context is the node root for the server being accessed.

You can use the previously mentioned keys in CORBA INS object URLs (corbaloc and corbaname) and as an argument to an ORB `resolve_initial_references` call. For examples, see CORBA and JNDI programming examples, which show how to get an initial context.

Default initial contexts

The default initial context depends on the type of client. Different categories of clients and the corresponding default initial context follow.

WebSphere Application Server v5 JNDI interface implementation

The JNDI interface is used by EJB applications to perform name space lookups. WebSphere Application Server clients by default use the WebSphere Application Server CosNaming JNDI plug-in implementation. The default initial context for clients of this type is the server root of the server specified by the provider URL. For more details, refer to the JNDI programming examples on getting initial contexts.

WebSphere Application Server JNDI interface implementation prior to v5

WebSphere Application Server clients running in releases prior to WebSphere Application Server v5 by default use WebSphere Application Server's v4.0 CosNaming JNDI plug-in implementation. The default initial context for clients of this type is the cell persistent root, also known as the *legacy root*.

Other JNDI implementation

Some applications can perform name space lookups with a non-WebSphere Application Server CosNaming JNDI plug-in implementation. Assuming the key `NamingContext` is used to obtain the initial context, the default initial context for clients of this type is the cell root.

CORBA

The standard CORBA client obtains an initial `org.omg.CosNaming.NamingContext` reference with the key `NamingContext`. The initial context in this case is the cell root.

Lookup names support in deployment descriptors and thin clients

Server objects, such as EJB homes, are bound relative to the server root context for the server in which the application is installed. Other objects, such as resources, can also be bound to a specific server root. The names used to look up these objects must be qualified so as to select the correct server root. This is a departure from previous versions of WebSphere Application Server, where these objects were all bound under a single root context. This section discusses what relative and qualified names are, when they can be used, and how you can construct them.

Relative names

All names are relative to a context. Therefore, a name that can be resolved from one context in the name space cannot necessarily be resolved from another context

in the name space. This point is significant because the system binds objects with names relative to the server root context of the server in which the application is installed. Each server has its own server root context. The initial JNDI context is by default the server root context for the server identified by the provider URL used to obtain the initial context. (Typically, the URL consists of a host and port.) For applications running in a server process, the default initial JNDI context is the server root for that server. A relative name will resolve successfully when the initial context is obtained from the server which contains the target object, but it will not resolve successfully from an initial context obtained from another server.

If all clients of a server application run in the same server process as the application, all objects associated with that application are bound to the same initial context as the clients' initial context. In this case, only names relative to the server's server root context are required to access these server objects. Frequently, however, a server application has clients that run outside the application's server process. The initial context for these clients can be different from the server application's initial context, and lookups on the relative names for server objects may fail. These clients need to use the qualified name for the server objects. This point must be considered when setting up the `jndiName` values in a J2EE client application deployment descriptors and when constructing lookup names in thin clients. Qualified names resolve successfully from any initial context in the cell.

Qualified names

All names are relative to a context. Here, the term qualified name refers to names that can be resolved from any initial context in a cell. This action is accomplished by using names that navigate to the same context, the cell root. The rest of the qualified name is then relative to the cell root and uniquely identifies an object throughout the cell. All initial contexts in a server (that is, all naming contexts in a server registered with the ORB as an initial reference) contain a binding with the name `cell`, which links back to the cell root context. All qualified names begin with the string `cell/` to navigate from the current initial context back to the cell root context.

A qualified name for an object is the same throughout the cell. The name can be topology-based, or some fixed name bound under the cell persistent root. Topology-based names, described in more detail below, navigate through the system name space to reach the target object. A fixed name bound under the cell persistent root has the same qualified name throughout the cell and is independent of the topology. Creating a fixed name under the cell persistent root for a server application object requires an extra step when the server application is installed, but this step eliminates impacts to clients when the application is moved to a different location in the cell topology. The process for creating a fixed name is described later in this section.

Generally speaking, you **must** use qualified names for EJB `jndiName` values in a J2EE client application deployment descriptors and for EJB lookup names in thin clients. The only exception is when the initial context is obtained from the server in which the target object resides. For example, a session bean which is a client to an entity bean can use a relative name if the two beans run in the same server. If the session bean and entity beans run in different servers, the `jndiName` for the entity bean must be qualified in the session bean's deployment descriptors. The same requirement may be true for resources as well, depending on the scope of the resource.

Topology-based names

The system name space partition in a cell's name space reflects the cell's topology. This structure can be navigated to reach any object bound into the cell's name space. Topology-based qualified names include elements from the topology which reflect the object's location within the cell. For a system-bound object, such as an EJB home, the form for a topology-based qualified name depends on whether the object is bound to a single server or cluster. Both forms are described below.

Single Server

An object bound in a single server has a topology-based qualified name of the following form:

```
cell/nodes/nodeName/servers/serverName/relativeJndiName
```

where *nodeName* and *serverName* are the node name and server name for the server where the object is bound, and *relativeJndiName* is the unqualified name of the object; that is, the object's name relative to its server's server root context.

Server Cluster

An object bound in a server cluster has a topology-based qualified name of the following form:

```
cell/clusters/clusterName/relativeJndiName
```

where *clusterName* is the name of the server cluster where the object is bound, and *relativeJndiName* is the unqualified name of the object; that is, the object's name relative to a cluster member's server root context.

Fixed names

It is possible to create a fixed name for a server object so that the qualified name is independent of the cell topology. This quality is desirable when clients of the application run in other server processes or as pure clients. Fixed names have the advantage of not changing if the object is moved to another server. The *jndiName* values in deployment descriptors for a J2EE client application can reference the qualified fixed name for a server object regardless of the cell topology on which the client or server application is being installed.

Defining a cell-wide fixed name for a server application object requires an extra step after the server application is installed. That is, a binding for the object must be created under the cell persistent root. A fixed name bound under the cell persistent root can be any name, but all names under the cell persistent root must be unique within the cell because the cell persistent root is global to the entire cell.

A qualified fixed name has the form:

```
cell/persistent/fixedName
```

where *fixedName* is an arbitrary fixed name.

The binding can be created programmatically (for example, using JNDI). However, it is probably more convenient to configure a cell-scoped binding for the server object.

You must keep the programmatic or configured binding up-to-date. Configured EJB bindings are based on the location of the enterprise bean within the cell topology, and moving the EJB application to another single server or to a server cluster, for example, requires the configured binding to be updated. Similar changes affect an EJB home reference programmatically bound so that the fixed

name would need to be rebound with a current reference. However, for J2EE clients, the `jndiName` value for the object, and for thin clients, the lookup name for the object, remains the same. In other words, clients that access objects by fixed names are not affected by changes to the configuration of server applications they access.

JNDI support in WebSphere Application Server

IBM WebSphere Application Server includes a name server to provide shared access to Java components, and an implementation of the `javax.naming` JNDI package which supports user access to the WebSphere Application Server name server through the JNDI naming interface.

WebSphere Application Server does *not* provide implementations for:

- `javax.naming.directory` or
- `javax.naming.ldap` packages

Also, WebSphere Application Server does *not* support interfaces defined in the `javax.naming.event` package.

However, to provide access to LDAP servers, the development kit shipped with WebSphere Application Server supports Sun's implementation of:

- `javax.naming.ldap` and
- `com.sun.jndi.ldap.LdapCtxFactory`

WebSphere Application Server's JNDI implementation is based on version 1.2 of the JNDI interface, and was tested with Version 1.2.1 of Sun's JNDI Service Provider Interface (SPI).

The default behavior of this JNDI implementation is adequate for most users. However, users with specific requirements can control certain aspects of JNDI behavior.

Developing applications that use JNDI

References to EJB homes and other artifacts such as data sources are bound to the WebSphere name space. These objects can be obtained through the JNDI interface. Before you can perform any JNDI operations, you need to get an initial context. You can use the initial context to look up objects bound to the WebSphere name space.

These examples describe how to get an initial context and how to perform lookup operations.

- Getting the default initial context
- Getting an initial context by setting the provider URL property
- Setting the provider URL property to select a different root context as the initial context
- Looking up an EJB home with JNDI
- Looking up a JavaMail session with JNDI

In these examples, the default behavior of features specific to WebSphere's JNDI Context implementation is used.

WebSphere Application Server's JNDI context implementation includes special features. JNDI caching enhances performance of repeated lookup operations on the same objects. Name syntax options offer a choice of a name syntaxes, one optimized for typical JNDI clients, and one optimized for interoperability with CosNaming applications. Most of the time, the default behavior of these features is the preferred behavior. However, sometimes you should modify the behavior for specific situations.

JNDI caching and name syntax options are associated with a `javax.naming.InitialContext` instance. To select options for these features, set properties that are recognized by the WebSphere Application Server's initial context factory. To set JNDI caching or name syntax properties which will be visible to WebSphere Application Server's initial context factory, follow the following steps.

Steps for this task

1. **(Optional)** Configure JNDI caches

JNDI caching can greatly increase performance of JNDI lookup operations. By default, JNDI caching is enabled. In most situations, this default is the desired behavior. However, in specific situations, use the other JNDI cache options.

Objects are cached locally as they are looked up. Subsequent lookups on cached objects are resolved locally. However, cache contents can become stale. This situation is not usually a problem, since most objects you look up do not change frequently. If you need to look up objects which change relatively frequently, change your JNDI cache options.

JNDI clients can use several properties to control cache behavior.

You can set properties:

- From the command line by entering the actual string value. For example:

```
java -Dcom.ibm.websphere.naming.jndicache.maxentrylife=1440
```

- In a `jndi.properties` file by creating a file named `jndi.properties` as a text file with the desired properties settings. For example:

```
...
com.ibm.websphere.naming.jndicache.cacheobject=none
...
```

Include the file as the beginning of the classpath, so that the classloader loads your copy of `jndi.properties` before any other copies.

- Within a Java program by using the **PROPS.JNDI_CACHE*** Java constants, defined in the `com.ibm.websphere.naming.PROPS` file. The constant definitions follow:

```
public static final String JNDI_CACHE_OBJECT =
    "com.ibm.websphere.naming.jndicache.cacheobject";
public static final String JNDI_CACHE_OBJECT_NONE = "none";
public static final String JNDI_CACHE_OBJECT_POPULATED = "populated";
public static final String JNDI_CACHE_OBJECT_CLEARED = "cleared";
public static final String JNDI_CACHE_OBJECT_DEFAULT =
    JNDI_CACHE_OBJECT_POPULATED;

public static final String JNDI_CACHE_NAME =
    "com.ibm.websphere.naming.jndicache.cachename";
public static final String JNDI_CACHE_NAME_DEFAULT = "providerURL";

public static final String JNDI_CACHE_MAX_LIFE =
    "com.ibm.websphere.naming.jndicache.maxcachelife";
public static final int JNDI_CACHE_MAX_LIFE_DEFAULT = 0;
```

```

public static final String JNDI_CACHE_MAX_ENTRY_LIFE =
    "com.ibm.websphere.naming.jndicache.maxentrylife";
public static final int JNDI_CACHE_MAX_ENTRY_LIFE_DEFAULT = 0;

```

To use the previous properties in a Java program, add the property setting to a hashtable and pass it to the InitialContext constructor as follows:

```

java.util.Hashtable env = new java.util.Hashtable();
...
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE);
// Disable caching
...
javax.naming.Context initialContext =
    new javax.naming.InitialContext(env);

```

2. (Optional) Specify the name syntax

Most WebSphere applications use JNDI to look up EJB objects and do not need to look up objects bound by CORBA applications. Therefore, the default name syntax used for JNDI names is the most convenient. If your application needs to look up objects bound by CORBA applications, you may need to change your name syntax so that all CORBA CosNaming names can be represented.

JNDI clients can set the name syntax by setting a property. The property setting is applied by the initial context factory when you instantiate a new `java.naming.InitialContext` object. Names specified in JNDI operations on the initial context are parsed according to the specified name syntax.

You can set the property:

- From the command line by entering the actual string value. For example:

```
java -Dcom.ibm.websphere.naming.name.syntax=ins
```

- In a `jndi.properties` file by creating a file named `jndi.properties` as a text file with the desired properties settings. For example:

```

...
com.ibm.websphere.naming.name.syntax=ins
...

```

Include the file as the beginning of the classpath, so that the classloader loads your copy of `jndi.properties` before any other copies.

- Within a Java program by using the `PROPS.NAME_SYNTAX*` Java constants, defined in the `com.ibm.websphere.naming.PROPS` file. The constant definitions follow:

```

public static final String NAME_SYNTAX =
    "com.ibm.websphere.naming.name.syntax";
public static final String NAME_SYNTAX_JNDI = "jndi";
public static final String NAME_SYNTAX_INS = "ins";

```

To use the previous properties in a Java program, add the property setting to a hashtable and pass it to the InitialContext constructor as follows:

```

java.util.Hashtable env = new java.util.Hashtable();
...
env.put(PROPS.NAME_SYNTAX, PROPS.NAME_SYNTAX_INS);
// Set name syntax to INS
...
javax.naming.Context initialContext =
    new javax.naming.InitialContext(env);

```

Example: Getting the default initial context

This example below gets the default initial context. That is, no provider URL is passed to the `javax.naming.InitialContext` constructor. The following section explains the process of determining the address of the bootstrap server to use to obtain the initial context.

Usage scenario

```
...
import javax.naming.Context;
import javax.naming.InitialContext;
...
Context initialContext = new InitialContext();
...
```

The default initial context returned depends the runtime environment of the JNDI client. The initial context returned in the various environments are listed below:

- Thin client: The server root context of the server running on the local host at port 2809.
- Pure client:
 - The context specified by the `java.naming.provider.url` property passed to `launchClient` command with the `-CCD` command line parameter. The context usually will be the server root context of the server at the address specified in the URL, although it is possible to construct a `corbaname` or `corbaloc` URL which resolves to some other context.
 - If no provider URL was specified, the server root context of the server running on the host and port specified by the `-CCBootstrapHost` `-CCBootstrapPort` command line parameters. The default host is the local host, and the default port is 2809.
- Server process: The server root context for that process.

Even though no provider URL is explicitly specified in the above example, the `InitialContext` may find a provider URL defined in other places that it searches for property settings.

Users of properties which affect ORB initialization should read the rest of this section for a deeper understanding of exactly how initial contexts are obtained, which has changed from previous releases.

Determining which server is used to obtain the initial context

WebSphere Application Server name servers are CORBA CosNaming name servers, and WebSphere Application Server provides a CosNaming JNDI plug-in implementation for JNDI clients to perform naming operations on WebSphere Application Server name spaces. The WebSphere Application Server CosNaming plug-in implementation is selected through a JNDI property that is passed to the `InitialContext` constructor. This property is `java.naming.factory.initial`, and it specifies the initial context factory implementation to use to obtain an initial context. The factory returns a `javax.naming.Context` instance, which is part of its implementation.

The WebSphere Application Server initial context factory, `com.ibm.websphere.naming.WsnInitialContextFactory`, is typically used by WebSphere Application Server applications to perform JNDI operations. The WebSphere Application Server run-time environment is set up to use this WebSphere Application Server initial context factory if one is not specified explicitly by the JNDI client. When the initial context factory is invoked, an *initial*

context is obtained. The following paragraphs explain how the WebSphere Application Server initial context factory obtains the initial context in client and server environments.

Understanding the registration of initial references in server processes

Every WebSphere Application Server has an ORB used to receive and dispatch invocations on objects running in that server. Services running in the server process can register initial references with the ORB. Each initial reference is registered under a key, which is a string value. An initial reference can be any CORBA object. WebSphere Application Server name servers register several initial contexts as initial references under predefined keys. Each name server initial reference is an instance of the interface `org.omg.CosNaming.NamingContext`.

Obtaining initial references in pure client processes

Pure JNDI clients, that is, JNDI clients which are not running in a WebSphere Application Server process, also have an ORB instance. This client ORB instance can be passed to the `InitialContext` constructor, but typically the initial context factory creates and initializes the client ORB instance transparently. A client ORB can be initialized with initial references, but the initial references most likely resolve to objects running in some server. The initial context factory does not define any default initial references when it initializes an ORB. If the `resolve_initial_references` method is invoked on the client ORB when no initial references have been configured, the method invocation fails. This condition is typical for pure client processes. To obtain an initial `NamingContext` reference, the initial context factory must invoke `string_to_object` with an IIOP type CORBA object URL, such as `corbaloc:iiop:myhost:2809`. The URL specifies the address of the server from which to obtain the initial context. The host and port information is extracted from the provider URL passed to the `InitialContext` constructor. If no provider URL is defined, the WebSphere Application Server initial context factory uses the default provider URL of `corbaloc:iiop:localhost:2809`. The `string_to_object` ORB method resolves the URL and communicates with the target server ORB to obtain the initial reference.

Obtaining initial references in server processes

If the JNDI client is running in a WebSphere Application Server process, the initial context factory obtains a reference to the server ORB instance if the JNDI client does not provide an ORB instance. Typically, JNDI clients running in server processes use the server ORB instance; that is, they do not pass an ORB instance to the `InitialContext` constructor. The name server which is running in the server process sets a provider URL as a `java.lang.System` property to serve as the default provider URL for all JNDI clients in the process. This default provider URL is `corbaloc:rir:/NameServiceServerRoot`. This URL resolves to the server root context for that server. (The URL is equivalent to invoking `resolve_initial_references` on the ORB with a key of `NameServiceServerRoot`. The name server registers the server root context as an initial reference under that key.)

Understanding the legacy ORB protocol

Previous versions of WebSphere Application Server used a different ORB implementation, which used a legacy protocol in contrast with the Interoperable Name Service (INS) protocol now used. This change has affected the implementation of the WebSphere Application Server initial context factory. **Certain types of pure clients can experience different behavior when getting initial JNDI**

contexts as compared to previous releases of WebSphere Application Server. This behavior is discussed in more detail below.

The following ORB properties are used with the legacy ORB protocol for ORB initialization and are now deprecated:

- `com.ibm.CORBA.BootstrapHost`
- `com.ibm.CORBA.BootstrapPort`

The new INS ORB is different in a major respect, in that it exhibits no default behavior if no initial references are defined. In the legacy ORB, the bootstrap host and port values defaulted to localhost and 900. All initial references were obtained from the server running on the bootstrap host and port. So, if the ORB user provided no bootstrap host and port, all initial references are resolved from the server running on the local host at port 900. The INS ORB has no concept of bootstrap host or bootstrap port. All initial references are defined independently. That is, different initial references could resolve to different servers. If `ORB.resolve_initial_references` is invoked with a key such that the ORB is not initialized with an initial reference having that key, the call fails.

In previous releases of WebSphere Application Server, the initial context factory invoked `resolve_initial_references` on the ORB in the absence of any provider URL. This action succeeded if a name server at the default bootstrap host and port was running. Today, with the INS ORB, this would fail. (Actually, the ORB would fall back to the legacy protocol during the deprecation period, but when the legacy protocol is no longer supported, the operation would fail.) The initial context factory now uses a default provider URL of `corbaloc:iiop:localhost:2809`, and invokes `string_to_object` with the provider URL. This operation preserves the behavior that pure clients in previous releases experienced when they set no ORB bootstrap properties or provider URL. **However, this different initial context factory implementation changes the behavior experienced by certain legacy pure clients, which do not specify a provider URL:**

- Clients which set the ORB bootstrap properties listed above when getting an initial context.
- Clients which supply their own ORB instance to the `InitialContext` constructor.

There are two ways to circumvent this change of behavior:

- Always specify an IIOP type provider URL. This approach does not depend on the bootstrap host and port properties and continues to work when support for the bootstrap host and port properties is removed. For example, you can express bootstrap host and port property values of `myHost` and `2809`, respectively, as `corbaloc:iiop:myHost:2809`.
- Use an `rir` type provider URL:
 - Specify `corbaloc:rir:/NameServiceServerRoot` if the ORB is initialized to use a WebSphere Application Server 5 server as the bootstrap server.
 - Specify `corbaname:rir:/NameService#domain/legacyRoot` if the ORB is initialized to use a WebSphere Application Server 4.0.x server as the bootstrap server.
 - Specify `corbaloc:rir:/NameService` if the ORB is initialized to use a server other than a WebSphere Application Server 5 or 4.0.x server as the bootstrap server.

URLs of this type are equivalent to invoking `resolve_initial_references` on the ORB with the specified key. If the bootstrap host and port properties are being used to initialize the ORB, this approach will not work when the bootstrap and host properties are no longer supported.

The `InitialContext` constructor search order for JNDI properties

If the code snippet shown at the beginning of this section is executed by an application, the bootstrap server depends on the value of the property, `java.naming.provider.url`. If the property is not set (in server processes the default value is set as a system property), the default host of `localhost` and default port of `2809` are used as the address of the server from which to obtain the initial context. The JNDI specification describes where the `InitialContext` constructor looks for `java.naming.provider.url` property settings, but briefly, the property is picked up from the following places in the order shown:

1. The `InitialContext` constructor. This does not apply to the above example since the example uses the empty `InitialContext` constructor.
2. System environment. You can add JNDI properties to the system environment as an option on the java command invocation and by program code. The recommended way to set the provider URL in the system environment is as an option supplied to the Java command invocation. Setting the provider URL in this manner is not temporal, so that getting a default initial context will always yield the same result. It is generally recommended that program code not set the provider URL property in the system environment because as a side-effect, this could adversely affect other, possibly unrelated, code running elsewhere in the same process.
3. `jndi.properties` file. There may be many `jndi.properties` files that are within the scope of the class loader in effect. All `jndi.properties` files are used for setting JNDI properties, but the provider URL setting is determined by the first `jndi.properties` file returned by the class loader.

Example: Getting an initial context by setting the provider URL property

In general, JNDI clients should assume the correct environment is already configured so there is no need to explicitly set property values and pass them to the `InitialContext` constructor. However, a JNDI client may need to access a name space other than the one identified in its environment. In this case, it is necessary to explicitly set the `java.naming.provider.url` (provider URL) property used by the `InitialContext` constructor. A provider URL contains bootstrap server information that the initial context factory can use to obtain an initial context. Any property values passed in directly to the `InitialContext` constructor take precedence over settings of those same properties found elsewhere in the environment.

You can use two different provider URL forms with WebSphere Application Server's initial context factory:

- A CORBA object URL (new for J2EE 1.3)
- An IIOP URL

CORBA object URLs are more flexible than IIOP URLs and are the recommended URL format to use. CORBA object URLs are part of the OMG CosNaming Interoperable Naming Specification. A `corbaname` URL, for example, can include initial context and lookup name information and can be used as a lookup name

without the need to explicitly obtain another initial context. The IIOP URLs are the legacy JNDI format, but are still supported by the WebSphere Application Server initial context factory.

The following examples illustrate the use of these URLs.

Using a CORBA object URL

This example shows a CORBA object URL.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
    "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
    "corbaloc:iiop:myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
...
```

Using a CORBA object URL with multiple name server addresses

CORBA object URLs can contain more than one bootstrap address. You can use this feature when attempting to obtain an initial context from a server cluster. You can specify the bootstrap addresses for all servers in the cluster in the URL. The operation succeeds if at least one of the servers is running, eliminating a single point of failure. There is no guarantee of any particular order in which the address list will be processed. For example, the second bootstrap address may be used to obtain the initial context even though the server at the first bootstrap address in the list is available.

Multiple-address provider URLs should only contain the bootstrap addresses of members of the same cluster. Otherwise, incorrect behavior may occur.

An example of a corbaloc URL with multiple addresses follows.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
    "com.ibm.websphere.naming.WsnInitialContextFactory");
// All of the servers in the provider URL below are
// members of the same cluster.
env.put(Context.PROVIDER_URL,
    "corbaloc::myhost1:9810,:myhost1:9811,:myhost2:9810");
Context initialContext = new InitialContext(env);
...
```

Using a CORBA object URL from a non-WebSphere Application Server JNDI

implementation

Initial context factories for CosNaming JNDI plug-in implementations other than the WebSphere Application Server initial context factory most likely obtain an

initial context using the object key, NameService. When you use such a context factory to obtain an initial context from a WebSphere Application Server name server, the initial context is the cell root context. Since system artifacts such as EJB homes associated with a server are bound under the server's server root context, names used in JNDI operations must be qualified. If you want to use relative names, ensure your initial context is the server root context under which the target object is bound. In order to make the server root context the initial context, specify a corbaloc provider URL with an object key of NameServiceServerRoot.

This example shows a CORBA object type URL from a non-WebSphere Application Server JNDI implementation. This example assumes full CORBA object URL support by the non-WebSphere Application Server JNDI implementation. The object key of NameServiceServerRoot is specified so that the initial context will be the specified server's server root context.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
"com.somecompany.naming.TheirInitialContextFactory");
env.put(Context.PROVIDER_URL,
"corbaname:iiop:myhost.mycompany.com:9810/
NameServiceServerRoot");
Context initialContext = new InitialContext(env);
...
```

If qualified names are used, you can use the default key of NameService.

Using an IIOP URL

The IIOP type of URL is a legacy format which is not as flexible as CORBA object URLs. However, URLs of this type are still supported. The following example shows an IIOP type URL as the provider URL.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
"com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
"iiop://myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
...
```

Example: Setting the provider URL property to select a different root context as the initial context

Each server contains its own server root context, and, when bootstrapping to a server, the server root is the default initial JNDI context. Most of the time, this default is the desired initial context, since system artifacts such as EJB homes are bound there. However, other root contexts exist, which can contain bindings of interest. It is possible to specify a provider URL to select other root contexts.

Selecting the initial root context with a CORBA object URL

There are several object keys registered with the bootstrap server that you can use to select the root context for the initial context. To select a particular root context with a CORBA object URL object key, set the object key to the corresponding value. The default object key is `NameService`. Using JNDI yields the server root context. A table that lists the different root contexts and their corresponding object key follows:

Root Context	CORBA Object URL Object Key
Server Root	<code>NameServiceServerRoot</code>
Cell Persistent Root	<code>NameServiceCellPersistentRoot</code>
Cell Root	<code>NameServiceCellRoot</code>
Node Root	<code>NameServiceNodeRoot</code>

The following example shows the use of a `corbaloc` URL with the object key set to select the cell persistent root context as the initial context.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
        "corbaloc:iiop:myhost.mycompany.com:2809/NameServiceCellPersistentRoot");
Context initialContext = new InitialContext(env);
...
```

Selecting the initial root context with the name space root property

You can also select the initial root context by passing a name space root property setting to the `InitialContext` constructor. Generally, the object key setting described above is sufficient. Sometimes a property setting is preferable. For example, you can set the root context property on the Java invocation to make which server root is being used as the initial context transparent to the application. The default server root property setting is `defaultroot`, which yields the server root context.

Root Context	Name Space Root Property Value
Server Root	<code>bootstrapserverroot</code>
Cell Persistent Root	<code>cellpersistentroot</code>
Cell Root	<code>cellroot</code>
Node Root	<code>bootstrapnoderoot</code>

The initial context factory ignores the name space root property if the provider URL contains an object key other than `NameService`.

The following example shows use of the name space root property to select the cell persistent root context as the initial context. Note that available constants are used instead of hardcoding the property name and value.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
```

```

import javax.naming.InitialContext;
import com.ibm.websphere.naming.PROPS;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "corbaloc:iiop:myhost.mycompany.com:2809");
env.put(PROPS.NAME_SPACE_ROOT, PROPS.NAME_SPACE_ROOT_CELL_PERSISTENT);
Context initialContext = new InitialContext(env);
...

