Borland AppServer™ 6.7

Developer's Guide

Borland Software Corporation
20450 Stevens Creek Blvd., Suite 800
Cupertino, CA 95014 USA
www.borland.com

Refer to the file deploy.html for a complete list of files that you can distribute in accordance with the License Statement and Limited Warranty.

Borland Software Corporation may have patents and/or pending patent applications covering subject matter in this document. Please refer to the product CD or the About dialog box for the list of applicable patents. The furnishing of this document does not give you any license to these patents.

Copyright 1999–2006 Borland Software Corporation. All rights reserved. All Borland brand and product names are trademarks or registered trademarks of Borland Software Corporation in the United States and other countries. All other marks are the property of their respective owners.

Microsoft, the .NET logo, and Visual Studio are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

For third-party conditions and disclaimers, see the Release Notes on your product CD.

BAS67DevGuide
December 2006

Borland®

Contents

Chapter 1
Introduction to Borland AppServer 1
AppServer features . . . 2
Borland AppServer Documentation . . . 2
Accessing AppServer online help topics . . . 3
Accessing AppServer online help topics from within a AppServer GUI tool . . . 3
Documentation conventions . . . 3
Platform conventions . . . 3
Contacting Borland support . . . 4
Online resources . . . 4
World Wide Web . . . 4
Borland newsgroups . . . 4

Chapter 2
Borland AppServer overview and architecture 5
AppServer architecture overview . . . 6
AppServer services overview . . . 6
Web Server . . . 7
JMS . . . 7
Smart Agent . . . 7
2PC Transaction Service . . . 7
The Partition and its services . . . 8
Connector Service . . . 8
EJB Container . . . 8
JDataStore Server . . . 8
Lifecycle Interceptor Manager . . . 8
Naming Service . . . 9
Session Storage Service . . . 9
Transaction Manager . . . 9
Web Container . . . 9
Borland AppServer and J2EE APIs . . . 10
JDBC . . . 10
Java Mail . . . 10
JTA . . . 10
JAXP . . . 10
JNDI . . . 11
RMI-IIOP . . . 11
Other Technologies . . . 11
Optimizeit Profiler and Optimizeit ServerTrace . . . 11

Chapter 3
Using Partitions 13
Creating, cloning, and deleting Partitions . . . 13
Creating a new Partition . . . 13
Cloning an existing Partition . . . 15
Deleting a Partition . . . 15
Deploying modules and libraries to a Partition . . . 15
Advanced options for deployment . . . 16
Hosting additional modules . . . 17
Configuring Partitions . . . 17
General properties . . . 17
Partition Settings properties . . . 19
Statistics properties . . . 19
JMX Agent properties . . . 19
JDK properties . . . 20
Advanced JDK options . . . 20
Performance Tuning Hints . . . 20
VisiBroker properties . . . 20
Advanced VisiBroker options . . . 20
Security properties . . . 21
Log Settings properties . . . 21
Time Rules properties . . . 21
Advanced properties . . . 21
Managed Object Settings . . . 22
Management Action Settings . . . 22
Viewing Partition information . . . 22
General tab . . . 22
General properties . . . 23
Partition properties . . . 23
Security properties . . . 23
Web Container Root Context . . . 23
Properties tab . . . 23
Virtual Machine . . . 23
Server Connection Manager Settings . . . 23
XML tab . . . 23
Class Loading tab . . . 24
Logs tab . . . 24
Status tab . . . 24
JDBC Pool States tab . . . 24
Tuning the Partition for performance . . . 24
Advanced performance tuning options . . . 26
Dumping the Partition stack trace to a log file . . . 26
Enabling a Partition to run with Optimizeit Profiler or ServerTrace . . . 26

Chapter 4
Web components 27
Apache web server implementation . . . 27
Apache configuration . . . 27
Apache configuration syntax . . . 28
Running Apache web server on a privileged port 28
Using the .htaccess files . . . 29
Apache directory structure . . . 29
Borland web container implementation . . . 29
Servlets and JavaServer Pages . . . 30
Typical web application development process . . . 31
Web application archive (WAR) file . . . 31
Borland-specific DTD . . . 31
Adding ENV variables for the web container . . . 31
Microsoft Internet Information Services (IIS) web server 32
IIS/IIOP redirector directory structure . . . 32
Smart Agent implementation . . . 32
Connecting an Apache web server to a Borland web container . . . 33
Connecting Borland web containers to Java Session Service . . . 33

Chapter 5
Web server to web container connectivity 35
Apache web server to Borland web container connectivity . . . 35
Modifying the Borland web container IIOP configuration . . . 35
Modifying the IIOP configuration in Apache . . . 37
Additional Apache IIOP directives . . . 39
Apache IIOP connector configuration . . . 39
Adding new clusters . . . 39
Adding new web applications . . . 41
Large data transfer . . . 41
Downloading large data . . . 42
Implementing chunked download . . . 42
Enabling chunked download . . . 42
Known content length versus unknown . . . 42
Chunked download with known content length 42
Chunked download with unknown content length 43
Browsers supporting only the HTTP 1.0 protocol 43
Implementing non-chunked download . . . 43
Uploading large data . . . 44
Implementing chunked upload . . . 44
Enabling chunked upload . . . 44
Changing the upload buffer size . . . 44
Known content length versus unknown . . . 45
Chunked upload with known content length . . 45
Chunked upload with unknown content length . 45
Implementing non-chunked upload . . . 45
IIS web server to Borland web container connectivity 46
Modifying the IIOP configuration in the Borland web container . . . 46
Microsoft Internet Information Services (IIS) server-specific IIOP configuration . . . 46
How to Configure your Windows 2003/XP/2000 system on which IIS is running . . . 46
IIS/IIOP redirector configuration . . . 48
Adding new clusters . . . 49
Adding new web applications . . . 50

Chapter 6
Java Session Service (JSS) configuration 51
Session management with JSS . . . 51
Managing and configuring the JSS . . . 54
Configuring the JSS Partition service . . . 54

Chapter 7
Clustering web components 55
Stateless and stateful connection services . . . 55
The Borland IIOP connector . . . 55
Load balancing support . . . 56
OSAgent based load balancing . . . 56
Corbaloc based load balancing . . . 56
Fault tolerance (failover) . . . 57
Smart session handling . . . 57
Setting up your web container with JSS . . . 58
Modifying a Borland web container for failover . . . 58
Session storage implementation . . . 58
Programmatic implementation . . . 58
Automatic implementation . . . 58
Using HTTP sessions . . . 59

Chapter 8
Apache web server to CORBA server connectivity 61
Web-enabling your CORBA server . . . 61
Determining the urls for your CORBA methods . . 61
Implementing the ReqProcessor IDL in your CORBA server . . . 62
The process() method . . . 63
Configuring your Apache web server to invoke a CORBA server . . . 63
Apache IIOP configuration . . . 64
Adding new CORBA servers (clusters) . . . 64
Mapping URIs to defined clusters . . . 66

Chapter 9
Borland AppServer Web Services 67
Web Services Overview . . . 67
Web Services Architecture . . . 67
Web Services and Partitions . . . 68
Web Service providers . . . 69
Specifying web service information in a deploy.wsdd file . . . 69
Java:RPC provider . . . 70
Java:EJB provider . . . 70
How Borland Web Services work . . . 71
Web Service Deployment Descriptors . . . 71
Creating a server-config.wsdd file . . . 71
Viewing and Editing WSDD Properties . . . 72
Packaging Web Service Application Archives . . . 72
Borland Web Services examples . . . 72
Using the Web Service provider examples . . . 72
Steps to build, deploy, and run the examples . 73
Apache Axis Web Service samples . . . 73
Tools Overview . . . 73
Apache ANT tool . . . 73
Java2WSDL tool . . . 74
WSDL2Java tool . . . 74
Axis Admin tool . . . 74

Chapter 10
Writing enterprise bean clients 75
Client view of an enterprise bean . . . 75
Initializing the client . . . 75
Locating the home interface . . . 75
Obtaining the remote interface . . . 76
Session beans . . . 76
Entity beans . . . 77
Find methods and primary key class . . . 77
Create and remove methods . . . 78
Invoking methods . . . 78
Removing bean instances . . . 79
Using a bean's handle . . . 79

Managing transactions 80
Getting information about an enterprise bean 81
Support for JNDI 81
EJB to CORBA mapping 81
Mapping for distribution 82
Mapping for naming 83
Mapping for transaction 83
Mapping for security 84

Chapter 11
The VisiClient Container 87
Application Client architecture 87
Packaging and deployment 88
Benefits of the VisiClient Container 88
Document Type Definitions (DTDs) 89
Example XML using the DTD 90
Support of references and links 91
Using the VisiClient Container 92
VisiClient Container usage example 92
Running a J2EE client application on machines not running AppServer 92
Embedding VisiClient Container functionality into an existing application 93
Use of Manifest files 94
Example of a Manifest file 94
Exception handling 95
Using resource-reference factory types 95
Other features 95
Using the Client Verify tool 95

Chapter 12
Caching of Stateful Session Beans 97
Passivating Session Beans 97
Simple Passivation 97
Aggressive Passivation 98
Sessions in secondary storage 99
Setting the keep alive timeout in Containers 99
Setting the keep alive timeout for a particular session bean 99

Chapter 13
Entity Beans and CMP 1.1 in Borland AppServer 101
Entity Beans 101
Container-managed persistence and Relationships 101
Implementing an entity bean 102
Packaging Requirements 102
Entity Bean Primary Keys 103
Generating primary keys from a user class . . 103
Generating primary keys from a custom class 103
Support for composite keys 103
Reentrancy 104
Container-Managed Persistence in AppServer . . . 104
AppServer CMP engine's CMP 1.1 implementation . 104
Providing CMP metadata to the Container . . 105
Constructing finder methods 105
Constructing the where clause 106
Parameter substitution 106
Compound parameters 106
Entity beans as parameters 107
Specifying relationships between entities . . . 107
Container-managed field names 109
Setting Properties 109
Using the Deployment Descriptor Editor 109
J2EE 1.2 Entity Bean using BMP or CMP 1.1 . 110
Container-managed data access support 110
Using SQL keywords 111
Using null values 111
Establishing a database connection 111
Container-created tables 112
Mapping Java types to SQL types 113
Automatic table mapping 113

Chapter 14
Entity Beans and Table Mapping for CMP 2.x 117
Entity Beans 117
Container-managed persistence and Relationships . 117
Packaging Requirements 118
A note on reentrancy 119
Container-Managed Persistence in AppServer . . . 119
About the Persistence Manager 119
Borland CMP engine's CMP 2.x implementation . 120
Optimistic Concurrency Behavior 120
Pessimistic Behavior 120
Optimistic Concurrency 121
Persistence Schema 122
Specifying tables and datasources 122
Basic Mapping of CMP fields to columns . . . 123
Mapping one field to multiple columns 123
Mapping CMP fields to multiple tables 124
Specifying relationships between tables . . . 125
Using cascade delete and database cascade delete . 128
Database cascade delete support 129

Chapter 15
Using Borland AppServer Properties for CMP 2.x 131
Setting Properties 131
Using the Deployment Descriptor Editor 131
The EJB Designer 132
J2EE 1.3 and 1.4 Entity Bean 132
Setting CMP 2.x Properties 133
Editing Entity properties 133
Editing Table and Column properties 134
Entity Properties 135
Table Properties 136
Column Properties 137
Security Properties 138

Chapter 16
EJB-QL and Data Access Support 139
Selecting a CMP Field or Collection of CMP Fields . 139
Selecting a ResultSet 140
Aggregate Functions in EJB-QL 140
Data Type Returns for Aggregate Functions . . . 140

Support for ORDER BY 141
Support for GROUP BY 143
Sub-Queries 143
Dynamic Queries 144
Overriding SQL generated from EJB-QL by the CMP engine 145
Container-managed data access support 146
Support for Oracle Large Objects (LOBs) 147
Container-created tables 148

Chapter 17
Generating Entity Bean Primary Keys 149
Generating primary keys from a user class. 150
Generating primary keys from a custom class . . . 150
Implementing primary key generation by the CMP engine 150
Oracle Sequences: using getPrimaryKeyBeforeInsertSql. 150
SQL Server: using getPrimaryKeyAfterInsertSql and ignoreOnInsert 150
JDataStore JDBC3: using useGetGeneratedKeys . 151
Automatic primary key generation using named sequence tables 151
Key cache size 152

Chapter 18
Transaction management 153
Understanding transactions 153
Characteristics of transactions 153
Transaction support 154
Transaction manager services 154
Distributed transactions and two-phase commit. 155
When to use two-phase commit transactions . . 155
Using multiple JDBC connections for access to multiple database resources from a single vendor in the same transaction 155
Using multiple JDBC connections to the same database resource in the same transaction 156
Using multiple disparate resources in a single transaction 156
EJBs and 2PC transactions 156
Example runtime scenarios 157
Declarative transaction management in Enterprise JavaBeans 159
Understanding bean-managed and container-managed transactions 160
Local and Global transactions 160
Transaction attributes 161
Programmatic transaction management using JTA APIs 162
JDBC API Modifications 162
Modifications to the behavior of the JDBC API . 163
Overridden JDBC methods 163
Java.sql.Connection.commit() 163
Java.sql.Connection.rollback() 163
Java.sql.Connection.close() 164
Java.sql.Connection.setAutoCommit(boolean) 164
Handling of EJB exceptions 164
System-level exceptions 164
Application-level exceptions 164
Handling application exceptions 165
Transaction rollback 165
Options for continuing a transaction 165

Chapter 19
Message-Driven Beans and JMS 169
JMS and EJB 169
EJB 2.0 Message-Driven Bean (MDB) 170
EJB 2.1 MDB 170
Client View of an MDB 170
MDB Configuration 170
Connecting to a JMS Server from EJB 2.0 MDBs 171
Connecting to message source from EJB 2.1 MDBs 171
Changes to ejb-jar.xml 172
Changes to ejb-borland.xml. 173
Clustering of MDBs 175
Error Recovery 176
Rebinding EJB 2.0 and EJB 2.1 MDBs configured with a JMS provider message source 176
Redelivered messages for EJB 2.0 and EJB 2.1 MDBs configured with a JMS provider message source 176
MDBs and Transactions 177

Chapter 20
Registering a Java Bean type object into JNDI 179
Adding a Java Bean Type Object into JNDI 179

Chapter 21
Connecting to Resources with Borland AppServer: using the Definitions Archive (DAR) 181
JNDI Definitions Module. 182
Migrating to DARs from previous versions of Borland AppServer 183
Creating and Deploying a DAR 183
Disabling and Enabling a Deployed DAR 184
Packaging DAR Modules in an Application EAR . . 184

Chapter 22
Using JDBC 185
Configuring JDBC Datasources 186
Deploying Driver Libraries 188
Defining the Connection Pool Properties for a JDBC Datasource 189
Getting debug output 193
Descriptions of AppServer's Pooled Connection States 194
Support for older JDBC 1.x drivers 194
Advanced Topics for Defining JDBC Datasources . 195
Connecting to JDBC Resources from J2EE Application Components 197

Chapter 23
Using JMS 199
JMS 1.1 Common APIs . . . 200
Configuring JMS Connection Factories and Destinations . . . 201
Defining Connection Pool Properties for JMS Connection Factories . . . 202
Defining Individual JMS Connection Factory Properties 204
Obtaining JMS Connection Factories and Destinations in J2EE Application Components . . . 204
J2EE 1.2 and J2EE 1.3 . . . 205
J2EE 1.4 . . . 207
JMS and Transactions . . . 209
Enabling the JMS services security . . . 212
Advanced Concepts for Configuring JMS Connection Factories and Destinations . . . 212

Chapter 24
Using Hibernate 213
Overview . . . 213
Creating a Hibernate Application . . . 214
Creating a Connection Data Source . . . 214
Working with Configuration Files and Session Factories . . . 214
Working with the Hibernate Configuration File 214
Configuring the Data Source . . . 214
Configuring Transactions . . . 215
Setting the Property for Automatic Session Flushing . . . 215
Setting the Property for Automatic Session Closing . . . 215
Setting the Property for Collecting Statistics Information . . . 215
Setting Entity Mapping . . . 215
Setting the Dialect . . . 216
Setting the Property to Drop and Recreate the Database Schema . . . 216
Setting the property to Log Executed SQL . . 216
Packaging . . . 217
EJB Module: EJB-JAR . . . 217
WEB Module: WAR . . . 218
Packaging the Mapping XMLs . . . 218
Packaging Multiple Session Factory configuration Files . . . 219
Working with EAR Modules . . . 219
Enabling and Disabling the Hibernate Service . . . 219
Working with Hibernate Statistics . . . 220

Chapter 25
JMS provider pluggability 221
Runtime pluggability . . . 221
Configuring JMS administered objects (connection factories, queues and topics) . . . 222
Setting Admin Objects Using Borland Deployment Descriptor . . . 222
Service Management for JMS Providers . . . 222
Tibco EMS 4.2 . . . 223
Added value for Tibco . . . 223
Configuring Admin Objects for Tibco . . . 223
Auto Queue Creation Feature in Tibco . . . 223
Tibco Admin Console . . . 223
Configuring clients for fault tolerant Tibco connections . . . 224
Enabling Security for Tibco . . . 225
Disabling security for Tibco . . . 225
OpenJMS . . . 225
Configuring JNDI objects for OpenJMS . . . 226
Connection Modes in OpenJMS . . . 228
Changing the Datasource for OpenJMS . . . 228
Creating Tables for OpenJMS . . . 229
Configuring Datasource to Achieve 2PC Optimization . . . 229
Configuring Security with OpenJMS . . . 229
Specifying Partition Level Properties for OpenJMS . 230
OpenJMS Topologies . . . 232
Using Message Driven Beans (MDB) with OpenJMS 232
Other JMS providers . . . 233

Chapter 26
Integrating SonicMQ into Borland AppServer 235
Installing SonicMQ . . . 235
Configuring SonicMQ Administered Objects in AppServer . . . 235
Resolving SonicMQ library modules in the AppServer environment . . . 236
Configuring Automatic Queue Creation for SonicMQ Queues deployed to AppServer . . . 236

Chapter 27
Integrating WebSphereMQ into Borland AppServer (BAS) 239
Supported Versions . . . 239
WebSphereMQ Configuration . . . 239
WebSphereMQ 5.3 . . . 239
WebSphereMQ 6.0 . . . 240
Configuring Admin Objects with WebSphereMQ . . 240
Locating WebSphereMQ Library modules at runtime . . 240
WebSphereMQ 6.0 . . . 241

Chapter 28
Using JACC 243
JACC Contracts . . . 243
Provider Configuration Subcontract . . . 243
Policy Configuration Subcontract . . . 243
Policy Decision and Enforcement Subcontract . . 243
How the JACC-based authorization works . . . 243
Configuring JACC provider in Borland AppServer . . 244
Configuring a JACC provider using AppServer Management Console . . . 244
Configuring a JACC provider through the configuration file . . . 245
Enabling/Disabling the JACC provider . . . 245
Configuring external JACC providers . . . 245

Chapter 29
Using ADLoginModule in BAS 247
How ADLoginModule works . . . 247
User Principal Name . . . 247
Authentication . . . 247
Configuring ADLoginModule . . . 248
Detailed Configuration Options . . . 248

Chapter 30
Using JAXR 251
Using JAXR in BAS . . . 251
System Property . . . 252
JAXR Connection Properties . . . 252
BAS JAXR Example code . . . 253

Chapter 31
Using the Scheduler Service 255
Configuring the Scheduler Service . . . 255
Using JDataStore to persist scheduler events . . . 256
Configuring other databases to persist scheduler events 257
Setting up for 2PC Optimization . . . 257
Partition Service properties for Scheduler Service . 258
Quartz properties used in AppServer . . . 259
Clustering support . . . 260

Chapter 32
VisiConnect overview 261
J2EE Connector Architecture . . . 261
Components . . . 261
System Contracts . . . 263
Connection Management . . . 264
Transaction Management . . . 265
One-Phase Commit Optimization . . . 266
Security Management . . . 266
Component-Managed Sign-on . . . 266
Container-Managed Sign-on . . . 266
EIS-Managed Sign-on . . . 266
Authentication Mechanisms . . . 266
Security Map . . . 267
Security Policy Processing . . . 268
Common Client Interface (CCI) . . . 268
Packaging and Deployment . . . 270
VisiConnect Features . . . 271
VisiConnect Partition Service . . . 271
Additional Classloading Support . . . 271
Secure Password Credential Storage . . . 271
Connection Leak Detection . . . 271
Security Policy Processing of ra.xml Specifications . . . 271
Resource Adapters . . . 272

Chapter 33
Implementing Partition Interceptors 273
Defining the Interceptor . . . 273
Creating the Interceptor Class . . . 274
Creating the JAR file . . . 276
Deploying the Interceptor . . . 276

Chapter 34
Using VisiConnect 277
VisiConnect service . . . 277
Service overview . . . 277
Connection management . . . 277
Configuring connection properties . . . 278
Security management with the Security Map . . . 279
Authorization domain . . . 280
Default roles . . . 280
Generating a resource vault . . . 280
Resource Adapter overview . . . 282
Development overview . . . 283
Editing existing Resource Adapters . . . 283
Resource Adapter Packaging . . . 284
Deployment Descriptors for the Resource Adapter . 284
Configuring ra.xml . . . 284
Configuring the transaction level type . . . 285
Configuring ra-borland.xml . . . 285
Changes to the Deployment Descriptors for Connectors 1.5 . . . 286
Resource Adapter Classloader Considerations . . . 286
Connection Factories and Connections . . . 287
Message Listeners . . . 287
Correcting ClassCastExceptions . . . 289
Developing the Resource Adapter . . . 289
Connection management . . . 289
Transaction management . . . 290
Security management . . . 290
Packaging and deployment . . . 290
Deploying the Resource Adapter . . . 291
Application development overview . . . 291
Developing application components . . . 291
Common Client Interface (CCI) . . . 291
Managed application scenario . . . 292
Non-managed application scenario . . . 292
Code excerpts—programming to the CCI . . . 293
Deployment Descriptors for Application Components 295
EJB 2.x example . . . 296
EJB 1.1 example . . . 297
Other Considerations . . . 299
Working with Poorly Implemented Resource Adapters . . . 299
Examples of Poorly Implemented Resource Adapters . . . 299
Working with a Poor Resource Adapter Implementation . . . 300

Chapter 35
Borland AppServer Ant tasks and running AppServer examples 305
General syntax and usage . . . 305
Name-value pair transformation . . . 306
Name-only argument transformation . . . 306
Multiple File Arguments . . . 306
Syntax and usage for iastool . . . 307
Omitting attributes . . . 309
Examples of iastool Ant tasks . . . 309
deploy . . . 309

merge . . . 309
ping . . . 309
restart . . . 309
Syntax and usage for java2iiop . . . 310
Example of java2iiop Ant task . . . 310
Syntax and usage for idl2java . . . 310
Example of idl2java Ant task . . . 311
Syntax and usage for appclient . . . 312
Building and running the Borland AppServer examples 312
Deploying the example . . . 312
Running the example . . . 312
Undeploying the example . . . 312
Troubleshooting . . . 313

Chapter 36
iastool command-line utility 315
Using the iastool command-line tools . . . 315
clonepartition . . . 316
compilejsp . . . 317
compress . . . 319
deploy . . . 320
dumpstack . . . 321
genclient . . . 322
gendeployable . . . 322
gendeployableverify . . . 323
genstubs . . . 324
info . . . 325
jndinamespace . . . 325
kill . . . 326
listpartitions . . . 327
listhubs . . . 328
listservices . . . 329
manage . . . 329
merge . . . 330
migrate . . . 331
newconfig . . . 332
patch . . . 333
ping . . . 333
pservice . . . 334
removestubs . . . 335
restart . . . 336
setmain . . . 337
start . . . 338
stop . . . 339
uncompress . . . 339
undeploy . . . 340
unmanage . . . 341
usage . . . 342
verify . . . 342
Executing iastool command-line tools from a script file . 344
Piping a file to the iastool utility . . . 344
Passing a file to the iastool utility . . . 344