```

Example: Looking up an EJB home with JNDI

Most applications which use JNDI run in a container. Some do not. The name used to look up an object depends on whether or not the application is running in a container. The examples below show lookups from each type of application. Sometimes it is more convenient for an application to use a corbaname URL as the lookup name. Container-based JNDI clients and thin Java clients can use a corbaname URL. An example of a lookup with a corbaname URL is also included in this section.

JNDI lookup from an application running in a container

Applications that run in a container can use `java:` lookup names. Lookup names of this form provide a level of indirection such that the lookup name used to look up an object is not dependent on the object's name as it is bound in the name server's name space. The deployment descriptors for the application provide the mapping from the `java:` name and the name server lookup name. The container sets up the `java:` name space based on the deployment descriptor information so that the `java:` name is correctly mapped to the corresponding object.

The following example shows a lookup of an EJB home. The actual home lookup name is determined by the application's deployment descriptors.

```

// Get the initial context as shown in a previous example
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup
        ("java:comp/env/com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}

```

JNDI lookup from an application that does not run in a container

Applications that do not run in a container cannot use `java:` lookup names because it is the container which sets the `java:` name space up for the application. Instead, an application of this type must look the object up directly from the name server. Each application server contains a name server. System artifacts such as EJB homes are bound relative to the server root context in that name server. The various name servers are federated by means of a system name space structure. The recommended way to look up objects on different servers is to qualify the name so that the name resolves from any initial context in the cell. If a relative name is used, the initial context must be the same server root context as the one under which the object is bound. The form of the qualified name depends on whether the qualified name is a topology-based name or a fixed name. A topology based name depends on whether the object resides in a single server or a server cluster. Examples of each form of qualified name follow.

Topology-based qualified names

Topology-based qualified names traverse through the system name space to the server root context under which the target object is bound. A topology-based qualified name resolves from any initial context in the cell. The topology-based qualified name depends on whether the object resides on a single server or server cluster. Examples of each lookup follow.

Single server

The following example shows a lookup of an EJB home that is running in the single server, `MyServer`, configured in the node, `Node1`.

```
// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup(
        "cell/nodes/Node1/servers/MyServer
        /com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)javad.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}
```

Server cluster

The example below shows a lookup of an EJB home which is running in the cluster, `MyCluster`. The name can be resolved if any of the cluster members is running.

```
// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup(
        "cell/clusters/MyCluster/com/mycompany/
        accounting/AccountEJB");
    accountHome = (AccountHome)javad.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}
```

Fixed qualified names

If the target object has a cell-scoped fixed name defined for it, you can use its qualified form instead of the topology-based qualified name. Even though the topology-based name works, the fixed name does not change with the specific cell topology or with the movement of the target object to a different server. An example lookup with a qualified fixed name is shown below.

```
// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup(
        "cell/persistent/com/mycompany/accounting/AccountEJB");
}
```

```

        accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
            (org.omg.CORBA.Object) ejbHome, AccountHome.class);
    }
    catch (NamingException e) { // Error getting the home interface
        ...
    }
}

```

JNDI lookup with a corbaname URL

A corbaname can be useful at times as a lookup name. If, for example, the target object is not a member of the federated name space and cannot be located with a qualified name, a corbaname can be a convenient way to look up the object. A lookup with a corbaname URL follows.

```

// Get the initial context as shown in a previous example
...
// Look up the home interface using a corbaname URL
try {
    java.lang.Object ejbHome = initialContext.lookup
        ("corbaname:iiop:someHost:2809#com/mycompany/accounting/
        AccountEJB");
    accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) {
// Error getting the home interface
...
}

```

Example: Looking up a JavaMail session with JNDI

The example below shows a lookup of a JavaMail resource. The actual lookup name is determined by the application's deployment descriptors.

Usage scenario

```

// Get the initial context as shown above
...
Session session = (Session)
    initialContext.lookup("java:comp/env/mail/MailSession");

```

JNDI interoperability considerations

This section explains considerations to take into account when interoperating with previous releases of WebSphere Application Server and with non-WebSphere Application Server JNDI clients. Also, the way resources from MQSeries must be bound to the name space has changed and is described below.

Interoperability with previous WebSphere Application Server Releases

EJB clients running on WebSphere Application Server v3.5 or v4.0 accessing EJB applications running on WebSphere Application Server v5

Applications migrated from previous versions of WebSphere Application Server may still have clients still running in a previous release. The default initial JNDI context for EJB clients running on previous versions of WebSphere Application Server is the cell persistent root (legacy root). The home for an enterprise bean deployed in version 5 is bound to its server's server root context. In order for the EJB lookup name for down-level clients to remain unchanged, configure a binding for the EJB home under the cell persistent root.

Note: EJB clients running in version 3.5 must be running in version 3.5.5 or above, or in version 3.5.3 or 3.5.4 with e-fix PQ51387 installed.

EJB clients running on WebSphere Application Server v5 accessing EJB applications running on WebSphere Application Server v3.5 or v4.0 servers

The default initial context for a WebSphere Application Server v3.5 or v4.0 server is the correct initial context. Simply look up the JNDI name under which the EJB home is bound.

Note: To enable WebSphere Application Server v5 clients to access version 3.5.x and 4.0.x servers, the down-level installations must have e-fix PQ60074 installed.

EJB clients running in an environment other than WebSphere Application Server accessing EJB applications running on WebSphere Application Server v5 servers

When an EJB application running in WebSphere Application Server v5 is accessed by a non-WebSphere Application Server EJB client, the JNDI initial context factory is presumed to be a non-WebSphere Application Server implementation. In this case, the default initial context will be the cell root. If the JNDI service provider being used supports CORBA object URLs, the corbaname format can be used to look up the EJB home. The construction of the stringified name depends on whether the object is installed on a single server or cluster, as shown below.

Single server

```
initialContext.lookup("corbaname:iiop:myHost:2809#cell/  
nodes/node1/servers/server1/myEJB");
```

According to the URL above, the bootstrap host and port are **myHost** and **2809**, and the enterprise bean is installed in a server **server1** in node **node1** and bound in that server under the name **myEJB**.

Server cluster

```
initialContext.lookup("corbaname:iiop:myHost:2809#cell/  
clusters/myCluster/myEJB");
```

According to the URL above, the bootstrap host and port are **myHost** and **2809**, and the enterprise bean is installed in a server cluster named **myCluster** and bound in that cluster under the name **myEJB**.

The above lookup will work with any name server bootstrap host and port configured in the same cell.

The above lookup will also work if the bootstrap host and port belongs to a member of the cluster itself. To avoid a single point of failure, the bootstrap server host and port for each cluster member could be listed in the URL as follows:

```
initialContext.lookup("corbaname:iiop:host1:9810,  
host2:9810#cell/clusters/myCluster/myEJB");
```

The name prefix **cell/clusters/myCluster/** is not necessary if bootstrapping to the cluster itself, but it will work. The prefix is needed, however, when looking up enterprise beans in other clusters. Name bindings under the **clusters** context are implemented on the name server to resolve to the server root of a running cluster member during a lookup; thus avoiding a single point of failure.

Without CORBA object URL support

If the JNDI initial context factory being used does not support CORBA object URLs, the initial context can be obtained from the server, and the lookup can be performed on the initial context as follows:

```
Hashtable env = new Hashtable();
env.put(CONTEXT.PROVIDER_URL, "iiop://myHost:2809");
Context ic = new InitialContext(env);
Object o = ic.lookup("cell/clusters/myCluster/myEJB");
```

Binding resources from MQSeries 5.2

In previous releases of WebSphere Application Server, the MQSeries jmsadmin tool could be used bind resources to the name space. When used with a WebSphere Application Server v5 name space, the resource will be bound within a transient partition in the name space and will not persist past the life of the server process. Instead of binding the MQSeries resources with the jmsadmin tool, bind them from the WebSphere Application Server administrative console, under Resources in the left panel on the console

JNDI caching

To increase the performance of JNDI operations, the WebSphere Application Server JNDI implementation employs caching to reduce the number of remote calls to the name server for lookup operations. For most cases, use the default cache setting.

When an InitialContext object is instantiated, an association is established between the InitialContext instance and a cache. The initial context and any contexts returned directly or indirectly from a lookup on the initial context are all associated with that same cache instance. By default, the association is based on the provider URL, in particular, the host name and port. The caller can specify the cache name to override this default behavior. A cache instance of a given name is shared by all instances of InitialContext configured to use a cache of that name which were created with the same context class loader in effect. Two EJB applications running in the same server will use their own cache instances, if they are using different context class loaders, even if the cache names are the same.

After an association between an InitialContext instance and cache is established, the association does not change. A `javax.naming.Context` object returned from a lookup operation inherits the cache association of the Context object on which the lookup was performed. Changing cache property values with the `Context.addToEnvironment()` or `Context.removeFromEnvironment()` method does not affect cache behavior. You can change properties affecting a given cache instance with each InitialContext instantiation.

A cache is restricted to a process and does not persist past the life of that process. A cached object is returned from lookup operations until either the max cache life for the cache is reached, or the max entry life for the object's cache entry is reached.

After this time, a lookup on the object causes the cache entry for the object to be refreshed. If a bind or rebind operation is executed on an object, the change is not reflected in any caches other than the one associated with the context from which the bind or rebind was issued. This scenario is most likely to happen when multiple processes are involved, since different processes do not share the same cache, and context objects in all threads in a process typically share the same cache instance for a given name service provider.

Usually, cached objects are relatively static entities, and objects becoming stale are not a problem. However, you can set timeout values on cache entries or on a cache so that cache contents are periodically refreshed.

JNDI cache settings

Various cache property settings follow. Ensure that all property values are string values.

com.ibm.websphere.naming.jndicache.cachename

The name of the cache to associate with an initial context instance can be specified with this property.

It is possible to create multiple InitialContext instances, each operating on the name space of a different name server. By default, objects from each bootstrap address are cached separately, since they each involve independent name spaces and name collisions could occur if they used the same cache. The provider URL specified when the initial context is created by default serves as the basis for the cache name. With this property, a JNDI client can specify a cache name. Valid options for cache names follow:

Valid options	Resulting cache behavior
providerURL (default)	Use the value for java.naming.provider.url property as the basis for the cache name. Cache names are based on the bootstrap host and port specified in the URL. The bootstrap host is normalized to a fully qualified name, if possible. For example, corbaname:iiop:server1:2809#some/starting/context and corbaloc:iiop://server1 are normalized to the same cache name. If no provider URL is specified, a default cache name is used.
Any string	Use the specified string as the cache name. You can use any arbitrary string with a value other than "providerURL" as a cache name.

com.ibm.websphere.naming.jndicache.cacheobject

Turn caching on or off and clear an existing cache with this property.

By default, when an InitialContext is instantiated, it is associated with an existing cache or, if one does not exist, a new one is created. An existing cache is used with its existing contents. In some circumstances, this behavior is not desirable. For example, when objects that are looked up change frequently, they can become stale in the cache. Other options are available. The following table lists these other options along with the corresponding property value.

Valid values	Resulting cache behavior
populated (default)	Use a cache with the specified name. If the cache already exists, leave existing cache entries in the cache; otherwise, create a new cache.
cleared	Use a cache with the specified name. If the cache already exists, clear all cache entries from the cache; otherwise, create a new cache.

none	Do not cache. If this option is specified, the cache name is irrelevant. Therefore, this option will not disable a cache that is already associated with other InitialContext instances. The InitialContext that is instantiated is not associated with any cache.
------	--

com.ibm.websphere.naming.jndicache.maxcachelife

Impose a limit to the age of a cache with this property.

By default, cached objects remain in the cache for the life of the process or until cleared with the `com.ibm.websphere.naming.jndicache.cacheobject` property set to "cleared". This property enables a JNDI client to set the maximum life of a cache. This property differs from the `maxentrylife` property (below) in that the entire cache is cleared when the cache lifetime is reached. The table below lists the various `maxcachelife` values and their affect on cache behavior:

Valid options	Resulting cache behavior
0 (default)	Make the cache lifetime unlimited.
Positive integer	Set the maximum lifetime of the entire cache, in minutes, to the specified value. When the maximum lifetime for the cache is reached, the next attempt to read any entry from the cache causes the cache to be cleared

com.ibm.websphere.naming.jndicache.maxentrylife

Impose a limit to the age of individual cache entries with this property.

By default, cached objects remain in the cache for the life of the process or until cleared with the `com.ibm.websphere.naming.jndicache.cacheobject` property set to `cleared`. This property enables a JNDI client to set the maximum lifetime of individual cache entries. This property differs from the `maxcachelife` property in that individual entries are refreshed individually as their maximum lifetime reached. This might avoid any noticeable change in performance that might occur if the whole cache is cleared at once. The table below lists the various `maxentrylife` values and their effect on cache behavior:

Valid options	Resulting cache behavior
0 (default)	Lifetime of cache entries is unlimited.
Positive integer	Set the maximum lifetime of individual cache entries, in minutes, to the specified value. When the maximum lifetime for an entry is reached, the next attempt to read the entry from the cache causes the individual cache entry to refresh.

Example: Controlling JNDI cache behavior from a program

Following are examples that illustrate how you can use JNDI cache properties to achieve the desired cache behavior. Cache properties take effect when an InitialContext object is constructed.

Usage scenario

```

import java.util.Hashtable;
import javax.naming.InitialContext;
import javax.naming.Context;
import com.ibm.websphere.naming.PROPS;

/*****
 Caching discussed in this section pertains to the
 WebSphere Application Server initial context factory.
 Assume the property, java.naming.factory.initial, is set
 to "com.ibm.websphere.naming.WsnInitialContextFactory"
 as a java.lang.System property.
 *****/

Hashtable env;
Context ctx;

// To clear a cache:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT,
        PROPS.JNDI_CACHE_OBJECT_CLEARED);
ctx = new InitialContext(env);

// To set a cache's maximum cache lifetime to 60 minutes:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_MAX_LIFE, "60");
ctx = new InitialContext(env);

// To turn caching off:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT,
        PROPS.JNDI_CACHE_OBJECT_NONE);
ctx = new InitialContext(env);

// To use caching and no caching:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT,
        PROPS.JNDI_CACHE_OBJECT_POPULATED);
ctx = new InitialContext(env);
env.put(PROPS.JNDI_CACHE_OBJECT,
        PROPS.JNDI_CACHE_OBJECT_NONE);
Context noCacheCtx = new InitialContext(env);

Object o;

// Use caching to look up home, since the home should
// rarely change.
o = ctx.lookup("com/mycom/MyEJBHome");
// Narrow, etc. ...

// Do not use cache if data is volatile.
o = noCacheCtx.lookup("com/mycom/VolatileObject");
//
...

```

JNDI name syntax

JNDI name syntax is the default syntax and is suitable for typical JNDI clients.

This syntax includes the following special characters: forward slash (/) and backslash (\). Components in a name are delimited by a forward slash. The

backslash is used as the escape character. A forward slash is interpreted literally if it is escaped, that is, preceded by a backslash. Similarly, a backslash is interpreted literally if it is escaped.

INS name syntax

INS syntax is designed for JNDI clients that need to interoperate with CORBA applications.

The INS syntax allows a JNDI client to make the proper mapping to and from a CORBA name. INS syntax is very similar to the JNDI syntax with the additional special character, dot (.). Dots are used to delimit the `id` and `kind` fields in a name component. A dot is interpreted literally when it is escaped. Only one unescaped dot is allowed in a name component. A name component with a non-empty `id` field and empty `kind` field is represented with only the `id` field value and must not end with an unescaped dot. An empty name component (empty `id` and empty `kind` field) is represented with a single unescaped dot. An empty string is not a valid name component representation.

JNDI to CORBA name mapping considerations

WebSphere Application Server name servers are an implementation of the CORBA CosNaming interface. WebSphere Application Server provides a JNDI implementation which you can use to access CosNaming name servers through the JNDI interface. Issues can exist when mapping JNDI name strings to and from CORBA names.

Each component in a CORBA name consists of an `id` and `kind` field, but a JNDI name component consists of no such fields. Each component in a JNDI name is atomic. Typical JNDI clients do not need to make a distinction between the `id` and `kind` fields of a name component, or know how JNDI name strings map to CORBA names. JNDI clients of this sort can use the JNDI syntax described below. When a name is parsed according to JNDI syntax, each name component is mapped to the `id` field of the corresponding CORBA name component. The `kind` field always has an empty value. This basic syntax is the least obtrusive to the JNDI client in that it has the fewest special characters. However, you cannot represent with this syntax a CORBA name with a non-empty `kind` field. This restriction can prevent EJB applications from interoperating with CORBA applications.

Some clients, however must interoperate with CORBA applications which use CORBA names with non-empty `kind` fields. These JNDI clients must make a distinction between `id` and `kind` so that JNDI names are correctly mapped to CORBA names, particularly when the CORBA names contain components with non-null `kind` fields. Such JNDI clients can use the INS name syntax. With its additional special character, you can use INS to represent any CORBA name. Use of this syntax is not recommended unless it is necessary, because this syntax is more restrictive from the JNDI client's perspective in that the JNDI client must be aware that name components with multiple unescaped dots are syntactically invalid. INS name syntax is part of the OMG CosNaming Interoperable Naming Specification.

Example: Setting the syntax used to parse name strings

JNDI clients which must interoperate with CORBA applications may need to use INS name syntax to represent names in string format. The name syntax property may be passed to the `InitialContext` constructor through its parameter, in the

System properties, or in a `jndi.properties` file. The initial context and any contexts looked up from that initial context will parse name strings based on the specified syntax.

The following example shows how to set the name syntax to make the initial context parse name strings according to INS syntax.

Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
import com.ibm.websphere.naming.PROPS;
// WebSphere naming constants
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, ...);
env.put(PROPS.NAME_SYNTAX, PROPS.NAME_SYNTAX_INS);
Context initialContext = new InitialContext(env);
// The following name maps to a CORBA name component
// as follows:
// id = "a.name", kind = "in.INS.format"
// The unescaped dot is used as the delimiter.
// Escaped dots are interpreted literally.
java.lang.Object o =
    initialContext.lookup("a\\.name.in\\.INS\\.format");
...
```

Developing applications that use CosNaming (CORBA Naming interface)

CORBA clients can perform naming operations on WebSphere name servers through the CosNaming interface. The following examples show how to obtain an ORB instance and an initial context as well as how to look up an EJB home.

Note: To enable WebSphere v5 clients to access Versions 3.5.x and 4.0.x servers, the earlier installations must have e-fix PQ60074 installed.

Steps for this task

1. Get an initial context
2. Perform desired CosNaming operations

Example: Getting an initial context with CosNaming

In the WebSphere Application Server, an initial context is obtained from a bootstrap server. The address for the bootstrap server consists of a host and port. To get an initial context, you must know the host and port for the server that is used as the bootstrap server.

Obtaining an initial context consists of two basic steps:

1. Obtain an ORB reference
2. Invoke a method on the ORB to obtain the initial reference

These steps are now explained in more detail.

Obtaining an ORB reference

Pure CosNaming clients, that is clients that are not running in a server process, must create and initialize an ORB instance with which to obtain the initial context. CosNaming clients which run in server processes can obtain a reference to the server ORB with a JNDI lookup. The following examples illustrate how to create and initialize a client ORB and how to obtain a server ORB reference.

Creating a client ORB instance

To create an ORB instance, invoke the static method, `org.omg.CORBA.ORB.init`. The `init` method requires a property set to the name of the ORB class you want to instantiate. An ORB implementation with the class name `com.ibm.CORBA.iiop.ORB` is included with the WebSphere Application Server. The WebSphere Application Server ORB recognizes additional properties with which you can specify initial references.

The basic steps for creating an ORB are as follows:

1. Create a Properties object.
2. Set the ORB class property to WebSphere Application Server's ORB class.
3. If the bootstrap server is INS-compliant, set the initial reference properties. If the bootstrap server is not INS-compliant (meaning, WebSphere Application Server v4.0.x or earlier), set bootstrap host and port for bootstrap server.
4. Invoke `ORB.init`, passing in the Properties object.

Usage scenario

```
...
import java.util.Properties;
import org.omg.CORBA.ORB;
...
Properties props = new Properties();
props.put("org.omg.CORBA.ORBClass",
"com.ibm.CORBA.iiop.ORB");
props.put("com.ibm.CORBA.ORBInitRef.NameService",
"corbaloc:iiop:myhost.mycompany.com:2809/NameService");
props.put("com.ibm.CORBA.ORBInitRef.NameServiceServerRoot",
"corbaloc:iiop:myhost.mycompany.com:2809/NameServiceServerRoot");
// props.put("com.ibm.CORBA.BootstrapHost",
// "myhost.mycompany.com");
// Use this if bootstrap server is WebSphere 4.0.x or before
// props.put("com.ibm.CORBA.BootstrapPort", "2809");
// Use this if bootstrap server is WebSphere 4.0.x or before
ORB _orb = ORB.init((String[])null, props);
...
```

Notice the initial reference definitions for `NameService` and `NameServiceServerRoot`. The initial context returned for `NameService` depends on the type of bootstrap server. The key `NameServiceServerRoot` is a key introduced in WebSphere Application Server v5. For more information on initial contexts, see the section [Initial Contexts](#).

Note: The properties `com.ibm.CORBA.BootstrapHost` and `com.ibm.CORBA.BootstrapPort` are deprecated. They are needed, however, to connect to WebSphere Application Servers of Version 4.0.x or earlier. The default bootstrap host is the local host and the default port is 2809.

Obtaining a reference to the server ORB

CosNaming clients which run in a server process can obtain a reference to the server ORB with a JNDI lookup on a java: name, shown as follows:

Usage scenario

```
...
import javax.naming.Context;
import javax.naming.InitialContext;
import org.omg.CORBA.ORB;
...
Context initialContext = new InitialContext();
ORB orb = (ORB) initialContext.lookup("java:comp/ORB");
...
```

Using an ORB reference to get an initial naming reference

There are two basic ways to get an initial CosNaming context. Both ways involve an ORB method invocation. The first way is to invoke the `resolve_initial_references` method on the ORB with an initial reference key. For this call to work, the ORB must be initialized with an initial reference for that key. The other way is to invoke the `string_to_object` method on the ORB, passing in a CORBA object URL with the host and port of the bootstrap server. The following examples illustrate both approaches.

Invoking `resolve_initial_references`

Once an ORB reference is obtained, invoke the `resolve_initial_references` method on the ORB to obtain a reference to the initial context. The following code example invokes `resolve_initial_references` on an ORB reference.

Usage scenario

```
...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Obtain ORB reference as shown in examples earlier
// in this section
...
org.omg.CORBA.Object obj =
    _orb.resolve_initial_references("NameService");
NamingContextExt initCtx =
    NamingContextExtHelper.narrow(obj);
...
```

Note that the key `NameService` is passed to the `resolve_initial_references` method. Other initial context keys are registered in WebSphere Application Servers. For example, `NameServiceServerRoot` can be used to obtain a reference to the server root context in the bootstrap name server. For more information on the initial contexts registered in server ORBs, please see the section `Initial Contexts`.

Invoking `string_to_object` with a CORBA object URL

You can use an INS-compliant ORB to obtain an initial context even if the ORB is not initialized with any initial references or bootstrap properties, or if those property settings are for a different server than the name server from which you want to obtain the initial context. To obtain an initial context by explicitly specifying the bootstrap name server, invoke the `string_to_object` method on the ORB, passing in a CORBA object URL which contains the bootstrap server host and port.

The code in the example below invokes the `string_to_object` method on an existing ORB reference, passing in a CORBA object URL which identifies the desired initial context.

Usage scenario

```
...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Obtain ORB reference as shown in examples earlier
// in this section
...
org.omg.CORBA.Object obj = orb.string_to_object
("corbaloc:iiop:myhost.mycompany.com:2809/NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...
```

Note that the key `NameService` is used in the `corbaloc` URL. Other initial context keys are registered in WebSphere Application Servers. For example, you can use `NameServiceServerRoot` to obtain a reference to the server root context in the bootstrap name server.

Using an existing ORB and invoking `string_to_object` with a CORBA object URL with multiple name server addresses to get an initial context

CORBA object URLs can contain more than one bootstrap server address. Use this feature when attempting to obtain an initial context from a server cluster. You can specify the bootstrap server addresses for all servers in the cluster in the URL. The operation will succeed if at least one of the servers is running, eliminating a single point of failure. There is no guarantee of any particular order in which the address list will be processed. For example, the second bootstrap server address may be used to obtain the initial context even though the first bootstrap server in the list is available. An example of a `corbaloc` URL with multiple addresses follows.

Usage scenario

```
...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Assume orb is an existing ORB instance
org.omg.CORBA.Object obj = orb.string_to_object(
"corbaloc:myhost1:9810,:myhost1:9811,
:myhost2:9810/NameService");
NamingContextExt initCtx=NamingContextExtHelper.narrow(obj);
...
```

Example: Looking up an EJB home with `CosNaming`

You can look up an EJB home or other CORBA object from a WebSphere Application Server name server through the CORBA `CosNaming` interface. You can invoke `resolve` or `resolve_str` on the initial context, or you can invoke `string_to_object` on the ORB. You can use a qualified name so that the name resolves regardless of which name server the lookup is executed on, or use an unqualified name that only resolves from the server root context on the name server that actually contains the object binding. (The qualified name traverses the federated system name space to the specified server root context.)

Qualified and unqualified names

Each application server contains a name server. System artifacts such as EJB homes are bound in that name server. The various name servers are federated by means of a system name space structure. The recommended way to look up objects on different servers is to use a qualified name. A qualified name can be a topology-based name, based on the name of the cluster or single server and node that contains the object. You can define fixed qualified names for objects. With qualified names, you can look up objects residing on different servers from the same initial context by traversing the system name space structure. Alternatively, you can use an unqualified name, but an unqualified name will only resolve using the name server associated with the object's application server.

CosNaming.resolve (and resolve_str) vs. ORB.string_to_object

If you have an initial context from any name server in a WebSphere Application Server cell, you can look up any CORBA object with a qualified name. You do not need additional host and port information for the target object's name server.

Alternatively, you can look up an object by invoking `string_to_object` on the ORB, passing in a corbaname URL. Typically, an IIOP type URL is specified, so the bootstrap address information required for an initial context must be contained in the URL. You can use a qualified or unqualified stringified name, but an unqualified name resolves only if the initial context is from the name server in which the object is bound.

The following examples show CosNaming resolve operations using qualified topology-based lookup names and an unqualified lookup name.

CosNaming resolve operation using a qualified name

The topology-based qualified name for an object depends on whether the object is bound in a single server or a server cluster. Examples of each follow.

Single Server

The following example shows the lookup of an EJB home that is running in a single server. The enterprise bean that is being looked up is running in the server, `MyServer`, on the node, `Node1`.