Chapter 37
Partition XML reference 345
<partition> element . . . 345
<jmx> element . . . 345
<mbean.server> element . . . 346
<mlet.service> element . . . 346
<http.adaptor> element . . . 346
<xslt.processor> element . . . 346
<rmi-iiop.adaptor> element . . . 347
<statistics.agent> element . . . 347
<security> element . . . 348
<container> element . . . 348
<user.orb> element . . . 348
<management.orb> element . . . 349
<shutdown> element . . . 349
<services> element . . . 350
<service> element . . . 350
<properties> element . . . 351
<archives> element . . . 351
<archive> element . . . 352

Chapter 38
EJB, JSS, and JTS Properties 353
EJB Container-level Properties . . . 353
EJB Customization Properties: Deployment Descriptor level . . . 356
Complete Index of EJB Properties . . . 357
Properties common for any kind of EJB . . . 357
Entity Bean Properties (applicable to all types of entities—BMP, CMP 1.1 and CMP 2) . . . 358
Message Driven Bean Properties . . . 361
Stateful Session Bean Properties . . . 363
EJB Security Properties . . . 364
Java Session Service (JSS) Properties . . . 364
Partition Transaction Service (Transaction Manager) . . 367

Chapter 39
Using LifeRay Portal with AppServer 6.7 369
Using Other Databases . . . 370
Deploying Portlet or J2EE modules to LifeRay module . 371

Chapter 40
Integrating Borland AppServer 6.7 with JBuilder 2006 373
Configuring JBuilder 2006 for Borland AppServer 6.7 . . 373
Displaying the Borland Management Console in JBuilder . . . 374
VisiBroker development with JBuilder . . . 375
Using the JBuilder Deployment Descriptor Editor to develop J2EE 1.4 applications . . . 375
Message Destinations page . . . 376
Message Destination Reference page . . . 378
Message-Driven Bean page . . . 379
Resource Environment References page . . . 380
Admin Object and Admin Object Properties page 381
Resource Adapter page . . . 381
BES Connection Definition page . . . 382
Creating a run configuration for Borland AppServer 6.7 targeted projects . . . 383

Changing the management port . . . 384
Launching the partition in JBuilder 2006 . . . 384
Deploying . . . 386
Remote debugging . . . 387
Preparing to remote debug partitions that are not managed in JBuilder . . . 387
Preparing to remote debug partitions with JBuilder . 387
Remote debugging from JBuilder . . . 388

Chapter 41
Managing the Borland AppServer 389
Understanding the SCU process . . . 389
Hub . . . 390
Agent . . . 390
Starting the Hub and Agents . . . 391
Starting a Hub and its Local Agent . . . 391
Starting an Agent . . . 391
Management Port . . . 391
Management Domain . . . 392
Configurations . . . 392
Managed Resources . . . 392
Managed Objects . . . 392
Borland Management Console . . . 393

Chapter 42
Configurations 395
Creating Configurations . . . 395
Configuration Templates . . . 395
configuration.xml files . . . 396
Implementing Configurations . . . 396
Starting Configurations . . . 396
Running Configurations . . . 396
Stopping Configurations . . . 396
Scheduling Configuration tasks . . . 397
Stopping the BAS infrastructure . . . 397
Stopping a Local Hub/Agent . . . 397
Restarting a Hub/Agent . . . 398

Chapter 43
Managed Objects 399
Creating and Adding Managed Objects . . . 399
Managed Object Templates . . . 399
Managed Object types . . . 400
Pure management Managed Object types . . . 400
Ordered Group . . . 401
Redundancy Group . . . 401
State Proxy . . . 403
Process Managed Object type . . . 403
UNIX: Managing Ownership of a Process Managed Object . . . 403
Java Process Managed Object type . . . 404
Extensibility Managed Object types . . . 404
Custom JavaScript Managed Object type . . . 404
Custom Executable Managed Object type . . . 404
AppServer-specific Managed Object types . . . 404
Managed Object actions . . . 405
Strategies . . . 405
Scheduling tasks for a Configuration . . . 405
Creating an availability schedule for a Managed Object 406

Chapter 44
Getting started with the Borland Management Shell 407
Using BMSH . . . 407
Running BMSH . . . 407
Using BMSH interactively (interactive mode) . . . 408
Using BMSH script files (batch mode) . . . 408
BMSH files and folders . . . 409
BMSH files and folders in the bin directory . . . 409
BMSH files and folders in the /bin/bscript/autoload subdirectory . . . 409
Files in /bin/bscript/scripts . . . 409
BMSH subdirectories in the examples directory . 409
Files in /examples/bmsh . . . 409
Files in /examples/bmsh/scripts . . . 410
Files in /examples/bmsh/autoload . . . 410
BMSH features . . . 410
Autoload feature . . . 410
Search path feature . . . 410
BMSH classpath add JARs feature . . . 411
BMSH Objects . . . 411
BMSH pre-instantiated objects . . . 411
BMSH user-instantiated objects . . . 411
BMSH factory objects . . . 411
Java LiveConnect . . . 412
Writing your own Rhino scriptable objects . . . 412
Environment variables . . . 412
JavaScript arrays . . . 412
Java string arrays . . . 413
JavaScript built-in functions . . . 413
BMSH interactive behavior . . . 413
BMSH Help . . . 413

Chapter 45
Configurations and configuration.xml files 415
configuration element . . . 415
configuration element attributes . . . 415
configuration element sub-elements . . . 416
configuration-id element . . . 417
main-root sub-element . . . 417

Chapter 46
Managed Object elements and attributes 419
Generic XML definition . . . 419
managed-object element attributes . . . 420
Managed Object type attribute . . . 422
ordered-group Managed Object type . . . 423
ordered-group sub-elements . . . 423
ordered-group sub-element . . . 423
ordered-group sub-element attributes . . . 425
member sub-element . . . 426
member sub-element attributes . . . 426
redundancy-group Managed Object type . . . 427

redundancy-group sub-elements 427
redundancy-group sub-element attributes 429
member sub-element 430
member sub-element attributes 430
state-proxy Managed Object type 431
state-proxy sub-elements 431
state-proxy sub-element attributes 432
process Managed Object type 432
process sub-elements 432
process sub-element attributes 434
process sub-elements 434
arguments sub-element 434
env-vars sub-element 435
env-vars sub-element attributes 435
library-path sub-element 436
path sub-element 436
stdin|stdout|stderr sub-elements 437
stdin|stdout|stderr sub-element attributes . . 437
platform-specific sub-element 438
java-process Managed Object type 439
java-process sub-elements 439
java-process sub-element attributes 440
java-process sub-elements 441
arguments sub-element 442
env-vars sub-element 442
env-vars sub-element attributes 443
library-path sub-element 443
path sub-element 443
stdin|stdout|stderr sub-elements 443
platform-specific sub-elements 444
VBJ-process sub-element attributes 444
VBJ-process sub-elements 445
options sub-element 446
arguments sub-element 447
env-vars sub-element 447
env-vars sub-element attributes 447
library-path sub-element 447
path sub-element 448
stdin|stdout|stderr sub-elements 448
platform-specific sub-elements 448
custom-javascript Managed Object type 449
custom-javascript sub-elements 449
control-overrides element sub-elements . . . 450
start sub-element attributes 452
ping sub-element attributes 453
stop sub-element attributes 453
kill sub-element attributes 454
jscript sub-element attributes 454
jscript sub-elements 455
run sub-element attributes 456
run sub-element 456
![CDATA[section 457
![CDATA[attribute 457
custom-executable Managed Object type 457
custom-executable sub-elements 458
control-overrides sub-elements 459
start sub-element attributes 461
ping sub-element attributes 462
stop sub-element attributes 463
kill sub-element attributes 463
process sub-element attributes 464
osagent Managed Object type 465
osagent sub-elements 465
process attributes and sub-elements 466
osagent sub-elements 467
apache-process Managed Object type 467
apache-process sub-elements 467
apache-data sub-element attributes 469
apache-data sub-elements 469
ots Managed Object type 471
ots sub-elements 471
tibco Managed Object type 473
tibco sub-elements 473
tibco-data sub-element attributes 475
process sub-element attributes 476
partition Managed Object type 476
partition-process sub-elements 476
partition-process sub-element attributes 477
partition-process sub-elements 477
optimizeit sub-element attributes 478
optimizeit element sub-elements 478
strace sub-element attributes 479
profiler sub-element attributes 479
strace element sub-elements 480
strace element classpath sub-element 480
strace element bootclasspath sub-element . . 480
strace element options sub-element 481
strace element path sub-element 481
profiler element sub-elements 481
profiler element classpath sub-element 482
profiler element bootclasspath sub-element . 482
profiler element options sub-element 482
profiler element path sub-element 483
jpda sub-element attributes 483
partition-services element sub-elements . . . 484
member sub-element attribute 485
jmx sub-element attributes 485
Partition Services Managed Object types 486

Chapter 47
Managed Object control-overrides element 487

control-overrides sub-elements 487
parameters sub-element attributes 489
start sub-element attribute 491
stop sub-element attribute 491
kill sub-element attribute 492
ping sub-element attribute 492
start sub-element strategy attribute 493
stop sub-element strategy attribute 494
ping sub-element strategy attribute 495
kill sub-element strategy attribute 496

Chapter 48
Managed Object time-rules element 499

time-rules Sub-element 499
time-rules Sub-element attributes 500
time-rule Sub-element 500
time-rule Sub-element attributes 501

time-rule Examples 503

Chapter 49
Working with Configuration properties 505
properties element 505
Defining a property 505
Referencing a defined property 506
Defining properties in terms of other properties . 506
Hub and Agent properties 506
Built-in, pre-defined properties 507
%platform% property 508
Using Property Expressions 509
if statements 509
Basic if statement examples 510
Nested if statements 511
switch statements 511

Chapter 50
Borland Management Shell (BMSH) API Specification 513

Index 529

1

Introduction to Borland AppServer

Borland AppServer (AppServer) is a set of services and tools that enable you to build, deploy, and manage distributed enterprise applications in your corporate environment.

The AppServer is a leading implementation of the J2EE 1.4 standard, and supports the latest industry standards such as EJB 2.1, JMS 1.1, Servlet 2.4, JSP 2.0, CORBA 2.6, XML, and SOAP. Borland provides two versions of AppServer, which include leading enterprise messaging solutions for Java Messaging Service (JMS) management (Tibco and OpenJMS). You can choose the degree of functionality and services you need in AppServer, and if your needs change, it is simple to upgrade your license.

The AppServer allows you to securely deploy and manage all aspects of your distributed Java and CORBA applications that implement the J2EE 1.4 platform standard.

With AppServer, the number of server instances per installation is unlimited, so the maximum of concurrent users is unlimited.

AppServer includes:

- Implementation of J2EE 1.4.
- Apache Web Server version 2.2.3
- Borland Security, which provides a framework for securing AppServer.
- Single-point management of leading JMS management solutions included with AppServer (Tibco, and OpenJMS).
- Strong management tools for distributed components, including applications developed outside of AppServer.

AppServer features

AppServer offers the following features:

- Support for BAS platforms (please refer to http://support.borland.com/kbcategory.jspa?categoryID=389 for a list of the platforms supported for AppServer).
- Full support for clustered topologies.
- Seamless integration with the VisiBroker ORB infrastructure.
- Integration with the Borland JBuilder integrated development environment.
- Enhanced integration with other Borland products including Borland Optimizeit Profiler and ServerTrace.
- AppServer allows existing applications to be exposed as Web Services and integrated with new applications or additional Web Services. Borland Web Services support is based on Apache Axis 1.2 technology, the next-generation Apache SOAP server that supports SOAP 1.2.

Borland AppServer Documentation

The AppServer documentation set includes the following:

- *Borland AppServer Installation Guide*—describes how to install AppServer on your network. It is written for system administrators who are familiar with Windows or UNIX operating systems.
- *Borland AppServer Developer's Guide*—provides detailed information about packaging, deployment, and management of distributed object-based applications in their operational environment.
- *Borland Management Console User's Guide*—provides information about using the Borland Management Console GUI.
- *Borland Security Guide*—describes Borland's framework for securing AppServer, including VisiSecure for VisiBroker for Java and VisiBroker for C++.
- Borland *VisiBroker for Java Developer's Guide*—describes how to develop VisiBroker applications in Java. It familiarizes you with configuration and management of the Visibroker ORB and how to use the programming tools. Also described is the IDL compiler, the Smart Agent, the Location, Naming and Event Services, the Object Activation Daemon (OAD), the Quality of Service (QoS), and the Interface Repository.
- Borland *VisiBroker VisiTransact Guide*—describes Borland's implementation of the OMG Object Transaction Service specification and the Borland Integrated Transaction Service components.

The documentation is typically accessed through the Help Viewer installed with your AppServer product. You can choose to view help from the standalone Help Viewer or from within a AppServer GUI tool. Both methods launch the Help Viewer in a separate window and give you access to the main Help Viewer toolbar for navigation and printing, as well as access to a navigation pane. The Help Viewer navigation pane includes a table of contents for all AppServer books and reference documentation, a thorough index, and a comprehensive search page.

The PDF books, Borland AppServer Developer's Guide and Borland Management Console User's Guide are available online at http://info.borland.com/techpubs/appserver.

Accessing AppServer online help topics

To access the online help, use one of the following methods:

Windows

Choose Start|Programs|Borland Deployment Platform|Help Topics

or, launch the Web browser and open <AppServer_Home>/doc/index.html.

UNIX

Launch a Web browser and open <AppServer_Home>/doc/index.html.

Accessing AppServer online help topics from within a AppServer GUI tool

To access the online help from within a AppServer GUI tool, use one of the following methods:

- From within the Borland Management Console, choose Help|Help Topics
- From within the Borland Deployment Descriptor Editor (DDEditor), choose Help|Help Topics

Documentation conventions

The documentation for AppServer uses the typefaces and symbols described below to indicate special text:

Convention	Used for
italics	Used for new terms and book titles.
`computer`	Information that the user or application provides, sample command lines and code.
`bold computer`	In text, bold indicates information the user types in. In code samples, bold highlights important statements.
[ ]	Optional items.
...	Previous argument that can be repeated.
|	Two mutually exclusive choices.

Platform conventions

The AppServer documentation uses the following symbols to indicate platform-specific information:

Symbol	Indicates
Windows	All supported Windows platforms.
Win2003	Windows 2003 only
WinXP	Windows XP only
Win2000	Windows 2000 only
UNIX	UNIX platforms
Solaris	Solaris only

Contacting Borland support

Borland offers a variety of support options. These include free services on the Internet where you can search our extensive information base and connect with other users of Borland products. In addition, you can choose from several categories of telephone support, ranging from support on installation of Borland products to fee-based, consultant-level support and detailed assistance.

For more information about Borland's support services or contacting Borland Technical Support, please see our web site at http://support.borland.com and select your geographic region.

When contacting Borland's support, be prepared to provide the following information:

- Name
- Company and site ID
- Telephone number
- Your Access ID number (U.S.A. only)
- Operating system and version
- Borland product name and version
- Any patches or service packs applied
- Client language and version (if applicable)
- Database and version (if applicable)
- Detailed description and history of the problem
- Any log files which indicate the problem
- Details of any error messages or exceptions raised

Online resources

You can get information from any of these online sources:

World Wide Web: http://www.borland.com

Online Support: http://support.borland.com (access ID required)

World Wide Web

Check http://www.borland.com regularly. The AppServer Product Team posts white papers, competitive analyses, answers to FAQs, sample applications, updated software, updated documentation, and information about new and existing products.

You may want to check these URLs in particular:

- http://www.borland.com/downloads/download_appserver.html (AppServer software and other files)
- http://support.borland.com (AppServer FAQs)

Borland newsgroups

You can participate in many threaded discussion groups devoted to the AppServer. Visit http://www.borland.com/newsgroups for information about joining user-supported newsgroups for Enterprise Server and other Borland products.

Note

These newsgroups are maintained by users and are not official Borland sites.

2

Borland AppServer overview and architecture

This section contains an overview of the Borland AppServer (AppServer).

AppServer architecture overview

The AppServer is a CORBA-based, J2EE server that utilizes distributed objects throughout its architecture. With the AppServer, you can establish connectivity to platforms from corporate mainframes to simpler systems with small-business applications and remote databases. The AppServer components process your enterprise application based on how it is packaged and how the deployment descriptors describe the application's modules.

In the following architectural diagram, your enterprise applications sit on top of the AppServer. An application server installation contains AppServer *core services* and *Partitions*.

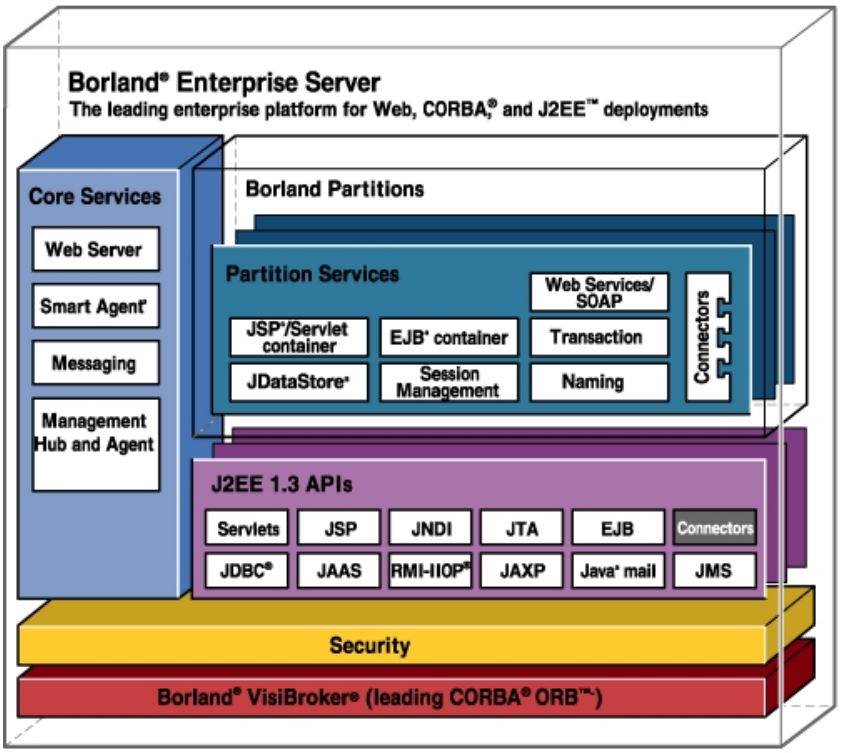

AppServer services overview

AppServer services are those services available to all applications being hosted on the AppServer. They are:

- Web Server
- Java Messaging (JMS)
- Smart Agent
- 2PC Transaction Service

Web Server

The AppServer includes the Apache Web Server version 2.2.3. The Apache web server is a robust, commercial grade reference implementation of the HTTP. protocol. The Apache web server is highly configurable and extensible through the addition of third-party modules. Apache supports clients with varying degrees of sophistication and supports content negotiation to this end. Apache also provides unlimited URL aliasing.

Borland has added an IIOP Plug-in to the Apache web server. The IIOP Plug-in allows Apache and the Borland web container to communicate via Internet Inter-ORB Protocol (IIOP), allowing users to add the power of CORBA with their Web applications in new ways. In addition, IIOP is the protocol of the VisiBroker ORB, allowing your Web applications to fully leverage the services provided by the object-request-broker provided by Borland.

JMS

The AppServer provides support for standard JMS pluggability, and currently bundles the Tibco messaging service. See “JMS provider pluggability” for vendor-specific information on JMS services.

Smart Agent

The Smart Agent is a distributed directory service provided by the VisiBroker ORB used in the AppServer. The Smart Agent provides facilities used by both client programs and object implementations, and must be started on at least one host within the local server network.

Note

Users of the Web Edition do not have to use the Smart Agent if they expect their Web server and Web containers to communicate through HTTP or another Web protocol. To leverage the IIOP Plug-in (and, by extension, the ORB provided with the Web Edition), however, the Smart Agent must be turned on.

More than one Smart Agent can be configured to run on your network. When a Smart Agent is started on more than one host, each Smart Agent will recognize a subset of the objects available and communicate with the other Smart Agents to locate objects it cannot find. In addition, if one of the Smart Agent processes should terminate unexpectedly, all implementations registered with that Smart Agent discover this event and they will automatically re-register with another Smart Agent. It should be noted that If a heavy services lookup load is necessary, it is advisable to use the Naming Service (VisiNaming). VisiNaming provides persistent storage capability and cluster load balancing whereas the Smart Agent only provides a simple round robin on a per osagent basis.

2PC Transaction Service

The Two-Phase Commit (2PC) Transaction Service exists provides a complete recoverable solution for distributed transactional CORBA applications. Implemented on top of the VisiBroker ORB, the 2PC Transaction Service simplifies the complexity of distributed transactions by providing an essential set of services, including a transaction service, recovery and logging, integration with databases, and administration facilities within one, integrated architecture.

The Partition and its services

A Partition is an application's deployment target. The Partition provides the J2EE server-side runtime environment required to support a complete J2EE 1.3 application. While a Partition is implemented as a single native process, its core implementation is Java. When a Partition starts, it creates an embedded Java Virtual Machine (JVM) within itself to run the Partition implementation and the J2EE application code.

Partitions are present in each AppServer Edition and product but they host less diverse archives in the Web Services, Team and VisiBroker Editions. This section describes the full-featured functional Partitions offered in the full Borland AppServer. Each Partition instance provides:

- Connector Service
- EJB Container
- JDataStore Server
- Lifecycle Interceptor Manager
- Naming Service
- Session Storage Service
- Transaction Manager
- Web Container

Connector Service

The Connector Service, also known as VisiConnect, is the Borland implementation of the Connectors 1.0 standard, which provides a simplified environment for integrating various EISs with the AppServer. The Connectors provide a solution for integrating J2EE-platform application servers and EISs, leveraging the strengths of the J2EE platform—connection, transaction and security infrastructure—to address the challenges of EIS integration. For more information see “VisiConnect overview.”

EJB Container

The AppServer provides integrated EJB container services. These services allow you to create and manage integrated EJB containers or EJB containers across multiple Partitions. Use this service to deploy, run, and monitor EJBs. Tools include a Deployment Descriptor Editor (DDEditor) and a set of task wizards for packaging and deploying EJBs and their related descriptor files. EJB containers can also make use of J2EE connector architecture, which enables J2EE applications to access Enterprise Information Systems (EISs).

JDataStore Server

Borland's JDataStore is a relational database service written entirely in Java. You can create and manage as many JDataStores as desired. For more information on JDataStore, see the JDatastore online documentation at `www.borland.com/techpubs/jdatastore/`.

Lifecycle Interceptor Manager

You can use Lifecycle Interceptors to further customize your implementation. Partition Lifecycle Interceptors allow you to perform operations at certain points in a Partition's lifecycle. For more information see “Implementing Partition Interceptors,”

Naming Service

The Naming Service is provided by the VisiBroker ORB. It allows developers, assemblers, and/or deployers to associate one or more logical names with an object reference and store those names in a VisiBroker namespace. It also allows application clients to obtain an object reference by using the logical name assigned to that object. Object implementations can bind a name to one of their objects within a namespace which client applications can then use to resolve a name using the `resolve()` method. The method returns an object reference to a naming context or an object.

Session Storage Service

The Java Session Service (JSS) is a service that stores information pertaining to a specific user session. The JSS provides a mechanism to easily store session information into a database. For example, in a shopping cart scenario, information about your session (your login name, the number of items in the shopping cart, and such) is polled and stored by the JSS. So when a session is interrupted by a Borland web container unexpectedly going down, the session information is recoverable by another Tomcat instance through the JSS. The JSS must be running on the local network. Any web container instance (in the cluster configuration) will find the JSS, connect to it, and continue session management. For more information, see “Java Session Service (JSS) configuration”.

Transaction Manager

A Partition Transaction Manager exists in each AppServer Partition. It is a Java implementation of the CORBA Transaction Service Specification. The Partition Transaction Manager supports transaction timeouts, one-phase commit protocol, and can be used in a two-phase commit protocol under special circumstances. For more information, see “Transaction management.”

Web Container

The Web Container is designed to support deployment of web applications or web components of other applications (for example, servlets and JSP files). The AppServer provides the Borland Web Container, which is based on Tomcat 5.5.17. Tomcat is a sophisticated and flexible open-source tool that provides support for servlets, JavaServer Pages, and HTTP. Borland has also provided an IIOP plug-in with its Web Container, enabling communication with application components and the web server over IIOP rather than strict HTTP. Other features of the Web Container are:

- EJB referencing
- DataSource referencing
- Environment referencing
- Integration into industry-standard web servers

For more information, see “Web components”.