```
// Get the initial context as shown in the previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the name under which
// the EJB home is bound
org.omg.CORBA.Object ejbHome =
    initialContext.resolve_str(
        "cell/nodes/Node1/servers/MyServer
        /mycompany/accounting/AccountEJB");
accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow
    (ejbHome, AccountHome.class);
```

Server Cluster

The following example shows a lookup of an EJB home that is running in a cluster. The enterprise bean being that is looked up is running in the cluster, `Cluster1`. The name can be resolved if any of the cluster members is running.

Usage scenario


```

// Get the initial context as shown in the previous example
// Using the form of lookup name below, it doesn't
// matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = initialContext.resolve_str(
"cell/clusters/Cluster1/
mycompany/accounting/AccountEJB");
accountHome=(AccountHome)javax.rmi.PortableRemoteObject.narrow
(ejbHome, AccountHome.class);

```

ORB string_to_object operation using an unqualified stringified name

If the resolve operation is being performed on the name server that contains the object, the system name space does not need to be traversed, and you can use an unqualified lookup name. Note that this name does not resolve on other name servers. If an unqualified name is provided, the object key must be NameServiceServerRoot so that the correct initial context is selected. If a qualified name is provided, you can use the default key of NameService.

The following example shows a lookup of an EJB home. The enterprise bean that is being looked up is bound on the name server running on the host myHost on port 2809. Note the object key of NameServiceServerRoot.

Usage scenario

```

// Assume orb is an existing ORB instance
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome=orb.string_to_object(
"corbaname:iiop:myHost:2809/NameServiceServerRoot
#mycompany/accounting");
accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow
(ejbHome, AccountHome.class);

```

Configured name bindings

Administrators can configure bindings into the name space. A configured binding is different from a programmatic binding in that the system creates the binding every time a server is started, even if the target context is in a transient partition.

Administrators can add name bindings to the name space through the configuration. Name servers add these configured bindings to the name space view, by reading the configuration data for the bindings. Configuring bindings is an alternative to creating the bindings from a program. Configured bindings have the advantage of being created each time a server starts, even when the binding is created in a transient partition of the name space. Cell-scoped configured bindings provide interoperability with JNDI clients running on previous versions of WebSphere Application Server. Additionally, you can configure cell-scoped bindings to create a fixed qualified name for server objects.

Scope

You can configure a binding at one of the following three scopes: cell, node, or server. Cell-scoped bindings are created under the cell persistent root context. Node-scoped bindings are created under the node persistent root context for the specified node. Server-scoped bindings are created under the server root context

for the selected server. If the target server of a server-scoped binding is a cluster, the binding is created under the server root context of each cluster member.

Note: The term *server* includes clusters and can be used interchangeably with the term *cluster* with respect to configured bindings. When applied to a cluster, a server-scoped binding is created in the server root for all member servers.

The scope you select for new bindings depends on how the binding is to be used. For example, if the binding is not specific to any particular node or server, or if you do not want the binding to be associated with any specific node or server, a cell-scoped binding is a suitable scope. Defining fixed names for enterprise beans to create fixed qualified names is just such an application. If a binding is to be used only by clients of an application running on a particular server, or if you want to configure a binding with the same name on different servers which resolve to different objects, a server-scoped binding would be appropriate. Note that two servers can have configured bindings with the same name but resolve to different objects. At the cell scope, only one binding with a given name can exist.

Intermediate Contexts

Intermediate contexts created with configured bindings are read-only. For example, if an EJB home binding is configured with the name `some/compound/name/ejbHome`, the intermediate contexts `some`, `some/compound`, and `some/compound/name` will be created as read-only contexts. You cannot add, update, or remove any read-only bindings.

The configured binding name cannot conflict with existing bindings. However, configured bindings can use the same intermediate context names. Therefore, a configured binding with the name `some/compound/name2/ejbHome2` does not conflict with the previous example name.

Configured binding types

Types of objects that you can bind follow:

The following data is required to configure an EJB home binding:

- JNDI name of the EJB server or server cluster where the enterprise bean is deployed
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

This type of binding is of special significance because you can use it to provide interoperability with WebSphere Application Server v3.5.x and v4.0.x JNDI clients. The default initial context for these earlier clients is the cell persistent root, which is different from the initial context of the server root for WebSphere Application Server v5 JNDI clients. If you migrate an application to the current release, you can configure an EJB binding at the cell scope so that the lookup names for the enterprise bean do not change for clients still running in an earlier WebSphere Application Server version.

A cell-scoped EJB binding is also useful for creating a fixed lookup name for an enterprise bean so that the qualified name is not dependent on the topology.

CORBA: CORBA object available from some CosNaming name server

You can identify any CORBA object bound into some INS compliant CosNaming server with a corbaname URL. The referenced object does not have to be available until the binding is actually referenced by some application.

The following data is required in order to configure a CORBA object binding:

- The corbaname URL of the CORBA object
- An indicator if the bound object is a context or leaf node object (to set the correct CORBA binding type of context or object).
- Target root for the configured binding
- The name of the configured binding, relative to the target root.

Indirect: Any object bound in WebSphere Application Server name space accessible with JNDI

Besides CORBA objects, this includes `javax.naming.Referenceable`, `javax.naming.Reference`, and `java.io.Serializable` objects. The target object itself is not bound to the name space. Only the information required to look up the object is bound. Therefore, the referenced name server does not have to be running until the binding is actually referenced by some application. The following data is required in order to configure an indirect JNDI lookup binding:

- JNDI provider URL of name server where object resides
- JNDI lookup name of object
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

A cell-scoped indirect binding is useful when creating a fixed lookup name for a resource so that the qualified name is not dependent on the topology. You can also achieve this topology by widening the scope of the resource definition.

Note: WebSphere Application Server v3.5.x clients cannot access this type of binding .

String: String constant

You can configure a binding of a string constant. The following data is required to configure a string constant binding:

- String constant value
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.


Name space federation

Federating name spaces involves binding contexts from one name space into another name space.

For example, assume that a name space, Name Space 1, contains a context under the name `a/b`. Also assume that a second name space, Name Space 2, contains a context under the name `x/y`. (See the following illustration.) If context `x/y` in Name Space 2 is bound into context `a/b` in Name Space 1 under the name `f2`, the two name spaces are federated. Binding `f2` is a federated binding because the context associated with that binding comes from another name space. From Name Space 1, a lookup of the name `a/b/f2` returns the context bound under the name `x/y` in Name Space 2. Furthermore, if context `x/y` contains an Enterprise JavaBeans (EJB)

home bound under the name `ejb1`, the EJB home could be looked up from Name Space 1 with the lookup name `a/b/f2/ejb1`. Notice that the name crosses name spaces. This fact is transparent to the naming client.

Federated Name Spaces



In a WebSphere Application Server name space, you can create federated bindings with the following restrictions:

- Federation is limited to CosNaming name servers. A WebSphere Application Server name server is a Common Object Request Broker Architecture (CORBA) CosNaming implementation. You can create federated bindings to other CosNaming contexts. You cannot, for example, bind contexts from an LDAP name server implementation.
- If you use JNDI to federate the name space, you must use WebSphere Application Server's initial context factory to obtain the reference to the federated context. If you use some other initial context factory implementation, you either may not be able to create the binding, or the level of transparency may be reduced.
- A federated binding to a non-WebSphere Application Server naming context has the following functional limitations:
 - JNDI operations are restricted to the use of CORBA objects. For example, you can look up EJB homes, but you cannot look up non-CORBA objects such as data sources.
 - JNDI caching is not supported for non-WebSphere Application Server name spaces. This restriction affects the performance of lookup operations only.
- Do not federate two WebSphere Application Server standalone server name spaces. Incorrect behavior may result. If you want to federate WebSphere Application Server name spaces, you should use servers running under the Network Deployment or Enterprise packages of WebSphere Application Server.

Name space bindings

Administrators can add name bindings to the name space through the configuration. Name servers add these configured bindings to the name space view by reading the configuration data for the bindings. Configuring bindings is an alternative to creating the bindings from a program.

Configured bindings are created each time a server starts, even when the binding is created in a transient partition of the name space. One major use of configured bindings is to provide interoperability with JNDI clients running on previous versions of the WebSphere Application Server.

There are four different kinds of bindings that you can configure:

- Enterprise JavaBeans (EJB)
- CORBA object
- Indirect Lookup
- String

Configuring and viewing name space bindings

To view or configure an EJB, CORBA, Indirect lookup or string name space binding, complete the following:

Steps for this task

1. Open the Administrative console.
2. Click **Environment**.
3. Click **Manage Name Space Bindings**.
4. Select the desired scope by entering in a node name for node-scoped bindings, or a node name and server name for server-scoped bindings, and click **Apply**.
5. To create a new binding, click **New** and follow the instructions. To edit a previously created binding, click the binding you want to edit and proceed to the next step.
6. Edit the Binding identifier, the Name in name space, and the String value fields as desired.
Note: All of these fields are required.
7. Click **Finish** to register the changes.

String binding settings

Use this page to configure a new string binding or to view or edit an existing string binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > *string_namespace_binding***.

Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

Binding Identifier

Specifies the name that uniquely identifies this configured binding.

Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

String Value

Specifies the string to be bound into the name space.

CORBA object binding settings

Use this page to configure a new name binding of a CORBA object binding, or to view or edit an existing CORBA object binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > CORBA_namespace_binding**.

Scope

Shows the scope of the configured binding. This value indicates the configuration location for the namebindings.xml file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

Binding Identifier

Specifies the name that uniquely identifies this configured binding.

Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

Corbaname URL

Specifies the CORBA name URL string identifying where the object is bound in a CosNaming server.

Federated Context

Specifies whether the target is a CosNaming context (true) or a leaf node object (false).

Value
true

Result

The target object is bound with a context CORBA binding type. If the corbaname URL does not resolve to a NamingContext, an error occurs when the binding is first used (which is when the URL is first resolved).

false

The target object is bound with an object CORBA binding type.

Indirect lookup binding settings

Use this page to configure a new indirect lookup name binding, or to view or edit an existing indirect lookup binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > *indirect_lookup_namespace_binding***.

Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

Binding Identifier

Specifies the name that uniquely identifies this configured binding.

Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

Provider URL

Specifies the provider URL string needed to obtain a JNDI initial context.

JNDI Name

Specifies the name used to look up the target object from the initial context.

EJB binding settings

Use this page to configure a new EJB binding, or to view or edit an existing EJB binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > *EJB_namespace_binding***.

Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

Binding Identifier

Specifies the name that uniquely identifies this configured binding.

Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

Enterprise Bean Location

Specifies whether the enterprise bean is running in a server cluster or a single server. If Single Server is specified, type the node name.

Server

Specifies the name of the cluster or non-clustered server in which the enterprise bean is configured.

JNDI Name

Specifies the JNDI name of the deployed enterprise bean (the bean's JNDI name that is in the enterprise bean bindings—not the java:comp name)

Name space binding collection

Use this page to configure a name binding of an EJB, a CORBA CosNaming NamingContext, a CORBA leaf node object, an object that you can look up using JNDI, or a constant string value.

Binding information for configured bindings is stored in the configuration and applied upon startup of the name server for each server within the scope of the binding.

To view the Manage Name Space Bindings Settings page, Click **Environment** > **Manage Namespace Bindings**.

Click the check boxes to select one or more of the users in your collection. Use the buttons to control the selected users.

Name space bindings

Shows the names given to uniquely identify these configured bindings.

Configuring name servers

To configure a name server, complete the following:

Steps for this task

1. Open the administrative console.
2. Click **Servers**.
3. Click **Application Servers**.
4. Click the application server you want to configure.
5. Click **Server Components**.
6. Click **Name Server**.
7. Edit the fields as desired.

Note: All of these fields are mandatory.

8. (Optional) To make other changes, click **Custom Properties**.
9. Click **OK** to register your changes.

Name server settings

Use this page to configure Naming Service Provider settings for the application server.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Server Components > Name Server**.

The **Configuration** tab on the page provides editable fields and the **Runtime** tab provides read-only information.

Name

Specifies the display name for the server.

Initial State

Specifies the execution state.

Troubleshooting name space problems

Many naming problems can be avoided by fully understanding the key underlying concepts of WebSphere Application Server naming.

Steps for this task

1. Review the key concepts of WebSphere Application Server naming, especially Name space logical view and Lookup names support in deployment descriptors and thin clients.
2. Review the programming examples that are included in the sections explaining the JNDI and CosNaming interfaces.
3. Read "Naming services component troubleshooting tips" (not in this document) for additional general information.
4. If you "Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client" (not in this document), read this article.

dumpNameSpace tool

You can use the name space dump utility to dump the contents of the name space accessed through a name server. This dump does not include local name spaces such java URL name spaces. The naming service for the WebSphere Application Server host must be active when this tool is invoked.

Note that the server root context for the server at the specified host and port is dumped (unless a non-default starting context which precludes it is specified). The server root contexts for other servers are not dumped.

Command line invocation descriptions of the name space dump utility follow. This section includes sample output.

You can also access this utility a through its program interface. Refer to the class `com.ibm.websphere.naming.DumpNameSpace` in the WebSphere Application Server API documentation.

Syntax

To invoke the tool through the command line, enter the following command from the *WebSphere/AppServer/bin* directory:

Platform	Command
UNIX	<code>dumpNameSpace.sh [[-keyword value]...]</code>
Windows NT	<code>dumpNameSpace [[-keyword value]...]</code>

Parameters

The keywords and associated values for the `dumpNameSpace` utility follow:

-host *myhost.austin.ibm.com*

Indicates the bootstrap host or the WebSphere Application Server host whose name space you want to dump. The value defaults to *localhost*.

-port *nnn*

Indicates the bootstrap port which, if not specified, defaults to **2809**.

-root {cell | server | node | host | legacy | tree | default}

Indicates the root context to use as the initial context for the dump. The applicable root options and default root context depend on the type of name server from which the dump is being obtained. This information is provided in the following tables.

For WebSphere Application Servers v5 or later:

cell	DumpNameSpace default. Dump the tree starting at the cell root context.
server	Dump the tree starting at the server root context.
node	Dump the tree starting at the node root context. (Synonymous with host.)

For WebSphere Application Servers v4.0 or later:

legacy	DumpNameSpace default. Dump the tree starting at the legacy root context.
host	Dump the tree starting at the bootstrap host root context. (Synonymous with node.)
tree	Dump the tree starting at the tree root context.

For all WebSphere Application Servers and other name servers:

default	Dump the tree starting at the initial context which JNDI returns by default for that server type. This is the only -root choice that is compatible with WebSphere Application Servers prior to v4.0 and with non-WebSphere Application Server name servers.
---------	---

-url *some provider URL*

Indicates the value for the `java.naming.provider.url` property used to get

the initial JNDI context. This option can be used in place of the `-host`, `-port`, and `-root` options. If the `-url` option is specified, the `-host`, `-port`, and `-root` options are ignored.

-factory *com.ibm.websphere.naming.WsnInitialContextFactory*
 Indicates the initial context factory to be used to get the JNDI initial context. The value defaults to:
com.ibm.websphere.naming.WsnInitialContextFactory The default value generally does not need to be changed.

-startAt *some/subcontext/in/the/tree*
 Indicates the path from the bootstrap host's root context to the top level context where the dump should begin. The utility recursively dumps subcontexts below this point. It defaults to an empty string, that is, the bootstrap host root context.

-format{jndi | ins}

Option	Description
jndi	The default. Displays name components as atomic strings.
ins	Shows name components parsed per INS rules (id.kind).

-report {short | long}

Option	Description
short	The default. Dumps the binding name and bound object type. This output is also provided by <code>JNDI.Context.list()</code> .
long	<p>Dumps the binding name, bound object type, local object type, and string representation of the local object (that is, the IORs, string values, and other values that are printed).</p> <p>For objects of user-defined classes to display correctly with the long report option, it may be necessary to add their containing directories to the list of directories searched. Set the environment variable WAS_USER_DIRS. The value can include one or more directories, as for example:</p> <p>Platform Command</p> <p>UNIX <code>WAS_USER_DIRS=/usr/classdir1:/usr/classdir2 export WAS_USER_DIRS</code></p> <p>Windows NT <code>set WAS_USER_DIRS=c:\classdir1;d:\classdir2</code></p> <p>All zip, jar, and class files in the specified directories can then be resolved by the class loader when running dumpNameSpace.</p>

-traceString *"some.package.name.to.trace.*=all=enabled"*
 Represents the trace string with the same format as that generated by the servers. The output is sent to the file, `DumpNameSpaceTrace.out`.

Example: Invoking the name space dump utility

It is often helpful to view a dump of the name space to understand why a naming operation is failing. You can invoke the name space dump utility from the command line or from a program. Examples of each option follow.

Invoking name space dump utility from a command line

Invoke the name space dump utility from the command line by entering the following command:

```
dumpNameSpace -host myhost.mycompany.com -port 901
```

OR

```
dumpNameSpace -url corbaloc:iiop:myhost.mycompany.com:901
```

There are several command line options to choose from. For detailed help, enter the following command:

```
dumpNameSpace -help
```

Invoking name space dump utility from a Java program

You can dump name spaces from a program with the `com.ibm.websphere.naming.DumpNameSpace` API. Refer to the WebSphere Application Server API documentation for details on the `DumpNameSpace` program interface.

The following example illustrates how to invoke the name space dump utility from a Java program:

```
{
    ...
    import javax.naming.Context;
    import javax.naming.InitialContext;
    import com.ibm.websphere.naming.DumpNameSpace;
    ...
    java.io.PrintStream filePrintStream = ...
    Context ctx = new InitialContext();
    ctx = (Context) ctx.lookup("cell/nodes/node1/servers/server1");
    // Starting context for dump
    DumpNameSpace dumpUtil = new DumpNameSpace(filePrintStream, DumpNameSpace.SHORT);
    dumpUtil.generateDump(ctx);
    ...
}
```

Name space dump utility for java: and local: name spaces

Sometimes it is helpful to dump the java: name space for a J2EE application. You cannot use the `dumpNameSpace` command line utility for this purpose because the application's java: name space is accessible only by that J2EE application. From the WebSphere Application Server scripting tool, you can invoke a `NameServer` MBean to dump the java: name space for any J2EE application running in that same server process.

There is another name space local to server processes which you cannot dump with the `dumpNameSpace` command line utility. This name space has the URL scheme of local: and is used by the container to bind objects locally instead of through the name server. There is only one local: name space in a server process. You can dump a local: name space by invoking the `NameServer` MBean associated with that server process.

Name space dump options

Name space dump options are specified in the MBean invocation as a parameter in character string format. The option descriptions follow.

-startAt *some/subcontext/in/the/tree*

Indicates the path from the name space root context to the top level context where the dump should begin. The utility recursively dumps subcontexts below this point. It defaults to an empty string, that is, the root context.

-report {short | long}

Option	Description
short	The default. Dumps the binding name and bound object type. This output is also provided by JNDI Context.list().
long	Dumps the binding name, bound object type, local object type, and string representation of the local object (that is, the IORs, string values, and other values that are printed).

NameServer MBean invocation

Enter the WebSphere Application Server scripting command prompt: Invoke a method on a NameServer MBean by using the WebSphere Application Server scripting tool. Enter the scripting command prompt by typing the following command:

Platform	Command
UNIX	wsadmin.sh
Windows NT	wsadmin

Use the `-help` option for help on using the `wsadmin` command.

Select the NameServer MBean instance to invoke: Execute the following script commands to select the NameServer instance you want to invoke. For example,

```
set mbean [$AdminControl completeObjectName WebSphere:*,
          type=NameServer,cell=cellName,node=nodeName,
          process=serverName]
```

where *cellName*, *nodeName*, and *serverName* are the names of the cell, node, and server for the MBean you want to invoke. The specified server must be running before you can invoke a method on the MBean.

You can see a list of all NameServer MBeans current running by issuing the following query:

```
$AdminControl queryNames {*:*,type=NameServer}
```

Invoke the NameServer MBean: `java: name space`

Dump a `java: name space` by invoking the `dumpJavaNameSpace` method on the NameServer MBean. Since each server application has its own `java: name space`, the application must be specified on the method invocation. An application is identified by the application name, module name, and component name. The method syntax follows:

```
$AdminControl invoke $mbean dumpJavaNameSpace
  {{appName}{modName}
  {compName}{opts}}
```

where *appName* is the application name, *modName* is the module name, and *compName* is the component name of the `java: name space` you want to dump. The value for *opts* is the list of name space dump options described earlier in this section. The list can be empty.

local: name space

Dump a java: name space by invoking the `dumpLocalNameSpace` method on the `NameServer` MBean. Since there is only one local: name space in a server process, you have to specify the name space dump options only.

```
$AdminControl invoke $mbean dumpLocalNameSpace {{opts}}
```

where *opts* is the list of name space dump options described earlier in this section. The list can be empty.

Name space dump output

Name space dump output is sent to the console. It is also written to the file `DumpNameSpace.log`, in the server's log directory.

Example: Invoking the name space dump utility for java: and local: name spaces

It is often helpful to view the dump of a java: or local: name space to understand why a naming operation is failing. The `NameServer` MBean running in the application's server process can be invoked from the WebSphere Application Server scripting tool to generate a dump of these name spaces. Examples of `NameServer` MBean calls to generate dumps of java: and local: name spaces follow.

Dumping a java: name space

Assume you want to dump the java: name space of an application component running in server `server1` on node `node1` of the cell `MyCell`. The application name is `AcctApp` in module `AcctApp.war`, and the component name is `Acct Servlet`. The following script commands generate a long format dump of the application's java: name space of that application:

```
set mbean [$AdminControl completeObjectName WebSphere:*,
           type=NameServer,cell=MyCell,node=node1,process=server1]
$AdminControl invoke $mbean dumpJavaNameSpace {{AcctApp}{AcctApp.war}
        {Acct Servlet}{-report long}}
```

Dumping a local: name space

Assume you want to dump the local: name space for the server `server1` on node `node1` of cell `MyCell`. The following script commands will generate a short format dump of that server's local name space:

```
set mbean [$AdminControl completeObjectName WebSphere:type=NameServer,
           cell=MyCell,node=node1,process=server1]
$AdminControl invoke $mbean dumpLocalNameSpace {{-report short}}
```

Name space dump sample output

Name space dump output looks like the following example, which is the **SHORT** dump format:

```
Getting the initial context
Getting the starting context
```

```
=====
Name Space Dump
Provider URL: corbaloc:iiop:localhost:9810
Context factory: com.ibm.websphere.naming.WsnInitialContextFactory
Requested root context: cell
Starting context: (top)=outpostNetwork
Formatting rules: jndi
Time of dump: Mon Sep 16 18:35:03 CDT 2002
=====
```

=====
Beginning of Name Space Dump
=====

```

1 (top)
2 (top)/domain                javax.naming.Context
2   Linked to context: outpostNetwork
3 (top)/cells                 javax.naming.Context
4 (top)/clusters              javax.naming.Context
5 (top)/clusters/Cluster1    javax.naming.Context
6 (top)/cellname              java.lang.String
7 (top)/cell                  javax.naming.Context
7   Linked to context: outpostNetwork
8 (top)/deploymentManager     javax.naming.Context
8   Linked to URL: corbaloc::outpost:9809/NameServiceServerRoot
9 (top)/nodes                 javax.naming.Context
10 (top)/nodes/will2          javax.naming.Context
11 (top)/nodes/will2/persistent javax.naming.Context
12 (top)/nodes/will2/persistent/SomeObject SomeClass
13 (top)/nodes/will2/nodename java.lang.String
14 (top)/nodes/will2/domain   javax.naming.Context
14   Linked to context: outpostNetwork
15 (top)/nodes/will2/cell     javax.naming.Context
15   Linked to context: outpostNetwork
16 (top)/nodes/will2/servers  javax.naming.Context
17 (top)/nodes/will2/servers/server1 javax.naming.Context
18 (top)/nodes/will2/servers/will2 javax.naming.Context
19 (top)/nodes/will2/servers/member2 javax.naming.Context
20 (top)/nodes/will2/node     javax.naming.Context
20   Linked to context: outpostNetwork/nodes/will2
21 (top)/nodes/will2/nodeAgent javax.naming.Context
22 (top)/nodes/outpost        javax.naming.Context
23 (top)/nodes/outpost/node   javax.naming.Context
23   Linked to context: outpostNetwork/nodes/outpost
24 (top)/nodes/outpost/nodeAgent javax.naming.Context
24   Linked to URL: corbaloc::outpost:2809/NameServiceServerRoot
25 (top)/nodes/outpost/persistent javax.naming.Context
26 (top)/nodes/outpost/nodename java.lang.String
27 (top)/nodes/outpost/domain javax.naming.Context
27   Linked to context: outpostNetwork
28 (top)/nodes/outpost/servers javax.naming.Context
29 (top)/nodes/outpost/servers/server1 javax.naming.Context
30 (top)/nodes/outpost/servers/server1/url javax.naming.Context
31 (top)/nodes/outpost/servers/server1/url/CatalogDAOURL
31   java.net.URL
32 (top)/nodes/outpost/servers/server1/mail javax.naming.Context
33 (top)/nodes/outpost/servers/server1/mail/PlantsByWebSphere
33   javax.mail.Session
34 (top)/nodes/outpost/servers/server1/TransactionFactory
34   com.ibm.ejs.jts.jts.
   ControlSet$LocalFactory
35 (top)/nodes/outpost/servers/server1/servername java.lang.String
36 (top)/nodes/outpost/servers/server1/WSsamples javax.naming.Context
37 (top)/nodes/outpost/servers/server1/WSsamples/TechSampDataSource
37   TechSamp
38 (top)/nodes/outpost/servers/server1/thisNode javax.naming.Context
38   Linked to context: outpostNetwork/nodes/outpost
39 (top)/nodes/outpost/servers/server1/cell javax.naming.Context
39   Linked to context: outpostNetwork
40 (top)/nodes/outpost/servers/server1/eis javax.naming.Context
41 (top)/nodes/outpost/servers/server1/eis/DefaultDataSource_CMP
41   Default_CF
42 (top)/nodes/outpost/servers/server1/eis/WSsamples javax.naming.Context
43 (top)/nodes/outpost/servers/server1/eis/WSsamples/TechSampDataSource_CMP
43   TechSamp_CF
44 (top)/nodes/outpost/servers/server1/eis/jdbc javax.naming.Context
45 (top)/nodes/outpost/servers/server1/eis/jdbc/PlantsByWebSphereDataSource_CMP

```



```

45                                     PLANTSDB_CF
46 (top)/nodes/outpost/servers/server1/eis/jdbc/petstore
46                                     javax.naming.Context
47 (top)/nodes/outpost/servers/server1/eis/jdbc/petstore/PetStoreDB_CMP
47                                     PetStore_CF
48 (top)/nodes/outpost/servers/server1/eis/jdbc/CatalogDB_CMP
48                                     Catalog_CF
49 (top)/nodes/outpost/servers/server1/jta
49                                     javax.naming.Context
50 (top)/nodes/outpost/servers/server1/jta/usertransaction
50                                     java.lang.Object
51 (top)/nodes/outpost/servers/server1/DefaultDataSource
51                                     Default DataSource
52 (top)/nodes/outpost/servers/server1/jdbc
52                                     javax.naming.Context
53 (top)/nodes/outpost/servers/server1/jdbc/CatalogDB
53                                     CatalogDB
54 (top)/nodes/outpost/servers/server1/jdbc/petstore
54                                     javax.naming.Context
55 (top)/nodes/outpost/servers/server1/jdbc/petstore/PetStoreDB
55                                     PetStoreDB
56 (top)/nodes/outpost/servers/server1/jdbc/PlantsByWebSphereDataSource
56                                     PLANTSDB
57 (top)/nodes/outpost/servers/outpost
57                                     javax.naming.Context
57   Linked to URL: corbaloc::outpost:2809/NameServiceServerRoot
58 (top)/nodes/outpost/servers/member1
58                                     javax.naming.Context
59 (top)/nodes/outpost/cell
59                                     javax.naming.Context
59   Linked to context: outpostNetwork
60 (top)/nodes/outpostManager
60                                     javax.naming.Context
61 (top)/nodes/outpostManager/domain
61                                     javax.naming.Context
61   Linked to context: outpostNetwork
62 (top)/nodes/outpostManager/cell
62                                     javax.naming.Context
62   Linked to context: outpostNetwork
63 (top)/nodes/outpostManager/servers
63                                     javax.naming.Context
64 (top)/nodes/outpostManager/servers/dmgr
64                                     javax.naming.Context
64   Linked to URL: corbaloc::outpost:9809/NameServiceServerRoot
65 (top)/nodes/outpostManager/node
65                                     javax.naming.Context
65   Linked to context: outpostNetwork/nodes/outpostManager
66 (top)/nodes/outpostManager/nodename
66                                     java.lang.String
67 (top)/persistent
67                                     javax.naming.Context
68 (top)/persistent/cell
68                                     javax.naming.Context
68   Linked to context: outpostNetwork
69 (top)/legacyRoot
69                                     javax.naming.Context
69   Linked to context: outpostNetwork/persistent
70 (top)/persistent/AnotherObject
70                                     AnotherClass

```

```

=====
End of Name Space Dump
=====

```

Naming and directories: Resources for learning

Use the following links to find relevant supplemental information about naming and directories. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming specifications

-  **IBM's Web Services architecture debuts**
(<http://www.ibm.com/developerworks/webservices/library/w-int.html?dwzone=webservices>)

Introducing IBM's Web Services, a distributed software architecture of service components. This brief overview and in-depth interview on IBM DeveloperWorks cover the fundamental concepts of Web Services architecture and what they mean for developers. The interview with IBM's Rod Smith, Vice President of Emerging Technologies, explores which types of developers Web Services targets, how Web services reduce development time, what developers could be doing with Web Services now, and takes a glance at the economics of dynamically discoverable services.

-  **OMG CosNaming Interoperable Naming Specification**(<http://www.omg.org/cgi-bin/doc?ptc/99-12-03>)

Chapter 14. Improving performance through the dynamic cache service

The dynamic cache works within an application server Java Virtual Machine (JVM), intercepting calls to cacheable objects, for example, through a servlet's `service()` method or a command's `execute()` method, and either stores the object's output to, or or serves the object's content from, the dynamic cache.

WebSphere Application Server, Version 4.0, supported the configuration of dynamic servlet caching through the use of a `servletcache.xml` file. For migration purposes, this file is still supported by this release. In order to utilize the new and improved functionality of the dynamic cache service in this release, you must configure your cache policy using the new `cachespec.xml` format.

The dynamic caching documentation provides you with the following tasks to enable and configure the dynamic cache service, as well as advanced features, such as controlling external caches and building user-defined drop-in components to customize the cache operation.

Steps for this task

1. Enable the dynamic cache service globally .
2. Configure servlet caching.
3. Configure Edge Side Include (ESI) caching.
4. Configure command caching.
5. Configure Web services caching.

Dynamic cache

Caching the output of servlets, commands and Java Server Pages (JSP)files, improves application performance. WebSphere Application Server consolidates several caching activities, including servlets, Web services, and WebSphere commands into one service called the dynamic cache. These caching activities work together to improve application performance, and share many configuration parameters, which are set in an application server's dynamic cache service.

The dynamic cache works within an application server Java Virtual Machine (JVM), intercepting calls to cacheable objects, for example, through a servlet's `service()` method or a command's `execute()` method, and either stores the object's output to, or serves the object's content from, the dynamic cache. Because J2EE applications have high read-write ratios and can tolerate small degrees of latency in the currency of their data, the dynamic cache creates an opportunity for significant gains in server response time, throughput, and scalability.

You can use the dynamic cache to improve the performance of servlet and JSP files by serving requests from an in-memory cache. Cache entries contain servlet output, results of servlet execution, and metadata.

Configuring globally the dynamic cache service

Steps for this task

1. Open the administrative console.

2. Click **Servers > Application Servers** in the administrative console navigation tree.
3. Click a *server*.
4. Click **Dynamic Cache Service**.
5. Select **Enable service at server startup** in the **Startup state** field.
6. Click **Apply** or **OK**.

DynamicCache service settings

Use this page to configure and manage the dynamic cache service settings.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service**.

Startup state

Specifies whether the dynamic cache is enabled.

Cache Size

Specifies a positive integer as the value for the maximum number of entries the cache holds.

Enter the cache size value in this field. Values are usually in the thousands, with no set maximum or minimum.

Default Priority

Specifies the default priority for cache entries, determining how long an entry stays in a full cache.

Default	1
---------	---

Disk offload

Specifies whether disk offload is enabled.

By default, the dynamic cache only maintains, at most, the number of entries configured in memory. If new entries are created while the cache is full, the priorities configured for each cache entry, along with a least recently used algorithm, are used to remove entries from the cache. As an alternative, you can configure disk offload, and rather than removing the entries from memory, have them copied onto the file system (the location is configurable) for later.

Cache replication

Specifies whether cache replication is enabled.

You can also configure advanced cache replication settings.

Configuring servlet caching

Before you begin

In order to enable servlet caching, the dynamic cache service must also be enabled.

Steps for this task

1. Open the administrative console.
2. Click **Servers > Application Servers** in the console navigation tree.
3. Click a *server*.
4. Click **Web Container**.

5. Select the **Enable servlet caching** check box under the Configuration tab.
6. Click **Apply** or **OK**.

Servlet caching

After a servlet is invoked once, generating the output to be cached, a cache entry is created containing not only the output, but also side effects of the invocation. For example, these side effects can include calls to other servlets or Java Server Pages (JSP) files, as well as metadata about the entry, including timeout and entry priority information.

Unique entries are distinguished by an ID string generated from the `HttpServletRequest` object for each invocation of the servlet. You can then base servlet caching on:

- Request parameters and attributes URI used to invoke the servlet
- Session information
- Other options, including cookies

Since JSP files are compiled by WebSphere Application Server into servlets, the dynamic cache function treats them the same, except in specifically documented situations.

Configuring the dynamic cache disk offload

By default, when the number of cache entries reaches the configured limit for a given WebSphere server, eviction of cache entries occurs, allowing new entries to enter the cache service. The dynamic cache includes an alternative feature named disk offload, that copies the evicted cache entries to disk for potential future access.

To configure disk offload:

Steps for this task

1. Open the administrative server.
2. Click **Server > Application Server** in the administrative console navigation tree.
3. Click *server*.
4. Click **Dynamic Cache Service**.
5. Click the **Enable disk offload** check box in the **Disk offload** field.
You can also set the disk offload location in this field.
6. Click **Apply** or **OK**.

Configuring Edge Side Include caching

The Web server plug-in contains a built-in Edge Side Include (ESI) processor. The ESI processor has the ability to cache whole pages, as well as fragments, providing a higher cache hit ratio. The cache implemented by the ESI processor is an in-memory cache, not a disk cache, therefore, the cache entries are not saved when the Web server is restarted.

The basic operation of the ESI processor is as follows: When a request is received by the Web server plug-in, it is sent to the ESI processor, unless the ESI processor is disabled. It is enabled by default. If a cache miss occurs, a `Surrogate-Capabilities` header is added to the request and the request is forwarded to the WebSphere Application Server. If the dynamic servlet cache is enabled in the application

server, and the response is edge cacheable, the application server returns a Surrogate-Control header in response to the WebSphere Application Server plug-in.

The value of the Surrogate-Control response header contains the list of rules which are used by the ESI processor in order to generate the cache ID. The response is then stored in the ESI cache, using the cache ID as the key. For each ESI include tag in the body of the response, a new request is processed such that each nested include results in either a cache hit or another request forwarded to the application server. When all nested includes have been processed, the page is assembled and returned to the client.

The ESI processor is configurable through the WebSphere Web server plug-in configuration file `plugin-cfg.xml`. The following is an example of the beginning of this file, which illustrates the ESI configuration options.

```
<?xml version="1.0"?>
<Config>
  <Property Name="esiEnable" Value="true"/>
  <Property Name="esiMaxCacheSize" Value="1024"/>
  <Property Name="esiInvalidationMonitor" Value="false"/>
</Config>
```

The first option, `esiEnable`, can be used to disable the ESI processor by setting the value to false. ESI is enabled by default. If ESI is disabled, then the other ESI options are ignored.

The second option, `esiMaxCacheSize`, is the maximum size of the cache in 1K byte units. The default maximum size of the cache is 1 megabyte. If the cache is full, the first entry to be evicted from the cache is the entry that is closest to expiration.

The third option, `esiInvalidationMonitor`, specifies whether or not the ESI processor should receive invalidations from the application server. There are three methods by which entries are removed from the ESI cache: first, an entry's expiration timeout could fire; second, an entry may be purged to make room for newer entries; or third, the application server could send an explicit invalidation for a group of entries. In order for the third mechanism to be enabled, the `esiInvalidationMonitor` property must be set to true and the `DynaCacheEsi` application must be installed on the application server. The `DynaCacheEsi` application is located in the `installableApps` directory and is named `DynaCacheEsi.ear`. If the `ESIInvalidationMonitor` property is set to true but the `DynaCacheEsi` application is not installed, then errors will occur in the webserver plugin and the request will fail.

The ESI processor's cache can be monitored through the `CacheMonitor` application. In order for ESI processor's cache to be visible in the `CacheMonitor`, the `DynaCacheEsi` application must be installed as described above and the `ESIInvalidationMonitor` property must be set to true in the `plugin-cfg.xml` file.

When WebSphere Application Server is used to serve static data, such as images and HTML on the application server, the URLs are also cached in the ESI processor. This data has a default timeout of 300 seconds. You can change the timeout value by adding the property `com.ibm.servlet.file.esi.timeOut` to your JVM's command line parameters. The following example shows how to set a one minute timeout on static data cached in the plug-in:

```
-Dcom.ibm.servlet.file.esi.timeOut=60
```

Configuring external cache groups

The dynamic cache has the ability to control caches outside of the application server, such as IBM Edge Server, IBM HTTP Server's FRCA cache, and the the WebSphere HTTP Server plug-in's ESI Fragment Processor. When external cache groups are defined, the dynamic cache matches externally cacheable cache entries with those groups, and pushes cache entries and invalidations out to those groups. This allows WebSphere to manage dynamic content beyond the application server. The content can then be served from the external cache, instead of the application server, improving savings in performance.

Steps for this task

1. Open the administrative console.
2. Enable the dynamic cache.
 - a. Click **Servers > Application Servers** in the administrative console navigation tree.
 - b. Click a *server*.
 - c. Click **Dynamic Cache Service**.
 - d. Select the check box in the **Startup state** field to enable the dynamic cache.
3. Define the external cache group in which WebSphere Application Server should control.
 - a. Click **External Caching Groups** from the Dynamic Cache administrative console page.
 - b. Click **New** or choose an external cache group from the list.
4. Configure cache group members.
 - a. Click **External cache groups** from the Dynamic Cache administrative console page. Then click **New** or choose an external cache group from the list.
 - b. Click **External cache group members > New** or choose an external cache group member from the list.
 - c. Type the configuration string in the **Address** field.
 - d. Type the adapter bean name in the **Adapter Bean Name** field.
 - e. **Save** the configuration.
 - f. Click **Apply** or **OK**.

External cache group collection

Use this page to define sets of external caches controlled by WebSphere Application Server on Web servers, such as IBM Edge Server and IBM HTTP Server.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache Groups**.

External Cache Group: Specifies the external cache group name.

The external cache group name needs to match the externalcache property as defined in the servlet or JSP cachespec.xml file.

When external caching is enabled, the cache matches pages with its URIs and pushes matching pages to the external cache. The entries can then serve from the external cache, instead of the application server.

Type: Specifies the external cache group type.

External cache group settings

Use this page to configure sets of external caches controlled by WebSphere Application Server on Web servers, such as IBM Edge Server and IBM HTTP Server.

To view this administrative console page, click **Servers > Application Server > server > Dynamic Cache Service > External Cache groups > external_cache_group**.

External Cache Group: Specifies the external cache group name.

The external cache group name needs to match the externalcache property as defined in the servlet or JSPs cachespec.xml file.

When external caching is enabled, the cache matches pages with its URIs and pushes matching pages to the external cache. The entries can then serve from the external cache, instead of the application server. This ability creates a significant savings in performance.

Type: Specifies the external cache group type.

External cache group member collection

Use this page to define specific caches that are members of a cache group.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache groups > external_cache_group > External cache group members**.

Address: Specifies a configuration string used by external cache adapter bean to connect to the external cache.

AdapterBeanName: Specifies the adapter bean name.

Example adapter bean names supported in WebSphere Application Server are:

- **AFPA:** com.ibm.ws.cache.servlet.Afpa
- **ESI:** com.ibm.websphere.servlet.cache.ESIInvalidatorServlet

External cache group member settings

Use this page to configure specific caches that are members of a cache group.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache groups > External_cache_group > External cache group members > External_cache_group_members**.

Address: Specifies a configuration string used by external cache adapter bean to connect to the external cache.

AdapterBeanName: Specifies the adapter bean name.

Example adapter bean names supported in WebSphere are:

- **AFPA:** com.ibm.ws.cache.servlet.Afpa
- **ESI:** com.ibm.websphere.servlet.cache.ESIInvalidatorServlet

Configuring high-speed external caching through the Web server

IBM HTTP Server for Windows NT and Windows 2000 operating systems contains a high speed cache referred to as the *Fast Response Cache Accelerator*, or *Cache Accelerator*.

The Fast Response Cache Accelerator is available on both Windows NT and Windows 2000 operating systems and AIX platforms. However, the dynamic cache is only available on Windows NT and Windows 2000 operating systems.

WebSphere Application Server fragment cache can use IBM HTTP Server as an external cache with the appropriate configuration.

To configure Fast Cache Response Accelerator:

Steps for this task

1. Install the WebSphere Application Server and IBM HTTP Server for Windows NT and Windows 2000 environments.
2. Configure the Fast Cache Accelerator on the Web server.
 - a. In the IBM HTTP Server conf directory, locate the `httpd.conf` configuration file, and add the following two lines at the end of the file:

```
LoadModule afpaplugin_module c:/WebSphere/AppServer/bin/afpaplugin.dll
AfpapluginHost 127.0.0.1:9081
```

The first line loads the IBM HTTP Server plug-in that connects the Fast Cache Accelerator to the WebSphere fragment cache. Ensure the `LoadModule` path points to your Websphere Application Server installation.

3. Configure an external cache group on the application server.
 - a. Define an external cache group named `afpa` for each application server that uses the Cache Accelerator.

To define an external cache group, click **New** on the External cache group administrative console page. Type `afpa` in the **External cache group** field. Click **Apply**.
 - b. Add a member to that group with an adapter bean name of `com.ibm.servlet.dynacache.Afpa`

To add a member, click **afpa > External cache group members**. Click **New** on the External cache group members administrative console page. Then type `com.ibm.servlet.dynacache.Afpa` in the **AdapterBean name** field.
4. Enter the assigned port number from the Web server `httpd.conf` file for the address.
5. Configure a cache policy using external cache.

Now you can use the external cache attribute when building your cache policies.

Configuring Fast Response Cache Accelerator cache size through the: Web server

In the default IBM HTTP Server configuration, the maximum Fast Cache Accelerator dynamic cache size is calculated as 1/8 of physical pin-able memory. On a machine with 384MG of RAM, it allows a maximum of approximately 50MG for the Fast Cache Accelerator dynamic cache. When this limit is reached, the Cache Accelerator then deletes older entries to cache new ones.

Follow these steps to configure the Cache Accelerator:

Update the following directives in the `httpd.conf` file of IBM HTTP Server:

```
AfpaEnable
AfpaCache on
AfpaLogFile "c:/Program Files/IBM HTTP Server/logs/afpalog" V-ECLF
AfpaDynaCacheMax 10
```


Steps for this task

1. Using the IBM HTTP Server directive `AfpaDynaCacheMax`, tune the maximum allowed cache size:
 - a. Place the directive in the global server configuration scope, along with the other default Fast Cache Accelerator directives.
 - b. Enable Fast Cache Accelerator.

To enable the Fast Cache Accelerator, update the following directives in the `http.conf` file of IBM HTTP Server:

```
AfpaEnable
AfpaCache on
AfpaLogFile "c:/Program Files/IBM HTTP Server/logs/afpalog" V-ECLF
AfpaDynaCacheMax 10
```

These above settings limit the dynamic cache size to 10MG. If you use these directives to increase cache size, do not make the cache so large that all the physical memory is consumed. Determine how much memory is available when all applications are running, by using the Windows Task Manager.

Assign no more than 50% of available physical memory to the dynamic cache. Specifying too large a cache not only decreases the performance of other applications, but also puts you at a risk for completely running out of memory.

The default configuration does not include the `AfpaDynaCacheMax` directive where the cache size is automatically calculated as 1/8 of physical memory.

Displaying cache information

The Dynamic Cache Monitor is an installable Web application that displays simple cache statistics, cache entries and cache policy information.

Steps for this task

1. Install the application from the `<<cache-entry>>install_root<>/installableApplications` directory.
The application is named `CacheMonitor.ear`.
2. Access the Web application using a Web browser and the URL `http://<<cache-entry>>your_hostname<>:<<cache-entry>>your_port_number<>/cachemonitor`.

Configuring cacheable objects with the `cachespec.xml` file

Before you begin

Define cacheable objects inside the `cachespec.xml`, found inside the Web application archive (WAR) `WEB-INF` or `enterprise bean META-INF` directory.

You can place a global `cachespec.xml` in the application server properties directory, but the recommended method is to place the cache configuration file with the deployment module. The root element of the `cachespec.xml` file is `<cache>`, which contains `<cache-entry>` elements.

Within a `<cache-entry>...</cache-entry>` element are parameters that allow you to complete the following tasks to enable the dynamic cache with the `cachespec.xml` file:

Steps for this task

1. Develop a `cachespec.xml` file.
 - a. Create a caching configuration file.

In the `<<<cache-entry>>install_root<>>/properties` directory, locate the `cachespec.sample.xml` file.
 - b. Copy the `cachespec.sample.xml` file to `cachespec.xml` in `WAR WEB-INF` or `enterprise bean META-INF` directory.
2. Define the cache-entry elements necessary to identify the servlet. See the article `Cachespec.xml` file for a list of elements.
3. Develop cache-ID rules.

To cache an object, WebSphere Application Server must know how to generate unique IDs for different invocations of that object. The `<cache-id>` element performs that task. Each cache entry can have multiple cache-ID rules that execute in order until either a rule returns non-empty cache-ID or no more rules remain to execute. If none of the cache-ID generation rules produce a valid cache ID, then the object is not cached. Develop these IDs in one of two ways:

- Use the `<component>` element defined in the cache policy of a cache entry (recommended)
- Write custom Java code to build the ID from input variables and system state

To configure the cache entry to use the `IdGenerator` specify your `IdGenerator` in the XML file, by using the `<idgenerator>` tag, for example:

```
<cache-entry>
<class>servlet</class>
<name>/servlet/CommandProcessor</name>
<cache-id>
<id-generator>com.mycompany.SampleIdGeneratorImpl
</id-generator>
<timeout>60</timeout>
</cache-id>
</cache-entry>
```

You can also use the Application Assembly Tool (AAT) to define the `IdGenerator` class in the cache policy's **Advanced** tab.

4. Specifying dependency ID rules.

Use dependency ID elements to specify additional cache group identifiers that associate multiple cache entries to the same group identifier.

The dependency ID is generated by concatenating the dependency ID base string with the values returned by its component elements. If a required component returns a null value, then the entire dependency ID does not generate and is not used. You can validate the dependency IDs explicitly through the WebSphere Dynamic Cache API, or use another cache-entry `<invalidation>` element. Multiple dependency ID rules can exist per cache-entry. All dependency ID rules separately execute. See the article `Cachespec.xml` file for a list of `<component>` elements.
5. Invalidate other cache entries as a side effect of this object execution, if relevant.

You can define invalidation rules in exactly the same manner as dependency IDs. However, the IDs that generate by invalidation rules are used to invalidate cache entries that have those same dependency IDs.

The invalidation ID is generated by concatenating the invalidation ID base string with the values returned by its component element. If a required component returns a null value, then the entire invalidation ID is not generated and no invalidation occurs. Multiple invalidation rules can exist per cache-entry. All invalidation rules separately execute.

6. Verify the cacheable page.

What to do next

Typically you declare several `<cache-entry>...</cache-entry>` elements inside a `cachespec.xml` file.

The dynamic cache responds to changes in this file. When new versions of the `cachespec.xml` are detected, the old policies are replaced. Objects cached through the old policy file are not automatically invalidated from the cache; they are either reused with the new policy or eliminated from the cache through its replacement algorithm.

For each of the three IDs (cache, dependency, invalidation) generated by cache entries, a `<cache-entry>` can contain multiple elements. The dynamic cache will execute the `<cache-id>` rules in order, and the first one that successfully generates an ID will be used to cache that output. If the object is to be cached, each one of the `<dependency-id>` elements will be executed to build a set of dependency IDs for that cache entry. Finally, each of the `<invalidation>` elements will be executed, building a list of IDs that the dynamic cache will invalidate, whether or not this object is cached.

Verifying the cacheable page

Verify the cacheable page by following these steps:

Steps for this task

1. View the snoop servlet in the default application by accessing the URI:
`/servlet/snoop`
2. Invoke and reload the URI several times.
This action returns the same output for the snoop servlet. The snoop servlet is now operating incorrectly, since it displays the request information from its first invocation rather than from the current request.
3. Inspect the entry in the cache with the Dynamic Cache Monitor.

Cachespec.xml file

The cache parses the `cachespec.xml` file on startup, and extracts from each `<cache-entry>` element a set of configuration parameters. Then, every time a new servlet or other cacheable object initializes, the cache attempts to match each of the different cache-entry elements, to find the configuration information for that servlet. Different cacheable objects have different `<class>` elements. You can define the specific object a cache policy refers to using the `<name>` element.

Location

The `cachespec.xml` file is found inside the META-INF directory of Web application archive (WAR) file or enterprise bean JAR file.

You can place a global `cachespec.xml` file in the application server properties directory, but the recommended method is to place the cache configuration file with the deployment module. The root element of the `cachespec.xml` file is `<<cache-entry>>cache<<`, which contains `cache-entry` elements.

The `cachespec.dtd` is available in the application server properties directory.

Usage notes

This section explains the function of each element of the `cachespec.xml` file including:

- Class
- Name
- Sharing-policy
- Property

Class

```
<class>command | servlet | webservice</class>
```

This element is required and governs how WebSphere will interpret the remaining cache policy definition. The value `servlet` refers to servlets and Java Server Pages (JSPs) deployed in the WebSphere Application Server servlet engine. The `webservice` class extends the `servlet` with special component types for Web services requests. Finally, the value `command` refers to classes using the WebSphere command programming model. The following is an example: `<class>command</class>`
`<class>servlet</class>`

Name

```
<name><<cache-entry>>name<</name> (name=The fully specified class name of the command/servlet/webservice)
```

There are two ways to use `<name>` to specify a cacheable object:

- Servlets and JSPs only: Use the relative Web path or servlet mapping of the servlet of JSP
- Use the object's full package and class name

For commands and Web services, this required element must include the full package and class name, including a trailing `.class` of the configured object. For servlets and JSP files, it can include the full URI, relative to the Web application root, of the JSP file or servlet to cache. Also, you can specify multiple `<name>` elements within a `<cache-entry>` if you have different mappings that refer to the same servlet.

The following are examples:

```
<name>com.mycompany.MyCommand.class</name>  
<name>/servlet/snoop/*</name>  
<name>com.mycompany.ControllerServlet.class</name> <name>myJsp.jsp</name>
```

Sharing-policy

<sharing-policy> not-shared | shared-push | shared-pull | shared-push-pull</sharing-policy>

When working within a cluster with a distributed cache, these values determine the sharing characteristics of entries created from this object. If this element is not present, a not-shared value is assumed. Also, in non-distributed environments, not-shared is the only valid value. This property does not affect distribution to Edge servers through the Edge fragment caching property.

Value	Description
not-share	Cache entries for this object will not be shared among different application servers. These entries can contain non-serializable data, for example, a cached servlet can place non-serializable objects into the request attributes, if the <class> type allows it.
shared-push	Cache entries for this object will automatically be distributed to the dynamic caches in other application servers or cooperating Java Virtual Machines (JVMs). Each cache then has a copy of the entry the time it is created. These entries can not store non-serializable data.
shared-pull	Cache entries for this object will be shared between application servers on demand. If an application server gets a cache miss for this object, it will first query the cooperating application servers to see if they have the object. If no application server has a cached copy of the object, the original application server will execute the request and generate the object. These entries can not store non-serializable data.
shared push-pull	Cache entries for this object will be shared between application servers on demand. When an application server generates a cache entry, it broadcasts the cache ID of the created entry to all cooperating application servers. Each server then knows whether an entry exists for any given cache ID. Thus, on a given request for that entry, the application server knows whether to generate the entry itself or pull it from somewhere else. These entries can not store non-serializable data.

An example of sharing policy is as follows:

```
<sharing-policy>not-shared</sharing-policy>
```

Property

```
<property name="<<cache-entry>>key<>"><<cache-entry>>value<></property>
```

(key= The name of the property being defined; value= The corresponding value)

You can set optional properties on a cacheable object, such as a description of the configured servlet. The class determines valid properties of the cache entry. At this time, the following properties are defined:

Property	Valid classes	Value
ApplicationName	All	Overrides the J2EEName application ID so that multiple applications can share a common cache ID namespace. For more information, see Cache ID Tuples.
EdgeCacheable	Servlet	True or False. Default is false. If the property is true, then the given servlet or JSP file is externally requested from an Edge Server. Whether or not the servlet or JSP file is cacheable, depends on the rest of the cache specification.
ExternalCache	Servlet	Specifies the external cache name. The external cache name needs to match the external cache group name.
consume-subfragments	Servlet or Web service	True or false. Default is false. When a servlet is cached, only the content of that servlet is stored, and includes placeholders for any other fragments to which it includes or forwards. Consume-subfragments (CSF) tells the cache not to stop saving content when it includes a child servlet. The parent entry (the one marked CSF) will include all the content from all fragments in its cache entry, resulting in one big cache entry that has no includes or forwards, but the content from the whole tree of entries. This can save a significant amount of application server processing, but is typically only useful when the external HTTP request contains all the information needed to determine the entire tree of included fragments.

Each cache-ID element defines a rule for caching an object and is composed of the sub-elements component, timeout, priority and property. Examples of is as follows:
`<cache-id>component*| timeout? | priority? | property* </cache-id>`

Component elements

Component is a generic term for some input variable with data that you want to use in caching a fragment. Because different classes of objects use different types of

inputs, such as, servlets use request parameters and commands use method values, the valid component types change for the different <class> values. Components are used in <cache-id>, <dependency-id> and <invalidation> elements.