Borland AppServer and J2EE APIs

Since the AppServer is fully J2EE 1.4 compliant, it supports the use of the following J2EE 1.4 APIs:

- JNDI: the Java Naming and Directory interface
- RMI-IIOP: remote method invocation (RMI) carried out via internet inter-ORB protocol (IIOP)
- JDBC: for getting connections to and modeling data from databases
- EJB 2.1: the Enterprise JavaBeans 2.1 APIs
- Servlets 1.0: the Sun Microsystems servlets APIs
- JSP: JavaServer Pages APIs
- JMS: Java Messaging Service
- JTA: the Java transactional APIs
- Java Mail: a Java email service
- Connectors 1.5: the J2EE Connector Architecture
- JAAS: the Java Authentication and Authorization Service
- JAXP: the Java API for XML parsing

JDBC

Borland implements the Java DataBase Connection APIs from Sun Microsystems. JDBC provides APIs for writing database drivers and a full Service Provider Interface (SPI) for those looking to develop their own drivers. JDBC also supports connection pooling and distributed transaction features. For more information, go to the Transaction management and JDBC, JDBC API Modifications section.

Java Mail

Java Mail is an implementation of Sun's Java Mail API. It is a set of abstract APIs that model a mail system. The API provides a platform independent and protocol independent framework to build Java-technology-based email client applications.

JTA

The Java Transactional API (JTA) defines the UserTransaction interface required by application components to start, stop, rollback, or commit transactions. EJBs establish transaction participation through the `getUserTransaction` method, while other components do so using JNDI lookups. JTA also specifies the interfaces needed by Connectors and resource managers to communicate with an application server's transaction manager.

JAXP

The Java APIs for XML Parsing (JAXP) enable the processing of XML documents using the DOM, SAX, and XSLT parsing implementations. Developers can easily use the parser provided with the reference implementation of the API to XML-enable their Java applications.

JNDI

The Java Naming and Directory Interface is used to allow developers to customize their application components at assembly and deployment without changes to the component's source code. The container implements the runtime environment for the components and provides the environment to the component as a JNDI naming context. The components' methods access the environment through JNDI interfaces. The JNDI naming context itself stores the application environment information and makes it available to all application components at runtime.

RMI-IIOP

The VisiBroker ORB supports RMI-over-IIOP protocol. When used in conjunction with the IIOP Connector Module for Apache and the Borland web container, it allows distributed web applications built on CORBA foundations. For more information, see “Using RMI over IIOP” in the *VisiBroker for Java Developer's Guide*.

Other Technologies

It is also possible to wrap other technologies, provide them as services, and run them in the AppServer.

Optimizeit Profiler and Optimizeit ServerTrace

Borland's Optimizeit Profiler (purchased separately) helps you track memory and CPU usage issues during the development of Java applications. Optimizeit ServerTrace provides a comprehensive, high-level application performance analysis and root-cause diagnostics that accelerate time-to-resolution of performance issues across complex, distributed J2EE and SOA-enabled systems. The AppServer runs Optimizeit Profiler and Optimizeit ServerTrace at the Partition level.

See the Sun Java Center for more information on these APIs.

3

Using Partitions

This section discusses how to use the Management Console to work with Partitions for Borland AppServer (AppServer). It describes the following tasks:

- Creating, cloning, and deleting Partitions
- Deploying modules to a Partition
- Configuring an existing Partition
- Viewing Partition information
- Tuning Partitions for Performance
- Dumping the Partition stack trace to a log file
- Enabling a Partition to run with Optimizeit Profiler or ServerTrace

Creating, cloning, and deleting Partitions

Using the Management Console, you can create, clone, and delete Partitions for your applications. A Partition can be created from a template or cloned from an existing Partition. Partitions are represented in the Navigation Pane of the Management Console's Hubs View as child nodes of their parent configurations.

Creating a new Partition

To create a new Partition:

1. Select the Configuration to which you want the new Partition to belong in the Navigation Pane.
2. Right-click the Configuration and select Add Managed Object.

   The Managed Object Template Gallery opens.
3. From the AppServer or OpenJMS Partition categories, choose from the following Partition templates:
   - **AppServer 6.7 Partition:** Produces a managed Partition using the default partition.config for AppServer 6.7.
   - **Standard Partition:** Produces a managed Partition.

- **Explicitly Pathed Partition:** Produces a path to an existing Partition. Use this template to migrate existing Partitions that you want under the management of the current configuration.
- **JBuilder Partition:** Produces an unmanaged Partition. **The JBuilder Partition is for use with JBuilder debugging of local servers. It is automatically created by JBuilder when used with AppServer for debugging purposes. You should not use the JBuilder Partition otherwise.**
- **Partition with Embedded OpenJMS:** Produces a Partition with embedded OpenJMS.

1 Select a Partition template and click Add.

The Add From Template dialog box appears.

Note

The information that appears on the Add From Template dialog box depends upon which template you chose in the previous step.

2 Enter the required information (shown in bold) in the dialog.

Different properties appear in this dialog, depending on which template you selected. Some of the common properties are:

- **Name:** Give the Partition a unique name in the Name field. This name will be used as the value of the string-replacement variable `${mo.name}` throughout the rest of the dialogs, and will be used as the display name of the Partition. If you want a display name different from the actual Partition name, change the value of the Display Name field.
- **Management Agent:** Select the host on which to create the Partition by selecting a valid Management Agent from the drop-down list. If you want to use the current Management Agent, represented by the string replacement variable `${hub.name}`, you do not need to change this field.
- **Display Name:** You can optionally enter a friendly name to be displayed in the Management Console. The string replacement variable, `${mo.name}`, indicates that the value in the Name field will be displayed.
- **Smart Agent Port:** You can modify the default osagent (Smart Agent) port number by entering a valid port number in this field.
- **HTTP Connector Port:** You can modify the default HTTP connector port number in this field.
- **Data Directory:** For an explicitly pathed Partition you must enter the path to an existing Partition.

1 Check the Show hidden properties check box to display some additional Partition configuration properties.

2 When you are finished, click OK.

Cloning an existing Partition

To clone an existing Partition:

1. Select the Partition you want to clone from the Navigation Pane.
2. Right-click the Partition and select Clone.

   The Clone Managed Object dialog appears.
3. Select the Configuration for which you want to target the new, cloned Partition from the Target Configuration drop-down list.
4. Enter the name you want to provide for the cloned Partition in the New Name field.

   This name will be used as the value of the string-replacement variable `${mo.name}` throughout the rest of the dialogs, and will be used as the display name of the cloned Partition. If you want a display name different from the actual Partition name, change the value of the New Display Name field.
5. Enter a Description, if desired. This is optional.
6. The Data Directory, `${config.path}/mos/${mo.name}`, identifies the location of the Partition relative to its Agent.
7. Choose a management agent to host the cloned Partition from the Target Agent drop-down list.
8. Choose the group to which the cloned Partition should belong from the Target Group drop-down list.
9. When you are finished, click OK.

Deleting a Partition

To delete a Partition:

1. Select the Partition you want to delete from the Navigation Pane.
2. Right-click the Partition and select Remove.

Deploying modules and libraries to a Partition

To deploy modules to a Partition:

1. Select the Partition to which you want to deploy in the Navigation Pane.
2. Right-click the Partition and select Deploy modules.

   The Module and Library Deployment Wizard appears.
3. To add a module, click Add.

   The Add J2EE Module dialog appears.

   a. Navigate to the module you wish to deploy to the Partition, and click OK.

   b. Repeat this step until all the application modules to be deployed have been added.

      If you make an error, you can highlight the module in the list and click Remove to delete it.

4 Select any additional options using the check boxes. The options are:
   - **Restart partitions on deploy (cold deploy):** Executes a “cold” deploy, and will restart the Partition after the deployment operation is complete. Leave this unchecked if you wish to “hot” deploy to a running Partition without restarting.
   - **Verify deployment descriptors:** Runs the verification tool which checks to make sure that all deployment descriptors have been constructed properly, including Borland-specific descriptors. This option is recommended.
   - **Generate stubs:** Check this box if you want your application stubs generated at deploy-time. This option is recommended.
5 Click Advanced Options to configure advanced properties for this deployment. See “Advanced options for deployment” for more information.
6 Click Next to continue.

   Step 2 of the wizard appears.
7 Select the Partitions that will be the deployment targets for your modules.

   The available Partitions appear automatically in the list. If no Partitions appear initially, click Refresh List. *Shift-click* or *Ctrl-click* to select multiple Partitions.
8 When you are done, click Finish.
9 While the module is deployed you can observe the progress in the Deploying Modules dialog.
10 When you are done, click Close.

Note

According to the specification for application module (EAR), the classpath entries of the submodule that is specified in the manifest file should be added to the application classpath. This feature is not used in most cases. If you want to use it, please specify the VM property `enable.add.classpath.entries` into the corresponding configuration file, for example, partition.config or iastool.config.

Advanced options for deployment

In Step 1 of the Deployment Wizard, you can choose to set Advanced Options for the Stub Generator and Verifier tools. To work with advanced settings:

1 In Step 1 of the Deployment Wizard, click Advanced Options.

   The Advanced Deployment Options dialog appears.
2 On the Stub Generator tab, you can set the following options:
   - **Generate stubs:** Check this box to enable stub generation on deployment.
   - **Classpath:** Select a Classpath from the drop-down list or click Edit to add archives to the list for addition to the Classpath.
   - **Edit:** Click Edit to open the Classpath Editor, where you can add and remove archives and paths using the editor.
   - **Java2IIOP arguments:** Include a sequence of `java2iiop` command-line arguments in this field. Click More Info for a list of valid command-line arguments, or see “Programmer tools for Java” in the *VisiBorker for Java Developer’s Gude.*
   - **More Info:** Click More Info to view the Stub Generator compiler flag usage information.
   - **Javac arguments:** Include a sequence of `javac` command-line arguments in this field.
1 On the Verifier tab, you can choose the level of verification using the following check boxes:

- **Verify deployment descriptors:** Check this box to verify all descriptors, both Borland and standard.
- **Show all warnings:** Check this box to receive logged information of all potential problems with your descriptors.
- **Use strict (pedantic) checks:** Perform pedantic checks on all descriptors.

1 When you are finished, click OK.

Hosting additional modules

You can also host “pre-deployed” modules that will not reside on the Partition footprint. Additionally, you have the option of hosting a path (that is, an application that is not converted to an archive form) called an “exploded archive.”

To have a Partition host an archive:

1 Select the Partition to which you want to deploy in the Navigation Pane.

2 Right-click the Partition and select Host additional module.

The Host Additional Module dialog appears.

3 Choose the Select File option to host an archive, or choose the Select Directory option to host an exploded archive.

4 Click Browse to locate the archive or exploded archive you want to host.

5 Optionally, you can provide a descriptive Module name for the hosted module.

6 When you are finished, click OK.

Configuring Partitions

You can set Partition properties, including all those that can be viewed in the Content Pane when a Partition is selected in the Navigation Pane. You can specify general information about the Partition, Partition properties, statistics gathering settings, JMX Agent settings, log settings, JMX client settings, time rules, and advanced options for management.

To edit a Partition's property settings:

1 Select the Partition in the Navigation Pane and right-click it.

2 Select Properties.

The Partition Properties dialog appears.

3 Use the tabs described in the following sections to edit the settings.

4 When you are finished making changes, click OK.

General properties

The properties on the General tab allow you to specify some general information about the Partition.

You can edit the following options:

- **Display name:** Enter the name of this Partition as displayed in the Management Console.
- **Data directory:** Enter the location of the Partition's footprint, relative to its Agent.
- **Version:** A version number created and tracked by the Management System.
- **Vendor:** The managed object vendor. The standard partition vendor is Borland Software Corporation.

- **Description:** An optional description for your Partition.

Partition Settings properties

The properties on the Partition Settings tab allow you to specify command-line arguments for the Partition and set up JPDA debugging.

You can set the following options:

- **Arguments:** You can either enter a space-separated list of command-line arguments for the Partition executable, or click the Edit button to store them in a list.
- **Enable JPDA remote debugging:** Flags whether or not debugging is enabled on the Partition.
- **JPDA debugging transport address:** The port used by the JPDA debugger to connect to the Partition. Leave this field blank for random port assignment.
- **Suspend partition until debugger attaches:** This is a flag that, when checked, does not mark the Partition as Running until the debugger successfully attaches.

Statistics properties

The properties on the Statistics tab allow you to gather statistical information which is displayed in the statistics tabs throughout the Management Console. Statistics gathering is enabled by default. Disable statistics gathering to improve the performance of the Partition.

When enabled you can configure the following statistics gathering settings:

- **Enable Agent Statistics:** Check this box to enable the Partition's statistics agent.

  You can set the logging level using the Statistics level drop-down list. Additionally you can set the polling interval by entering a value in the Snapshot period field.

- **Enable Agent Statistics Reaping:** Check this box to periodically delete stored statistics information and preserve disk space, known as *reaping*.

  You determine how long to keep statistics by entering a value in Reap older than. You can set how often the statistics logs are reaped by entering a value in Reap period.

JMX Agent properties

The properties on the JMX Agent tab allow you to configure some aspects of the JMX MBean Server, the RMI-IIOP and HTTP adaptors, and the MLet service implemented for the Partition.

You can edit the following options:

- **Enable JMX:** Check this box to enable the JMX MBean Server.

  The MBean Server is an interface and a factory object defined by the agent specification level of the JMX. This option must be enabled to launch the JMX Console (see “Using the JMX console”).

- **Enable HTTP Adaptor:** Check this box to enable the HTTP adaptor.

  The HTTP adaptor is the adaptor for the HTTP protocol through which the Partition can be managed through any HTML 3.2 compliant browser or application.

  You can configure the port number at which the HTTP adaptor is listening (default 8082) and enable the XSLT processor in this configuration tab. If you are using a Web browser to monitor the Partition, the XSLT processor must be enabled to transform the HTTP adaptor output from raw XML to HTML. To configure the host name (default `localhost`) you must edit partition.xml (see “Partition XML reference” in *the AppServer Developer's Guide* for more details).

- **Enable RMI-IIOP adaptor:** Check this box to enable the RMI-IIOP adaptor.

The RMI-IIOP adaptor is based on the JMX client framework, which helps managing applications to communicate with the MBean Server through RMI.

– **Enable mlet service:** Check this box to enable the MLet service.

The MLet service is useful for loading MBean classes and resources inside an MBean Server's JVM from a remote host and registering the MBeans in a single action.

For more information about configuring the JMX agent properties see “Partition XML reference” in *AppServer Developer's Guide*. For information about using the JMX console, a JMX client available to monitor MBeans associated with the Partition, see “Using the JMX console.”

JDK properties

The properties on the JDK tab allow you to set JDK properties on the Partition.

You can edit the following options:

– **Select the JDK to be used by this partition:** Select the appropriate JDK from the list by clicking on it.

– **Heap and Thread Stack Sizes:** Enter a value in the Initial heap size field to modify the starting heap size.

You can control the maximum heap size by entering a value in the Maximum heap size field. Use the Java thread stack size field to control how threads are managed within the VM.

– **Java VM Type:** Select the Java VM type by selecting the appropriate radio button. The values are Server, Client, or Other (which you can customize from the drop-down list).

Advanced JDK options

Click Advanced Configuration to open the `partition_server.config` file in an editing window. Add entries to the file as needed. When you are finished, click OK.

Performance Tuning Hints

Click Performance Tuning Hints to get hints on optimizing the performance of your Partition. When you are finished, click OK. For more information about performance tuning, see “Tuning the Partition for performance”.

VisiBroker properties

The properties on the VisiBroker tab allow you to tweak some of the VisiBroker ORB properties used by the Partition.

You can edit the following options:

– **Select a server connection manager:** Select a server connection manager from the drop-down list.

– **Listener port:** Set the server connection manager listener port.

– **Connection Pool:** Set the maximum number of allowable connections and the maximum connection idle time in these fields.

– **Thread Dispatcher Pool:** Set the range of allowable threads and the maximum thread idle time in these fields.

Advanced VisiBroker options

Click Advanced to open the `vbroker.properties` file in an editing window. Edit the file as needed. When you are finished, click OK.

Security properties

The properties on the Security tab allow you to set security information for your Partition.

To enable security, select a Security profile from the drop-down list. To enable SSL on this Partition select the ssl_enabled option. Additionally, if you selected `ssl_enabled` you can configure the following settings:

- **SSL Listener Settings:** Set the SSL listener port and enable trust in a client connected through the port.
- **SSL Connection Pool:** Set the maximum number of allowable connections and the maximum connection idle time in these fields.
- **SSL Thread Dispatcher Pool:** Set the range of allowable threads and the maximum thread idle time in these fields.

Log Settings properties

The properties on the Log Settings tab allow you to set logging properties on the Partition.

You can edit the following options:

- **Partition log format:** Choose a log format (XML or text) from the drop-down list.
- **Trace Level Settings:** Select the trace level (minimal or verbose) for each of the services where trace is provided.

Time Rules properties

The Time Rules tab allows you to configure rules for when the Partition should be running and when it should be stopped.

You can edit the following properties:

- **Default state:** Indicate whether the default state of the Partition is Running or Stopped with this drop-down list.
- **Rules:** Configure the Managed Object Availability Timer rules for the Partition here.

Advanced properties

The properties on the Advanced tab allow you to configure advanced information about how the Partition is managed.

You can set the following two types of management information: Managed Object Settings, and Management Action Settings.

Managed Object Settings

You can edit the following Managed Object Settings:

- **Local restart:** When checked, instructs the local Agent to carry out restarts on the Partition rather than having a remote Hub send the request.
- **Escalate stop:** When checked, the stop operation on the Partition will be escalated to a kill if the stop operation times out.
- **Ping policy:** When set to Always, the Partition is always pinged for its status whether it is running or not. When set to Not-when-stopped, the Partition is only pinged when it is known to be in a Running state.
- **Ping strategy:** The AppServer Action Strategy to use for this operation.
- **Ping interval:** Amount of time in seconds between state checks on the Partition.
- **Max failure retries:** Maximum number of times the Managed Object will attempt to retry an operation, such as Start or Stop.
- **Failure retry interval:** Amount of time between Managed Object attempts to retry an operation.

Management Action Settings

This section contains three tabs for configuring starting, stopping, and killing the Partition. For each operation (Start, Stop, and Kill) you can edit the following:

- **Strategy:** the AppServer Action Strategy to use for this operation.
- **First ping delay:** Amount of time before any pinging begins. Leave this field blank for no delay before the first Ping operation.
- **Ping interval:** Amount of time between state checks on the Partition to determine the success of the operation.
- **Retry interval:** (Stop and Kill tabs only) Amount of time between retry attempts.
- **Max retries:** Maximum number of times the Managed Object will attempt to retry this operation if it fails the first time.
- **Timeout:** Amount of time to allow the operation to proceed without success.

Viewing Partition information

When you select a Partition in the Navigation Pane, a variety of information is displayed in the Content Pane on the right side of the Management Console. This section details the tabs that appear in the Content Pane and what information they display.

General tab

The General tab displays basic information about the Partition. It shows three different categories of information: General Properties, Partition Properties, Security Settings, and the Web Container Root Context.

General properties

- **Display Name:** The name of the Partition as displayed in the Management Console.
- **Name:** The logical name of the Partition.
- **Agent name:** The hosting Agent for the Partition.
- **Data directory:** The location of the Partition's footprint, relative to its Agent.
- **Version:** An optional version number for the Partition, provided by the user.
- **Vendor:** The Partition's vendor, provided by the user.
- **Description:** An optional description for the Partition, provided by the user.

Partition properties

- **Path of partition on server:** The physical location of the Partition's footprint.
- **Verify all modules when loaded:** Flags whether or not the verification tool runs at Partition startup.

Security properties

- **Security profile:** Flags whether or not basic security (authentication/authorization) is enabled on the Partition.
- **Secure transport enabled:** Flags whether or not RPCs implement SSL.

Web Container Root Context

This section of the General properties tab displays the Web Container's root context.

Properties tab

The Properties tab displays information about JPDA remote debugging and the Server Connection Manager settings.

Virtual Machine

- **Enable JPDA remote debugging:** Flags whether or not debugging is enabled on the Partition.
- **JPDA debugging transport address:** The port used by the JPDA debugger to connect to the Partition.
- **Suspend partition until debugger attaches:** Flag that, when true, does not mark the Partition as “Running” until the debugger successfully attaches.

Server Connection Manager Settings

- **Server connection manager name:** The server connection manager servicing the Partition.
- **Listener port:** The port on which the server connection manager listens for incoming connections.
- **Connection Pool:** Shows the range of minimum and maximum allowable connection to the Partition.
- **Dispatcher Pool:** Shows the range of minimum and maximum allowable threads within the Partition. Additionally displays the Maximum Thread Idle time.

XML tab

The XML tab displays the XML data block defining the Partition's management properties. The data block is excerpted from the `configuration.xml` file.

Class Loading tab

The Class Loading tab displays the Partition's classloading policy and the classloader hierarchy.

Logs tab

The Logs tab displays the log4j, stderr, and stdout logs. Use the drop-down list in the upper left corner of the tab to view the contents of each log. The log4j log has filtering capabilities.

Status tab

If the Partition's statistics agent is enabled, you can view “real-time” performance information about your Partition on the Status tab. To enable statistics gathering see “Statistics properties”.

JDBC Pool States tab

If the Partition's statistics agent is enabled, you can view “real-time” performance information about the JDBC pool on the JDBC Pool States tab. To enable statistics gathering see “Statistics properties”.

Tuning the Partition for performance

The Partition has a Performance Tuning Wizard that allows you to set particular preferences about a Partition's operation. You can also configure some advanced properties for performance tuning.

To tune a Partition:

1. Select the Partition you want to tune from the Navigation Pane.
2. Right-click the Partition and select Performance Tuning. The Performance Tuning Wizard Step 1 appears.

   In this step you can select the usage model for statistics tuning. The choices here are reflected in the statistics gathering settings in the next Wizard step. Choose from among the following models:

   - **Developer:** Enables a maximum level of statistics gathering.
   - **System test:** Enables a medium level of statistics gathering.
   - **Production:** Enables a minimum level of statistics gathering.
   - **Best performance:** Selecting this option disables statistics gathering, improving the performance of your Partition.
   - **Custom:** Selecting this option means that your desired settings do not match any of the predefined usage models listed above. You can customize any of the other usage models in which case your usage model will be described as *custom*.

1. Click Next when you are finished. The Performance Tuning Wizard Step 2 appears.

   In this step you can determine how the Partition collects statistics about its operation. Statistics information is written to disk and displayed in the Partition's Statistics tab in the Content Pane. For more information on Partition statistics, see “Partition XML reference” in the *AppServer Developer's Guide*. For maximum performance, disable statistics gathering. In this step you can:

- **Enable Agent Statistics:** Check this box to enable the Partition's statistics agent. You can set the logging level using the Statistics level drop-down list. Additionally you can set the polling interval by entering a value in the Snapshot period field.
- **Enable Agent Statistics Reaping:** Check this box to periodically delete stored statistics information and preserve disk space, known as *reaping*. You determine how long to keep statistics by entering a value in Reap older than. You can set how often the statistics logs are reaped by entering a value in Reap period.

1 To continue, click Next. Step 3 appears.

Having unneeded or unused deployed modules enabled in a Partition adds to Partition start up time, increases it's memory footprint, and adds to the time taken by the garbage collector. Choose the modules you want to disable by selecting them in the list. You can *Shift-click* or *Ctrl-click* to select multiple modules.

2 To continue, click Next. Step 4 appears.

You can disable those Partition services not needed by your applications. To disable a Partition service, uncheck its check box. To enable it, check the check box.

3 To continue, click Next. Step 5 appears.

You can improve performance by tuning the Java VM parameters appropriately for your application. In this step, you can configure the following settings:

- **Select the JDK to be used by this partition:** Select the JDK from the list by clicking on it.
- **Heap and Thread Stack Sizes:** Enter a value in the Initial heap size field to modify the starting heap size. You can control the maximum heap size by entering a value in the Maximum heap size field. Use the Java thread stack size field to control how threads are managed within the VM.
- **Java VM Type:** Select the appropriate option. Valid values are either Server, Client, or Other (which you can select from the drop-down list).

In addition you can do some advanced configuration (see “Advanced performance tuning options”) or look at the performance tuning hints.