```
<component id="id"
type="type"
ignore-value="true false">...</component>
(id=string,
type=string)
```

Use the component element to generate a portion of the cache ID. Each component element consists of the following attributes:

- **ID:** Use the ID attribute to identify the component.
- **Type:** Use the type attribute to identify the type of component, for example, method, field, request parameter, session value. An example is as follows:

```
<component id="getPrice"
type="method">
```

Type value	Valid classes	Meaning
method	command	Calls the indicated method on the command or object
field	command	Retrieves the named field in the command or object
parameter	servlet	Retrieves the named parameter from the request object
session	servlet	Retrieves the named value from the HttpSession
cookie	servlet	Retrieves the named cookie value
attribute	servlet	Retrieves the named request attribute
header	servlet and webservice	Retrieves the named request header
pathInfo	servlet	Retrieves the pathInfo from the request
servletpath	servlet	Retrieves the servlet path
locale	servlet	Retrieves the request locale
SOAPEnvelope	Web service	Retrieves the SOAPEnvelope from a Web services request. An ID attribute of Hash uses a Hash of the SOAPEnvelope, while Literal uses the SOAPEnvelope as received.
SOAPAction	Web service	Retrieves the SOAPAction header, (if available, for a Web services request
serviceOperation	Web service	Retrieves the service operation for a Web service request
serviceOperationParameter	Web service	Retrieves the specified parameter from a Web services request

- **Ignore-value** (optional, defaults to false): If true, this attribute specifies that you should not use the value returned by this component in cache ID formation. Only the ID of the component is used when creating a cache-ID, or no output is used when creating a dependency or invalidation ID. An example is as follows:

```
<component id="isValid" type="method"
ignore-value="true"
```

and sub elements:

- **Method:** Use method elements to call a void method on a returned object. Method and field objects can be infinitely nested in any combination. The method must be public. Not valid for edge cacheable components. An example is as follows:

```
<component id="getColor" type="method">
<method>getUserInfo<method>getName
</method></method></component>
```

This is equivalent to `getUser().getUserInfo().getName()`

- **Field:** Use field elements to access a field in a returned object. Method and field objects can be infinitely nested in any combination. The field must be public. Not valid for edge cacheable components. An example is as follows:

```
<component id="getColor" type="method">
<method>getUserInfo<field>name
</field></method></component>
```

This element is equivalent to `getUser().getUserInfo().name`

- **Required:** If true, then this component must return a non-null value for this cache-ID to represent a valid cache-ID. If false, the default, then a non-null value is used in the formation of the cache-ID and a null value means that this component is not used at all in the ID formation. An example is as follows:

```
<required>true</required>
```

- **Value:** Use value elements to specify values that must match to use this component in cache-ID formation. An example is as follows:

```
<component id="getColor" type="method">
<value>blue</value>
<value>red</value> </component>
```

- **Not-value:** Use not-value elements to specify values that must not match to use this component in cache-ID formation. This element is similar to the `<value>`, but instead proscribes the defined values from caching. You can use multiple `<not-value>` elements when there is more than one invalid value. An example is as follows:

```
<component id="getColor" type="method">
<required>true</required>
<not-value>blue</not-value>
<not-value>red</not-value>
</component>
```

IdGenerator and MetadataGenerator elements

The `IdGenerator` element is used to specify the classname loaded for cache ID's ID generation. The `IdGenerator` must implement the `com.ibm.websphere.servlet.cache.IdGenerator` interface. The `IdGenerator` must build and set cache IDs, group IDs and invalidation IDs. An example is as follows:

```
<idgenerator>
classname classname </idgenerator>
```

(classname= Fully qualified name of the class to be used)

The `MetadataGenerator` element is used to specify the classname loaded for the cache ID's metadata generation. The `MetadataGenerator` class must implement the

com.ibm.websphere.servlet.cache.MetadataGenerator interface. The MetadataGenerator defines properties like timeout, external caching or generic properties.

Configuring command caching

Cacheable commands are stored in the cache for re-use with a similar mechanism for servlets and Java Server Pages (JSP) files. However, in this case, the unique cache IDs are generated based on methods and fields present in the command as input parameters. For example, a **GetStockQuote** command can have a symbol as its input parameter.

A unique cache ID can generate from the name of the command, plus the value of the symbol.

To use command caching you must:

Steps for this task

1. Create a command.

- a. Define an interface.

The Command interface specifies the most basic aspects of a command.

You must define the interface that extends one or more of the interfaces in the command package. The command package consists of three interfaces:

- TargetableCommand
- CompensableCommand
- CacheableCommand

In practice, most commands implement the TargetableCommand interface, which allows the command to execute remotely. The code structure of a command interface for a targetable command follows:

```
...
import com.ibm.websphere.command.*;
public interface MyCommand extends CacheableCommand {
    // Declare application methods here
}
```

- a. Provide an implementation class for the interface.

Write an interface that extends the CacheableCommandImpl class and implements your command interface. This class contains the code for the methods in your interface, the methods inherited from extended interfaces like the CacheableCommand interface, and the required or abstract methods in the CacheableCommandImpl class.

You can also override the default implementations of other methods provided in the CacheableCommandImpl class.

Command class

To write a command interface, extend one or more of the three interfaces included in the command package. The base interface for all commands is the Command interface. This interface provides only the client-side interface for generic commands and declares three basic methods:

- **isReadyToCallExecute.** This method is called on the client side before the command passes to the server for execution.
- **execute.** This method passes the command to the target and returns any data.

- **reset.** This method reverts any output properties to the values they had before the execute method was called so that you can reuse the object.

The implementation class for your interface must contain implementations for the `isReadyToCallExecute` and `reset` methods.

CacheableCommandImpl class

Commands are implemented by extending the class `CacheableCommandImpl`, which implements the `CacheableCommand` interface.

The `CacheableCommandImpl` class is an abstract class that provides implementations for some of the methods in the `CacheableCommand` interface, for example, setting return values. This class declares additional methods that the application must implement, for example, how to execute the command.

The code structure of an implementation class for the `CacheableCommand` interface follows:

```
...
import com.ibm.websphere.command.*;
public class MyCommandImpl extends CacheableCommandImpl
implements MyCommand {
    // Set instance variables here    ...
    // Implement methods in the MyCommand
    // interface
    ...
    // Implement abstract methods in the CacheableCommandImpl class
    ...
}
```

Example: Caching a command object

This example of command caching is a simple stock quote command.

Examples

The following is a stock quote command bean. It accepts a ticker as an input parameter and produces a price as its output parameter.

```
public class QuoteCommand extends CacheableCommandImpl
{
    private String ticker;
    private double price;
    // called to validate that command input
    // parameters have been set
    public boolean isReadyToCallExecute() {
        return (ticker!=null);
    }
    // called by a cache-hit to copy output properties
    // to this object
    public void setOutputProperties(TargetableCommand fromCommand) {
        QuoteCommand f = (QuoteCommand)fromCommand;
        this.price = f.price;
    }

    // business logic method called when the stock price
    // must be retrieved
    public void performExecute()throws Exception {...}

    //input parameters for the command
    public void setTicker(String ticker) { this.ticker=ticker;}
    public String getTicker() { return ticker;}
}
```

```

    //output parameters for the command
    public double getPrice() { return price;};
}

```

Examples

To cache the above command object using the stock ticker as the cache key and using a 60 second time-to-live, use the following cache policy:

```

<cache>
<cache-entry>
<class>command</class>
<sharing-policy>not-shared</sharing-policy>
<name>QuoteCommand</name>
<cache-id>
<component type="method" id="getTicker">
<required>true</required>
</component>
<priority>3</priority>
<timeout>60</timeout>
</cache-id>
</cache-entry>
</cache>

```

Example: Caching Web services

The following is an example of building a set of cache policies for a simple Web services application. The application in this example stores stock quotes, and has operations to read, update the price of, and buy a given stock symbol. .

Following are two SOAP message examples that the application can receive, with accompanying HTTP Request headers.

The first message sample contains a SOAP message for a GetQuote operation, requesting a quote for IBM. This is a read-only operation that gets its data from the back-end, and is very cacheable. In this example the SOAP message is cached and a timeout is placed on its entries to guarantee the quotes it returns are not too out of date.

Message example 1

```

POST /soap/servlet/soaprouter
HTTP/1.1
Host: www.myhost.com
Content-Type: text/xml; charset="utf-8"
SOAPAction: urn:stockquote-lookup
<SOAP-ENV:Envelope xmlns:SOAP-ENV=
"http://schemas.xmlsoap.org/soap/envelope/"
SOAP-ENV:encodingStyle=
"http://schemas.xmlsoap.org/soap/encoding/">
<SOAP-ENV:Body>
<m:getQuote xmlns:m="urn:stockquote:>
<symbol>IBM</symbol>
</m:getQuote>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```


The SOAPAction HTTP header in the request is defined in the SOAP specification and is used by HTTP proxy servers to dispatch requests to particular HTTP servers. WebSphere Application Server dynamic cache can use this header in its cache policies to build IDs without having to parse the SOAP message.

Message example 2 illustrates a SOAP message for a BuyQuote operation. While message 1 is cacheable, this message is not, because it updates the back-end database.

Message example 2

```
POST /soap/servlet/soaprouter
HTTP/1.1
Host: www.myhost.com
Content-Type: text/xml; charset="utf-8"
SOAPAction: urn:stockquote-update
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
<SOAP-ENV:Body>
<m:buyStock xmlns:m="urn:stockquote:>
<symbol>IBM</symbol>
</m:buyStock>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

The graphic illustrates how to invoke methods with the SOAP messages. In Web services terms, especially Web Service Definition Language (WSDL), a service is a collection of operations such as getQuote and buyStock. A body element namespace (urn:stockquote in our example) defines a service, and the name of the first body element indicates the operation.



The following is an example of WSDL for the getQuote operation:

```
<?xml version="1.0"?>
<definitions name="StockQuoteService-interface"
targetNamespace="http://www.getquote.com/StockQuoteService-interface"
xmlns:tns="http://www.getquote.com/StockQuoteService-interface"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns="http://schemas.xmlsoap.org/wsdl/"
<message name="SymbolRequest">
<part name="return" type="xsd:string"/>
</message>
<portType name="StockQuoteService">
<operation name="getQuote">
<input message="tns:SymbolRequest"/>
<output message="tns:QuoteResponse"/>
</operation>
</portType>
<binding name="StockQuoteServiceBinding"
type="tns:StockQuoteService">
<soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
<operation name="getQuote">
<soap:operation soapAction="urn:stockquote-lookup"/>
<input>
<soap:body use="encoded" namespace="urn:stockquote"
```

```

encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
</input>
<output>
<soap:body use="encoded" namespace="urn:stockquotes"
encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
</output>
</operation>>
</binding>
</definition>

```

To build a set of cache policies for a Web services application configure WebSphere Application Server dynamic cache to recognize cacheable service operation of the operation.

WebSphere Application Server inspects the HTTP request to determine whether or not an incoming message can cache based on the cache policies defined for an application. In the case of this example, buyStock and stock-update are not be cached, but stockquote-lookup is. In the cachespec.xml file for this Web application, the cache policies need defining for these services so that the dynamic cache can handle both SOAPAction and service operation.

WebSphere Application Server uses the operation and the message body in Web services cache IDs, each of which has a component associated with them. Therefore, each Web services <cache-id> rule contains only two components. The first is for the operation. Because you can perform the stockquote-lookup operation by either using a SOAPAction header or a service operation in the body, you must define two different <cache-id> elements, one for each method. The second component is of type "body", and defines how WebSphere Application Server should incorporate the message body into the cache ID. You can use a hash of the body, although it is legal to use the literal incoming message in the ID.

The incoming HTTP request is analyzed by WebSphere Application Server to determine which of the <cache-id> rules match. Then, the rules are applied to form cache or invalidation IDs.

The following is sample code of a cachespec.xml file defining SOAPAction and servicesOperation rules:

```

<cache>
<cache-entry>
<class>webservice</class>
<name>/soap/servlet/soaprouter</name>
<sharing-policy>not-shared</sharing-policy>
<cache-id>
<component id="" type=SOAPAction>
<value>urn:stockquote-lookup</value>
</component>
<component id="Hash" type="SOAPEnvelope"/>
<timeout>3600</timeout>
<priority>1</priority>
</cache-id>
<cache-id>
<component id="" type="serviceOperation">
<value>urn:stockquote:getQuote</value>
</component>
<component id="Hash" type="SOAPEnvelope"/>
<timeout>3600</timeout>
<priority>1</priority>
</cache-id>
</cache-entry>
</cache>

```

Example: Configuring the dynamic cache

This example puts all the steps together for configuring the dynamic cache with the `cachespec.xml` file, showing the use of the cache ID generation rules, dependency IDs, and invalidation rules.

Suppose we have a servlet which is used to manage a simple news site. This servlet uses the query parameter "action" to determine whether the request is being used to "view" news or "update" news (used by the administrator). Further, another query parameter "category" is used to select the news category. Further, suppose that this site supports an optional customized layout, which is stored in the user's session using the attribute name "layout". Here are example URL requests to this servlet:

`http://yourhost/yourwebapp/newscontroller?action=view&category=sports` (Returns a news page for the sports category)

`http://yourhost/yourwebapp/newscontroller?action=view&category=money` (Returns a news page for the money category)

`http://yourhost/yourwebapp/newscontroller?action=update&category=fashion` (Allows the administrator to update news in the fashion category)

Here are the steps for configuring dynamic cache with `cachespec.xml`, using the information provided to you:

1. Define the cache-entry elements necessary to identify the servlet. In this case, the servlet's URI is "newscontroller" so this will be our cache-entry's name element. Also, since we are caching a servlet/JavaServer Page (JSP), the cache-entry class is "servlet".

```
<cache-entry>
<name> /newscontroller </name>
<class>servlet </class>
</cache-entry>
```

2. Define cache ID generation rules. For this servlet, we only want to cache when `action=view`, so one component of the cache ID will be the parameter "action" when the value equals "view". The news category is also an essential part of the cache ID. Finally, the optional session attribute for the user's layout is included in the cache ID. The cache-entry now looks like this:

```
<cache-entry>
<name> /newscontroller </name>
<class>servlet </class>
<cache-id>
<component id="action" type="parameter">
<value>view</value>
<required>>true</required>
</component>
<component id="category" type="parameter">
<required>>true</required>
</component>
<component id="layout" type="session">
<required>>false</required>
</component>
</cache-id>
</cache-entry>
```

3. Define dependency ID rules. For this servlet, a dependency ID will be added for the category. Later, when the category is invalidated due to an update event, all views of that news category will be invalidated. After adding our dependency-id, the cache-entry now looks like this:

```

<cache-entry>
  <name>newscontroller </name>
  <class>servlet </class>
  <cache-id>
    <component id="action" type="parameter">
      <value>view</value>
      <required>true</required>
    </component>
    <component id="category" type="parameter">
      <required>true</required>
    </component>
    <component id="layout" type="session">
      <required>>false</required>
    </component>
  </cache-id>
  <dependency-id>category
    <component id="category" type="parameter">
      <required>true</required>
    </component>
  </dependency-id>
</cache-entry>
<cache-entry>
  <name>newscontroller </name>
  <class>servlet </class>
  <cache-id>
    <component id="action" type="parameter">
      <value>view</value>
      <required>true</required>
    </component>
    <component id="category" type="parameter">
      <required>true</required>
    </component>
    <component id="layout" type="session">
      <required>>false</required>
    </component>
  </cache-id>
  <dependency-id>category
    <component id="category" type="parameter">
      <required>true</required>
    </component>
  </dependency-id>
</cache-entry>

```

4. Define invalidation rules. Since we defined a category dependency ID, we will now define an invalidation rule to invalidate the category when action=update. To incorporate the conditional logic, we will add "ignore-value" components into the invalidation rule. These components will not add to the output of the invalidation ID, but will only determine whether or not the invalidation ID is created and executed. The final cache-entry now looks like this:

```

<cache-entry>
  <name>newscontroller </name>
  <class>servlet </class>
  <cache-id>
    <component id="action" type="parameter">
      <value>view</value>
      <required>true</required>
    </component>
    <component id="category" type="parameter">
      <required>true</required>
    </component>
    <component id="layout" type="session">
      <required>>false</required>
    </component>
  </cache-id>
  <dependency-id>category
    <component id="category" type="parameter">
      <required>true</required>
    </component>
  </dependency-id>
</cache-entry>

```

```
</component
</dependency-id
<invalidation>category<
  <component id="action" type="parameter" ignore-value="true">
    <value>update</value>
    <required>true</required>
  </component>
  <component id="category" type="parameter">
    <required>true</required>
  </invalidation>
</cache-entry>
```

Chapter 15. Managing user profiles

Note: User Profile Manager API is deprecated in the current release, and there is no replacement available.

IBM WebSphere Application Server provides a service for processing user profiles, called the *User Profile Manager*. The service is provided in the form of an EJB entity bean that servlets can call whenever they are required to access a user profile.

The key activities for implementing user profiles are summarized.

Steps for this task

1. Customize the user profile support as necessary.

Options include:

- Using the data representation class with the name-value pairs it currently supports (no action required)
- Extending the data representation class to support additional, arbitrary name-value pairs
- Adding columns to the base user profile representation
- Extending the User Profile enterprise bean to import existing databases

Evaluate whether the user profile representation provided by IBM represents the kind of data you want to keep about your users. You might find it desirable to customize the IBM user profile support.

2. Create or modify servlets to use the User Profile Manager and related user profile support classes to maintain user profiles on behalf of Web applications.
3. (Assemble your application.)
4. (Deploy your application.)
5. Ensure the administrator appropriately configures User Profile Managers using `userprofile.xml` file.

If the programmer and administrator are not the same person, the programmer might need to provide settings information to the administrator, based on how the programmer implemented user profiles.

User profile

Some applications collect data about the users with which they interact. The data is stored in a database. The next time the user interacts with the application, the application recalls the data.

Because the application already knows something about the user, it can provide the user with a more personalized experience.

User profiles provides a means by which a company can maintain and manage database tables containing fields for demographic data, and use those tables to interact with a database of individual customers or other users on the company system.

For example, when a repeat user logs onto a Web site that supports user profiles, the Web site can display headlines and advertising tailored to the shopping preferences of that user. The site can address the user by logon name.

An application implementing user profiles requires database access for storing the user profile data it gathers.

UserProfileManager class

Servlets and other application building blocks requiring user profile support should make calls to the class:

```
com.ibm.websphere.userprofile.UserProfileManager
```

The class supports the following functions:

- Creating and deleting user profiles
- Getting and updating (cached and immediate) to and from the database
- Getting user profiles for read-only tasks
- Performing queries on database columns

User profile development options

The application developer has a few options for customizing the user profile support provided by IBM WebSphere Application Server. The Related information provides instructions and additional details about each option.

Extending the data represented in user profiles

Web applications can maintain several pieces of data about users. You can extend the data representation to allow the collection of arbitrary name-value pairs.

Use the following interface with the `com.ibm.websphere.userprofile.UserProfileExtender` class to extend a user profile hash table:

```
com.ibm.websphere.userprofile.UserProfileProperties
```

This action enables you to place arbitrary name-value pairs in the user profile. Extending the hash table is similar to using the `java.util.Dictionary` class in the base JDK 1.x, or any of the classes that extend it.

Adding columns to the base user profile implementation

Application developers can customize user profiles by adding columns to the base user profile implementation. Adding new columns is accomplished by implementing the interface:

```
com.ibm.websphere.userprofile.UserProfileExtender
```

and extending the base class:

```
com.ibm.servlet.personalization.userprofile.UserProfile
```

The application developer can add columns to but not delete columns from the base implementation.

Adding columns is a two-step process, as follows:

1. Extend the `UserProfile` class.
2. Modify your existing servlets to use the new columns.

Several examples are available to demonstrate how to extend the base user profile implementation and utilize the extension with a servlet.

Example	Description
UPServletExample.java	Demonstrates how a servlet opens a user profile and prints the fields contained within.
UserProfileExtendedSample.java	Shows how to extend the UserProfile class to add a column to the user profile for a cellular phone number. The WebSphere Application Server administrator configures the User Profile Manager to point to the extended class.
UPServletExampleExtended.java	Shows how to modify the UPServletExample servlet to include the cellular phone number in the output.
UserProfileExtended.java	Shows how to extend a hash table to place arbitrary name-value pairs into the user profile.
UPServletExtended.java	Shows how to extend the servlet. When any of the newly added columns are removed or replaced, look for the table named "USERPROFILE" in the database to which the user profile is configured and drop that table.

The examples are encoded in HTML for viewing in a browser. The documentation directory also contains nonHTML versions (.java files) that are ready for use.

Extending the User Profile enterprise bean and importing legacy databases

Application developers can extend the User Profile enterprise bean itself and import legacy databases into the user profile. The main advantage in extending the User Profile enterprise bean is to gain the ability to import existing databases into the user profile. You can also extend this enterprise bean to add columns to the base user profile implementation.

UPServletExample.java

```
import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;

//Creates a Userprofile using the new API

public class UPServlet_ReadWrite extends HttpServlet {

    public void doGet (HttpServletRequest req,
        HttpServletResponse res)
```

```

throws ServletException, IOException {

    PrintWriter out;
    res.setContentType("text/html");
    out = res.getWriter();

    UserProfileManager manager =
        UserProfileManager.getUserProfileManager();
    UserProfile userprofile;

    try {

        //Try creating the UserProfile
        userprofile = manager.addUserProfile("bpink");

    } catch(UserProfileCreateException e1) {

        try {

            //Try finding the existing in readWrite mode.
            //Second argument indicates whether we want
            //to get userprofile
            //in read only mode or read write mode.

            userprofile = manager.getUserProfile("bpink",true);

        } catch(UserProfileFinderException e) {
            e.printStackTrace();
            return;
        }
    }

    //Set the properties

    userprofile.setAddress1("myaddress1");
    userprofile.setAddress2("myaddress2");
    userprofile.setFirstName("Pinkowski");
    userprofile.setSurName("Ben");
    userprofile.setDayPhone("555-6677");
    userprofile.setNightPhone("556-6765");
    userprofile.setCity("MYCITY");
    userprofile.setNation("myCountry");
    userprofile.setEmployer("MyEmployer");
    userprofile.setFax("7823470");

    userprofile.setLanguage("mylanguage");
    userprofile.setEmail("MyEmail@email");
    userprofile.setStateOrProvince("myState");
    userprofile.setPostalCode("xxxxx");

    //Freeing resources held by userprofile
    manager.releaseResources(userprofile);
    userprofile=null;

    //Checking whether it updated the info

    try {

        //Getting the existing userprofile in ReadOnly mode.

        userprofile = manager.getUserProfile("bpink",false);

    } catch(UserProfileFinderException e1) {

```

```

        out.println("Error finding ");
        e1.printStackTrace();
        return;
    }

    //Displaying the properties of userprofile

    out.println(userprofile.getAddress1()+"<br>");
    out.println(userprofile.getAddress2()+"<br>");;
    out.println(userprofile.getFirstName()+"<br>");;
    out.println(userprofile.getSurName()+"<br>");;
    out.println(userprofile.getDayPhone()+"<br>");;
    out.println(userprofile.getNightPhone()+"<br>");;
    out.println(userprofile.getCity()+"<br>");;
    out.println(userprofile.getNation()+"<br>");;
    out.println(userprofile.getEmployer()+"<br>");;
    out.println(userprofile.getFax()+"<br>");;
    out.println(userprofile.getLanguage()+"<br>");;
    out.println(userprofile.getEmail()+"<br>");;
    out.println(userprofile.getStateOrProvince()+"<br>");;
    out.println(userprofile.getPostalCode()+"<br>");;

    //Freeing resources held by userprofile
    manager.releaseResources(userprofile);
}
}

```

UserProfileExtendedSample.java