1 When you are finished, click Next. Step 6 appears.

You can set some of the VisiBroker parameters which affect performance. In this step, you can configure the following settings:

- **Connection Pool:** Set the maximum number of allowable connections and the maximum connection idle time in the appropriate dialog boxes.
- **Dispatcher Pool:** Set the range of allowable threads in the Minimum number of threads and Maximum number of threads fields. Set the Maximum thread idle time in its field.
- **Security:** Flags whether or not basic security (authentication/authorization) is enabled on the Partition.

1 To continue, click Next. The last step appears.

In this step you can tune some J2EE specific properties. You can configure the following settings:

- **EJB Container Tuning:** Set the Session passivation timeout field. Check the check box to use PBV for intra-bean calls.
- **Message-driven bean thread pool:** Set the thread pool range by setting values in the Minimum number of threads and Maximum number of threads fields. Set the Maximum thread idle time in its field.
- **Web Container Tuning:** Select the HTTP Service Connector from the drop down list, and set the range of processors and connection timeout settings.

1 When you are done, click Finish.

Advanced performance tuning options

To access the advanced options for Performance Tuning:

1. In Step 4 of the Performance Tuning Wizard, click Advanced Configuration.
2. The `partition_server.config` file opens in an editing window.
3. Edit the file as needed. When you are finished, click OK.

Dumping the Partition stack trace to a log file.

You can dump the Partition's stack trace to a log file for troubleshooting.

To dump the Partition's stack trace to a log file:

1. Select the Partition in the Navigation Pane.
2. Right-click the Partition and select Dump stack to log.
3. The stack trace is dumped to the stdout.log, located in:

```
<install_dir>/var/domains/<domain-name>/configurations/<configuration-name>/
    mos/<partition-name>/adm/logs/<partition_name>.stdout.log
```

Enabling a Partition to run with Optimizeit Profiler or ServerTrace

See "Using Optimizeit Profiler and ServerTrace" for information about configuring and using the Borland Management Console integration with Optimizeit ServerTrace and Optimizeit Profiler.

4

Web components

This section provides information about the web components which are included in the Borland AppServer (AppServer). For more information, see “Installing Borland AppServer on Windows” or “Installing Borland AppServer on Solaris and HP-UX” in the *Installation Guide*.

Apache web server implementation

The AppServer implementation of the open-source Apache web server version 2.2.3 (an httpd server) is HTTP 1.1-compliant and is highly customizable through the Apache modules.

Apache configuration

The Apache web server comes pre-configured and ready-to-use when it is initially started. Many modules are dynamically loaded during the Apache startup. You can later customize its configuration for the IIOP connector, clustering, failover, and load balancing with one or more web container(s). You can use the Management Console to modify the configuration file, or you can use the directives in the plain text configuration file: `httpd.conf`.

By default, the Apache `httpd.conf` file is located in the following directory:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
  mos/<apache_managedobject_name>/conf
```

Otherwise, for the location of the `httpd.conf` file, go to the `configuration.xml` file located in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>
```

and search for the Apache Managed Object, `apache-process` `httpd-conf` sub-element attribute:

```
httpd-conf=
```

For information about configuring the `httpd.conf` file for the IIOP connector/redirector, see, “Modifying the IIOP configuration in Apache” .

Apache configuration syntax

When you edit the `httpd.conf` file, you must adhere to the following configuration syntax guidelines:

- The `httpd.conf` files contain one directive per line.
- To indicate that a directive continues onto the next line, use a back-slash "\" as the last character on a line.
- No other characters or white space must appear between the back-slash "\" and the end of the line.
- Arguments to directives are often case-sensitive, but directives are not case-sensitive.
- Lines which begin with the hash character "#" are considered comments.
- Comments cannot be included on a line after a configuration directive.
- Blank lines and white space occurring before a directive are ignored, so you can indent directives for clarity.

Running Apache web server on a privileged port

Processes which access privileged ports on UNIX hosts must have appropriate permissions; the process must be started with the permissions of user `root`. Typically, only some processes in a particular configuration need to be started with root permissions. The `setuser` script helps configure AppServer to allow Apache web server start as root or with root permissions. See "Using the setuser tool to manage ownership" in the *Installation Guide* for information about using the `setuser` tool and multi-user mode.

Before starting the configuration, gather the following information about the system on which Apache is installed:

1. AppServer installation directory
2. User and group names for the account that the Agent should become after shedding its root UID, that is, the installation owner.
3. User and group names for the system root account (typically root/sys)

The following steps describe the procedure for configuring the Apache web server to run on port 80.

1. Ensure that the Management Hub and the configuration containing the Apache web server are not running.
2. Enable multi-user mode on the AppServer installation.
   a. Edit the property

      ```
      agent.mum.enabled.root.mo=true
      ```

      in

      ```
      <install_dir>/var/domains/base/adm/properties/agent.config
      ```

   b. Become `root`.
   c. Run the setuser script:

      ```
      setuser -u <user> -g <group> +m
      ```

      where <user> and <group> are the attributes of the installation owner account (as described in B above).
3. Start the Management Hub.
4. Edit the Apache web server properties in the Management Console.
   a. Right click the Apache web server MO, and select Properties.

b In the Properties dialog, select the Apache Process Settings tab.

c Click More settings to open the Advanced Process Settings dialog.

d Select the Platform Specific Settings tab.

e In the Unix Settings group enter the user and group names for the system root account in the Start as user and Start as group fields (as described in C above).

f Click OK to close the Advanced Process Settings dialog.

g In the Properties dialog, select the Files tab, and select the httpd.conf file.

h Change the User and Group directives to the user and group name values for the account that owns the AppServer installation (as described in B above).

i Change the Listen directive to 80.

j Click OK to close the Apache Properties dialog.

5 Start the configuration.

Using the .htaccess files

The Apache web server allows for decentralized management of configuration through the `.htaccess` files placed inside the web tree. These files are specified in the `AccessFileName` directive.

Directives placed in `.htaccess` files apply to the directory where you place the file, and all sub-directories. The `.htaccess` files follow the same syntax as the main configuration files. Since `.htaccess` files are read on every request, changes made in these files take immediate effect. To find which directives can be placed in `.htaccess` files, check the Context of the directive. You can control which directives can be placed in `.htaccess` files by configuring the `AllowOverride` directive in the main configuration files.

Apache directory structure

After installing the Apache web server, by default, the following Apache-specific directory structure appears in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
  mos/<apache_managedobject_name>/
```

Apache-specific Directory Name	Description
`cgi-bin`	Contains all CGI scripts.
`conf`	Contains all configuration files.
`error`	Contains all error html documents.
`htdocs`	Contains all HTML documents and web pages.
`icons`	Contains the icon images in `.gif` format.
`logs`	Contains all log files.
`proxy`	Contains the proxies for your web application.

Borland web container implementation

The Borland web container supports development and deployment of web applications. The Borland web container, which is based on Tomcat 5.5.17., is included in the AppServer. For more information, see “Installing Borland AppServer on Windows” or “Installing Borland AppServer on Solaris and HP-UX” in the *Installation Guide*.

The Borland web container is a sophisticated and flexible tool that provides support for Servlets 2.4 and JSP 2.0 specifications.

As a “Partition service”, all the Borland web container configuration files are located in each of your Partitions' data directory under:

```
adm/tomcat/conf/
```

By default, a Partition's data directory is located in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
  mos/<partition_name>/
```

For example, for a Partition named “standard”, by default the Borland web container configuration files are located in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
  mos/standard/adm/tomcat/conf/
```

Otherwise, for the location of a Partition data directory, go to the `configuration.xml` file located in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
```

and search for the Partition Managed Object, `partition-process` sub-element `directory` attribute:

```
<partition-process directory=
```

Servlets and JavaServer Pages

A *servlet* is a Java program that extends the functionality of a web server, generating dynamic content and interacting with web clients using a request-response paradigm.

JavaServer Pages (JSP) are a further abstraction to the servlet model. JSPs are an extensible web technology that uses template data, custom elements, scripting languages, and server-side Java objects to return dynamic content to a client. Typically the template data is HTML or XML elements, and in many cases the client is a web browser.

Servlets and JSPs are server components that normally run within a web server. Servlets are written as web server extensions separate from the HTML page, while JSP embeds the Java code directly in the HTML. At runtime, the JSP Java code is automatically converted into a servlet.

Servlets process web requests, pass them into the back-end enterprise application systems, and dynamically render the results as HTML or XML client interfaces. Servlets also manage the client session information, so that users do not need to repeatedly input the same information.

Typical web application development process

In a typical development phase for a web application:

1. The web designer writes the JSP components, and the software developer creates the servlets for handling presentation logic.
2. In conjunction, other software engineers write Java source code for servlets and the `.jsp` and `.html` for processing client request to the server-side components (EJB application tier, CORBA object, JDBC object).
3. The Java class files, `.jsp` files, and the `.html` files are bundled with a deployment descriptor as a Web ARchive (WAR) file.
4. The WAR file (or web module) is deployed in the Borland web container as a web application.

For more information about using the AppServer Deployment Descriptor Editor (DDE) to create a Web ARchive (WAR) file, see “Adding WAR information” in the *Management Console User's Guide*.

Web application archive (WAR) file

In order for the Borland web container to deploy a web application, the web application must be packaged into a Web ARchive (WAR) file. This is achieved by using the standard Java Archive tool `jar` command.

The WAR file includes the `WEB-INF` directory. This directory contains files that relate to the web application. Unlike the document root directory of the web application, the files in the `WEB-INF` directory do not have direct interaction with the client. The `WEB-INF` directory contains the following:

Directory/File name	Contents
`/WEB-INF/web.xml`	the deployment descriptor
`/WEB-INF/web-borland.xml`	the deployment descriptor with Borland-specific extensions.
`/WEB-INF/classes/*`	the servlets and utility classes. The application class loader loads any class in this directory.
`/WEB-INF/lib/*.jar`	the Java ARchive (JAR) files which contain servlets, beans, and other utility classes useful to the web application. All JAR files are used by the web application class loader to load classes from.

Borland-specific DTD

For the Borland-specific DTD information, see the DTD documentation.

Adding ENV variables for the web container

You add web container ENV variables for a Partition the same way you set any ENV variables for any Partition service; you use the `<env-vars>` element and insert the xml code within the `partition-process` sub-element.

Note

When adding web container ENV variables, be sure to type space-separated, value pairs.

By default, all `configuration.xml` files are located in the following directory:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
```

To add web container ENV variables for a Partition Managed Object, use the `env-vars` element and `env-var` sub-element and the following syntax:

```
<managed-object name="standard"> ...>
    <partition-process ...>
        <env-vars ...>
          <env-var name="name" value="value"/>
        </env-vars>
        ⋮
</managed-object>
```

where `<name>` is the ENV variable name and `<value>` is the value you want to set for the named ENV variable.

For example:

```
<managed-object name="standard"> ...>
    <partition-process ...>
        <env-vars ...>
          <env-var name="ABC" value="val_abc"/>
        </env-vars>
        ⋮
</managed-object>
```

Microsoft Internet Information Services (IIS) web server

The Microsoft Internet Information Services (IIS) web server is not included with any AppServer product offerings. However, AppServer does include the IIOP redirector which provides connectivity from the Borland Tomcat-based web container to the IIS web server, and from the IIS web server to a CORBA server. The IIOP redirector is supported for the following IIS versions:

- Microsoft Windows 2000/IIS version 5.0
- Microsoft Windows XP/IIS version 5.1
- Microsoft Windows 2003/IIS version 6.0

For more information, see “IIS web server to Borland web container connectivity”.

IIS/IIOP redirector directory structure

After installing any of the AppServer products, by default, the following IIS/IIOP redirector-specific directory structure appears in:

```
<install_dir>/etc/iisredir2/
```

IIS/IIOP redirector-specific directory name	Description
`conf`	Contains all configuration files.
`logs`	Contains all log files.

Smart Agent implementation

The Smart Agent is a service that helps in locating and mapping client programs and object implementation. The Smart Agent is automatically started with default properties. For information on configuring the Smart Agent, see “Using the Smart Agent” in the *VisiBroker for Java Developer's Guide*.

The Smart Agent is a dynamic, distributed directory service that provides facilities for both the client programs and object implementation. The Smart Agent maps client programs to the appropriate object implementation by correlating the object or service name used by the client program to bind to an object implementation. The object implementation is an object reference provided by a server, such as the Borland web container.

The Smart Agent must be started on at least one host within your local network. When your client program invokes an object (using the `bind` method), the Smart Agent is automatically consulted. The Smart Agent locates the specified object implementation so that a connection can be established between the client and the object implementation. The communication with the Smart Agent is transparent to the client program.

The following are examples of how the Smart Agent is used by the AppServer web components:

- Connecting Apache web server to a Borland web container.
- Connecting Borland web containers to Java Session Service (JSS).

Connecting an Apache web server to a Borland web container

As a distributed directory service, the Smart Agent registers an active ID of an object reference for the client programs to use. The following diagram shows the interaction between the client program binding to an object through the Smart Agent. In this example, the Apache web server is acting as a client and the Borland web container is acting as a server (and provides the object reference).

Figure 4.1 Client program binding to an object reference

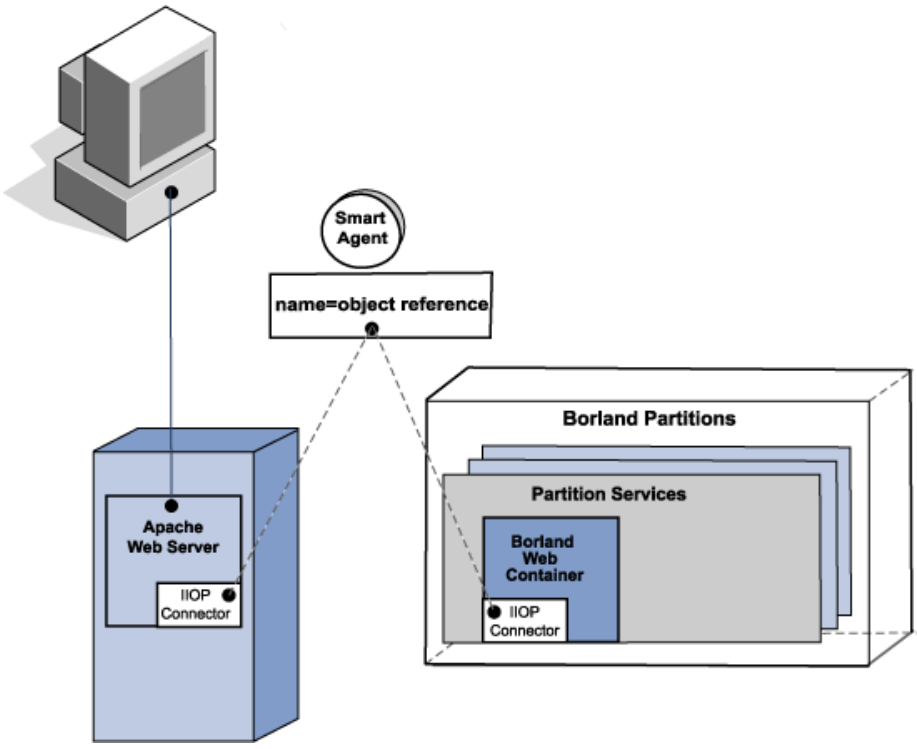

Connecting Borland web containers to Java Session Service

In this scenario, there are multiple web containers that need to connect to a Java Session Service during start up. The Smart Agent is used to make a client/server connection. The following diagram shows multiple instances of the Borland web container. Each web container is acting as a client. During start up, the Smart Agent is consulted as a directory service to find and connect a JSS object reference. For more information about the Java Session Service (JSS), see “Java Session Service (JSS) configuration.”

Figure 4.2 Connecting multiple web containers to a single JSS

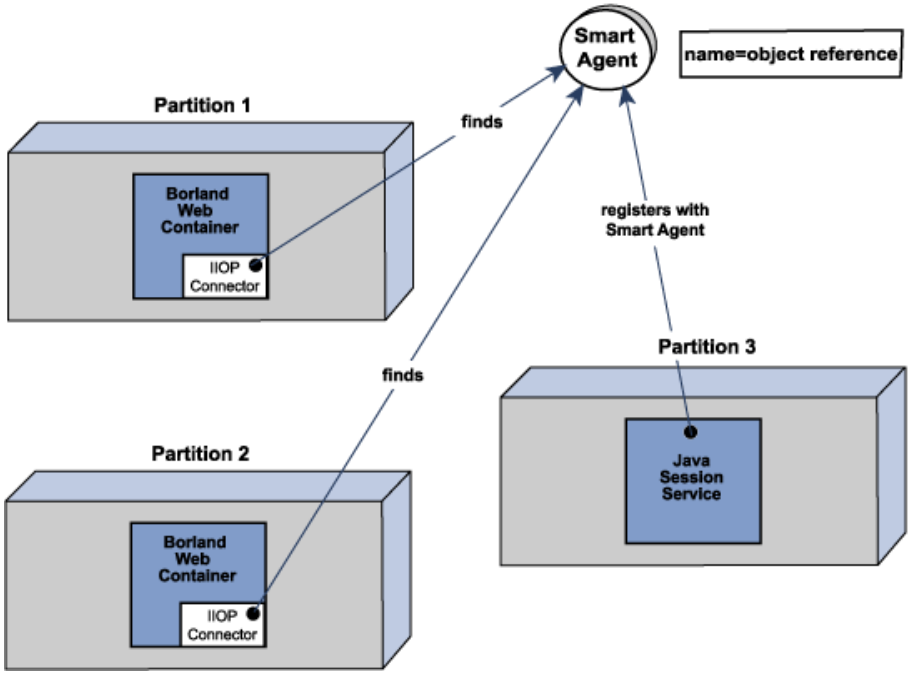

5

Web server to web container connectivity

This section describes the web server to web container IIOP connectivity provided in the Borland AppServer (AppServer). For more information, see “Installing Borland AppServer on Windows” or “Installing Borland AppServer on Solaris and HP-UX” in the *Installation Guide*.

For information about Apache to CORBA connectivity, see “Apache web server to CORBA server connectivity.”

Apache web server to Borland web container connectivity

This section provides information about the following web components:

- an implementation of the open-sourced Apache web server version 2.2.3,
- the Tomcat-based Borland web container, and
- the IIOP connector, which provides connectivity from the Apache web server to the Tomcat-based Borland web container.

These web components are included in the AppServer. For more information, see “Installing Borland AppServer on Windows” or “Installing Borland AppServer on Solaris and HP-UX” in the *Installation Guide*.

Modifying the Borland web container IIOP configuration

The `server.xml` is the main configuration file for the Borland web container and is stored in your Partition's data directory:

```
adm/tomcat/conf/
```

For more information, see “Borland web container implementation”.

Within the `server.xml` file are the following lines of code that pertain to the IIOP connector configuration.

```
<Connector className="com.borland.catalina.connector.iiop.IiopConnector"
 name="tc_inst1 debug="0" minProcessors="5"
 maxProcessors="75" enableChunking="false" port="0"
 canBufferHttp10Data="true" downloadBufferSize="4096" shortSessionId="false" />
```

Use these lines of code and the following attributes to configure the Borland web container IIOP connector.

Attribute	Default	Description
`name`	`tc_inst1`	The name by which this connector can be reached by Apache and IIS web servers.
`debug`	0 (zero)	Integer that sets the level of debug information. When set to `0` (zero), the default, debug is turned off. To turn debug on, set to `1`. For very detailed debug messages, set to `99`.
`minProcessors`	`5`	The number of minimum threads previously created to service requests on this connector.
`maxProcessors`	`75`	The number of maximum threads that will be created on this connector to service requests.
`enableChunking`	`false`	Enables chunking behavior on the connector. To enable chunking, set this attribute to `true`. **Important:** To enable chunking, you must also set the servlet response header `Transfer-Encoding` value to `chunked`. For more information, see “Downloading large data”.
`downloadBufferSize`	`4096`	Defines the “chunked” buffer size employed when `enableChunking` is set to `true`. This directive accepts a numeric value >`0`. Essentially, the larger the number of bytes you set this directive to, the less the number of CORBA RPCs that are required to send the data to Apache or IIS. However, the larger you set this directive, the more memory will be consumed in servicing the transaction. Tuning this parameter allows you to fine-tune the performance characteristics. This enables the administrator to weigh the RPC costs against memory resource usage to optimize uploading on their system. **Note:** If an invalid value is presented (non numeric/negative number) then the default `4096` value is employed. For more information, see “Downloading large data”.
`port`	`0` (zero)	The IIOP connector port. If set to `0` (zero), the default, a random port gets picked. **Note:** If the corbaloc mechanism must be used to locate this connector from Apache or IIS, then `port` must be set to a value other than `0` (zero).

Attribute	Default	Description
canBufferHttp10Data	true	When the HTTP protocol is less than 1.1 and the content length is not set on a servlet, the following two choices are available for the web container. It can buffer up the data, compute the content length, and then send the response or it can raise an error message. To avoid buffering the data and consuming memory, set this attribute to false.For more information, see "Browsers supporting only the HTTP 1.0 protocol".
shortSessionId	false	To enable the use of the Borland "collapsed locator" feature that reduces the size of the IIOP connectors session id, set to true. By default, this attribute is disabled; set to false. The IIOP connector uses a stringified object reference (IOR) as a means to associate a client with a given Borland web container (service affinity). Some network routers and browser implementations are unable to cope with the large payload stored within the Session Id Cookie. To resolve this problem, Borland has developed the collapsed locator string that allows service affinity to be implemented using a much shorter session id string. Setting shortSessionId to true enables this feature, setting it to false or omitting the parameter entirely, defaults to the conventional IOR-based solution. **Note:** The "collapsed locator" is essentially an encoded CORBALOC string. Resolving a CORBALOC string to an object-reference is a more costly operation than the equivalent operation on an IOR. To alleviate this extra cost, the Apache web server will maintain a look-aside list of encoded CORBALOC strings to their resolved object references. This will incur a slight increase in memory use per Apache web server. **Important:** In order for the shortSessionID feature to work correctly, you **must** specify a port value for the IIOP connector; you cannot use the port default value of 0 (zero).

Modifying the IIOP configuration in Apache

The httpd.conf file is the global configuration file for the Apache web server. Within the httpd.conf file are the following lines which pertain to the IIOP connector.

Windows

```
LoadModule iiop2_module <install_dir>/lib/<apache_managedobject_name>/
     mod_iiop2.dll
   IIopLogFile <install_dir>/var/domains/<domain_name>/
configurations/<configuration_name>/mos/<apache_managedobject_name>/
logs/mod_iiop.log
   IIopLogLevel error
   IIopClusterConfig <install_dir>/var/domains/<domain_name>/
configurations/<configuration_name>/mos/<apache_managedobject_name>/
conf/WebClusters.properties
   IIopMapFile <install_dir>/var/domains/<domain_name>/
configurations/<configuration_name>/mos/<apache_managedobject_name>/
conf/UriMapFile.properties
```

Use these lines of code to configure the Apache web server IIOP connector.

Directive	default	Description
LoadModule	<install_dir>/lib/<apache_managedobject_name>/mod_iiop2.dll	Enables Apache 2.2.3 to load the IIOP connector. This directive instructs the Apache web server to load the Apache mod_iiop2 module from the location specified. Once the module is loaded, the following four directives are required to enable the IIOP connector to locate the web container(s) or CORBA server(s) it must communicate with and perform other functions.
IIopLogFile	<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/mos/<apache_managedobject_name>/logs/mod_iiop.log	Specifies the location where the IIOP connector writes log output.
IIopLogLevel	error	Specifies the level of log information to write. This directive can take one of the following: debug|warn|info|error.
IIopClusterConfig	<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/mos/<apache_managedobject_name>/conf/WebClusters.properties	Specifies the location of the "cluster" instance file. For CORBA servers, identifies the file that contains the "cluster" name by which they are known to the IIOP connector[1].
IIopMapFile	<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/mos/<apache_managedobject_name>/conf/UriMapFile.properties	Specifies the location of the URI-to-Instance mapping file. For CORBA servers, maps HTTP URIs to a specific "cluster" known to the IIOP connector.

[1] "cluster" is used to represent a CORBA server instance that is known to the system by a single name or URI. The IIOP connector is able to load-balance across multiple instances, hence the term "cluster" is used.

The following are examples of typical configurations of these 5 lines for the IIOP connector for Apache 2.2.3:

Windows example

```
  LoadModule iiop2_module C:/Borland/BDP/lib/myapache/mod_iiop2.dll
  IIopLogFile C:/Borland/BDP/var/domains/base/configurations/j2ee/mos/
myapache/logs/mod_iiop.log
  IIopLogLevel error
  IIopClusterConfig C:/Borland/BDP/var/domains/base/configurations/j2ee/
mos/myapache/conf/WebClusters.properties
  IIopMapFile C:/Borland/BDP/var/domains/base/configurations/j2ee/mos/
myapache/conf/UriMapFile.properties
```

Solaris example

```
  LoadModule iiop2_module /opt/Borland/BDP/lib/myapache/mod_iiop2.so
  IIopLogFile /opt/Borland/BDP/var/domains/base/configurations/j2ee/mos/
myapache/logs/mod_iiop.log
  IIopLogLevel error
  IIopClusterConfig /opt/Borland/BDP/var/domains/base/configurations/j2ee/
mos/myapache/conf/WebClusters.properties
  IIopMapFile /opt/Borland/BDP/var/domains/base/configurations/j2ee/mos/
myapache/conf/UriMapFile.properties
```

Additional Apache IIOP directives

The following optional additional directives are available for you to use to further customize your Apache IIOP configuration.

Directive	default	Description
`IIopChunkedUploading`	(commented out) `true`	Controls whether Apache attempts “chunked” uploads to the Borland web container IIOP connector. To enable Apache to “chunk” large size data that is greater than the `IIopUploadBufferSize` value, uncomment and make sure it is set to `true`. **Note:** “Chunked” upload must also be enabled on the web container by setting the `server.xml` attribute `enablechunking="true"`. If you want Apache to wait until it has collected all data before invoking the CORBA RPC to send the large size data to the Borland web container, leave commented out or set to `false`. For more information, see “Implementing chunked upload”.
`IIopUploadBufferSize`	(commented out) `4096`	Defines the “chunked” buffer size employed when `IIopChunkedUploading` is set to `true`. This directive accepts a numeric value >0. Essentially, the larger the number of bytes you set this directive to, the less the number of CORBA RPCs that are required to send the data to the Borland web container. However, the larger you set this directive, the more memory will be consumed in servicing the transaction. Tuning this parameter allows you to fine-tune the performance characteristics. This enables the administrator to weigh the RPC costs against memory resource usage to optimize uploading on their system. **Note:** If an invalid value is presented (non numeric/negative number) then the default `4096` value is employed. For more information, see “Implementing chunked upload”.

Apache IIOP connector configuration

The Apache IIOP connector has a set of configuration files that you must update with web server cluster information. By default, these IIOP connector configuration files are located in:

```
<install_dir>/var/domains/<domain_name>/configurations/
<configuration_name>/mos/<apache_managedobject_name>/conf
```

The two configuration files are:

IIOP configuration file	Description
WebClusters.properties	Specifies the cluster(s) and the corresponding web container(s) for each cluster.
UriMapFile.properties	Maps URI references to the clusters defined in the `WebClusters.properties` file.

Note

Modifying either of these configuration files can be done so without starting up or shutting down the Apache web server(s) or CORBA server(s) because the file is automatically loaded by the IIOP connector.

Adding new clusters

The `WebClusters.properties` file tells the IIOP connector:

- The name of each available cluster (`ClusterList`).
- The web container identification.
- Whether to provide automatic load balancing (`enable_loadbalancing`) for a particular cluster.

To add a new cluster, in the `WebClusters.properties` file:

1. add the name of the configured cluster to the `ClusterList`. For example:

```
ClusterList=cluster1,cluster2,cluster3
```

2. define each cluster by adding a line in the following format specifying the cluster name, the required `webcontainer_id` attribute, and any additional attributes (see the following Cluster definition attributes table). For example:

```
<clustername>.webcontainer_id = <id> <attribute>
```

Note

Failover and smart session are always enabled, for more information see "Clustering web components".

Attribute	Required	Definition
`webcontainer_id`	yes	the object "bind" name or corbaloc string identifying the web container implementing the cluster.
`enable_loadbalancing`	no	To enable load balancing, do not include this attribute **or** include and set to `true`; load balancing is enabled by default. To disable load balancing, set to `false` indicating that this cluster instance should **not** employ load-balancing techniques. **Warning:** Ensure that when entering the `enable_loadbalancing` attribute you give it a legal value (`true` or `false`).

For example:

```
   ClusterList=cluster1,cluster2,cluster3
cluster1.webcontainer_id = tc_inst1
cluster2.webcontainer_id = corbaloc::127.20.20.2:20202,:127.20.20.3:20202/
        tc_inst2
   cluster2.enable_loadbalancing = true
   cluster3.webcontainer_id = tc_inst3
   cluster3.enable_loadbalancing = false
```

In the above example, the following three clusters are defined:

1. The first, uses the osagent naming scheme and is enabled for load balancing.
2. The second cluster employs the corbaloc naming scheme, and is also enabled for load balancing.
3. The third uses the osagent naming scheme, but has the load balancing features disabled.

Note

To disable use of a particular cluster, simply remove the cluster name from the `ClusterList` list. However, we recommend you do not remove clusters with active http sessions attached to the web server (attached users), because requests to these "live" sessions will fail.

Note

Modifications you make to the `WebClusters.properties` file automatically take effect on the next request. You do not need to restart your server(s).

Adding new web applications

Important

By default, your web application is not made available through Apache. In order to make it available through Apache, you must add some information to the web application descriptor. For step-by-step instructions on how to do so, see “Web Deploy Paths” in the *Management Console User's Guide*.

For new applications that you have deployed to the Borland web container, you need to do the following to make them available through the Apache web server. Use the `UriMapFile.properties` file to map `HTTP URI` strings to web cluster names configured in the `WebClusters.properties` file (see “Adding new clusters”).

- In the `UriMapFile.properties` file, type:

  ```
  <uri-mapping> = <clustername>
  ```

  where `<uri-mapping>` is a standard URI string or a wild-card string, and `<clustername>` is the cluster name as it appears in the `ClusterList` entry in the `WebClusters.properties` file.

For example:

```
/examples = cluster1
/examples/* = cluster1

/petstore/index.jsp = cluster2
/petstore/servlet/* = cluster2
```

In this example:

- Any URI that starts with `/examples` will be forwarded to a web container running in the “cluster1” web cluster.
- URIs matching either `/petstore/index.jsp` or starting with `/petstore/servlet` will be routed to “cluster2”.

Note

With the URI mappings, the wild-card “*” is only valid in the **last** term of the URI and may represent the follow cases:

- the whole term (and all inferior references) as in `/examples/*`.
- the filename part of a file specification as in `/examples/*.jsp`.

Note

Modifications you make to the `UriMapFile.properties` file automatically take effect on the next request. You do not need to restart your server(s).

If the `WebCluster.properties` or `UriMapFile.properties` is altered, then it is automatically loaded by the IIOP connector. This means that the adding and removing of web applications and the altering of cluster configurations may be done so without starting up or shutting down the Apache web server(s) or Borland web container(s).

Large data transfer

This section details the AppServer options available to you for handling large data transfers between a client and the Borland web container with Apache 2.2.3 in between. The data to be transferred may be either:

- static content obtained from a file, or
- dynamically generated content

Typically, the content length is known in advance for static content, but is not known for dynamic content.

Downloading large data

The following modes are available for downloading large data from the Borland web container to the browser:

- Chunked download
- Non-chunked download

Implementing chunked download

In the *chunked* download mode, the Borland web container does not wait until it has all the data to send. As soon as servlet generates the data, the web container starts sending the data to the browser via Apache in fixed size buffers.

Because the data is flushed as soon as it is available, the chunked download mode of transfer has low memory requirements both on Apache and the Borland web container. The browser user sees data as it arrives rather than as one large lump at the end of the full transfer.

Enabling chunked download

To enable chunked download mode, you update the Borland web container `server.xml` file which is stored in your Partition's data directory:

```
adm/tomcat/conf/
```

For more information, see “Borland web container implementation”.

To enable the chunked download:

1. In the Borland web container `server.xml`, locate the `<Service name="IIOP">` section of the code.
2. For the IIOP service you want to enable chunked download, set `enableChunking="true"`
3. By default, the download buffer size is set to `4096`. To change it for an IIOP service, use the `downloadBufferSize` attribute as follows:

   ```
   downloadBufferSize=<value>
   ```

   Where `<value>` is a numeric value >`0`.

   Note

   If an invalid value is presented (non numeric/negative number) then the default `4096` value is employed.

The chunked download mode of transfer has an overhead of an extra thread per each request.

Known content length versus unknown

Based on whether content length is known in advance or not, chunked download mode can take one of the following two paths:

- chunked download with known content length
- chunked download with unknown content length

Chunked download with known content length

In this case, a servlet or JSP knows the content length of the data in advance of the transfer. The servlet sets the `Content-Length` HTTP header before writing out the data.

The Borland web container writes out a single response header followed by multiple chunks of data. Apache receives this from the web container and sends data in the same fashion to the browser.

The response header contains the following header:

```
Content-Length=<actual data size>
```

Chunked download with unknown content length

HTTP protocol version 1.1 adds a new feature to handle the case of data transfer when data length is **not** known in advance. This feature is called *HTTP chunking*. In this case, a servlet does not know the content length of the data in advance of a transfer. The servlet does **not** set the `Content-Length` HTTP header.

The Borland web container sends the data to the Apache web server in exactly the same way as in the case of the chunked download where the content length is known in advance; a single response header is sent followed by multiple data chunks. The response header contains the following header:

```
Transfer-Encoding="chunked"
```

If the browser protocol is HTTP 1.1 and the `Content-Length` header is not set by the servlet, the Borland web container automatically adds the `Transfer-Encoding="chunked"` header.

When an Apache web server sees this `Transfer-Encoding` header, it starts sending the data as “HTTP chunks”—a response header followed by multiple combinations of “chunked” header, “chunked” data, and “chunked” trailers.

Note

Per the HTTP 1.1 specification, if a servlet sets both the `Content-Length` and `Transfer-Encoding` headers, the `Content-Length` header is dropped by the Borland web container.

Browsers supporting only the HTTP 1.0 protocol

If the browser only supports the HTTP 1.0 protocol or less and a servlet does not set the `Content-Length` header, the Borland web container can not automatically add the `Transfer-Encoding` header. The reason being that to the HTTP 1.0 protocol, the `Transfer-Encoding` header has no meaning. In this case, the Borland web container:

1. buffers all the data until the data is finished,
2. calculates the content length, and
3. sets the `Content-Length` header itself.

If you do not want the Borland web container to perform this buffering behavior, you can set the IIOP connector attribute `canBufferHttp10Data="false"`. By default, this attribute is set to `true`.

Note

When the `canBufferHttp10Data` attribute is set to `false`, the following error message is sent to the browser:

```
Servlet did not set the Content-Length
```

Implementing non-chunked download

This is the default transfer of data mode for the IIOP connector. In the *non-chunked* download mode, the Borland web container waits until it has all the data to send. Then it calculates the content length and sets the `Content-Length` header to the actual content length. The Borland web container then sends the response header followed by a single huge data block.

This mode of transfer has high memory requirements both on the Apache web server and the Borland web container, because the data is cached until all of it is available. Only when all the data is transferred does the browser user see the data.

The non-chunked download mode of transfer has no overhead of extra thread per each request. This download mode works well under both the HTTP protocol versions 1.0 and 1.1, because the `Transfer-Encoding` header is never set in this mode.

Uploading large data

The following modes are available for uploading large data initiated by a client (which can be either a browser or a non-browser (such as Java) client that speaks HTTP):

- Chunked upload
- Non-chunked upload

The browser always sends the data to an Apache web server in a “chunked” fashion. *Chunked* and *non-chunked* upload refers to the data transfer mode between an Apache web server and a Borland web container.

Implementing chunked upload

By default, Apache will try to upload large size data in “chunks”. In this mode, Apache does not wait until it has all the data from the browser before it starts sending data to a Borland web container. Apache sends the data in fixed size buffers as the data becomes available from the browser.

Because the data is flushed as soon as possible, the chunked mode of upload transfer has low memory requirements both on Apache and the Borland web container.

The chunked mode of transfer has an overhead of an extra thread per each request on the Borland web container.

Enabling chunked upload

To enable chunked upload mode, you must update **both** of the following:

- the Borland web container `server.xml` file, which is stored in your Partition's data directory: `adm/tomcat/conf`

  For more information see “Borland web container implementation”.

- the Apache `httpd.conf` file, which by default is located in the following directory:

  ```
  <install_dir>/var/domains/<domain_name>/configurations/
  <configuration_name>/mos/<apache_managedobject_name>/conf
  ```

  For more information see “Apache configuration”.

To enable the chunked upload:

1. In the Borland web container `server.xml`, locate the `<Service name="IIOP">` section of the code.
2. By default, the `enableChunking` attribute is set to `false`. Change this value to `enableChunking="true"`
3. In the Apache `httpd.conf` file, locate and uncomment the following IIOP directive:

   ```
   #IIopChunkedUploading true
   ```

Note

The chunked upload mode of transfer has an overhead of an extra thread per each request for the Borland web container.

Changing the upload buffer size

By default, `IIopUploadBufferSize` is set to `4096` bytes. To change this value:

1. In the Apache `httpd.conf`, locate the following commented out directive:

   ```
   #IIopUploadBufferSize 4096
   ```

2. Uncomment this directive and set as follows:

```
IIopUploadBufferSize <value>
```

where <value> is a numeric value >0 (greater than zero).

Note

If you specify an invalid value (non numeric/negative number) then the default 4096 value is employed.

Known content length versus unknown

Based on whether content length is known in advance or not, chunked upload mode can take one of the following two paths:

- chunked upload with known content length
- chunked upload with unknown content length

Chunked upload with known content length

In this case, the client knows the content length of the data in advance of the transfer. The client sets the Content-Length HTTP header before writing out the data. The client writes out a single response header followed by multiple chunks of data. Apache receives this from the browser and sends data in the same fashion to the Borland web container.

The response header contains the following header:

```
Content-Length=<actual data size>
```

Chunked upload with unknown content length

HTTP protocol version 1.1 adds a new feature to handle the case of data transfer when data length is **not** known in advance. This feature is called *HTTP chunking*.

In this case, a client does not know the content length of the data in advance of a transfer. The client does **not** set the Content-Length HTTP request header. Instead, the client sets the Transfer-Encoding HTTP request header to a value of chunked as follows:

```
Transfer-Encoding="chunked"
```

The client sends the data to the Apache web server as “HTTP chunks”; a single request header followed by multiple combinations of *chunk header*, *chunk data*, and *chunk trailer*.

When the Apache web server sees this Transfer-Encoding header, it strips out the chunk header and chunk trailers and sends the data as normal data chunks to the Borland web container.

At this time, no major browsers support uploading data without knowing the content length. In other words, browsers never add a Transfer-Encoding="chunked″ header to an HTTP request. However, a non-browser client can add this header to an HTTP request.

Implementing non-chunked upload

This is the default transfer of data mode for the IIOP connector. In the *non-chunked* upload mode, the Apache web server waits until it has all the data to send. Then it calculates the content length and sets the Content-Length header to the actual content length. Apache then sends the request header followed by a single huge data block.

This mode of transfer has high memory requirements both on the Apache web server and the Borland web container, because the data is cached until all of it is available.

The non-chunked upload mode of transfer has no overhead of extra thread per each request (in the Borland web container). This download mode works well under both the HTTP protocol versions 1.0 and 1.1, because the Transfer-Encoding header is never set in this mode.

IIS web server to Borland web container connectivity

This section describes the Tomcat-based Borland web container, its IIOP connector, as well as the IIS/IIOP redirector which provides connectivity from the Microsoft Internet Information Services (IIS) web server (not included with AppServer) to the Borland web container. These features are provided in AppServer. For more information, see “Installing Borland AppServer on Windows” or “Installing Borland AppServer on Solaris and HP-UX” in the *Installation Guide*.

Modifying the IIOP configuration in the Borland web container

The `server.xml` is the main configuration file for the Borland web container and is stored in your Partition's data directory:

```
adm/tomcat/conf/
```

Within the `server.xml` file is a section that pertains to the IIOP connector configuration. For detailed configuration information, see “Modifying the Borland web container IIOP configuration”.

Microsoft Internet Information Services (IIS) server-specific IIOP configuration

Before the IIS/IIOP redirector can be used on your system, you need to complete the following:

- Configure your Windows 2003/XP/2000 system on which IIS is running
- Configure the IIS/IIOP redirector

How to Configure your Windows 2003/XP/2000 system on which IIS is running

1 Add the `OSAGENT_PORT` required environment variable to the SYSTEM environment.

The IISredirectory relies on VisiBroker to provide the IIOP communication layer between IIS and the Borland web container. In order to function, VisiBroker requires that the following environment variable be defined as follows:

Environment Variable	Value	Description
`OSAGENT_PORT`	`14000`	The numeric value of the OSAGENT port number used by your AppServer.

Important

After setting the `OSAGENT_PORT` environment variable, in order for IIS to detect it, you must reboot the Windows system.

2 Add the IIS/IIOP redirector as an ISAPI filter.

a Right-click My Computer and choose Manage.

The Computer Management dialog appears.

b Expand the tree, expand the Services and Applications node.

c Expand the Internet Information Services node.

d Right-click the Default Web Site node and choose Properties.

The Default Web Site Properties dialog appears.

e Go to the ISAPI Filters tab.

f Click Add.

g In the Filter Properties dialog, type a Filter Name and the path for the Executable in the corresponding entry boxes.

By convention, the name of the filter should reflect its task, for example:

```
iisredir2
```

Also, the executable should point to the `iisredir2.dll` in the `<install_dir>\bin`. For example:

```
C:\borland\BDP\bin\iisredir2.dll
```

h Click OK.

Your new ISAPI filter appears on the list. You do not need to change the filter Priority.

i Click OK.

3 Add a “borland” virtual directory to your IIS web site.

a In the Computer Management dialog, right-click Default Web Site and choose New|Virtual Directory.

The Virtual Directory Creation Wizard appears.

b Click Next.

c For the Alias, enter “borland”.

The `borland` virtual directory is required to allow the IIS/IIOP redirector extension to be located by the IIS web server when it responds to a URI of:

```
http://localhost/borland/iisredir2.dll
```

d For the Directory, browse to `<install_dir>\bin`.

e Click Next to proceed.

f For Access Permissions, select “Execute” in addition to “Read” and “Run scripts” which are selected by default.

g Click Next.

h Click Finish.

4 **Windows 2003 only:** Configure ISAPI extension permissions for Windows 2003.

The IIS version limits which application extensions may be loaded into IIS. You have the choice of enabling all extensions or selectively picking which ISAPI extensions that may be run in your IIS installation. The following procedure enables just the `iisredir2` extension.

a In the Computer Management dialog, open “Services and Applications”.

b Open “Internet Information Service”.

c Open “Web Service Extensions” and click Add a new Service Extension.

d Name the extension “iisredir2.dll”.

e Browse (using the add button) to <install>\bin\iisredir2.dll

f Select this file.

g Check the Extension allowed checkbox.

h Click OK.

5 Restart IIS by stopping **then** starting the IIS Service:

a In the Computer Management dialog, right-click the Internet Information Services node and choose Restart IIS.

b In the Stop/Start/Reboot dialog, from the drop-down choose “Stop Internet Services on <name of your IIS web server>"

c Click OK.

The web service unloads any dlls loaded by the IIS administrator.

d After shut down of the server is complete, right-click the Internet Information Services node and choose Restart IIS.

e In the Stop/Start/Reboot dialog, choose “Start Internet Services on <your IIS web server name>”.

f Click OK.

The web service reloads any dlls loaded by the IIS administrator.

6 Make sure the `iisredir2` filter is active.

a In the Computer Management dialog, right-click the Default Web Site node and choose Properties.

b In the Default Web Site Properties dialog, go to the ISAPI Filters tab.

c The `iisredir2` filter should be marked with a green up-pointing arrow indicating that it has been activated.

If not, then check the `iisredir2.log` file for details of why it may not have loaded correctly. This file can be found in:

```
<install_dir>\etc\iisredir2\logs
```

d To exit, click OK.

7 Attempt to access the `\examples` context via the IIS web-server.

If you have followed the previous steps, the `\examples` context should be accessible following a restart of your IIS Server.

Note

In the example the port number of the web server should match that configured for your site. For instance, if your IIS administrator has configured IIS to listen on port `6060`, then a valid URL is:

```
http://localhost:6060/examples
```

Of course, if your IIS is configured as per Microsoft defaults, then it listens on port `80`, in which case you may dispense with a port number. For example:

```
http://localhost/examples
```

IIS/IIOP redirector configuration

The IIS/IIOP redirector has a set of configuration files that you must update with web server cluster information. By default, these IIOP redirector configuration files are located in the following directory:

```
<install_dir>/etc/iisredir2/conf
```

The configuration files are:

IIOP configuration file	Description
WebClusters.properties	Specifies the cluster(s) and the corresponding web container(s) for each cluster.
UriMapFile.properties	Maps URI references to the clusters defined in the `WebClusters.properties` file.

Note

Modifying either of these configuration files can be done so without starting up or shutting down the IIS web server(s) or Borland web container(s) because the file is automatically loaded by the IIOP redirector.

Adding new clusters

The `WebClusters.properties` file tells the IIOP redirector:

- the name of each available cluster: (`ClusterList`).
- the web container identification.
- whether to provide automatic load balancing (`enable_loadbalancing`) for a particular cluster.

To add a new cluster, in the `WebClusters.properties` file:

1 add the name of the configured cluster to the `ClusterList`. For example:

```
ClusterList=cluster1,cluster2,cluster3
```

2 define each cluster by adding a line in the following format specifying the cluster name, the required `webcontainer_id` attribute, and any additional attributes (see the following Cluster definition attributes table). For example:

```
<clustername>.webcontainer_id = <id> <attribute>
```

Note

Failover and smart session are always enabled, for more information see “Clustering web components”.

Attribute	Required	Definition
`webcontainer_id`	yes	the object “bind” name or corbaloc string identifying the web container(s) implementing the cluster.
`enable_loadbalancing = true|false`	no	To enable load balancing, do not include this attribute **or** include and set to `true`; load balancing is enabled by default. To disable load balancing, set to `false` indicating that this cluster instance should **not** employ load-balancing techniques. **Warning:** Ensure that when entering the `enable_loadbalancing` attribute you give it a legal value (`true` or `false`).

For example:

```
   ClusterList=cluster1,cluster2,cluster3
cluster1.webcontainer_id = tc_inst1
cluster2.webcontainer_id = corbaloc::127.20.20.2:20202,:127.20.20.3:20202/
     tc_inst2
   cluster2.enable_loadbalancing = true
   cluster3.webcontainer_id = tc_inst3
   cluster3.enable_loadbalancing = false
```

In the above example, the following three clusters are defined:

1 The first, uses the osagent naming scheme and is enabled for load balancing.

2 The second cluster employs the corbaloc naming scheme, and is also enabled for load balancing.

3 The third uses the osagent naming scheme, but has the load balancing features disabled.

Note

To disable use of a particular cluster, simply remove the cluster name from the `ClusterList` list. However, we recommend you do not remove clusters with active http sessions attached to the web server (attached users), because requests to these “live” sessions will fail.

Note

Modifications you make to the `WebClusters.properties` file automatically take effect on the next request. You do not need to restart your server(s).

Adding new web applications

Important

By default, your web applications are not made available through IIS. In order to make a web application available through IIS, you must add some information to the web application descriptor. For step-by-step instructions on how to do so, see “Web Deploy Paths” in the *Management Console User's Guide*.

The `\examples` context is useful for verifying your IIS/IIOP installation configuration, however, for new applications that you have deployed to the Borland web container, you need to do the following to make them available through the IIS web server. Use the `UriMapFile.properties` file to map `HTTP URI` strings to web cluster names configured in the `WebClusters.properties` file (see “Adding new clusters”).

– In the `UriMapFile.properties` file, type:

```
<uri-mapping> = <clustername>
```

where `<uri-mapping>` is a standard URI string or a wild-card string, and `<clustername>` is the cluster name as it appears in the `ClusterList` entry in the `WebClusters.properties` file.

For example:

```
/examples = cluster1
/examples/* = cluster1

/petstore/index.jsp = cluster2
/petstore/servlet/* = cluster2
```

In this example:

– Any URI that starts with `/examples` will be forwarded to a web container running in the “cluster1” web cluster.

– URIs matching either `/petstore/index.jsp` or starting with `/petstore/servlet` will be routed to “cluster2”.

Note

With the URI mappings, the wild-card “*” is only valid in the last term of the URI and may represent the follow cases:

– the whole term (and all inferior references) as in `/examples/*`.

– the filename part of a file specification as in `/examples/*.jsp`.

Note

Modifications you make to the `UriMapFile.properties` file automatically take effect on the next request. You do not need to restart your server(s).

If the `WebCluster.properties` or `UriMapFile.properties` is altered, then it is automatically loaded by the IIOP redirector. This means that the adding and removing of web applications and the altering of cluster configurations may be done so without starting up or shutting down the IIS web server(s) or Borland web container(s).

6

Java Session Service (JSS) configuration

The Java Session Service (JSS) is a service that stores information pertaining to a specific user session. JSS is used to store session information for recovery in case of container failure.

Borland provides an Interface Definition Language (IDL) interface for the use of JSS. Two implementations are bundled, one using DataExpress and another with any JDBC capable database.

JSS provides a mechanism to easily store session information in a database. For example, in a shopping cart scenario, information about your session (the number of items in the shopping cart, and such) is stored by the JSS. So, if a session is interrupted by a Borland web container unexpectedly going down, the session information is recoverable by another Borland web container instance through the JSS. The JSS must be running on the local network. Any web container (within the cluster configuration) finds the JSS and connects to it and continues session management.

For more information about the Borland web container, see “Borland web container implementation” .

Session management with JSS

The following diagrams show typical landscapes of web components and how session information is managed by the JSS. The JSS session management is completely transparent to the client.

In the JSS Management with a Centralized JSS and Two Web Containers diagram, there are four virtual machines:

- The first machine hosts the Apache web server,
- two other machines contain an instance of the Borland web container,
- and the fourth machine hosts the JSS and relational database (JDataStore or a JDBC datasource).

If an interruption occurs between the Apache web server (Machine 1) which is passing a client request to the first web container instance (Machine 2), then the second web container instance (Machine 3) can continue processing the client request by retrieving the session information from the JSS (Machine 4). The items in the Shopping Cart are retained and the client request continues to be processed.

Figure 6.1 JSS Management with a Centralized JSS and Two Web Containers

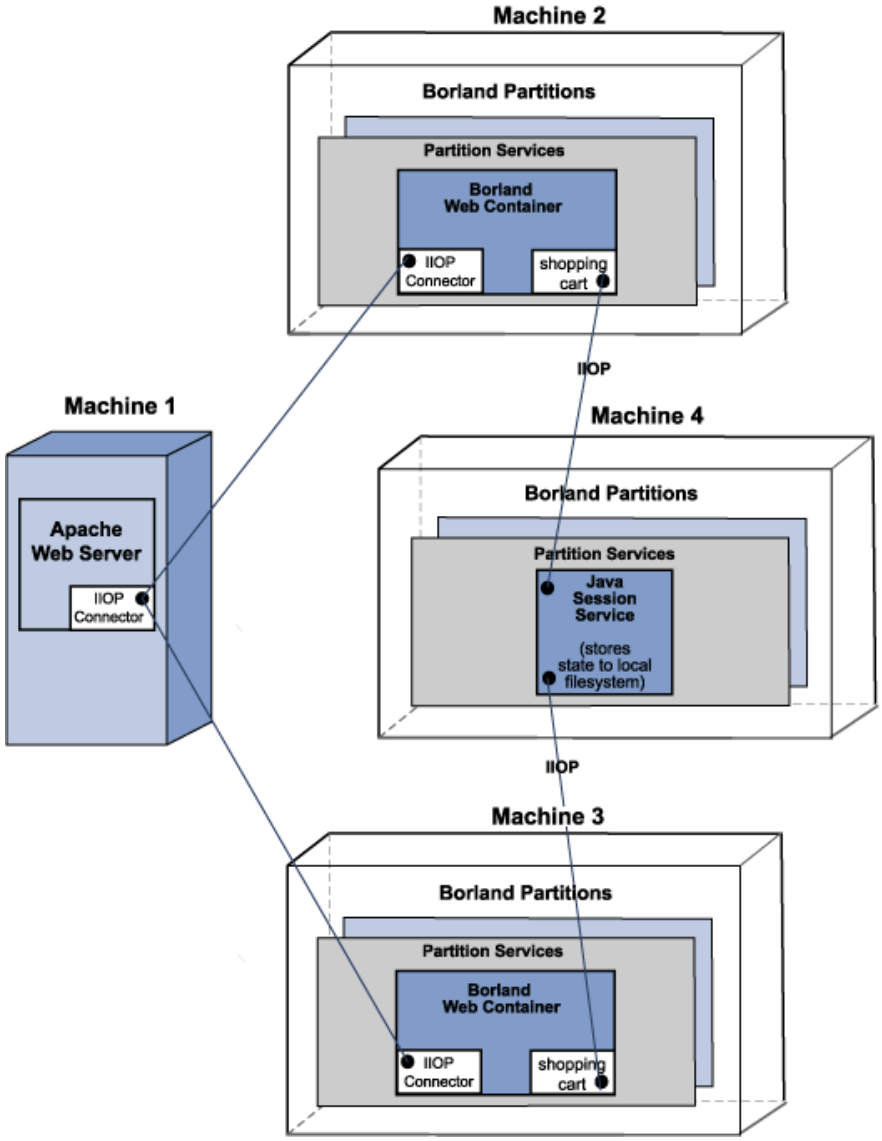

In the JSS Management with Two Web Containers and a Centralized Backend Datastore diagram, are the following four virtual machines:

- The first machine hosts the Apache web server,
- the two other machines contain an instance of the Borland web container as well as each hosting the JSS,
- and the fourth machine hosts the relational database (JDataStore or a JDBC datasource).

If an interruption occurs between the Apache web server (Machine 1) which is passing a client request to the first web container instance (Machine 2), then the second web container instance (Machine 3) can continue processing the client request by retrieving the session information from the JSS (Machine 4). The items in the Shopping Cart are retained and the client request continues to be processed.

Figure 6.2 JSS Management with Two Web Containers and a Centralized Backend Datastore

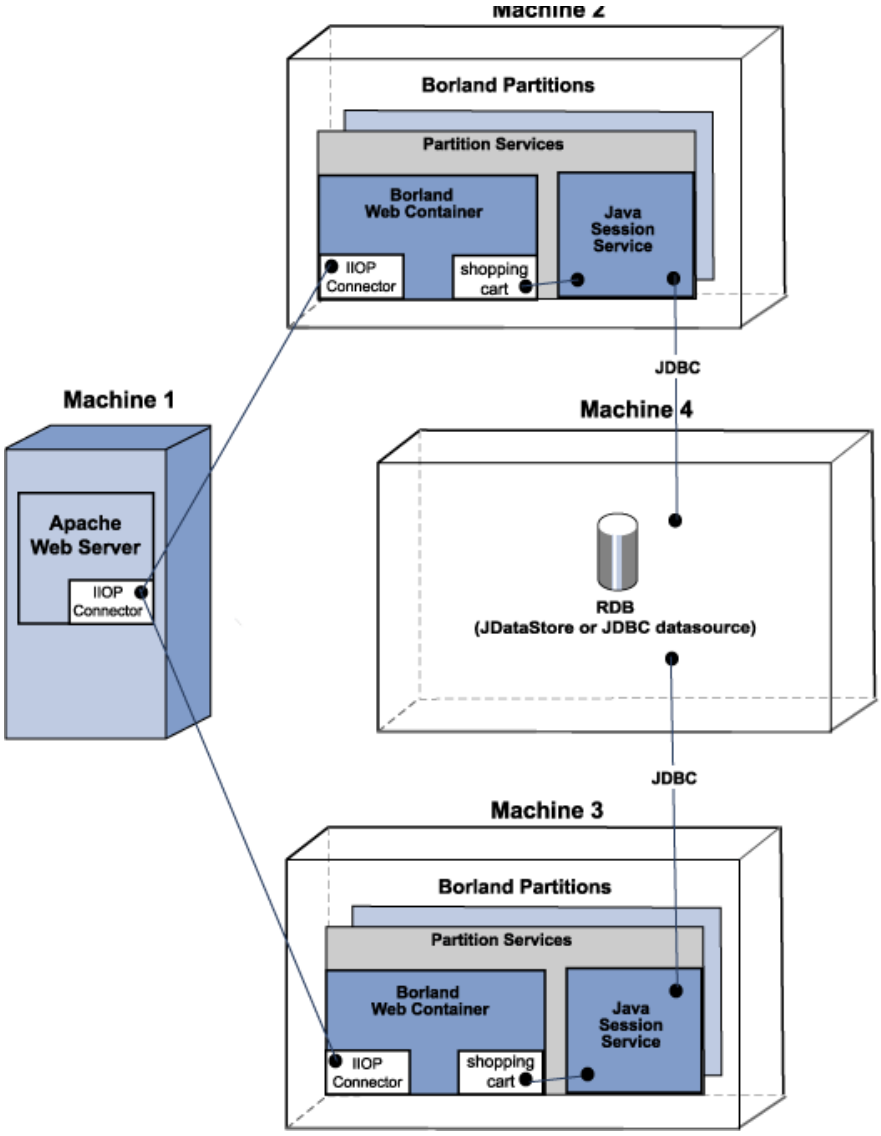

Managing and configuring the JSS

The JSS configuration is defined through its properties. The Borland AppServer (AppServer) is designed to work with any J2EE certified JDBC 2 driver; however, AppServer is only certified with and supported on JDataStore and Oracle.

- If JSS is configured to use a JDataStore file, the database tables are automatically created by JSS.
- If JSS is configured to use a JDBC datasource, three database tables needs to be pre-created in the backend database by your system administrator using the following SQL statements:

```
CREATE TABLE "JSS_KEYS" ("STORAGE_NAME" java_string primary key, "KEY_BASE"
java_float);
CREATE TABLE "JSS_WEB" ("KEY" java_string primary key, "VALUE"
java_serializable, "EXPIRATION" java_float);
CREATE TABLE "JSS_EJB" ("KEY" java_string primary key, "VALUE"
java_serializable, "EXPIRATION" java_float);
```

Note

When using the above SQL statements, make sure to substitute the equivalent data types for your database.

The JSS can run as part of the Partition side-by-side with other Partition services.

Configuring the JSS Partition service

As a “Partition service”, JSS configuration information is located in each Partition's data directory in the `partition.xml` file. By default, this file is located in the following directory:

```
<install_dir>/var/domains/base/configurations/<configuration_name>/mos/
<partition_name>/adm/properties.
```

For example, for a Partition named “MyPartition”, by default the JSS configuration information is located in:

```
<install_dir>/var/domains/base/configurations/<configuration_name>/mos/
mypartition/adm/properties/partition.xml
```

For more information see “<service> element”.

Otherwise, for the location of a Partition data directory, go to the `configuration.xml` file located in:

```
<install_dir>/var/domains/base/configurations/<configuration_name>/
```

and search for the Partition Managed Object directory attribute:

```
<partition-process directory=
```

For a listing and description of the session service (JSS) level properties, see “Java Session Service (JSS) Properties”.

7

Clustering web components

This section discusses the clustering of multiple web components which includes Apache web servers and the Tomcat-based Borland web containers. In a typical deployment scenario, you can use multiple Borland Partitions to work together in providing a scalable *n*-tier solution.

Each Borland Partition can have the same or different services. Depending on your clustering scheme, these services can be turned off or on. In any case, leveraging these resources together or clustering, makes deployment of your web application more efficient. Clustering of the web components involves session management, load balancing and fault tolerance (failover).

Stateless and stateful connection services

Interaction between the client and server involves two types of services: stateless and stateful. A *stateless service* does not maintain a state between the client and the server. There is no “conversation” between the server and the client while processing a client request. In a *stateful service*, the client and server maintains a dialog of information.

For information about the location of the Borland web container configuration files, see “Borland web container implementation”.

The Borland IIOP connector

The IIOP connector is software that is designed to allow an http web server to redirect requests to the Borland web container. The Borland AppServer (AppServer) includes the IIOP connector for the Apache 2.2.3 and Microsoft Internet Information Server(IIS) versions 5.0, 5.1 and 6.0 web servers. The job of handling the redirection of http requests is split between two components:

- a native library running on the web server.
- a jar file running of the web container.

The AppServer supports clustering of web components. The Borland IIOP connector uses the IIOP protocol. The following unique features are provided:

- Load Balancing
- Fault tolerance (failover)

– Smart Session handling

Load balancing support

Load balancing is the ability to direct http requests across a set of web containers. This enables the system administrator to spread the load of the http traffic across multiple web containers. Load balancing techniques can significantly improve the scalability of a given system. The Borland IIOP connector can be configured to offer load balancing in the following two ways:

– OSAgent based load balancing

– Corbaloc based load balancing

OSAgent based load balancing

This is simple to achieve and requires the least amount of configuration. In this setup, you start a number of Borland web container instances and name the IIOP connector in those Borland web container with the same name.

For more information about setting the `name` attribute, see “Modifying the Borland web container IIOP configuration”.

Apache does load balancing across Borland web container instances for each request. Essentially, Apache does a new bind for each request. The newly started Borland web container containers can be dynamically discovered.

Important

All Borland web containers and Apache must be running in the same ORB domain; osagent based load balancing is not possible in cases where you are using different Partitions on different ORB domains.

Corbaloc based load balancing

This approach uses a static configuration of the web containers that make up the cluster. However, it can span ORB domains. In this case you specify the locations where the web containers are running using the CORBA corbaloc semantics. For example:

```
corbaloc::172.20.20.28:30303,:172.20.20.29:30304/tc_inst1
```

In the above corbaloc example string:

– two TCP/IP endpoints are configured for a web container named “`tc_inst1`”

– a web container with an object name of “`tc_inst1`” is running on host `172.20.20.28` with its IIOP connector at port `30303`

– there is another web container running with the same object name on host `172.20.20.29` with it's IIOP connector listening on port `30304`.

For more information about setting the `port` attribute, see “Modifying the Borland web container IIOP configuration”.

The web server side IIOP connector converts this corbaloc string into CORBA objects using `orb.string_to_object` and uses the underlying features of VisiBroker to load balance across these “endpoints” specified in the corbaloc string. There can be any number of endpoints.

Note

All of the listed web containers do not have to be running for the load balancing to function. The ORB simply moves on to the next endpoint until a valid connection is obtained.

However, corbaloc based load balancing does require the web container's IIOP connector be started at a known port and be available for corbaloc kind of object naming. The following is a snippet of the web container IIOP connector configuration that is required:

```
<Connector className="org.apache.catalina.connector.iiop.IiopConnector"
name="tc_inst1"  port="30303"/>
```

This snippet starts the IIOP connector at port `30303` and names the Borland web container object "`tc_inst1`". The `port` attribute is optional. However, if you do not specify the port, a random port gets picked up by the ORB and you will be unable to use the corbaloc scheme to locate the object.

Your organization can impose policies on how to name web containers and the IIOP port or port ranges used.

Fault tolerance (failover)

Failover using osagent bind naming and corbaloc naming is automatic in both cases. In corbaloc naming, the next configured endpoint in the corbaloc name string is tried and so on in a cyclic fashion until all endpoints in the corbaloc string are tried.

For osagent bind naming, the osagent automatically redirects the client to an alternative (but equivalent) object instance.

Note

If there is no object available to the osagent, or none of the endpoints specified in the corbaloc name string are running, then the http request fails.

Smart session handling

When there is no session involved, the IIOP connector can do indiscriminate round robining. However, when sessions are involved, it is important that Apache routes its session requests to the web container that initiated the session.

In other http-to-servlet redirectors (and in the earlier version of the IIOP connector) this is achieved by maintaining a list of sessions-ids-to-web-container-id's in the web server's cache. This presents numerous issues with maintaining the state of this list. This list can be very large and wasteful of system resources. It can become out of date, for example, sessions can timeout and, in general, is an inefficient and problematic facet of the web server to web container redirection paradigm.

The IIOP connector resolves this by utilizing a technique called "smart session ids". This is where the IOR of the web container is embedded within the session id returned by the web container as part of the session cookie (or URL in the case of url-rewriting).

When the web container generates the session ID, it first determines if the request originated from the IIOP connector. If so, it obtains the stringified IOR of the IIOP connector through which the request is received. The web container generates the normal session ID as it always generates, but pre-fixes the stringified IOR in front of it. For example:

```
Stringified IOR: IOR:xyz
Normal session ID: abc
The new session ID: xyz_abc
```

In the case where the original web container has stopped running, failover is employed to locate another instance of an equivalent web container.

In the case of corbaloc identified web containers, where automatic osagent failover is not guaranteed, the IIOP connector performs a manual "rebind" to obtain a valid reference to the running equivalent web container.

Obviously, if there are no other running instances of the web container, then the http request fails.

The new web container obtains the old state from the session database and continues to service the request. When returning the response the new web container changes the session ID to reflect its IOR. This should be transparent to Apache as it does not look at the session ID on the way back to the browser client.

Setting up your web container with JSS

To properly failover when sessions are involved, you must set up the web containers with the same JSS backend.

Modifying a Borland web container for failover

In the Borland web container configuration file, `server.xml`, you need to add an entry similar to the following code sample for each web application. For more information about the `server.xml` file, see "Modifying the Borland web container IIOP configuration".

```
<Manager className="org.apache.catalina.session.PersistentManager">
          <Store className="org.apache.catalina.session.BorlandStore"
          storeName="jss_factory"/>
     </Manager>
```

The preceding code specifies the use of a `PersistentManager` with a storage class `BorlandStore`. It also specifies the connection to a `BorlandStore` factory named `jss_factory`. There must be a JSS with that factory name running in the local network.

For a description of `jss.factoryName`, see "Java Session Service (JSS) Properties".

Session storage implementation

There are two methods of implementing session storage for your clustered web components:

- Programmatic implementation
- Automatic implementation

Programmatic implementation

The *Programmatic implementation* assumes that each time you change the session attributes, you call `session.SetAttribute()` to notify the Borland web container that you have changed the session attributes.

This is a common operation in servlet development and when executed, there is no need to modify the `server.xml` file. Each time you change the session data, it is immediately written to the database through the JSS. Then if your web container instance unexpectedly goes down, the next web container instance designated to pick up the session accesses the session data. In essence, the Programmatic implementation guarantees to save changes immediately.

Automatic implementation

The *Automatic implementation* lets you store the session data periodically to JSS, regardless of whether the data has changed. By using this implementation, you do not need to notify the web container that the session attribute has changed.

For example, you can change state without calling `setAttribute ()` as depicted in the following code example:

```
Object myState = session.getAttribute("myState");

// Modify mystate here and do not call setAttribute ()
```

Your configuration file, `server.xml`, will have the following code snippet:

```
<Manager className=
"org.apache.catalina.session.PersistentManager"
         maxIdleBackup="xxx">
<Store className=
"org.apache.catalina.session.BorlandStore"
storeName="jss_factory">
</Manager>
```

where `xxx` is the time interval in seconds that you want the session data to be stored.

For more information about the `server.xml` file, see “Modifying the Borland web container IIOP configuration”.

Note

When using the Automatic implementation, you need to consider the following limitations:

1. If the web container goes down between two save intervals, the latest changes are not visible for the next web container instance. This is an important concern when defining the time interval value for the heartbeat.
2. The data is saved at the specified time interval no matter if the data is changed or not. This can be wasteful if a session frequently does not change and the defined time interval value is set too low.

Using HTTP sessions

The HyperText Transfer Protocol (HTTP) is a stateless protocol. In the client/server paradigm, it means that all client requests that the Apache web server receives are viewed as independent transactions. There is no relationship between each client request. This is a typical stateless connection between the client and the server.

However, there are times when the client deems it necessary to have a session concept for transaction completeness. A *session concept* typically means having a stateful interaction between the client and server. An example of the session concept is shopping online with an interactive shopping cart. Every time you add a new item into your shopping cart, you expect to see that new item added to a list of previously added items. HTTP is not usually regarded for handling client request in a stateful manner. But it can.

AppServer supports the HTTP sessions through two methods of implementations:

- Cookies: The web server send a cookie to identify a session. The web browser keeps sending back the same cookie with future requests. This cookie helps the server-side components to determine how to handle the transaction for a given session.
- URL rewriting: The URL that the user clicks on is dynamically rewritten to have session information.

8

Apache web server to CORBA server connectivity

The Apache IIOP connector can be configured to enable your web server to communicate with any standalone CORBA server implementing the `ReqProcessor` Interface Definition Language (IDL). This means you can easily put a web-based front-end on almost any CORBA server.

For more information, see “Installing Borland AppServer on Windows” or “Installing Borland AppServer on Solaris and HP-UX” in the *Installation Guide*.

Web-enabling your CORBA server

The following steps are required to make your CORBA server accessible over the internet:

- Determine the urls for your CORBA methods
- Implement the ReqProcessor IDL

Determining the urls for your CORBA methods

In order to make your CORBA server accessible over the internet, you need to:

1. Decide what business operations you want to expose.
2. Provide a url for those business operations (CORBA methods).

For example, your banking company's CORBA server is implementing the methods: `debit()`, `credit()`, and `balance()` and you want to expose these business methods to users through the internet. You need to map each of the CORBA server operations to what the user types in a browser.

Your bank company web site is `http://www.bank.com`.

To provide a url for each of the business operations you want to expose to the internet users:

1. Append the web application name to the company root url.

   For example:

```
http://www.bank.com/accounts
```

where `accounts` is the web application name.

Important

By default, your web application is not made available through the web server. In order to make it available through Apache, you must add some information to the web application descriptor. For step-by-step instructions on how to do so, see “Web Deploy Paths” in the *Management Console User's Guide*.

2 Append a name that is meaningful to users for the method in the web application that you want to expose.

For example:

```
http://www.bank.com/accounts/balance
```

where `balance` is the meaningful name for the `balance()` method.

Implementing the ReqProcessor IDL in your CORBA server

The `ReqProcessor` IDL allows communication between a web server and a CORBA server using IIOP. Once you implement the `ReqProcessor` IDL in your CORBA server, http requests can be passed from your web server to your CORBA server.

In implementing this IDL, you must expect the request url as part of the `HttpRequest` and invoke the appropriate CORBA method in response to that url.

IDL Specification for ReqProcessor Interface

```
*/
module apache {
        struct NameValue {
                string name;
                string value;
    };
        typedef sequence<NameValue> NVList;
    typedef sequence<octet> OctetSequence_t;

    struct HttpRequest {
                string  authType;     // auth type (BASIC,FORM etc)
                string  userid;       // username associated with request
                string  appName;      // application name (context path)
                string  httpMethod;   // PUT, GET etc,
                string  httpProtocol; // protocol HTTP/1.0, HTTP/1.1 etc
                string  uri;          // URI associated with request
                string  args;      // query string associated with this request
                string  postData;  // POST (form) data associated with request
                boolean isSecure;  // whether client specified https or http
                string  serverHostname; // server hostname specified with URI
  string  serverAddr; // [optionally] server IP address specified with URI
                long    serverPort; // server port number specified with URI
               NVList headers; // headers associated with this request format:
                   header-name:value
        };



        struct HttpResponse {
                long    status;                 // HTTP status, OK etc.
                boolean isCommit; // server intends to commit this request
                NVList  headers;                // header array
```