```

/* -----
** Copyright 1997-99 IBM Corporation. All rights reserved.
**
** -----
*/
package com.ibm.servlet.personalization.userprofile;

import com.ibm.servlet.personalization.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileExtender;

//Extensions of UserProfile to add new Columns should
//implement UserProfileExtender
public class UserProfileExtendedSample
    extends com.ibm.servlet.personalization.userprofile.UserProfile
    implements UserProfileExtender {

    //New column that is being added by this
    //derived class.
    public String cellPhone;

    //Manager Class will call this method to append new
    //Column types.
    //If UserProfile class is extended to append new columns
    //TOTAL COLUMNS: Base Class columns + columns returned
    //by this class

    public String[] getNewColumns() {
        //If variable name is "cellPhone," you need to
        //return "cellPhone" in array format. JDBC equivalent will
        //be generated automatically. You can add multiple columns.
        //For multiple columns: String newCol=
        //{"fieldName1","fieldName2",...};
        String[] newCol={"cellPhone"};
        return newCol;
    }
}

```

```

public String getCellPhone() {
    // Need to call this method to
    // get the things from persistence store.
    return(String)getByType("cellPhone");
}

public void setCellPhone(String value) {
    cellPhone = value;
    //Call this method to store the
    //things in persistence store
    setByType("cellPhone", value);
}
}

```

UPServletExampleExtended.java

```

import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;
import com.ibm.servlet.personalization.userprofile.UserProfileExtendedSample;

public class UPServletExtendedSample extends HttpServlet {

    public void doGet (HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException
    {
        UserProfileManager manager =
            UserProfileManager.getUserProfileManager();
        UserProfile userprofile;

        PrintWriter out;

        res.setContentType("text/html");
        out = res.getWriter();

        try {

            //try Creating the UserProfile
            userprofile = manager.addUserProfile("bpink");
            } catch( UserProfileCreateException e1) {

                try { //try finding the existing in readWrite mode

                    userprofile = manager.getUserProfile("bpink",true);
                } catch(UserProfileFinderException e) {
                    e.printStackTrace();
                    return;
                }
            }

            userprofile.setAddress1("myaddress1");
            userprofile.setAddress2("myaddress2");

```

```

        userprofile.setFirstName("Pinkowski");
        userprofile.setSurName("Ben");
        userprofile.setDayPhone("555-6677");
        userprofile.setNightPhone("556-6765");
        userprofile.setCity("MYCITY");
        userprofile.setNation("myCountry");
        userprofile.setEmployer("MyEmployer");
        userprofile.setFax("7823470");
        userprofile.setLanguage("mylanguage");
        userprofile.setEmail("MyEmail@email");
        userprofile.setStateOrProvince("myState");
        userprofile.setPostalCode("xxxxx");

//calling setCellPhone
((com.ibm.servlet.personalization.userprofile.
    UserProfileExtendedSample)
    userprofile).setCellPhone("346-4588");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info
try {

    //Getting the existing userprofile
    userprofile = manager.getUserProfile("bpink",false);

} catch( UserProfileFinderException e1) {

    out.println("Error finding ");
    e1.printStackTrace();
    return;
}

out.println(userprofile.getAddress1()+"<br>");
out.println(userprofile.getAddress2()+"<br>");
out.println(userprofile.getFirstName()+"<br>");
out.println(userprofile.getSurName()+"<br>");
out.println(userprofile.getDayPhone()+"<br>");
out.println(userprofile.getNightPhone()+"<br>");
out.println(userprofile.getCity()+"<br>");
out.println(userprofile.getNation()+"<br>");
out.println(userprofile.getEmployer()+"<br>");
out.println(userprofile.getFax()+"<br>");
out.println(userprofile.getLanguage()+"<br>");
out.println(userprofile.getEmail()+"<br>");
out.println(userprofile.getStateOrProvince()+"<br>");
out.println(userprofile.getPostalCode()+"<br>");

//Calling getCellPhone
out.println(((UserProfileExtendedSample)userprofile).
    getCellPhone()+"<br>");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//For getting values by cellPhone
out.println("<br><br>Retreiving by Cell Phone <br>");

```



```

Enumeration enum = manager.findUserProfiles
    ("cellPhone","346-4588");
while(enum.hasMoreElements()) {

    com.ibm.websphere.userprofile.UserProfile up =
        (com.ibm.websphere.userprofile.UserProfile)
            enum.nextElement();
    out.println("first name :"+up.getFirstName()+
        "<br>");

        //Freeing resources held by userprofile
        manager.releaseResources(up);
    }
}

```

UserProfileExtended.java

```

package com.ibm.servlet.personalization.userprofile;
/* -----
** Copyright 1997-99 IBM Corporation. All rights reserved.
** -----
*/
import java.util.*;

import com.ibm.servlet.personalization.userprofile.UserProfile;

import com.ibm.websphere.userprofile.UserProfileExtender;
import com.ibm.websphere.userprofile.UserProfileProperties;

public class UserProfileExtended extends UserProfile
    implements UserProfileExtender,
        UserProfileProperties {
    //New column that is being added by this
    //derived class.
    public Hashtable properties;

    static String propCol ="properties";

    //Manager Class will call this method to append
    //new Column types
    //to SQL Strings. If UserProfile class is extended
    //to append new columns
    //it should implement UserProfileExtender.
    //COLUMNS: Base Class columns + columns returned by this class

    public String[] getNewColumns() {
        //if variable name is properties, you need to
        //return "properties" . JDBC equivalent will be
        //generated automatically.
        String[] newCol={propCol};
        return newCol;
    }

    public Object getValue(String key) {
        // Need to call this method to
        // get the things from persistent store
        properties = (Hashtable) getByType(propCol);

        if(properties != null)
            return properties.get(key);

        else return null;
    }
}

```

```

public void putValue(String key, Object value) {
    properties =(Hashtable) getByType(propCol);
    if(properties == null)
        properties = new Hashtable();
    properties.put(key,value);
    //store in persistent store
    setByType(propCol, properties);
}

public void removeValue(String key) {
    properties = (Hashtable) getByType(propCol);
    if(properties == null)
        return;
    properties.remove(key);
    //store in persistent store
    setByType(propCol, properties);
}
}

```

UPServletExtended.java

```

import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;
import com.ibm.websphere.userprofile.UserProfileProperties;

public class UPServletExtended extends HttpServlet {

    public void doGet (HttpServletRequest req,
        HttpServletResponse res)
        throws ServletException, IOException
    {
        UserProfileManager manager =
            UserProfileManager.getUserProfileManager();
        UserProfile userprofile;

        PrintWriter out;

        res.setContentType("text/html");
        out = res.getWriter();

        try {

            //try Creating the UserProfile

            userprofile = manager.addUserProfile("bpink");

        } catch( UserProfileCreateException e1) {

            try { //try finding the existing in readWrite mode

```

```

        userprofile = manager.getUserProfile("bpink",true);
    } catch(UserProfileFinderException e) {
        e.printStackTrace();
        return;
    }
}

userprofile.setAddress1("myaddress1");
userprofile.setAddress2("myaddress2");
userprofile.setFirstName("Pinkowski");
userprofile.setSurName("Ben");
userprofile.setDayPhone("555-6677");
userprofile.setNightPhone("556-6765");
userprofile.setCity("MYCITY");
userprofile.setNation("myCountry");
userprofile.setEmployer("MyEmployer");
userprofile.setFax("7823470");
userprofile.setLanguage("mylanguage");
userprofile.setEmail("MyEmail@email");
userprofile.setStateOrProvince("myState");
userprofile.setPostalCode("xxxxx");

//calling putValue
((UserProfileProperties)userprofile).putValue
("name","HHHHHHH");
((UserProfileProperties)userprofile).putValue
("Date",new java.util.Date());

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info
try {

    //Getting the existing userprofile

    userprofile = manager.getUserProfile("bpink",false);

    out.println(userprofile.getAddress1()+"<br>");
    out.println(userprofile.getAddress2()+"<br>");
    out.println(userprofile.getFirstName()+"<br>");
    out.println(userprofile.getSurName()+"<br>");
    out.println(userprofile.getDayPhone()+"<br>");
    out.println(userprofile.getNightPhone()+"<br>");
    out.println(userprofile.getCity()+"<br>");
    out.println(userprofile.getNation()+"<br>");
    out.println(userprofile.getEmployer()+"<br>");
    out.println(userprofile.getFax()+"<br>");
    out.println(userprofile.getLanguage()+"<br>");
    out.println(userprofile.getEmail()+"<br>");
    out.println(userprofile.getStateOrProvince()+"<br>");
    out.println(userprofile.getPostalCode()+"<br>");

    //Getting the values

    out.println(((UserProfileProperties)
        userprofile).getValue("name")+"<br>");
    out.println(((UserProfileProperties)
        userprofile).getValue("Date")+"<br>");
    out.println("Removing Values ");
    ((UserProfileProperties)

```

```

        userprofile).removeValue("name");
        ((UserProfileProperties)
        userprofile).removeValue("Date");
        out.println(((UserProfileProperties)
        userprofile).getValue("name")+
        "<br>");
        out.println(((UserProfileProperties)
        userprofile).getValue("Date")+"<br>");

        //Freeing resources held by userprofile
        manager.releaseResources(userprofile);

    } catch( UserProfileFinderException e1) {

        out.println("Error finding ");
        e1.printStackTrace();
        return;
    }
}
}
}

```

userprofile.xml

To *installation_root*/properties, add a file named `userprofile.xml` in the following format. Specify enterprise bean class names; data wrapper class name; and JNDI names for the read-only bean, read/write bean, and data source (from step 2). You must also add user ID and password information for the JNDI data source

The following example file contains class names as provided in WebSphere Application Server. If data wrapper and enterprise bean class names are extended programmatically, change them accordingly.

```

<?xml version="1.0"?>

<userprofile>
<userprofile-enabled>true</userprofile-enabled>

<userprofile-wrapper-class>
<classname>
    com.ibm.servlet.personalization.userprofile.UserProfile
</classname>
</userprofile-wrapper-class>
<userprofile-manager-name>
    User Profile Manager
</userprofile-manager-name>
<userprofile-bean>
<readonly-interface>
    com.ibm.servlet.personalization.userprofile.UP_ReadOnly
</readonly-interface>
<readwrite-interface>
    com.ibm.servlet.personalization.userprofile.UP_ReadWrite
</readwrite-interface>
<readonlyhome-interface>
    com.ibm.servlet.personalization.userprofile.UP_ReadOnlyHome
</readonlyhome-interface>
<readwritehome-interface>
    com.ibm.servlet.personalization.userprofile.UP_ReadWriteHome
</readwritehome-interface>
<readonly-JNDI-lookupName>UP_ReadOnlyHome</
    readonly-JNDI-lookupName>
<readwrite-JNDI-lookupName>UP_ReadWriteHome</
    readwrite-JNDI-lookupName>
</userprofile-bean>

<userprofile-store>

```

```
<database-userid></database-userid>  
<database-password></database-password>  
<database-datasource></database-datasource>  
</userprofile-store>  
  
</userprofile>
```

Chapter 16. Assembling applications

Before you begin

Assemble application modules (known as EAR files) from new or existing J2EE 1.3 modules, including these archives: Web application archives (WAR), resource adapter archives (RAR), enterprise beans (EJB JAR), and application client archives (JAR). This packaging and configuration of code artifacts into application modules or standalone Web modules is necessary for deploying the applications onto the application server.

Gather the code artifacts that you want to package into one or more assembled modules. Code artifacts include these items that you have created and unit tested in your favorite integrated development environment:

- Enterprise beans
- Servlets, JavaServer Pages (JSP) files and other Web components
- Resource adapter (*connector*) implementations
- Application clients
- Other supporting classes and files

Steps for this task

1. Start the Application Assembly Tool (AAT).
2. **(Optional)** Migrate existing J2EE 1.2 modules to J2EE 1.3.

The Application Assembly Tool (AAT) has an option for migrating J2EE 1.2 application modules to J2EE 1.3. The J2EE 1.2 module is kept intact, with a new 1.3 module created. See also the `earconvert` tool documentation.

You must migrate J2EE 1.2 application modules to which you want to add J2EE 1.3 level WAR, RAR, EJB and client modules. This tool migrates only the application modules. J2EE 1.2-level modules inside a J2EE 1.2 application module must be migrated by other means.

3. (Assemble new EJB modules (enterprise beans JAR files)) as needed.

Assemble an EJB module to contain enterprise beans and related code artifacts. (Group Web components, client code, and resource adapter code in separate modules.)

You can install an EJB module as a standalone application or you can combine it with other modules into an enterprise application.

4. (Assemble new Web modules (WAR files)) as needed.

Assemble a Web module to contain servlets, JSP files, and related code artifacts. (Group enterprise beans, client code, and resource adapter code in separate modules.)

You can install a Web module as a standalone application or combine it with other modules into an enterprise application.

5. (Assemble new application client modules (client JAR files)) as needed.
6. Assemble new resource adapter archives (RAR files) as needed.

Assemble a resource adapter archive module to contain the library implementation code that your application uses to connect to enterprise information systems (EIS). (Group enterprise beans, Web components, and client code in separate modules.)

7. Assemble an application module from other module types.
You are ready to combine your new or migrated modules into an application module (EAR file).
For applications containing only Web modules, this step is optional. It is feasible to deploy Web modules without assembling them into application modules.
8. Verify your archive files.
Verify your archive files and correct any problems so that generation of deployment code is successful. During verification, the Application Assembly Tool (AAT) checks that an archive file is complete, and that deployment descriptor properties and references contain appropriate values.
9. Remember to save your application one last time.
10. Generate code for deployment for applications containing EJB modules.
If the application modules contain EJB modules, you must generate deployment code for the enterprise beans in the application before you deploy applications on the server. The Application Assembly Tool (AAT) provides this ability, or you can use the `ejbdeploy` command line tool.
11. Open existing modules (**File > Open**) in the AAT to modify them as needed.
For example, you can add or remove modules and edit deployment descriptor properties.

What to do next

After assembling your applications, use a systems management tool to deploy the EAR or WAR files onto the application server.

The systems management tool follows the security and deployment instructions defined in the deployment descriptor, and enables you to modify bindings specified within the AAT. The tool locates the required external resources that the application uses, such as enterprise beans and databases.

Select a tool to use:

- Deploying and managing applications with the GUI
- Deploying and managing applications using programming
- Deploying and managing applications using scripting

If you are uncertain of which systems management tool to use, try using the administrative console.

Application assembly and J2EE applications

Application assembly is the process of creating an Enterprise Archive (EAR) file containing all files related to an application, as well as an XML deployment descriptor for the application. This configuration and packaging prepares the application for deployment onto an application server.

EAR files are comprised of the following archives:

- Enterprise bean (JAR) files (known as EJB modules)
- Web application (WAR) files (known as Web modules)
- Application client (JAR) files (known as client modules)
- Resource adapter (RAR) files (known as resource adapter modules)

Ensure that modules are contained in an EAR file so that they may be deployed onto the server. The exceptions are WAR modules, which you can deploy individually. Although WAR modules can contain regular JAR files, they cannot contain the other module types described previously.

The assembly process includes the following:

- Selecting all of the files to include in the module
- Creating a deployment descriptor containing instructions for module deployment on the application server.

As you configure properties using the Application Assembly Tool (AAT), the tool generates the deployment descriptor for you. While the AAT graphical interface is recommended, you can also edit descriptors directly in your favorite XML editor.

- Packaging modules into a single Enterprise archive (EAR) file, which contains one or more files in a compressed format

Archive support in Version 5.0

These archives and Web components are supported:

- J2EE 1.3 Enterprise application (EAR) files
- EJB 2.0 (JAR) files
- Servlet 2.3 Web application WAR files
- Application Client 1.3 JAR files
- Connector 1.0 RAR files

These archive files and Web components are back-level and may be read but not created or changed:

- J2EE 1.2 EAR files
- EJB 1.1 JAR files
- Servlet 2.2 WAR files
- Application Client 1.2 JAR files

Starting the Application Assembly Tool (AAT)

Before you begin

A graphical interface is available for packaging code artifacts into various archives (modules) and configuring their J2EE 1.3 compliant deployment descriptors.

The Application Assembly Tool (AAT) is available from the Windows Start menu, or you can invoke the tool from a command line as described.

Steps for this task

1. Change the directory at a system command prompt to the location of the `assembly.bat|sh` file, typically `install_root/bin`.
2. Run the assembly script to launch the graphical interface.
3. Select whether to work with an existing module or create a new one.

Results

The navigation tree displays a hierarchical structure used to build the contents of a new module, or to work with the contents of an existing module. Icons in the tree

represent the components, assembly properties, and files for the module. The assembly properties appear in the AAT workspace.

What to do next

Consider whether you have any existing J2EE 1.2 application modules that you would like to migrate to J2EE 1.3.

You can create new modules of the following types, to assemble into an application module later:

- (Assembling EJB modules)
- (Assembling Web modules)
- (Assembling application client modules)
- Assembling resource adapter modules

Rather than create new modules to assemble an application, you can proceed directly to assembling a new application module. While assembling an application module, you can create any new modules that you need.

Migrating application modules from J2EE 1.2 to J2EE 1.3

The Application Assembly Tool (AAT) has an option for migrating J2EE 1.2 application modules to J2EE 1.3. The J2EE 1.2 module is kept intact, with a new 1.3 module created. See also the earconvert tool documentation.

Migrate J2EE 1.2 application modules to which you want to add J2EE 1.3 level Web application (WAR) modules, Resource adapter (RAR) modules, Entity bean (EJB) modules, and application client modules. This tool migrates only the application modules. Migrate J2EE 1.2-level modules inside a J2EE 1.2 application module by other means.

Note: When Entity beans are moved from a J2EE 1.2 module to a J2EE 1.3 module, the EJB container will then apply rules defined in the EJB 2.0 specification to these beans. The EJB 2.0 specification mandates that when a `findBy` method is called on a bean home (except for `findByPrimaryKey`), the EJB container must cause other Entity beans enlisted in the same transaction to write out their current state to the persistent store. This is to ensure that the `findBy` operation is performed on the most current data. Application developers should plan for and be aware of any changes to the application behavior as a result of this rule.

Steps for this task

1. Start the AAT.
2. Use it to open the J2EE 1.2 application module you want to migrate.
3. Click **Convert EAR** from the file menu.
4. Save the new J2EE 1.3 application.

What to do next

Assemble zero or more new modules of your choice:

- (Assembling EJB modules)
- (Assembling Web modules)
- (Assembling application client modules)
- Assembling resource adapter modules

Another option is to proceed directly to assembling a new application module. You can create any new modules that you need, while assembling an application module.

earconvert tool

A command line tool is provided for migrating J2EE 1.2 application modules to J2EE 1.3. This migration enables you to add J2EE 1.3 modules to the migrated application module. See also the Application Assembly Tool (AAT) for information on performing this task.

Migrate J2EE 1.2 application modules to which you want to add J2EE 1.3 level Web application (WAR), Resource adapter (RAR), Enterprise beans (EJB), and client modules. This tool migrates only the application modules. Ensure that you migrate J2EE 1.2-level modules inside a J2EE 1.2 application to prevent working with back-level files.

Syntax

You must locate the file that contains this tool (`earconvert.bat|sh`), which must be located in the `bin` directory of the product installation root.

earconvert

j2ee_1.2_file_name

j2ee_1.3_file_name

Parameters

Supported arguments include:

"j2ee_1.2_file_name"

Specifies the actual name of the existing J2EE 1.2 application file. (In this and other arguments, use quotation marks to allow for path names that contain spaces.)

"j2ee_1.3_file_name"

Specifies what you would like to name the new J2EE 1.3 application file.

Examples

The following command creates a new J2EE 1.3 archive, `new_application.ear`, based on the J2EE 1.2 archive, `existing_application.ear`.

```
earconvert existing_application.ear new_application.ear
```

Assembling new or modifying existing modules

Before you begin

Ensure code artifacts, such as servlets, JSP files, enterprise beans, and application clients are assembled into their respective modules.

If you want to use existing J2EE 1.2 modules in your J2EE 1.3 application, migrate these modules to J2EE 1.3 first. Also migrate any J2EE 1.2 application modules to which you want to add J2EE 1.3 modules.

You are now ready to combine your new or migrated modules into an application module Enterprise application (EAR file).

The Application Assembly Tool (AAT) provides flexibility in assembling applications from various Web application (WAR), Resource adapter (RAR), Enterprise beans (EJB JAR), and application client (JAR) files. Options described below include:

- Importing an existing module (JAR, RAR or WAR file)
- Creating a new module while you create the new application
- Copying code artifacts, such as servlets, from one module to another of the same type, to reside in the new application

Steps for this task

1. Start the AAT.
2. Click **File > New > Application** if you did not already specify to create a new application module.
Each of the next three steps is optional, but you must perform at least one of them.
3. **(Optional)** Import existing modules into the application module.
 - a. Right-click the folder for the type of module you want to import, such as an EJB module, in the navigation tree.
 - b. Click **Import** from its right-click menu.
 - c. Use the file browser to locate and select the archive file for the module.
 - d. Click **Open**. The archive file appears under the appropriate folder in the navigation tree.
 - e. Click the plus sign (+) next to the icon for the archive, to view the module contents and edit its properties if needed.
 - f. Save the application module.
4. **(Optional)** Create a new archive file to include in the application.
 - a. Right-click the folder for the type of module to create (such as enterprise beans (EJB) modules, Web application modules (WAR), resource adapter (RAR) files, or application client modules) in the navigation tree.
 - b. Click **New** from its right-click menu.
 - c. Configure properties of the new module when it displays.
 - d. Click **OK**. The archive file displays under the appropriate folder.
 - e. Click the plus sign (+) to verify file contents and enter assembly properties.
 - f. Add enterprise beans, if this is an EJB module.
 - g. Right-click the folder corresponding to the type of bean to create (session bean or entity bean), and click **New** or **Import**.
 - h. Configure properties of the enterprise bean when it displays.
 - i. Click **OK**. The enterprise bean appears in the navigation pane.
 - j. Click the plus sign (+) to verify file contents and enter assembly properties.
 - k. Save the application module.
5. **(Optional)** Copy code artifacts, such as servlets, from one module to another of the same type, to reside in the new application.
 - a. Identify the code artifact to copy, and the type of module in which it resides. Make sure you already have the same kind of module (such as a Web module) created in the new application module.

- b. Open a separate, existing module in the AAT by selecting **File > Open** from the menu bar.
 - c. Arrange the AAT workspace so that you can see both the new application module and the source archive containing the code artifact.
 - d. Copy and paste the code artifact from the source module to the same module type in the new application.
For example, copy a container-managed persistence (CMP) bean from the source EJB module into the new EJB application module.
 - e. Save the application module.
6. Continue to add desired modules to the application module.
 7. Define security properties for the application.
 - a. Right-click the **Security Roles** icon in the navigation tree.
 - b. Click **New**.
 - c. Configure the security properties.
 - d. Click **OK**.
 8. Add supplementary files needed by the application.
 - a. Right-click the **Files** icon in the navigation tree, and select **Add Files**.
 - b. Add files, using the Add Files dialog.
 9. Save the application module.

Results

You are performing application assembly results in a J2EE 1.3 compliant EAR file containing one or more WAR, RAR, or JAR files.

Note: If you use the Application Assembly Tool to create application client modules, you must also use the Application Client Resource Configuration Tool. Using this tool, you can define references to resources (other than enterprise beans) on the machine where the application client resides.

What to do next

"Verifying archive files". "Generating code for deployment".

After an application is assembled and you generate your application code for deployment, use the administrative console to install the application in an application server.

Use the administrative console at installation time to carry out the security instructions defined in the deployment descriptor and to locate required external resources, such as enterprise beans and databases. You can add configuration properties and redefine binding properties defined in the Application Assembly Tool.

After the application deploys, you can use the Application Assembly Tool to modify the application by adding or removing modules, editing deployment descriptor properties and regenerating code for deployment.

Note: If your application has a large number of modules, it might not install successfully onto a server. Package your application so the .ear file has as few modules as are necessary. Modules can include metadata for the modules such as information on deployment descriptors, bindings and IBM extensions.

Adding files to assembled modules

Before you begin

Review the usage scenario (as follows) to become familiar with the **Add Files** dialog.

Use the **Add Files** dialog box of the Application Assembly Tool (AAT) to import files into assembled modules including Enterprise application (EAR), Web application (WAR), Resource adapter (RAR) and Application client (JAR) files.

This task assumes that you are performing another task, such as assembling a Web or EJB module, when the **Add Files** dialog is presented to you.

Steps for this task

1. Click **Browse**. Locate the files to add.
 - To add specific individual files, select the directory or archive (WAR, JAR, RAR, ZIP, for example) containing the files.
 - To add an entire directory of files, select its parent directory.
2. Click **OK**.


The selected directory or archive appears in the top left part of the dialog box, in an expandable tree.

The top right part of the dialog box shows the contents of the directory, subdirectory, or archive that is selected on the left-hand side.
3. Select one or more items to add from the top right part of the dialog, then click **Add**.

As you add files, they will be displayed in the lower half of the dialog box.
4. Change your left-hand selection to gain access to other files that you want to add, as needed.
5. Click **OK** when all of the files that you want to add appear in the lower half of the dialog box.

Usage scenario

The following example refers to the main areas of the Add Files dialog box. Details such as clicking OK are omitted. Refer to the detailed task steps above for this information.



Suppose you are constructing a new application module and want to add `myFile.txt` file to the archive as a supplementary file. The `myFile.txt` currently is contained within the `myFiles` subdirectory of a JAR file that resides somewhere on your directory system.

1. Browse for the JAR file.
2. Select the JAR file.
3. Exit the browse dialog. At this point:
 - Area 1 of the **Add Files** dialog contains the path to the JAR file.
 - Area 2 displays the JAR file name as the root directory of an expandable tree showing the directories in the JAR file — including `myFiles`.
 - Area 3 shows the root contents of the JAR file, as well as any subdirectories visible from the root.
4. Select the `myFiles` directory from area 2, causing `myFile.txt` to become visible in area 3.
5. Click `myFiles.txt` from area 3 and specify to **Add** the file. Now this file is listed in area 4, the lower half of the dialog, which indicates it is the file that you want to add to the new application module.
6. Exit the **Add Files** dialog.

Enterprise application assembly properties

Use the assembly properties to change the settings for an enterprise application.

File name

Specifies the file name of the application.

Data type String

Display name

(Required) Specifies a short name intended to be displayed by GUIs. It identifies the application at deployment time.

Data type String

Description

Describes the application.

Data type String

Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels). The image provides an icon to represent the application in a GUI.

Data type String

Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels). The image provides an icon to represent the application in a GUI.

Data type String

Reload interval

Specifies the number of seconds that should elapse before reloading EJB modules.

Data type	Integer
Units	Seconds

Enterprise application name

Specifies the global JNDI name of the application.

Data type	String
-----------	--------

Run-as bindings - Authorization Data - User ID

Specifies the security principal used for the running of the application's enterprise bean methods. This security principal is granted the named role.

Data type	String
-----------	--------

Run-as bindings - Authorization Data - Password

Specifies the password of the security principal.

Data type	String
-----------	--------

Run-as bindings - Role name

Specifies a security role name. This name appears as the value of the Use Identity Assigned to Specified Role property of the enterprise bean.

Data type	String
-----------	--------

Resource environment reference assembly settings

Resource environment reference elements contain declarations of an enterprise bean's reference to an administered object associated with a resource in the enterprise bean's environment.

Name

Specifies the name of the resource environment reference.

Its value is the environment entry name used in the enterprise bean code.

Data type	String
-----------	--------

Description

Contains the information that the EJB jar file producer wants to provide to the EJB jar file consumer.

Data type	String
-----------	--------

Type

Specifies the type of a resource environment reference.

Data type	String
-----------	--------

Resource Adapter Archive file assembly settings

Use this page to set the resource adapter archive file properties.

File name

Specifies the file name of the Resource Adapter Archive.

Data type String

Display name

Specifies a short name that is intended to be displayed by the GUI.

Data type String

Description

Specifies a description that should include any information that the component file producer wants to provide to the consumer of the component file (that is, to the deployer).

Data type String

EIS type

This helps in identifying EIS instances that can be used with this resource adapter.

Data type String

Vendor name

Specifies a string-based version of the resource adapter from the resource adapter provider.

Data type String

Version

Specifies a string-based version of the resource adapter from the resource adapter provider.

Data type String

Specification

Specifies the version of the connector architecture specification that is supported by this resource adapter.

Data type String

License required

Specifies if a license is or is not required.

Description

If a license is required, this field specifies the licensing requirements for the resource adapter module. For example, duration of license, number of connection restrictions, and so forth.

Data type String

Implementation

The element (managedconnectionfactory class) that specifies the fully qualified name of the Java class that implements the javax.resource.spi.ManagedConnectionFactory interface.

Data type

Class

Interface

The element (credential-interface) that specifies the interface that the resource adapter implementation supports for the representation of the credentials.

The possible values are:

```
<credential-interface> javax.resource.spi.security.PasswordCredential
</credential-interface>
<credential-interface> javax.resource.spi.security.GenericCredential
</credential-interface>
```

Data type

Class

Implementation

The element (connectionfactory class) that specifies the fully-qualified name of the ConnectionFactory class that implements the resource adapter specific ConnectionFactory interface.

Data type

Class

Interface

The element (connection-interface) that specifies the fully-qualified name of the Connection interface supported by the resource adapter.

Data type

Class

Implementation

The element (connection class) that specifies the fully-qualified name of the Connection class that implements the resource adapter specific Connection interface.

Data type

Class

Support Reauthentication

Specifies whether the resource adapter implementation supports re-authentication of existing ManagedConnection instances.

The values are either True or False.

Data type

String

Transaction

Specifies the level of transaction support provided by the resource adapter.

The three possible values are:

- NoTransaction
- LocalTransaction
- XATransaction

Data type

String

Small Icon

The image is used as an icon to represent the module in a GUI.

Specifies a JPEG or GIF file containing a small image (16x16 pixels).

Data type Image

Large Icon

The image is used as an icon to represent the module in a GUI.

Specifies a JPEG or GIF file containing a small image (32x32 pixels).

Data type Image

Basic Password

The basic user password authentication mechanism that is specific to an EIS.

Credential Interface

Specifies the interface that the resource adapter implementation supports for the representation of the credentials. For Basic Password the credential value is *javax.resource.spi.security.PasswordCredential*.

Description

Any information that describes Basic Password selection.

Data type String

Kerberos V5

Specifies a Kerberos version 5 authentication mechanism.

Credential Interface

Specifies the interface that the resource adapter implementation supports for the representation of the credentials. For Kerberos version 5, the credential value is *javax.resource.spi.security.GenericCredential*.

Description

Any information that describes the Kerberos V5 selection.

Data type String

Property Name

Specifies the name of a configuration property.

The possible values are:

```
<config-property-name>ServerName</config-property-name>  
<config-property-name>PortNumber</config-property-name>  
<config-property-name>UserName</config-property-name>  
<config-property-name>Password</config-property-name>  
<config-property-name>ConnectionURL</config-property-name>
```

Data type String

Property Type

Contains the fully-qualified Java type of a configuration property as required by the *ManagedConnectionFactory* instance.

Data type String

Property Value

Contains the value of a configuration entry.

Data type String

Description

Describes the parent element.

Data type String

Permission Specification

Specifies a security permission that is required by the resource adapter code.

Data type String

Saving applications after assembly

Periodically save modules that you assemble with the Application Assembly Tool (AAT). Save any changes right before you close the module with which you are working.

Before you begin

This task assumes you have started the AAT and are working with a particular module.

Steps for this task

1. Save the archive file by clicking **File > Save As**.
 - If you are saving an existing archive file or application, click **File > Save**.
2. Name the new archive file or application whatever you like.

This step is optional if you are working with an existing archive file or application.

What to do next

Now that you have saved your assembled application, you can verify your archives and generate code for deployment.

Verifying archive files

Verify your archive files and correct any problems so that generation of deployment code is successful. During verification, the Application Assembly Tool (AAT) checks that an archive file is complete, and that deployment descriptor properties and references contain appropriate values.

This task assumes you have previously assembled and saved one or more modules.

Steps for this task

1. Start the Application Assembly Tool (AAT).
2. Click **File > Open** and select the module to verify.

3. Right-click the name of the module at the top of the navigation pane and click **Verify**.
4. Click **Verify** in the **Verify** window. The tool displays a scrolling window for viewing status messages as the verification proceeds.
5. Save the application.

Results

Archive files have been verified. The following list includes, but is not limited to, areas that the verification process has checked:

- Required deployment properties contain values.
- Values specified for environment entries match their associated Java types.
- In both Enterprise application (EAR) and Web application (WAR) files:
 - The target enterprise bean of the link exists for EJB references.
 - The target role exists for security role references.
 - Security roles are unique.
- Each module listed in the deployment descriptor exists in the archive for EAR files.
- Files for icons, servlets, error and welcome pages listed in the deployment descriptor have corresponding files in the archive for WAR files.
- For EJB modules:
 - All class files referenced in the deployment descriptor exist in the JAR file.
 - Method signatures for enterprise bean home, remote and implementation classes are compliant with the EJB 2.0 specification.

What to do next

If your application module contains EJB modules, generate code for deployment.

Otherwise, you are ready to deploy this application module (or standalone Web module) onto the application server.

Generating code for deployment

Before deploying applications on the server, if the application modules contain EJB modules, you must generate deployment code for the enterprise beans in the application. The Application Assembly Tool (AAT) provides this ability, or you can use the `ejbdeploy` command line tool.

Before you begin

This task assumes you have already assembled an EJB module, added it to an application module, saved the application module, and verified the application module.

Before installing your application in WebSphere Application Server, you must generate deployment code for the application. This step is required for EJB modules and for any Enterprise application (EAR) files that contain EJB modules. During code generation, the Application Assembly Tool invokes the `EJBDeploy` tool to prepare entity bean (JAR) files for deployment in run time environment. To deploy a J2EE application, you can install the application in the administrative console.

The following steps assume that you are using the Application Assembly Tool to generate code for deployment.

Steps for this task

1. Start the Application Assembly Tool (AAT).
2. Open the EAR or JAR file for which you want to generate code for deployment.
3. Click **File > Generate code for deployment** from the menu bar.
4. Specify the options for the server to use for generating code for the application deployment.

For more information on ejbdeploy options you can select through the AAT, see the Usage Scenario section.

Note: For Container managed persistence (CMP) entity beans, if the JAR file that you opened (inputJar file) contains a map and schema document, that schema is used. If the JAR file does not contain a map and schema document, the Application Assembly Tool uses a top-down mapping to generate files that contain mapping and database schema information.

5. Click **Generate Now**.

Review the messaging box for details of any error that might occur.

Note: Do not change the default output file name to be the same as the input filename, as the AAT cannot read and write to the same file name, and therefore, an error will occur.

Results

After deployment code is generated for an application, the deployable archive is renamed with the prefix Deployed_.

What to do next

Install the application on your server machine.

Note: Before deploying the application in your run time environment, you might need to set classpaths.

ejbdeploy tool

Abstract:

You can generate code for deployment by either using the Application Assembly Tool (AAT) or by using the Deployment Tool for Enterprise Java Beans (ejbdeploy) from a command prompt. For example, the options that you are able to set in AAT correspond with commands that the EJBDeploy tool uses to generate code for deploying an application.

View the Deployment Tool for Enterprise Java Beans documentation.

Syntax

For a detailed list of available options in the EJBDeploy tool, enter ejbdeploy from a command prompt.

ejbdeploy syntax -- relationship to Application Assembly Tool options

Abstract:

Many options for generating code for deployment in the Application Assembly Tool (AAT) directly correlate with the options in the command line EJBDeploy tool.

Application Assembly Tool options	EJBDeploy tool options
Deployed module location	outputJar
Working Directory	workingDirectory
Dependent classpath	cp
Code generation only	codegen
Verify archive (unchecked)	novalidate
RMIC options	rmic options
Database type	dbvendor
Database name	dbname
Schema name	dbschema

Chapter 17. Deploying and managing applications

After you develop an enterprise application and configure an application server, you can use the administrative console to install application files on the server and manage the activity of deployed applications.

Steps for this task

1. Install your application on your application server.
2. Start and stop applications.
3. Edit the administrative configuration for an application. Go to the settings page for an application, change the values for settings as needed, and click **OK**.
4. **(Optional)** Export applications.
5. **(Optional)** Export DDL files.
6. Update application binary files.
7. Uninstall applications.

After making changes to administrative configurations of your applications, ensure that you click **Save** on the administrative console taskbar to save the changes.

Enterprise applications

Enterprise applications (or J2EE applications) are applications that conform to the Java 2 Platform, Enterprise Edition, specification.

Enterprise applications can consist of the following:

- Zero or more EJB modules
- Zero or more Web modules
- Zero or more connector modules (packaged in RAR files)
- Zero or more application client modules
- Optionally, additional JAR files containing dependent classes or other components required by the application
- Any combination of the above

A J2EE application is represented by, and packaged in, an enterprise archive (EAR) file.

Installing a new application

To install an enterprise application to a WebSphere Application Server configuration, you can use the administrative console or the wsadmin tool. The steps below describe how to use the administrative console to install an application, EJB component, or Web module.

Note: Once you start performing the steps below, click **Cancel** to exit if you decide not to install the application. Do not simply move to another administrative console page without first clicking **Cancel** on an application installation page.

Steps for this task

1. Click **Applications > Install New Application** in the console navigation tree. The first of two Preparing for application install pages is shown.
2. On the first Preparing for application install page:
 - a. Specify the full path name of the source application file (.ear file otherwise known as an EAR file). The EAR file that you are installing can be either on the client machine (the machine that runs the Web browser) or on the server machine (the machine to which the client is connected). If you specify an EAR file on the client machine, then the administrative console uploads the EAR file to the machine on which the console is running and proceeds with application installation. You can also specify a standalone WAR or JAR file for installation.
 - b. If you are installing a standalone WAR file, specify the context root.
 - c. Click **Next**.
3. On the second Preparing for application install page:
 - a. Select whether to generate default bindings. Using the default bindings causes any incomplete bindings in the application to be filled in with default values. Existing bindings are not altered. You can customize default values used in generating default bindings. For example, you can specify JNDI prefix for all the EJB files in EJB modules, default data source and connection factory settings for EJB modules, virtual host for web modules, and so on. "Preparing for application install settings" describes available customizations and provides sample bindings.
 - b. Click **Next**. The Install New Application pages are now shown. If you chose to generate default bindings, you can proceed to the Summary step (step 23 below). ""Example: Installing an EAR file using the default bindings"" provides sample steps.
4. On the **Step: Provide options to perform the installation** panel, provide values for the following settings specific to WebSphere Application Server. Default values are used if you do not specify a value.
 - a. For **Pre-compile JSP**, specify whether to precompile JSP files as a part of installation. The default is not to precompile JSP files.
 - b. For **Directory to Install Application**, specify the directory to which the application EAR file will be installed. The default value is the value of `APP_INSTALL_ROOT/cell_name`, where the `APP_INSTALL_ROOT` variable is `install_root/installedApps`; for example, `C:\WebSphere\AppServer\installedApps\cell_name`.
You can specify an absolute path or use a pathmap variable such as `#{MY_APPS}`. You can use a pathmap variable in any installation. A WebSphere Application Server variable `#{CELL}` that denotes the current cell name can also be in the pathmap variable; for example, `#{MY_APP}/#{CELL}`.
 - c. For **Distribute Application**, specify whether WebSphere Application Server expands or deletes application binaries in the installation destination. The default is to enable application distribution. As a result, when you save changes in the console, application binaries for newly installed applications are expanded to the directory specified. The binaries are also deleted when you uninstall and save changes to the configuration. If you disable this option, then you must ensure that the application binaries are expanded appropriately in the destination directories of all nodes where the application is expected to run.
 - d. For **Use Binary Configuration**, specify whether the application server uses the binding, extensions, and deployment descriptors located with the

application deployment document, the deployment.xml file (default), or those located in the EAR file. The default is not to use the binary configuration.

- e. For **Deploy EJBs**, specify whether the EJBDeploy tool runs during application installation. The tool generates code needed to run EJB files. The default is not to run the EJBDeploy tool. You must enable this setting if the EAR file was assembled using the (Application Assembly (AAT) tool) and the EJBDeploy tool was not run during assembly, if the EAR file was not assembled using the AAT tool, or if the EAR file was assembled using versions of the AAT tool previous to Version 5. Note that enabling this setting might cause the installation program to run for several minutes.
- f. For **Application Name**, name the application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.
- g. For **Create MBeans for Resources**, specify whether to create MBeans for various resources (such as servlets or JSP files) within an application when the application is started. The default is to create MBean instances.
- h. For **Enable class reloading**, specify whether to enable class reloading when application files are updated. The default is not to enable class reloading.
- i. For **Reload Interval**, specify the number of seconds to scan the application's file system for updated files. The default is the value of the reload interval attribute in the IBM extension (META-INF/ibm-application-ext.xmi) file of the EAR file. This setting takes effect only if class reloading is enabled.

The reload interval specified here overrides the value specified in the IBM extensions for each Web module in the EAR file (which in turn overrides the reload interval specified in the IBM extensions for the application in the EAR file).

- 5. If your application uses EJB modules, on the **Step: Provide JNDI Names for Beans** panel, specify a JNDI name for each enterprise bean in every EJB module. You must specify a JNDI name for every enterprise bean defined in the application. For example, for the EJB module MyBean.jar, specify MyBean.
- 6. If your application uses EJB modules that contain Container Managed Persistence (CMP) beans that are based on the EJB 1.x specification, for **Step: Provide default datasource mapping for modules containing 1.x entity beans**, specify a JNDI name for the default data source for the EJB modules. The default data source for the EJB modules is optional if data sources are specified for individual CMP beans.
- 7. If your application has CMP beans that are based on the EJB 1.x specification, for **Step: Map datasources for all 1.x CMP**, specify a JNDI name for data sources to be used for each of the 1.x CMP beans. The data source attribute is optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, then a validation error displays after you click **Finish** (step 23) and the installation is cancelled.
- 8. If your application defines EJB references, for **Step: Map EJB references to beans**, specify JNDI names for enterprise beans that represent the logical names specified in EJB references. Each EJB reference defined in the application must be bound to an EJB file before clicking **Finish** on the Summary panel.

9. If your application defines resource references, for **Step: Map resource references to resources**, specify JNDI names for the resources that represent the logical names defined in resource references. Each resource reference defined in the application must be bound to a resource defined in your WebSphere Application Server configuration before clicking on **Finish** on the Summary panel.
10. If your application uses Web modules, for **Step: Map virtual hosts for web modules**, select a virtual host from the list that should map to a Web module defined in the application. The port number specified in the virtual host definition is used in the URL that is used to access artifacts such as servlets and JSP files in the Web module. Each Web module must have a virtual host to which it maps. Not specifying all needed virtual hosts will result in a validation error displaying after you click **Finish** on the Summary panel.
11. On the **Step: Map modules to application servers** panel, for every module select a target server or a cluster from the **Clusters and Servers** list. Place a checkmark in the check box beside **Module** to select all of the application modules or select individual modules.
12. If the application has security roles defined in its deployment descriptor then, for **Step: Map security roles to users/groups**, specify users and groups that are mapped to each of the security roles. Place a checkmark in the check box beside **Role** to select all of the roles or select individual roles. For each role, you can specify if predefined users such as **Everyone** or **All Authenticated users** are mapped to it. To select specific users or groups from the user registry:
 - a. Select a role and click **Lookup users** or **Lookup groups**.
 - b. On the Lookup users/groups panel shown, enter search criteria to extract a list of users or groups from the user registry.
 - c. Select individual users or groups from the results displayed.
 - d. Click **OK** to map the selected users or groups to the role selected on the **Step: Map security roles to users/groups** panel.
13. If the application has RunAs roles defined in its deployment descriptor, for **Step: Map RunAs roles to user**, specify the RunAs user name and password for every RunAs role. RunAs roles are used by enterprise beans that must run as a particular role while interacting with another enterprise bean. Place a checkmark in the check box beside **Role** to select all of the roles or select individual roles. After selecting a role, enter values for the user name, password, and verify password and click **Apply**.
14. If your application contains EJB 1.x CMP beans that do not have method permissions defined for some of the EJB methods, for **Step: Ensure all unprotected 1.x methods have the correct level of protection**, specify if you want to leave such methods unprotected or assign protection with deny all access.
15. If your application contains message driven enterprise beans, for **Step: Provide Listener Ports for messaging beans**, provide a listener port name for every message driven bean. If a name is not specified for each bean, then a validation error displays after you click on **Finish** on the Summary panel.
16. If your application uses EJB modules that contain CMP beans that are based on the EJB 2.0 specification, for **Step: Provide default datasource mapping for modules containing 2.0 entity beans**, specify a JNDI name for the default data source and the type of resource authorization to be used for the default data source for the EJB modules. The default data source for EJB modules is optional if data sources are specified for individual CMP beans.

17. If your application has CMP beans that are based on the EJB 2.0 specification, on the **Step: Map datasources for all 2.0 CMP** panel, for each of the 2.0 CMP beans specify a JNDI name and the type of resource authorization for data sources to be used. The data source attribute is optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, then a validation error is shown after you click **Finish** and installation is cancelled.
18. If your application contains EJB 2.0 CMP beans that do not have method permissions defined in the deployment descriptors for some of the EJB methods, on the **Step: Ensure all unprotected 2.0 methods have the correct level of protection** panel, specify whether you want to assign a specific role to the unprotected methods, add the methods to the exclude list, or mark them as unchecked. Methods added to the exclude list are marked as uncallable. For methods marked unchecked no authorization check is performed prior to their invocation.
19. If the **Deploy EJBs** setting is enabled on the **Provide options to perform the installation** panel, then you can specify options for the EJBDeploy tool on the **Step: Provide options to perform the EJB Deploy** panel. On this panel, you can specify extra classpath, rmic options, database types, and database schema names to be used while running the EJBDeploy tool. The tool is run on the EAR file during installation after you click **Finish**.
20. If your application contains resource environment references, for **Step: Mapping Resource Environment References to Resources**, specify JNDI names of resources that map to the logical names defined in resource environment references. If each resource environment reference does not have a resource associated with it, a validation error is shown after you click **Finish**.
21. If your application defines **Run-As Identity** as *System Identity*, for **Step: Replacing RunAs System to RunAs Roles**, you can optionally change it to *Run-As role* and specify a user name and password for the RunAs role specified. Selecting *System Identity* implies that the invocation is done using the WebSphere Application Server security server ID and should be used with caution as this ID has more privileges.
22. If your application has resource references that map to resources that have an Oracle database doing back-end processing, for **Step: Specify the isolation level for Oracle type provider**, specify or correct the isolation level to be used for such resources when used by the application. Oracle databases support ReadCommitted and Serializable isolation levels only.
23. On the Summary panel, verify the cell, node, and server onto which the application modules will install. Beside the **Cell/Node/Server** option, click **Click here** and verify the settings. Then click **Finish**.

Note: After clicking **Finish**, if you receive an OutOfMemory exception and the source application file does not install, your system might not have enough memory or your application might have too many modules in it to install successfully onto the server. If lack of system memory is not the cause of the exception, package your application again so the .ear file has fewer modules. If lack of system memory and the number of modules are not the cause of the exception, check the options you specified on the Java Virtual Machine page of the application server running the administrative console. Then, try installing the application file again.
24. Associate any shared libraries that the application needs to the application.
25. Click **Save** on the administrative console taskbar to save the changes to your configuration. The application is registered with the administrative

configuration and application files are copied to the target directory, which is *install_root/installedApps/cell_name* by default or the directory that you designate. For the single-server (base) installation, application files are copied to the destination directory when you click **Save**; for the Network Deployment installation, files are copied to remote nodes when the configuration on the deployment manager synchronizes with the configuration on individual nodes.

26. Test the application. For example, point a Web browser at the URL for the deployed application and examine the performance of the application. If necessary, update the application.

Preparing for application install settings

Use this page to install an application (EAR file) or module (JAR or WAR file).

To view this administrative console page, click **Applications > Install New Application**.

Follow the steps on this page to install an application or module. You must complete, at minimum, the first step; you must complete some or all of the later steps, depending on whether you are installing an application, EJB module, or Web module.

Path

Specifies the fully qualified path to the .ear, .jar, or .war file for the enterprise application.

Use **Local path** if the browser and application files are on the same machine (whether or not the server is on that machine, too).

Use **Server path** if the application file resides on any node in the current cell context. You can browse the entire file system of a node if the node agent or deployment manager is running on that selected node. Only .ear, .jar, or .war files are shown during the browsing.

During application installation, application files are typically uploaded from a client machine running the browser to the server machine running the administrative console, where they are deployed. In such cases, the Web browser running the administrative console is used to select EAR, WAR, or JAR modules to upload to the server machine.

In some cases, however, the application files reside on the file system of any of the nodes in a cell. To have the application server install these files, use the **Server path** option.

You can also use this option to specify an application file already residing on the machine running the application server. For example, the field value on Windows NT might be C:\WebSphere\AppServer\installableApps\test.ear. If you are installing a standalone WAR module, then you also must specify the context root.

Context Root

Specifies the context root of the Web application (WAR).

This field is used only to install a standalone WAR file. The context root is combined with the defined servlet mapping (from the WAR file) to compose the full URL that users type to access the servlet. For example, if the context root is

/gettingstarted and the servlet mapping is MySession, then the URL is `http://host:port/gettingstarted/MySession`.

Generate Default Bindings

Specifies whether to generate default bindings. If you place a checkmark in the check box, then any incomplete bindings in the application are filled in with default values. Existing bindings are not altered.

By choosing this option, you can directly jump to the Summary step and install the application if none of the steps have a red asterisk (*) next to them. A red asterisk denotes that the step has incomplete data and requires a valid value. On the Summary panel, verify the cell, node and server on which the application is installed.

Bindings are generated as follows:

- EJB JNDI names are generated of the form *prefix/ejb-name*. The default prefix is `ejb`, but can be overridden. The *ejb-name* is as specified in the deployment descriptors `<ejb-name>` tag.
- EJB references are bound as follows: If an `<ejb-link>` is found, it is honored. Otherwise, if a unique enterprise bean is found with a matching home (or local home) interface as the referenced bean, the reference is resolved automatically.
- Resource reference bindings are derived from the `<res-ref-name>` tag. Note that this action assumes that the `java:comp/env` name is the same as the resource global JNDI name.
- Connection factory bindings (for EJB 2.0 JAR files) are generated based on the JNDI name and authorization information provided. This action results in default connection factory settings for each EJB 2.0 JAR file in the application being installed. No bean-level connection factory bindings are generated.
- Data source bindings (for EJB 1.1 JAR files) are generated based on the JNDI name, data source user name password options. This results in default data source settings for each EJB JAR file. No bean-level data source bindings are generated.
- Message Driven Bean (MDB) listener ports are derived from the MDB `<ejb-name>` tag with the string `Port` appended.
- For `.war` files, the virtual host is set as `default_host` unless otherwise specified.

The default strategy suffices for most applications or at least for most bindings in most applications. However, it does not work if:

- You want to explicitly control the global JNDI names of one or more EJB files.
- You need tighter control of data source bindings for CMPs. That is, you have multiple data sources and need more than one global data source.
- You must map resource references to global resource JNDI names that are different from the `java:comp/env` name.

In such cases, you can alter the behavior with an XML document (a custom strategy). Use the **Specific bindings file** field to specify a custom strategy and see the field's help for examples.

Prefixes

Specifies prefixes to use for generated JNDI names.

Override

Specifies whether to override existing bindings.

If this check box is checked, the existing bindings are overridden by the generated ones.

EJB 1.1 CMP bindings

Specifies the default data source JNDI name.

If the **Default Bindings for EJB 1.1 CMPs** radio button is selected, specify the JNDI name for the default data source to be used with the CMP 1.1 beans. Also specify the user ID and password for this default data source.

Connection Factory Bindings

Specifies the default data source JNDI name.

If the **Default connection factory bindings** radio button is selected, specify the JNDI name for the default data source to be used with the bindings. Also specify the resource authorization.

Virtual Host

Specifies the virtual host for WAR modules.

Specific bindings file

Specifies a bindings file that overrides the default binding.

Alter the behavior of the default binding with an XML document (*aka* custom strategy). Custom strategies extend the default strategy so you only need to customize those areas where the default strategy is insufficient. That is, you only need to describe how you want to change the bindings generated by the default strategy; you do not have to define bindings for the entire application.

Brief examples of how to override various aspects of the default bindings generator follow:

Controlling an EJB JNDI name

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
<module-bindings>
<ejb-jar-binding>
<jar-name>helloEjb.jar</jar-name>
<!--this name must match the module name in the .ear file-->
  <ejb-bindings>
    <ejb-binding>
      <ejb-name>HelloEjb</ejb-name>
      <!--this must match the <ejb-name> entry in the EJB jar DD-->
      <jndi-name>com/acme/ejb/HelloHome</jndi-name>
    </ejb-binding>
  </ejb-bindings>
</ejb-jar-binding>
</module-bindings>
</dfltbndngs>
```

Setting the connection factory binding for an EJB JAR file

```
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
<module-bindings>
<ejb-jar-binding>
  <jar-name>yourEjb20.jar</jar-name>
  <connection-factory>
  <jndi-name>eis/jdbc/YourData_CMP</jndi-name>
  <res-auth>Container</res-auth>
```

```

        </connection-factory>
    </ejb-jar-binding>
</module-bindings>
</dfltbndngs>

```

Setting the connection factory binding for an EJB file

```

<?xml version="1.0">
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>yourEjb20.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>YourCmp20</ejb-name>
<!-- this matches the ejb-name tag in the DD -->
          <connection-factory>
            <jndi-name>eis/jdbc/YourData_CMP</jndi-name>
            <res-auth>PerConnFact</res-auth>
          </connection-factory>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</dfltbndngs>

```

Overriding a Resource Ref Binding from a WAR, EJB JAR file, or J2EE client JAR file

Example code for overriding a Resource Ref Binding from a WAR file follows. Use similar code to override a Resource Ref Binding from an EJB JAR file or a J2EE client JAR file.

```

<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <war-binding>
      <jar-name>hello.war</jar-name>
      <resource-ref-bindings>
        <resource-ref-binding>
<!-- the following must match the resource-ref in the DD -->
          <resource-ref-name>jdbc/MyDataSrc</resource-ref-name>
          <jndi-name>war/override/dataSource</jndi-name>
        </resource-ref-binding>
      </resource-ref-bindings>
    </war-binding>
  </module-bindings>
</dfltbndngs>

```

Overriding MDB JMS listener ports

```

<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>YourEjbJar.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>YourMDB</ejb-name>
          <listener-port>yourMdbListPort</listener-port>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</dfltbndngs>

```



```
</ejb-bindings>
</ejb-jar-binding>
</module-bindings>
</df1tbndngs>
```

Example: Installing an EAR file using the default bindings

An example of a simple .ear file installation using the default bindings follows:

1. Go to the Preparing for application install pages. Click **Applications > Install an Application** in the console navigation tree.
2. For **Path**, specify the full path name of the .ear file. For this example, the base file name is my_appl.ear and the file resides on a server at C:\sample_apps.
3. Now that a value is given for **Path**, on the first Preparing for application install page, click **Next**.
4. On the second Preparing for application install page, place a checkmark beside the **Generate Default Bindings** check box and click **Next**. Using the default bindings causes any incomplete bindings in the application to be filled in with default values. Existing bindings are not changed. By choosing this option, you can directly jump to the Summary step.
5. On the Install New Application page, click on **Summary**, the last step.
6. On the Summary panel, verify the cell, node, and server onto which the application files will install.
 - a. Beside the **Cell/Node/Server** option, click **Click here**.
 - b. On the **Map modules to application servers** panel, select the server onto which the application files will install from the **Clusters and Servers** list, place a checkmark in the check box beside **Module** to select all of the application modules, and click **Next**.