```
                OctetSequence_t data;    // data buffer
        };

interface ReqProcessor {
        HttpResponse process(in HttpRequest req);
  };
};
```

The process() method

The `ReqProcessor` IDL includes the `process()` method which your Apache web server calls for internet requests. The web server passes the user's request as an argument to the `process()` method. Basically, the input for the `process()` method is a request from a browser: `HttpRequest`, and the output for the `process()` method is an html page contained in: `HttpResponse`.

Configuring your Apache web server to invoke a CORBA server

Before an Apache web server can invoke a CORBA server, you must modify the lines of code that pertain to the IIOP connector in the `httpd.conf` file. For detailed information, see “Modifying the IIOP configuration in Apache”.

Figure 8.1 Connecting from Apache to a CORBA server

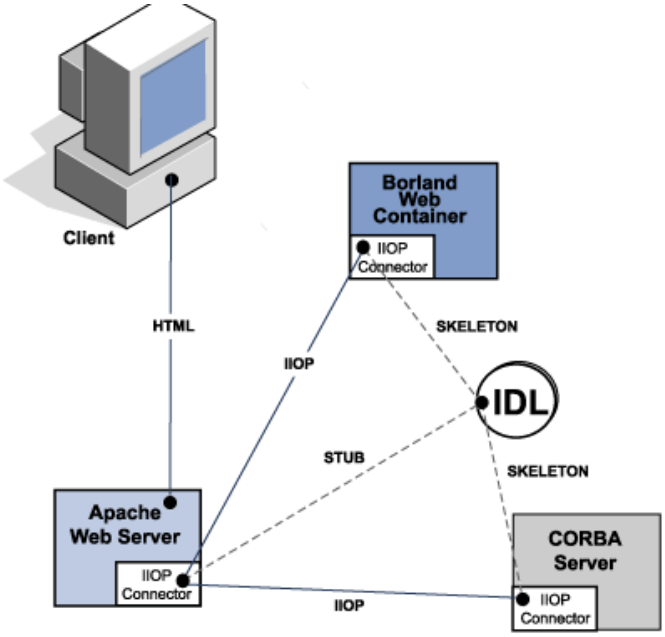

Apache IIOP configuration

The Apache IIOP connector has a set of configuration files that you must update with web server cluster information. By default, these IIOP connector configuration files are located in:

```
<install_dir>/var/domains/<domain_name>/configurations/
<configuration_name>/mos/<apache_managedobject_name>/conf
```

Note

"Cluster" is used to represent a CORBA object instance(s) that is known to the system by a single name or URI. The IIOP connector is able to load-balance across multiple instances, hence the term "cluster" is used.

The two configuration files are:

IIOP configuration file	Description
WebClusters.properties	Specifies the cluster(s) and the corresponding CORBA server(s) for each cluster.
UriMapFile.properties	Maps URI references to the clusters defined in the `WebClusters.properties` file.

Modifying either of these configuration files can be done so without starting up or shutting down the Apache web server(s) or CORBA server(s) because the file is automatically loaded by the IIOP connector.

Adding new CORBA servers (clusters)

CORBA servers are known as "clusters" to the IIOP connector. To configure your CORBA server for use with the IIOP connector, you need to define and add a "cluster" to the `WebClusters.properties` file.

The `WebClusters.properties` file tells the IIOP connector:

- The name of each available cluster (`ClusterList`).
- The web container identification.
- Whether to provide automatic load balancing (`enable_loadbalancing`) for a particular cluster.

To add a new cluster:

- In the `WebClusters.properties` file:
  - **a** add the name of the configured cluster to the `ClusterList`. For example:

    ```
    ClusterList=cluster1,cluster2,cluster3
    ```

  - **b** define each cluster by adding a line in the following format specifying the cluster name, the required `webcontainer_id` attribute, and any additional attributes (see the following Cluster definition attributes table). For example:

    ```
    <clustername>.webcontainer_id = <id> <attribute>
    ```

Note

Failover and smart session are always enabled, see “Clustering web components”.

Attribute	Required	Definition
`webcontainer_id`	yes	the object “bind” name or corbaloc string identifying the web container implementing the cluster.
`enable_loadbalancing`	no	Load balancing is enabled by default; to enable load balancing, do not include this attribute **or** include and set to `true`. To disable load balancing, set to `false` indicating that this cluster instance should **not** employ load-balancing techniques. **Warning:** Ensure that when entering the `enable_loadbalancing` attribute you give it a legal value (`true` or `false`).

For example:

```
ClusterList=cluster1,cluster2,cluster3
cluster1.webcontainer_id = tc_inst1
cluster2.webcontainer_id = corbaloc::127.20.20.2:20202,:127.20.20.3:20202/
tc_inst2
   cluster2.enable_loadbalancing = true
   cluster3.webcontainer_id = tc_inst3
   cluster3.enable_loadbalancing = false
```

In the above example, the following three clusters are defined:

1. The first, uses the osagent naming scheme and is enabled for load balancing.
2. The second cluster employs the corbaloc naming scheme, and is also enabled for load balancing.
3. The third uses the osagent naming scheme, but has the load balancing features disabled.

Note

To disable use of a particular cluster, simply remove the cluster name from the `ClusterList` list. However, we recommend you do not remove clusters with active http sessions attached to the CORBA server (attached users), because requests to these “live” sessions will fail.

Note

Modifications you make to the `WebClusters.properties` file automatically take effect on the next request. You do not need to restart your server(s).

Mapping URIs to defined clusters

Once the cluster entry is defined, all that remains is to identify which HTTP requests received by the web server need to be forwarded to your CORBA server. Use the `UriMapFile.properties` file to map http uri strings to web cluster names (CORBA instances) configured in the `WebClusters.properties` file.

– In the `UriMapFile.properties` file, type:

```
<uri-mapping> = <clustername>
```

where `<uri-mapping>` is a standard URI string or a wild-card string, and `<clustername>` is the cluster name as it appears in the `ClusterList` entry in the `WebClusters.properties` file.

For example:

```
/examples = cluster1
/examples/* = cluster1

/petstore/index.jsp = cluster2
/petstore/servlet/* = cluster2
```

In this example:

– Any URI that starts with `/examples` will be forwarded to a CORBA server running in the “cluster1” web cluster.

– URIs matching either `/petstore/index.jsp` or starting with `/petstore/servlet` will be routed to “cluster2”.

Note

With the URI mappings, the wild-card “*” is only valid in the **last** term of the URI and may represent the follow cases:

– the whole term (and all inferior references) as in `/examples/*`.

– the filename part of a file specification as in `/examples/*.jsp`.

Note

Modifications you make to the `UriMapFile.properties` file automatically take effect on the next request. You do not need to restart your server(s).

If the `WebCluster.properties` or `UriMapFile.properties` is altered, then it is automatically loaded by the IIOP connector. This means that modifications to either of these files can be done so without starting up or shutting down the web server(s) or CORBA server(s).

9

Borland AppServer Web Services

The Borland AppServer (AppServer) provides an out-of-the-box web services capability in all Borland Partitions.

Web Services Overview

A *Web Service* is an application component that you can describe, publish, locate, and invoke over a network using standardized XML messaging. Defined by new technologies like Simple Object Access Protocol (SOAP), Web Services Description Language (WSDL), and Universal Discovery, Description and Integration (UDDI), this is a new model for creating e-business applications from reusable software modules that are accessed on the World Wide Web.

Web Services Architecture

The standard Web Service architecture consists of the three roles that perform the web services publish, find, and bind operations:

- The *Service Provider* registers all available web services with the Service Broker.
- The *Service Broker* publishes the web services for the Service Requestor to access. The information published describes the web service and its location.
- The *Service Requestor* interacts with the Service Broker to find the web services. The Service Requestor can then bind or invoke the web services.

The Service Provider hosts the web service and makes it available to clients via the Web. The Service Provider publishes the web service definition and binding information to the Universal Description, Discovery, and Integration (UDDI) registry. The Web Service Description Language (WSDL) documents contain the information about the web service, including its incoming message and returning response messages.

The Service Requestor is a client program that consumes the web service. The Service Requestor finds web services by using UDDI or through other means, such as email. It then binds or invokes the web service.

The Service Broker manages the interaction between the Service Provider and Service Requestor. The Service Broker makes available all service definitions and binding information. Currently, SOAP (an XML-based, messaging and encoding protocol format for exchange of information in a decentralized, distributed environment) is the standard for communication between the Service Requestor and Service Broker.

Figure 9.1 Standard Web Services Architecture

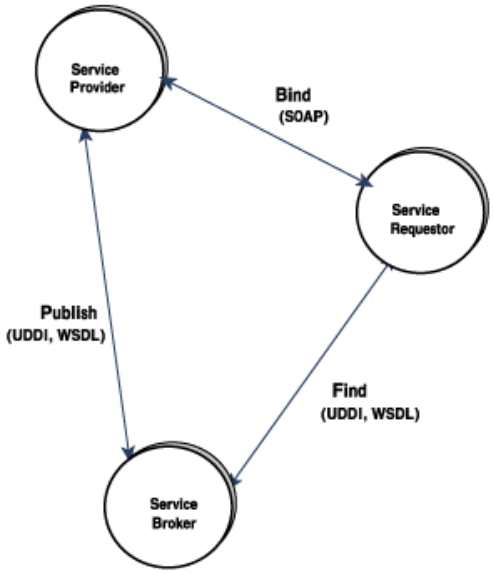

Web Services and Partitions

All AppServer Partitions are configured to support web services. You simply need to start a Partition and deploy WARs (or EARs containing WARs) containing web services.

Additionally, you can expose a previously deployed stateless session bean as a web service. For more information, see “Export EJB as a Web Service Wizard” in the *Management Console User's Guide*.

The Borland web services is based on the Apache Axis technology and supports dispatch of incoming SOAP web services requests to the following “Web Service providers”:

- EJB providers
- RPC/Java providers
- MDB/Java providers

Figure 9.2 Borland Web Services Architecture

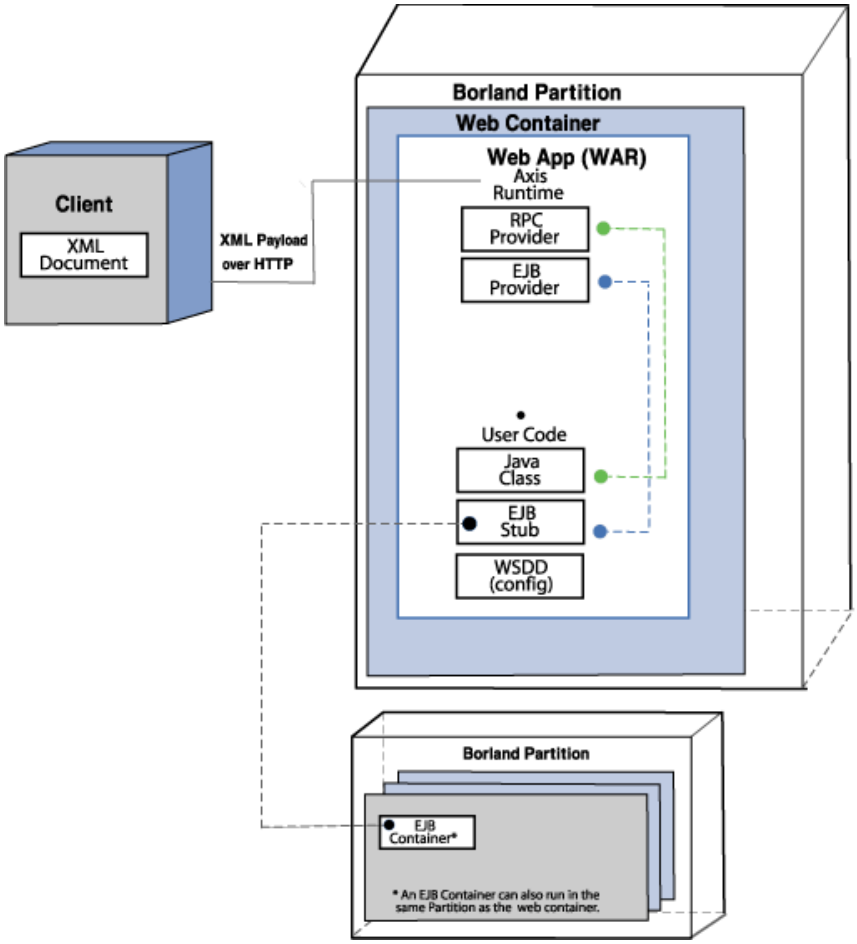

Web Service providers

The Borland web services engine includes a number of providers. A *provider* is the link that connects a client web service request to the user's class on the server side.

All providers do the following:

- Create an instance of an object on which they can invoke methods. The exact way of creating this object differs from provider to provider.
- Invoke the methods on that object and pass all the parameters that the XML client sent.
- Pass the return value to the Axis Runtime engine, which then converts it to XML and sends it back to the client.

Specifying web service information in a deploy.wsdd file

When installing a new web service, you must name the web service and specify which provider the service is going to use. Each provider takes different parameters. The following describes the service providers and the required parameters for each.

Java:RPC provider

This provider assumes that the class serving the web service is in the application archive (WAR). When a web service request arrives, the RPC provider:

1 Loads the Java class associated with the service.

2 Creates a new instance of the object.

3 Invokes the specified method using reflection.

The parameters are:

– `className`: The name of the class that is loaded when a request arrives on this service.

– `allowedMethods`: The methods that are allowed to be invoked on this class. The class can have more methods than listed here; the methods listed here are available for remote invocation.

Example:

```
<service name="Animal" provider="java:RPC">
 <parameter name="className" value="com.borland.examples.web
services.java.Animal"/>
 <parameter name="allowedMethods" value="talk sleep"/>
</service>
```

Java:EJB provider

This provider assumes that the class serving the web service is an EJB.

Note

You can expose a previously deployed stateless session bean as a web service. For more information, see “Export EJB as a Web Service Wizard” in the *Management Console User's Guide*.

When a web service request arrives:

1 The EJB provider looks up the bean name in JNDI initial context.

2 Locates the home class and creates a bean.

3 Invokes the specified method using reflection on the EJB stub.

The actual EJB itself must be deployed to any Partition before a client can access it.

The essential parameters are:

– `beanJndiName`: The name of the bean in JNDI.

– `homeInterfaceName`: The fully specified class of the home interface. This class must be present in the WAR.

– `className`: The name of the EJB remote interface.

– `allowedMethods`: The methods that are allowed to be invoked on this EJB, separated by spaces. The EJB can have more methods than listed here; the methods listed here are available for remote invocation.

Example:

```
<service name="Animal" provider="java:EJB">
   <parameter name="beanJndiName" value="Animal"/>
   <parameter name="homeInterfaceName"
value="com.borland.examples.webservices.ejb.AnimalHome"/>
   <parameter name="className"
value="com.borland.examples.webservices.ejb.Animal"/>
   <parameter name="allowedMethods" value="talk sleep"/>
     </service>
```

How Borland Web Services work

1. The web services server receives an XML SOAP message from a client.
2. It then:
   a. Interprets the SOAP message.
   b. Extracts the SOAP service name.
   c. Determines the appropriate provider who can respond to this service.
3. The mapping between the SOAP service and the type of provider is obtained from the Web Service Deployment Descriptor (WSDD) as part of WAR deployment.
4. The message is then passed onto the right provider. For information about the different ways in which each provider deals with the message, see “Java:RPC provider” and “Java:EJB provider”.

Web Service Deployment Descriptors

Web services are deployed as part of a WAR. A single WAR can contain multiple web services. You can also deploy multiple WARs with each containing many web services.

The difference between a normal WAR and a WAR containing web services is the presence of an extra descriptor named `server-config.wsdd` in the `WEB-INF` directory. The `server-config.wsdd` file provides configuration information (the name of the web service, the provider, any corresponding Java classes and allowed methods).

There is one WSDD file per WAR and it contains information about all available web services within that WAR.

The typical component structure of a WAR containing web services has the following elements:

- `WEB-INF/web.xml`
- `WEB-INF/server-config.wsdd`
- `WEB-INF/classes/<classes corresponding to your web services are located here>`
- `WEB-INF/lib/<classes corresponding to your web services are located here in the packed JAR form>`

The `WEB-INF/lib` also contains some standard JARs that are necessary for the Axis Runtime engine.

To publish your Java classes as a web service, use the WSDD format to define the items that you want to deploy to the Partition. For example, an entry corresponding to a service named “BankService” can be:

```
<service name="BankService" provider="java:RPC">
   <parameter name="allowedMethods" value="create_account query_account"/>
   <parameter name="className" value="com.fidelity.Bank"/>
</service>
```

In this case, the `com.fidelity.Bank` Java class links to web service `BankService`. The class `com.fidelity.Bank` can have a number of public methods, but only the methods `create_account` and `query_account` are available through the web service.

Creating a server-config.wsdd file

To create the `server-config.wsdd`:

- Use JBuilder to generate the deployment descriptor as part of your WAR.

or

1. Use a text editor to write a `deploy.wsdd` file. Refer to the `deploy.wsdd` file in `<install_dir>/examples/webservices/java/server`.

2 Run the Tools Overview with the `deploy.wsdd` file by typing:

```
prompt>java org.apache.axis.utils.Admin server deploy.wsdd
```

The `server-config.wsdd` file is packaged as part of the WAR.

Viewing and Editing WSDD Properties

You can view and edit the properties of any web service deployment descriptor (WSDD) (`server-config.wsdd` file) that is packaged in a WAR file using either the Borland Management Console or the DDEditor. For more information, see “Viewing Web Services WSDD properties” or “Web Services” in the *Management Console User's Guide*.

Packaging Web Service Application Archives

To Create a WAR that can be deployed to the web services archive:

1 Make sure your web service classes are in `WEB-INF/classes` or `WEB-INF/lib`.

2 Copy the Axis toolkit libraries to `WEB-INF/lib`. The Axis libraries can be found in: `<install_dir>/lib/axis`

3 Copy the `web.xml` necessary for the Axis tool kit to WEB-INF directory. The `web.xml` can be found in: `<install_dir>/etc/axis`

4 Create a `deploy.wsdd` that has deployment information about your web services.

5 Run the Axis Admin tool on this `deploy.wsdd` to generate the `server-config.wsdd` as follows:

```
java org.apache.axis.utils.Admin server deploy.wsdd
```

6 Copy this `server-config.wsdd` to `WEB-INF`

7 JAR your web application into a WAR file.

Borland Web Services examples

To help you get started with developing and deploying web services, we provide samples and examples for the Borland web services engine. These examples are included in your AppServer installation at:

```
<install_dir>/examples/webservices/
```

The examples that illustrate the different web service providers are located in the web services examples directory in the `Java`, `EJB`, `MDB` or `VisiBroker` folder.

Your AppServer installation also includes several Apache Axis samples in:

```
<install_dir>/examples/webservices/axis/samples/
```

Using the Web Service provider examples

The AppServer examples must be built before they are deployed and deployed before they are run. Building the examples involves generating the necessary WSDL files and packaging the application's code and descriptors into a deployable unit, in this case a WAR. This WAR can then be deployed to a Borland Partition. The application is run by invoking its client from a command-line. Building and running the examples is automated through the use of the Apache Ant utility, while deployment is performed using tools provided with AppServer.

Steps to build, deploy, and run the examples

1 **Build.** You can build all of the examples simultaneously or build each one individually. To build them all simultaneously, navigate to the:

```
/examples/webservices/
```

directory and execute the Ant command. For example:

```
C:/BDP/examples/webservices>Ant
```

builds all the examples.

To build an individual example, navigate to its specific directory and execute the Ant command.

For example:

```
C:/BDP/examples/webservices/java>Ant
```

builds only the Java Provider example.

2 **Deploy.** You deploy the examples to a running instance of AppServer. You can use the `ant deploy` target, or any of the following to deploy your WAR and JAR:

- iastool command-line utility; for more information see “iastool command-line utility”.
- Deployment Wizard; see “Deployment Wizard” in the *Management Console User's Guide*.

1 **Run.** To run an example, navigate to its directory and use the `ant run-client` command.

For example, to run the Java Provider client:

```
C:/BDP/examples/webservices/java>Ant run-client
```

Apache Axis Web Service samples

The Apache Axis web service samples are already deployed in the `axis-samples.war` file present in the Borland Partition. Since these are already pre-deployed, you do not need to use the Apache Axis deploy commands suggested in the Apache Axis User's Guide.

The Apache Axis User's Guide is also provided with the AppServer installation and is located in:

```
<install_dir>/doc/axis/user-guide.html
```

These samples illustrate the capabilities of Axis. They are unmodified from the original Apache Axis implementation and are not guaranteed to run.

Tools Overview

This section describes the various tools used in examples.

Apache ANT tool

The Apache ANT utility is a platform-independent, java-based build tool used to build the examples.

The XML build script `build.xml` is used to drive the tool. This `build.xml` file describes the various targets available for a project and the commands executed in response to those targets. The AppServer conveniently provides all necessary JARs and scripts to run the Apache Ant tool.

Java2WSDL tool

The Java2WSDL is an Apache Axis utility class that generates WSDL corresponding to a Java class. This class can accept a number of command line arguments. You can get the full usage help by running this utility without arguments as follows:

```
java org.apache.axis.wsdl.Java2WSDL
```

Note

You must set your CLASSPATH to include all jar files in the `<install-dir>\lib\axis` directory, before you run the following command.

WSDL2Java tool

The WSDL2Java is an Apache Axis utility class that generates Java classes from a WSDL file. This tool can generate java stubs (used on the client side), or java skeletons (used on the server side). The generated files make it easy to develop your client or server for a given WSDL.

This class can accept a number of command line arguments. You can get the full usage help by running this utility without arguments as follows:

```
java org.apache.axis.wsdl.WSDL2Java
```

Note

You must set your CLASSPATH to include all jar files in the `<install-dir>\lib\axis` directory, before you run the following command.

Axis Admin tool

The Apache Admin tool is a utility class that generates WAR level global configuration files based on deployment information specific to some web services.

The input to this utility is a XML file (typically named `deploy.wsdd`) containing deployment information about one or more web services. The Apache Admin utility adds some global definitions that are necessary and writes an output file. Use this tool as follows:

```
java org.apache.axis.utils.Admin server|client deployment-file
```

Note

You must set your CLASSPATH to include all jar files in the `<install-dir>\lib\axis` directory, before you run the command.

This tool generates `server-config.wsdd` or `client-config.wsdd` based on what option you choose.

10

Writing enterprise bean clients

Client view of an enterprise bean

A client of an enterprise bean is an application—a stand-alone application, an application client container, servlet, or applet—or another enterprise bean. In all cases, the client must do the following things to use an enterprise bean:

- Locate the bean's home interface. The EJB specification states that the client should use the JNDI (Java Naming and Directory Interface) API to locate home interfaces.
- Obtain a reference to an enterprise bean object's remote interface. This involves using methods defined on the bean's home interface. You can either create a session bean, or you can create or find an entity bean.
- Invoke one or more methods defined by the enterprise bean. A client does not directly invoke the methods defined by the enterprise bean. Instead, the client invokes the methods on the enterprise bean object's remote interface. The methods defined in the remote interface are the methods that the enterprise bean has exposed to clients.

Initializing the client

The SortClient application imports the necessary JNDI classes and the SortBean home and remote interfaces. The client uses the JNDI API to locate an enterprise bean's home interface.

A client application can also use logical names (as recommended in the various J2EE specifications) to access resources such as database connections, remote enterprise beans, and environment variables. The container, per the J2EE specification, exposes these resources as administered objects in the local JNDI name space (that is, `java:comp/env`).

Locating the home interface

A client locates an enterprise bean's home interface using JNDI, as shown in the code sample below. The client first needs to obtain a JNDI initial naming context. The code instantiates a new `javax.naming.Context` object, which in our example it calls initialContext. Then, the client uses the context `lookup()` method to resolve the name to

a home interface. Note that the initialization of the initial naming context factory is EJB container/server specific.

A client application can also use logical names to access a resource such as the home interface. See “Initializing the client” for more information.

The context's `lookup()` method returns an object of type `java.lang.Object`. Your code must cast this returned object to the expected type. The following code sample shows a portion of the client code for the sort example. The `main()` routine begins by using the JNDI naming service and its context lookup method to locate the home interface. You pass the name of the remote interface, which in this case is `sort`, to the `context.lookup()` method. Notice that the program eventually casts the results of the `context.lookup()` method to SortHome, the type of the home interface.