Because my_appl.ear does not require any additional settings to complete an installation, the Summary panel displays again.

7. On the Summary panel, click **Finish**.

Enterprise application collection

Use this page to view and manage enterprise applications.

To view this administrative console page, click **Applications > Enterprise Applications**.

Name

Specifies the name of the installed (or deployed) application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.

Status

Indicates whether the application deployed on the application server is started, stopped, or unavailable.

An application with an unavailable status might, in fact, be running but have an unavailable status because the server running the administrative console cannot communicate with the server running the application.

Enterprise application settings

Use this page to configure an enterprise application.

To view this administrative console page, click **Applications > Enterprise Applications > *application_name***.

Name

Specifies a logical name for the application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.

Data type	String
-----------	--------

Starting Weight

Specifies the order in which applications are started when the server starts. The application with the lowest starting weight is started first.

Data type	Integer
Default	1
Range	0 to 100

Application Binaries

Specifies the directory to which the application EAR file will be installed. The default value is the value of `APP_INSTALL_ROOT/cell_name`, where the `APP_INSTALL_ROOT` variable is `install_root/installedApps`; for example, `C:\WebSphere\AppServer\installedApps\cell_name`.

You can specify an absolute path or use a pathmap variable such as `${MY_APPS}`. You can use a pathmap variable in any installation. A WebSphere Application Server variable `${CELL}` that denotes the current cell name can also be in the pathmap variable; for example, `${MY_APP}/${CELL}`.

Data type	String
Units	Full path name

Use Metadata From Binaries

Specifies whether the application server uses the binding, extensions, and deployment descriptors located with the application deployment document, the `deployment.xml` file (default), or those located in the enterprise application resource (EAR) file.

Data type	Boolean
Default	true

Enable Distribution

Specifies whether WebSphere Application Server expands or deletes application binaries in the installation destination. The default is to enable application distribution. Application binaries for installed applications are expanded to the directory specified. The binaries are also deleted when you uninstall and save changes to the configuration. If you disable this option, then you must ensure that the application binaries are expanded appropriately in the destination directories of all nodes where the application runs.

Data type	Boolean
Default	true

ClassLoader Mode

Specifies whether the classloader searches in the parent classloader or in the application classloader first to load a class. The standard for JDK classloaders and WebSphere Application Server classloaders is PARENT_FIRST. By specifying PARENT_LAST, your application can override classes contained in the parent classloader, but this action can potentially result in ClassCastException or LinkageErrors if you have mixed use of overridden classes and non-overridden classes.

The options are PARENT_FIRST and PARENT_LAST. The default is to search in the parent classloader before searching in the application classloader to load a class.

Data type	String
Default	PARENT_FIRST

WAR Classloader Policy

Specifies whether to use a single classloader to load all WAR files of this application or to use a different classloader for each WAR file.

The options are APPLICATION and MODULE. The default is to use a separate classloader to load each WAR file.

Data type	String
Default	MODULE

Create MBeans for Resources

Specifies whether to create MBean files for various resources (such as servlets or JSP files) within an application.

Data type	Boolean
Default	true

Reload Enabled

Specifies whether to enable class reloading when application files are updated.

Data type	Boolean
Default	true

Reload Interval

Specifies the number of seconds to scan the application's file system for updated files. The default is the value of the reload interval attribute in the IBM extension (META-INF/ibm-application-ext.xmi) file of the EAR file. This setting takes effect only if class reloading is enabled.

The reload interval specified here overrides the value specified in the IBM extensions for each Web module in the EAR file (which in turn overrides the reload interval specified in the IBM extensions for the application in the EAR file).

Data type	Integer
Units	Seconds

Target mapping collection

Use this page to manage mappings of deployed applications or modules to servers or clusters.

To view this administrative console page, click **Applications > Enterprise Applications > *application_name* > Target Mappings**.

Target: States the name of the target server or cluster to which the application or module maps. You specify the target on the Map modules to application servers page accessed from the settings for an application.

Node: Specifies the node name if the target is a server.

Status: Indicates whether the status of the application running on the target server is started, stopped or unavailable.

Target mapping settings: Use this page to map a deployed application or module to a server or cluster.

To view this administrative console page, click **Applications > Enterprise Applications > *application_name* > Target Mappings > *target_name***.

Target: States the name of the target server or cluster to which the application or module maps. You specify the target on the Map modules to application servers page accessed from the settings for an application.

Data type	String
-----------	--------

Enabled: Indicates whether the application modules installed on the target server are started (or enabled) when the server starts. A true value indicates that the corresponding modules are enabled and thus are accessible when the server starts. A false value indicates that the corresponding modules are not enabled and thus are not accessible when the server starts.

Data type	Boolean
Default	true

Starting and stopping applications

You can start an application that is not running (has a status of *Stopped*) or stop an application that is running (has a status of *Started*).

Steps for this task

1. Go to the Enterprise Applications page. Click **Applications > Enterprise Applications** in the console navigation tree.
2. Check the check box for the application you want started or stopped.
3. Click a button:

Start Runs the application and changes the state of the application from *Stopped* to *Started*.

Stop Stops the processing of the application and changes the state of the application from *Started* to *Stopped*.

To restart a running application, place a checkmark in the check box for the application you want to restart, click **Stop** and then click **Start**.

Results

The status of the application changes and a message stating that the application started or stopped displays at the top the page.

Exporting applications

You can export an enterprise application to a location of your choice. Exporting applications enables you to back up your applications and preserve binding information for the applications. You might export your applications before updating installed applications or migrating to a later version of the WebSphere Application Server product.

Steps for this task

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Place a checkmark in the check box beside the application and click **Export**.
3. On the Export Application EAR Files page, click on the link to download the exported EAR file.
4. Use the browser dialogue to specify a location at which to save the exported EAR file and click **OK**.

The file containing binding information is exported to the specified node and directory, and has the name *enterprise_application_name.ear*.

Exporting DDL files

You can export the DDL files (Table.ddl) in the EJB modules of the application to a location of your choice.

Steps for this task

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Place a checkmark in the check box beside the application and click **Export DDL**. If the application has no DDL files in any of its EJB modules, then the message *No DDL files were found* displays at the top of the page. If the application has DDL files in its EJB modules, then a page displays listing DDL files in the format *appname.ear/_module.jar_Table.ddl*.
3. Click on a file in the list to download the file to your machine.

Updating applications

You can update an application deployed on a server. The steps below describe how to update a deployed application using the administrative console.

Steps for this task

1. Update the contents of the application and reassemble it, using the Application Assembly Tool. Typical tasks include adding or editing assembly properties, adding or importing modules into an application, and adding enterprise beans, Web components, and files.

2. Go to the Applications page of the administrative console. Click **Applications > Enterprise Applications** in the console navigation tree.
3. **(Optional)** Back up the application. Place a checkmark in the check box beside the application you want uninstalled and click **Export** to export the application to an EAR file and preserve the binding information.
4. With a checkmark beside the application, click **Update**.
The binding information of the updated (new) version of the application merges with the binding information from the installed (old) version. Then, the older version uninstalls from the configuration and the new version installs.
5. Complete the steps in the Preparing for application install page and the pages that follow it. See information on installing applications and on the settings page for application installation for guidance.
Note that the installation steps have the merged binding information from the new version and the old version. If the new version has bindings for application artifacts such as EJB JNDI names, EJB references or resource references, then those bindings will be part of the merged binding information. If new bindings are not present, then bindings are taken from the installed (old) version. If bindings are not present in the old version and if the default binding generation option is enabled, then the default bindings will be part of the merged binding information.
You can select whether to ignore bindings in the old version or ones in the new version.
6. **(Optional)** Map the installed application or module to servers. Use the Map modules to application servers page of the Install New Application pages displayed during updating the application. Or, after updating the application, use the Map modules to application servers page accessed from the Enterprise Applications page.
 - a. Go to the Map modules to application servers page. Click **Applications > Enterprise Applications** in the console navigation tree, click the application name, and then click **Map modules to application servers**.
 - b. Specify the application server where you want to install modules contained in your application and click **OK**.
7. Click **Save** on the administrative console taskbar to save the changes to your configuration. In the single server (base) product, after you click **Save** the old version of the application is uninstalled and the new version is installed into the configuration. The application binaries for the old version are deleted from the destination directory and the new binaries are copied to the directory.
8. Restart the application so the changes take effect.
 - a. Click **Applications > Enterprise Applications** in the console navigation tree to go to the Enterprise Applications page.
 - b. Check the check box for the updated application.
 - c. Click **Start**.

Note that you can also update applications using the wsadmin tool, which provides updating capabilities identical to that provided by clicking **Update** on the Enterprise Applications page. Further, you can update applications without stopping and restarting the server.

Hot deployment and dynamic reloading

You can make various changes to applications and their contents without having to stop the server and start it again. Making these types of changes is known as *hot deployment and dynamic reloading*.

Hot deployment is the process of adding new components (such as WAR files, EJB Jar files, enterprise Java beans, servlets, and JSP files) to a running server without having to stop the application server process and start it again.

Dynamic reloading is the ability to change an existing component without needing to restart the server in order for the change to take effect. Dynamic reloading involves:

- Changes to the implementation of a component of an application, such as changing the implementation of a servlet
- Changes to the settings of the application, such as changing the deployment descriptor for a Web module

Steps for this task

1. Locate your expanded application files. The application files are in the directory you specified when installing the application or, if you did not specify a custom target directory, are in the default target directory, *install_root/installedApps/cell_name*. Your EAR file, *install_root/installedApps/cell_name/application_name.ear*, points to the target directory. The *variables.xml* file for the node defines *install_root*.

It is important to locate the expanded application files because, as part of installing applications, a WebSphere application server unjars portions of the EAR file onto the file system of the computer that will run the application. These expanded files are what the server looks at when running your application.

If you cannot locate the expanded application files, look at the *binariesURL* attribute in the *deployment.xml* file for your application. The attribute designates the location the run time uses to find the application files.

For the remainder of this information on hot deployment and dynamic reloading, *application_root* represents the root directory of the expanded application files.

2. Locate application metadata files. The metadata files include the deployment descriptors (*web.xml*, *application.xml*, *ejb-jar.xml*, and the like), the bindings files (*ibm-web-bnd.xmi*, *ibm-app-bnd.xmi*, and the like), and the extensions files (*ibm-web-ext.xmi*, *ibm-app-ext.xmi*, and the like).

Metadata XML files for an application can be loaded from one of two locations. The metadata files can be loaded from the same location as the application binary files (such as *application_root/META-INF*) or they can be loaded from the WebSphere configuration tree, *install_root/cells/cell_name/applications/application_EAR_name/deployments/application_name/*. The value of the *useMetadataFromBinary* flag specified during application installation controls which location is used. If specified, the metadata files are loaded from the same location as the application binary files. If not specified, the metadata files are loaded from the application deployment folder in the configuration tree.

For the remainder of this information, *metadata_root* represents the location of the metadata files for the specified application or module.

3. Change or add the following components or modules as needed:
 - Application files
 - WAR files
 - EJB Jar files
 - HTTP plug-in configuration files
4. For changes to take effect, you might need to start, stop, or restart an application. "Starting and stopping applications" provides information on using

the administrative console to start, stop, or restart an application. "Example: Starting an application using wsadmin" (not in this document) and "Example: Stopping running applications on a server using wsadmin" (not in this document) provide information on using the wsadmin scripting tool.

Changing or adding application files

You can change or add application files on application servers without having to stop the server and start it again. This file describes—

- Updating an existing application on a running server (providing a new EAR file)
- Adding a new application to a running server
- Removing an existing application from a running server
- Adding a new EJB or Web module to an existing, running application
- Changing the application.xml file for an application
- Changing the ibm-app-ext.xmi file for an application
- Changing the ibm-app-bnd.xmi file for an application
- Changing a non-module Jar file contained in the EAR file

Updating an existing application on a running server (providing a new EAR file)

Reinstall an updated application using the administrative console or the wsadmin \$AdminApp install command with the -update option.

Both reinstallation methods enable you to update an existing application using any of the other steps listed in this file, including changing classes, adding modules, removing modules, changing modules, or changing metadata files. The application reinstallation methods detect the changes in your application and prompt you for additional binding data that might be needed to install the application. The reinstallation process automatically stops and restarts your application on the appropriate servers.

Hot deployment:	Yes
Dynamic reloading:	Yes

Adding a new application to a running server

Install an application using the administrative console or the wsadmin install command.

Hot deployment:	Yes
Dynamic reloading:	No

Removing an existing application from a running server

Stop the application and then uninstall it from the server. Use the administrative console to stop the application and then uninstall it. Or run the wasadmin stopApplication command and then the uninstall command.

Hot deployment:	Yes
Dynamic reloading:	No

Adding a new EJB or Web module to an existing, running application

1. Update the application files in the *application_root* location.

- Restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Yes
Dynamic reloading:	No

Changing the application.xml file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing the ibm-app-ext.xmi file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing the ibm-app-bnd.xmi file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing a non-module Jar file contained in the EAR file

- Update the non-module Jar file in the *application_root* location.
- If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.
If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment:	Yes
Dynamic reloading:	Yes

Changing or adding WAR files

You can change WAR files on application servers without having to stop the server and start it again. This file describes—

- Changing an existing JSP file
- Adding a new JSP file to an existing application
- Changing an existing servlet class (editing and recompiling)
- Changing a dependent class of an existing servlet class

- Adding a new servlet using the Invoker (Serve Servlets by class name) facility or adding a dependent class to an existing application
- Adding a new servlet, including a new definition of the servlet in the web.xml deployment descriptor for the application
- Changing the web.xml file of a WAR file
- Changing the ibm-web-ext.xmi file of a WAR file
- Changing the ibm-web-bnd.xmi file of a WAR file

Changing an existing JSP file

Place the changed JSP file directly in the *application_root/module_name* directory or the appropriate subdirectory. The change will be automatically detected and the JSP will be recompiled and reloaded.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Adding a new JSP file to an existing application

Place the new JSP file directly in the *application_root/module_name* directory or the appropriate subdirectory. The new file will be automatically detected and compiled on the first request to the page.

Hot deployment:	Yes
Dynamic reloading:	Yes

Changing an existing servlet class (editing and recompiling)

1. Place the new version of the servlet .class file directly in the *application_root/module_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application_root/module_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.
If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing a dependent class of an existing servlet class

1. Place the new version of the dependent .class file directly in the *application_root/module_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application_root/module_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Adding a new servlet using the Invoker (Serve Servlets by class name) facility or adding a dependent class to an existing application

1. Place the new .class file directly in the *application_root/module_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application_root/module_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.

This case is treated the same as changing an existing class. The difference is that adding the servlet or class does not immediately cause the Web application to reload because the class has never been loaded before. The class simply becomes available for execution.

2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment:	Yes
Dynamic reloading:	Not applicable

Adding a new servlet, including a new definition of the servlet in the web.xml deployment descriptor for the application

1. Place the new .class file directly in the *application_root/module_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application_root/module_name/WEB-INF/lib*.

You can edit the web.xml file in place or copy it into the *application_root/module_name/WEB-INF/classes* directory. The new .class file will not trigger a reloading of the application.

2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands. After the application restarts, the new servlet is available for service.

Hot deployment:	Yes
Dynamic reloading:	Not applicable

Changing the web.xml file of a WAR file

1. Edit the web.xml file in place or copy it into the *metadata_root/module_name/WEB-INF* directory.
2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment:	Yes
Dynamic reloading:	Yes

Changing the ibm-web-ext.xmi file of a WAR file

Edit the extension settings as needed. You can change all of the extension settings. The only warning is if you set the `reloadInterval` property to zero (0) or the `reloadEnabled` property to false, the application will no longer automatically detect changes to class files. Both of these changes disable the automatic reloading function. The only way to re-enable automatic reloading is to change the appropriate property and restart the application. See other task descriptions in this file for information on restarting an application.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing the `ibm-web-bnd.xml` file of a WAR file

1. Edit the bindings as needed. You can change all of the values but ensure that the entities you are binding to are present in the configuration of the server.
2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing or adding EJB Jar files

You can change EJB Jar files on application servers without having to stop the server and start it again. This file describes—

- Changing the `ejb-jar.xml` file of an EJB Jar file
- Changing the `ibm-ejb-jar-ext.xml` or `ibm-ejb-jar-bnd.xml` file of an EJB Jar file
- Changing the `Table.ddl` file for an EJB Jar file
- Changing the `Map.map.xml` or `Schema.db.xml` file for an EJB Jar file
- Updating the implementation class for an EJB file or a dependent class of the implementation class for an EJB file
- Updating the Home/Remote interface class for an EJB file
- Adding a new EJB file to an existing EJB Jar file

Changing the `ejb-jar.xml` file of an EJB Jar file

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing the `ibm-ejb-jar-ext.xml` or `ibm-ejb-jar-bnd.xml` file of an EJB Jar file

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Changing the `Table.ddl` file for an EJB Jar file

Rerun the DDL file on the user database server. Changing the Table.ddl file has no effect on the application server and is a change to the database table schema for the EJB files.

Hot deployment:	Not applicable
Dynamic reloading:	Not applicable

Changing the Map.mapxmi or Schema.dbxmi file for an EJB Jar file

1. Change the Map.mapxmi or Schema.dbxmi file for an EJB Jar file.
2. Regenerate the deployed code artifacts for the EJB file.
3. Apply the new EJB Jar file to the server.
4. Restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Updating the implementation class for an EJB file or a dependent class of the implementation class for an EJB file

1. Update the class file in the *application_root/module_name.jar* file.
2. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

If automatic reloading is not enabled, restart the application of which the EJB file is a member. If the updated module is used by other modules in other applications, restart those applications as well. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Updating the Home/Remote interface class for an EJB file

1. Update the interface class of the EJB file.
2. Regenerate the deployed code artifacts for the EJB file.
3. Apply the new EJB Jar file to the server.
4. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

If automatic reloading is not enabled, restart the application of which the EJB file is a member. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:	Not applicable
Dynamic reloading:	Yes

Adding a new EJB file to an existing EJB Jar file

1. Apply the new or updated Jar file to the *application_root* location.
2. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment:	Yes
Dynamic reloading:	Yes

Changing the HTTP plug-in configuration

You can change the HTTP plug-in configuration without having to stop the server and start it again. This file describes—

- Changing the `application.xml` file to change the context root of a WAR file
- Changing the `web.xml` file to add, remove, or modify a servlet mapping
- Changing the `server.xml` file to add, remove, or modify an HTTP transport or changing the `virtualhost.xml` file to add or remove a virtual host or to add, remove, or modify a virtual host alias

Changing the `application.xml` file to change the context root of a WAR file

1. Change the `application.xml` file.
2. Regenerate the plug-in configuration file using the administrative console or by running the `GenPluginCfg.bat/sh` script.

Hot deployment:	Yes
Dynamic reloading:	Yes

Changing the `web.xml` file to add, remove, or modify a servlet mapping

1. Change the `web.xml` file.
2. Regenerate the plug-in configuration file using the administrative console or by running the `GenPluginCfg.bat/sh` script.

If the Web application has file serving enabled or has a servlet mapping of `/`, you do not have to regenerate the plug-in configuration. In all other cases the regeneration is required.

Hot deployment:	Yes
Dynamic reloading:	Yes

Changing the `server.xml` file to add, remove, or modify an HTTP transport or changing the `virtualhost.xml` file to add or remove a virtual host or to add, remove, or modify a virtual host alias

1. Change the `server.xml` file to add, remove, or modify an HTTP transport or change the `virtualhost.xml` file to add or remove a virtual host or to add, remove, or modify a virtual host alias.
2. Regenerate the plug-in configuration file using the administrative console, by running the `GenPluginCfg.bat/sh` script, or by running a `wasadmin` command.

Hot deployment:	Yes
Dynamic reloading:	Yes

Uninstalling applications

After an application no longer is needed, you can uninstall it. Uninstalling an application deletes the application from the WebSphere Application Server configuration repository and it deletes the application binaries from the file system of all nodes where the application modules are installed.

Steps for this task

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. **(Optional)** Stop the application. Place a checkmark in the check box beside the application you want uninstalled and click **Stop**. Note that it is recommended you stop the application before uninstalling. However, even if the application is running when uninstallation commences, the uninstallation program stops the application before deleting application binaries from the file system.
3. **(Optional)** Back up the application. Place a checkmark in the check box beside the application you want uninstalled and click **Export** to export the application to an EAR file and preserve the binding information.
4. With a checkmark in the check box beside the application you want uninstalled, click **Uninstall**.
5. Confirm the uninstallation operation.
6. Click **Save** on the console taskbar to save changes made to the administrative configuration.




In the single-server (base) product, application binaries are deleted after you click **Save**.

Deploying and managing applications: Resources for learning

Use the following links to find relevant supplemental information about deploying and managing applications using the administrative console. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Administration
-  **The J2EETM Tutorial: The Duke's Bank Application**
http://java.sun.com/j2ee/tutorial/1_3-fcs/doc/Ebank.html
-  **Best Practices in WebSphere Application: Separating the developers from the administrators**
<http://www.sys-con.com/websphere/articleprint.cfm?id=26>
-  **Designing Enterprise Applications with the Java™ 2 Platform, Enterprise Edition, Second Edition**

[http://java.sun.com/blueprints/guidelines/
designing_enterprise_applications_2e/](http://java.sun.com/blueprints/guidelines/designing_enterprise_applications_2e/)

-  **Designing Enterprise Applications, Second Edition**
<http://developer.java.sun.com/developer/Books/j2ee/designingenterprise/>
-  **Building Java™ Enterprise Applications Volume I: Architecture**
<http://developer.java.sun.com/developer/Books/j2ee/bjeapps/>
- ✓ **WebSphere Application Server education**
<http://www.ibm.com/software/webservers/learn/>
- ✓ **Developing and Testing a Complete 'Hello World' J2EE Application with IBM WebSphere Studio Application Developer for Linux**
[http://www7b.software.ibm.com/wsdd/library/tutorials/
0206_wosnick/wosnick_reg.html?open&l=937,t=gr](http://www7b.software.ibm.com/wsdd/library/tutorials/0206_wosnick/wosnick_reg.html?open&l=937,t=gr)
- ✓ **IBM WebSphere Application Server: The Complete Reference**
<http://www.mcgraw-hill.co.uk/html/0072223944.html>
- ✓ **Writing Enterprise Applications with Java™ 2 Platform, Enterprise Edition**
<http://developer.java.sun.com/developer/onlineTraining/J2EE/Intro/>
- ✓ **WebSphere Application Server Bible**
<http://www.wiley.com/cda/product/0,,0764548964%7Ctoc%7C2948,00.html>
- ✓ **Listing of all IBM WebSphere Application Server Redbooks**
<http://publib-b.boulder.ibm.com/Redbooks.nsf/Portals/WebSphere>
- ✓ **IBM WebSphere V4.0 Advanced Edition Handbook**
<http://www.redbooks.ibm.com/abstracts/sg246176.html>
- ✓ **Redbook on WebSphere 4.0 Installation and Configuration on the IBM iSeries Server**
[http://publib-b.boulder.ibm.com/Redbooks.nsf/
9445fa5b416f6e32852569ae006bb65f/
7b1a07251256f08b85256b750067aee1?OpenDocument](http://publib-b.boulder.ibm.com/Redbooks.nsf/9445fa5b416f6e32852569ae006bb65f/7b1a07251256f08b85256b750067aee1?OpenDocument)
- ✓ **Redbook on WebSphere Application Server V4 for Linux, Implementation and Deployment Guide**
[http://publib-
b.boulder.ibm.com/Redbooks.nsf/RedpaperAbstracts/REDP0405.html](http://publib-b.boulder.ibm.com/Redbooks.nsf/RedpaperAbstracts/REDP0405.html)