```
// SortClient.java
import javax.naming.InitialContext;
import SortHome; // import the bean's home interface
import Sort; // import the bean's remote interface
public class SortClient {
   ⋮
   public static void main(String[] args) throws Exception {
      javax.naming.Context context;

      // preferred JNDI context lookup
      // get a JNDI context using a logical JNDI name in the local JNDI context,
         i.e.,ejb-ref
      javax.naming.Context context = new javax.naming.InitialContext();
      Object ref = context.lookup("java:comp/env/ejb/Sort");
      SortHome home = (SortHome) javax.rmi.PortableRemoteObject.narrow
         (ref, SortHome.class);
      Sort sort = home.create();
      ... //do the sort and merge work
      sort.remove();
   }
}
```

The `main()` routine of the client program throws the generic exception coded this way, the SortClient program does not have to catch any exceptions that might occur, though if an exception occurs it will terminate the program.

Obtaining the remote interface

Now that we have obtained the home interface of an enterprise bean we can get a reference to the enterprise bean's remote interface. To do this, we use the home interface's create or finder methods. The exact method to invoke depends on the type of the enterprise bean and the methods the enterprise bean provider has defined in the home interface.

For example, the first code sample shows how SortClient obtains a reference to the Sort remote interface. Once SortClient obtains the reference to the home interface and casts it to its proper type (SortHome), then the code can create an instance of the bean and call its methods. It calls the home interface's `create()` method, which returns a reference to the bean's remote interface, Sort. (Because SortBean is a stateless session bean, its home interface has only one `create()` method and that method by definition takes no parameters.) SortClient can then call the methods defined on the remote interface—`sort()` and `merge()`—to do its sorting work. When the work finishes, the client calls the remote interface's `remove()` method to remove the instance of the enterprise bean.

Session beans

A client obtains a reference to a session bean's remote interface by calling one of the create methods on the home interface.

All session beans must have at least one `create()` method. A stateless session bean must have only one `create()` method, and that method must have no arguments. A stateful session bean can have one `create()` method, and may have additional

`create()` methods whose parameters vary. If a `create()` method does have parameters, the values of these parameters are used to initialize the session bean.

The default `create()` method has no parameters. For example, the sort example uses a stateless session bean. It has, by definition, one `create()` method that takes no parameters:

```
Sort sort = home.create();
```

The cart example, on the other hand, uses a stateful session bean, and its home interface, CartHome, implements more than one `create()` method. One of its `create()` methods takes three parameters, which together identify the purchaser of the cart contents, and returns a reference to the Cart remote interface. The CartClient sets values for the three parameters—`cardHolderName`, `creditCardNumber`, and `expirationDate`—then calls the `create()` method. This is shown in the code sample below:

```
Cart cart;
   {
      String cardHolderName = "Jack B. Quick";
      String creditCardNumber = "1234-5678-9012-3456";
      Date expirationDate = new GregorianCalendar(2001,
         Calendar.JULY, 1).getTime();
      cart = home.create(cardHolderName, creditCardNumber, expirationDate);
   }
```

Session beans do not have finder methods.

Entity beans

A client obtains a reference to an entity object either through a find operation or a create operation. Recall that an entity object represents some underlying data stored in a database. Because the entity bean represents persistent data, entity beans typically exist for quite a long time; certainly for much longer than the client applications that call them. Thus, a client most often needs to find the entity bean that represents the piece of persistent data of interest, rather than creating a new entity object, which would create and store new data in the underlying database.

A client uses a find operation to locate an existing entity object, such as a specific row within a relational database table. That is, find operations locate data entities that have previously been inserted into data storage. The data may have been added to the data store by an entity bean or it may have been added outside of the EJB context, such as directly from within the database management system (DBMS). Or, in the case of legacy systems, the data may have existed prior to the installation of the EJB container.

A client uses an entity bean object's `create()` method to create a new data entity that will be stored in the underlying database. An entity bean's `create()` method inserts the entity state into the database, initializing the entity's variables according to the values in the `create()` method's parameters. A `create()` method for an entity bean always returns the remote interface, but the corresponding `ejbCreate()` method returns primary key of the entity instance.

Every entity bean instance must have a primary key that uniquely identifies it. An entity bean instance can also have secondary keys that can be used to locate a particular entity object.

Find methods and primary key class

The default find method for an entity bean is the `findByPrimaryKey()` method, which locates the entity object using its primary key value. Its signature is as follows:

```
<remote interface> findByPrimaryKey( <key type> primaryKey )
```

Every entity bean must implement a `findByPrimaryKey()` method. The `primaryKey` parameter is a separate primary key class that is defined in the deployment descriptor. The key type is the type for the primary key, and it must be a legal value type in RMI-IIOP. The primary key class can be any class—a Java class or a class you've written yourself.

For example, you have an Account entity bean that defines the primary key class AccountPK. AccountPK is a String type, and it holds the identifier for the Account bean. You can obtain a reference to a specific Account entity bean instance by setting the AccountPK to the account identifier and invoking the `findByPrimaryKey()` method, as shown in the following code sample.

```
AccountPK accountPK = new AccountPK("1234-56-789");
Account source = accountHome.findByPrimaryKey( accountPK );
```

The bean provider can define additional finder methods that a client can use.

Create and remove methods

A client can also create entity beans using create methods defined in the home interface. When a client invokes a `create()` method for an entity bean, the new instance of the entity object is saved in the data store. The new entity object always has a primary key value that is its identifier. Its state may be initialized to values passed as parameters to the `create()` method.

Keep in mind that an entity bean exists for as long as data is present in the database. The life of the entity bean is not bound by the client's session. The entity bean can be removed by invoking one of the bean's `remove()` methods—these methods remove the bean and the underlying representation of the entity data from the database. It is also possible to directly delete an entity object, such as by deleting a database record using the DBMS or with a legacy application.

Invoking methods

Once the client has obtained a reference to the bean's remote interface, the client can invoke the methods defined in the remote interface for this enterprise bean. The methods pertaining to the bean's business logic are of most interest to the client. There are also methods for getting information about the bean and its interfaces, getting the bean object's handle, testing if one bean is identical to another bean, and methods for removing the bean instance.

The next code sample illustrates how a client calls methods of an enterprise bean, in this case, a cart session bean. We pick up the client code from the point where it has created a new session bean instance for a card holder and retrieved a Cart reference to the remote interface. At this point, the client is ready to invoke the bean methods.

First, the client creates a new book object, setting its title and price parameters. Then, it invokes the enterprise bean business method `addItem()` to add the book object to a shopping cart. The `addItem()` method is defined on the CartBean session bean, and is made public through the Cart remote interface. The client adds other items (not shown here), then calls its own `summarize()` method to list the items in the shopping cart. This is followed by the `remove()` method to remove the bean instance. Notice that a client calls the enterprise bean methods in the same way that it invokes any method, such as its own method `summarize()`.

```
⋮
Cart cart;
{
   ⋮
   // obtain a reference to the bean's remote interface
   cart = home.create(cardHolderName, creditCardNumber, expirationDate);
}
// create a new book object
Book knuthBook = new Book("The Art of Computer Programming", 49.95f);
// add the new book item to the cart
cart.addItem(knuthBook);

⋮
// list the items currently in the cart
summarize(cart);
cart.removeItem(knuthBook);
⋮
```

Removing bean instances

The `remove()` method operates differently for session beans than for entity beans. Because a session object exists for one client and is not persistent, a client of a session bean should call the `remove()` method when finished with a session object. There are two `remove()` methods available to the client: the client can remove the session object with the `javax.ejb.EJBObject.remove()` method, or the client can remove the session handle with the `javax.ejb.EJBHome.remove(Handle handle)` method. See "Using a bean's handle" for more information on handles.

While it is not required that a client remove a session object, it is considered to be good programming practice. If a client does not remove a stateful session bean object, the container eventually removes the object after a certain time, specified by a timeout value. The timeout value is a deployment property. However, a client can also keep a handle to the session for future reference.

Clients of entity beans do not have to deal with this problem as entity beans are only associated with a client for the duration of a transaction and the container is in charge of their life cycles, including their activation and passivation. A client of an entity bean calls the bean's `remove()` method only when the entity object is to be deleted from the underlying database.

Using a bean's handle

A handle is an another way to reference an enterprise bean. A handle is a serializable reference to a bean. You can obtain a handle from the bean's remote interface. Once you have the handle, you can write it to a file (or other persistent storage). Later, you can retrieve the handle from storage and use it to reestablish a reference to the enterprise bean.

However, you can only use the remote interface handle to recreate the reference to the bean; you cannot use it to recreate the bean itself. If another process has removed the bean, or the system crashed or shutdown and removed the bean instance, then an exception is thrown when the client application tries to use the handle to reestablish its reference to the bean.

When you are not sure that the bean instance will still be in existence, rather than using a handle to the remote interface, you can store the bean's home handle and recreate the bean object later by invoking the bean's create or find methods.

After the client creates a bean instance, it can use the `getHandle()` method to obtain a handle to this instance. Once it has the handle, it can write it to a serialized file. Later, the client program can read the serialized file, casting the object that it reads in to a Handle type. Then, it calls the `getEJBObject()` method on the handle to obtain the bean reference, casting the results of `getEJBObject()` to the correct type for the bean.

To illustrate, the CartClient program might do the following to utilize a handle to the CartBean session bean:

```
import java.io;
import javax.ejb.Handle;
⋮
Cart cart;
⋮
cart = home.create(cardHolderName, creditCardNumber, expirationDate);
// call getHandle on the cart object to get its handle
cartHandle = cart.getHandle();
// write the handle to serialized file
FileOutputStream f = new FileOutputStream ("carthandle.ser");
ObjectOutputStream o = new ObjectOutputStream(f);
o.writeObject(myHandle);
o.flush();
o.close();
⋮
// read handle from file at later time
FileInputStream fi = new FileInputStream ("carthandle.ser");
ObjectInputStream oi = new ObjectInputStream(fi);
```

```
//read the object from the file and cast it to a Handle
cartHandle = (Handle)oi.readObject();
oi.close();
⋮
// Use the handle to reference the bean instance
try {
   Object ref = context.lookup("cart");
   Cart cart1 = (Cart) javax.rmi.PortableRemoteObject.narrow(ref, Cart.class);
   ⋮
} catch (RemoteException e) {
   ⋮
}
⋮
```

When finished with the session bean handle, the client can remove it with the `javax.ejb.EJBHome.remove(Handle handle)` method.

Managing transactions

A client program can manage its own transactions rather than letting the enterprise bean (or container) manage the transaction. A client that lmanages its own transaction does so in exactly the same manner as a session bean than manages its own transaction.

When a client manages its own transactions, it is responsible for delimiting the transaction boundaries. That is, it must explicitly start the transaction and end (commit or roll back) the transaction.

A client uses the `javax.transaction.UserTransaction` interface to manage its own transactions. It must first obtain a reference to the `UserTransaction` interface, using JNDI to do so. Once it has the `UserTransaction` context, the client uses the `UserTransaction.begin()` method to start the transaction, followed later by the `UserTransaction.commit()` method to commit and end the transaction (or `UserTransaction.rollback()` to rollback and end the transaction). In between, the client does its queries and updates.

This code sample shows the code that a client would implement to manage its own transactions. The parts that pertain specifically to client-managed transactions are highlighted in bold.

```
⋮
import javax.naming.InitialContext;
import javax.transaction.UserTransaction;
⋮
public class clientTransaction {
   public static void main (String[] argv) {
      UserTransaction ut = null;
      InitialContext initContext = new InitialContext();
      ⋮
      ut = (UserTransaction)initContext.lookup("java:comp/UserTransaction");
      // start a transaction
      ut.begin();
      // do some transaction work
      ⋮
      // commit or rollback the transaction
      ut.commit(); // or ut.rollback();
      ⋮
   ]
]
```

Getting information about an enterprise bean

Information about an enterprise bean is referred to as metadata. A client can obtain metadata about a bean using the enterprise bean's home interface `getMetaData()` method.

The `getMetaData()` method is most often used by development environments and tool builders that need to discover information about an enterprise bean, such as for linking together beans that have already been installed. Scripting clients might also want to obtain metadata on the bean.

Once the client retrieves the home interface reference, it can call the `getEJBMetaData()` method on the home interface. Then, the client can call the `EJBMetaData` interface methods to extract such information as:

- The bean's `EJBHome` home interface, using `EJBMetaData.getEJBHome()`.
- The bean's home interface class object, including its interfaces, classes, fields, and methods, using `EJBMetaData.getHomeInterfaceClass()`.
- The bean's remote interface class object, including all class information, using `EJBMetaData.getRemoteInterfaceClass()`.
- The bean's primary key class object, using `EJBMetaData.getPrimaryKeyClass()`.
- Whether the bean is a session bean or an entity bean, using `EJBMetaData.isSession()`. The method returns true if this is a session bean.
- Whether a session bean is stateless or stateful, using `EJBMetaData.isStatelessSession()`. The method returns true if the session bean is stateless.

Support for JNDI

The EJB specification defines the JNDI API for locating home interfaces. JNDI is implemented on top of other services, including CORBA's Naming Service, LDAP/ X.500, flat files, and proprietary directory services. The diagram below illustrates the different implementation choices. Typically, the EJB server provider selects a particular implementation of JNDI.

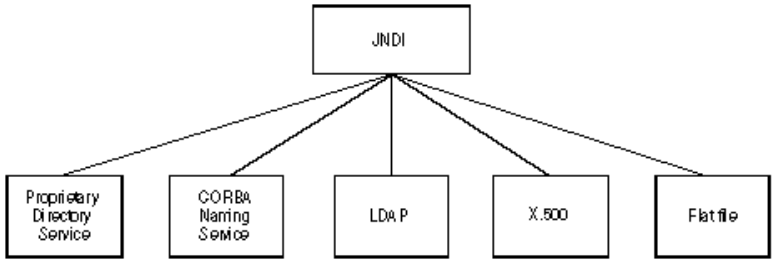

The technology implemented beneath JNDI is of no concern to the client. The client needs to use only the JNDI API.

EJB to CORBA mapping

There are a number of aspects to the relationship between CORBA and Enterprise JavaBeans. Three important ones are the implementation of an EJB container/server with an ORB, the integration of legacy systems into an EJB middle tier, and the access of enterprise beans from non-Java components, specifically clients. The EJB specification is currently only concerned with the third aspect.

CORBA is a very suitable and natural platform on which to implement an EJB infrastructure. CORBA addresses all of the concerns of the EJB specification with the CORBA Core specification or the CORBA Services:

- **Support for distribution.** CORBA Core and CORBA Naming Service
- **Support for transactions.** CORBA Object Transaction Service
- **Support for security.** CORBA Security Specification, including IIOP-over-SSL

Additionally, CORBA allows the integration of non-Java components into an application. These components can be legacy systems and applications, plus different kinds of clients. Back-end systems can be easily integrated using OTS and any programming language for which an IDL mapping exists. This requires an EJB container to provide OTS and IIOP APIs.

The EJB specification is concerned with the accessibility of enterprise beans from non-Java clients and provides an EJB to CORBA mapping. The goals of the EJB/CORBA mapping are:

- Supporting interoperability between clients written in any CORBA-supported programming language and enterprise beans running on a CORBA-based EJB server.
- Enabling client programs to mix and match calls to CORBA objects and enterprise beans within the same transaction.
- Supporting distributed transactions involving multiple enterprise beans running on CORBA-based EJB servers provided by different vendors.

The mapping is based on the Java-to-IDL mapping. The specification includes the following parts: mapping of distribution-related aspects, the mapping of naming conventions, the mapping of transactions, and the mapping of security. We explain each of these aspects in the following sections. Since the mapping uses new IDL features introduced by the OMG's Object-by-Value specification, interoperability with other programming languages requires CORBA 2.3-compliant ORBs.

Mapping for distribution

An enterprise bean has two interfaces that are remotely accessible: the remote interface and the home interface. Applying the Java/IDL mapping to these interfaces results in corresponding IDL specifications. The base classes defined in the EJB specification are mapped to IDL in the same manner.

For example, look at the IDL interface for an ATM enterprise session bean that has methods to transfer funds between accounts and throws an insufficient funds exception. By applying the Java/IDL mapping to the home and the remote interface, you get the following IDL interface.

```
module transaction {
   module ejb {
   valuetype InsufficientFundsException : ::java::lang::Exception {};
   exception InsufficientFundsEx {
      ::transaction::ejb::InsufficientFundsException value;
   };
   interface Atm : ::javax::ejb::EJBObject{
   void transfer (in string arg0, in string arg1, in float arg2)
      raises (::transaction::ejb::InsufficientFundsEx);
   };
   interface AtmHome : ::javax::ejb::EJBHome {
      ::transaction::ejb::Atm create ()
      raises (::javax::ejb::CreateEx);
   };
};};};};
```

Mapping for naming

A CORBA-based EJB runtime environment that wants to enable any CORBA clients to access enterprise beans must use the CORBA Naming Service for publishing and resolving the home interfaces of the enterprise beans. The runtime can use the CORBA Naming Service directly or indirectly via JNDI and its standard mapping to the CORBA Naming Service.

JNDI names have a string representation of the following form "`directory1/directory2/.../directoryN/objectName`". The CORBA Naming Service defines names as a sequence of name components.

```
typedef string Istring;
   struct NameComponent {
      Istring id;
      Istring kind;
   };
typedef sequence<NameComponent> Name;
```

Each "/" separated name of a JNDI string name is mapped to a name component; the leftmost component is the first entry in the CORBA Naming Service name.

A JNDI string name is relative to some naming context, which calls the JNDI root context. The JNDI root context corresponds to a CORBA Naming Service initial context. CORBA Naming Service names are relative to the CORBA initial context.

A CORBA program obtains an initial CORBA Naming Service naming context by calling `resolve_initial_references("NameService")` on the ORB (pseudo) object. The CORBA Naming Service does not prescribe a rooted graph for organizing naming context and, hence, the notion of a root context does not apply. The initialization of the ORB determines the context returned by `resolve_initial_references()`.

For example, a C++ Client can locate the home interface to the ATMSession bean, which has been registered with a JNDI string name "transaction/corbaEjb/atm". You first obtain the initial naming context.

```
Object_ptr obj = orb->resolve_initial_references("NameService");
NamingContext initialNamingContext= NamingContext.narrow( obj );
if( initialNamingContext == NULL ) {
   cerr << "Couldn't initial naming context" << endl;
   exit( 1 );
}
```

Then you create a CORBA Naming Service name and initialize it according to the mapping explained previously.

```
Name name = new Name( 1 );
name[0].id = "atm";
name[0].kind = "";
```

Now resolve the name on the initial naming context. Assume that you have successfully performed the initialization and that you have the context of the naming domain of the enterprise bean. We narrow the resulting CORBA object to the expected type and make sure that the narrow was successful.

```
Object_ptr obj = initialNamingContext->resolve( name );
ATMSessionHome_ptr atmSessionHome = ATMSessionHome.narrow( obj );
if( atmSessionHome == NULL ) {
   cerr << "Couldn't narrow to ATMSessionHome" << endl;
   exit( 1 );
}
```

Mapping for transaction

A CORBA-based enterprise bean runtime environment that wants to enable a CORBA client to participate in a transaction involving enterprise beans must use the CORBA Object Transaction Service for transaction control.

When an enterprise bean is deployed it can be installed with different transaction policies. The policy is defined in the enterprise bean's deployment descriptor.

The following rules have been defined for transactional enterprise beans: A CORBA client invokes an enterprise through stubs generated from the IDL interfaces for the enterprise bean's remote and home interface. If the client is involved in a transaction, it uses the interfaces provided by CORBA Object Transaction Service. For example, a C++ client could invoke the ATMSession bean from the previous example as follows:

```
try {
   ⋮
   // obtain transaction current
   Object_ptr obj = orb->resolve_initial_refernces("Current");
   Current current = Current.narrow( obj );
   if( current == NULL ) {
      cerr << "Couldn't resolve current" << endl;
      exit( 1 );
   }
// execute transaction
   try {
      current->begin();
      atmSession->transfer("checking", "saving", 100.00 );
      current->commit( 0 );
   }   catch( ... ) {
      current->rollback();
   }
}
catch( ... ) {
   ⋮
}
```

Mapping for security

Security aspects of the EJB specification focuses on controlling access to enterprise beans. CORBA defines a number of ways to define the identities, including the following cases:

- **Plain IIOP.** CORBA's principal interface was deprecated in early 1998. The principal interface was intended for determining the identity of a client. However, the authors of the CORBA security services implemented a different approach, GIOP.
- The GIOP specification contains a component called service context, which is an array of value pairs. The identifier is a CORBA long and the value is a sequence of octet. Among other purposes, entries in the service context can be used to identify a caller.
- **Secure IIOP.** The CORBA security specification defines an opaque data type for the identity. The real type of the identity is determined by the chosen security mechanism; for example, GSS Kerberos, SPKM, or CSI-ECMA.
- **IIOP over SSL.** SSL uses X.509 certificates to identify servers and, optionally, clients. When a server requests a client certificate, the server can use the certificate as a client identity.

11

The VisiClient Container

VisiClient is a container that provides a J2EE environment for services for application clients.

Containers are an integral part of J2EE applications. Most applications provide containers for each application type. Application clients depend on their containers to supply system services to all J2EE components.

Application Client architecture

J2EE application clients are first tier client programs that execute in their own Java virtual machines. Application clients obey the model for Java technology-based applications, in that they are invoked at their main method and run until the virtual machine is terminated. Like other J2EE application components, application clients depend on a container to provide system services; though in the case of application clients, these services are limited.

Figure 11.1 VisiClient architecture

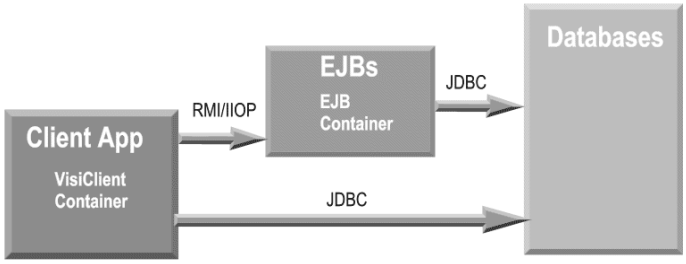

Packaging and deployment

Deploying the application client components into a VisiClient container requires the specification of deployment descriptors using XML. (Refer to J2EE 1.3 Specification for more information about application clients, and their deployment into a J2EE 1.3 compliant container.)

Application clients are packaged in JAR files and include a deployment descriptor (similar to other J2EE application components). The deployment descriptor defines the EJB and the external resources referenced by the application. You can use the Borland AppServer (AppServer) Deployment Descriptor Editor for packaging and editing application client components. For more information, see “Using the Deployment Descriptor Editor” in the *Management Console User’s Guide*.

The deployment descriptor is necessary because there are a number of functions that must be configured at deployment time, such as assigning names to EJBs and their resources. The minimum requirements for deployment of an application client into a VisiClient container are:

- All the client-side classes are packaged into a JAR. See below section on required client JARs and files. A well-formed JAR should have the following:
  - Application specific classes including the class containing the application entry point (main class)
  - The JAR file must have a META-INF subdirectory with the following:
    - A manifest file
    - A standard XML file (application-client.xml), as required by J2EE 1.3 specifications
    - A vendor-specific XML file (application-client-borland.xml)
- RMI-IIOP stubs which can also be packaged separately. In this case, the file needs the classpath attribute of the manifest file set to the appropriate value. The JAR formed in this manner is deployable to a standalone container or to an EAR file. The following sections in this chapter describe this process in detail.

Benefits of the VisiClient Container

VisiClient offers users a range of benefits from the use of J2EE applications. These include:

- **Client code portability:** Applications can use logical names (as recommended in the J2EE specifications) to access resources such as database connections, remote EJBs and environment variables. The container, per the J2EE specification, exposes the resources as administered objects in the local JNDI namespace (java:comp/env).
- **JDBC Connection Pooling:** Client applications in Borland AppServer can use JDBC 2-based datasources (factories). VisiClient Container provides connection pooling to client applications in the Server that employ a JDBC 2-based datasource. For example, the VisiClient container's application uses java.net.URL, JMS, and Mail factories.

Datasource and URL factories are deployed in the in-process local JNDI subcontext that resides in the client container virtual machine on startup. Other res-ref-types (such as JMS and Mail) are configured and deployed using the relevant tools from the vendor of these products. Refer to the Deployment, Datasources and Transaction chapters of the Borland AppServer *Developer's Guide* for more information about configuration and deployment.

Document Type Definitions (DTDs)

There are two deployment descriptors for each J2EE compliant application client module. One is a J2EE standard deployment descriptor, and the other is a vendor specific file.

The XML grammar for a J2EE application client deployment descriptor is defined in the J2EE application-client Document Type Definition (DTD). The root element of the deployment descriptor for an application client is the application-client.

Note

The content of XML elements are generally case sensitive. All valid application client deployment descriptors must contain the following DOCTYPE declaration:

```
<!DOCTYPE application-client PUBLIC '-//Sun Microsystems, Inc.//DTD J2EE
Application Client
 1.3//EN';';http://java.sun.com/j2ee/dtds/application-client_1_3.dtd'>
```

The vendor-specific deployment descriptor for an application client must contain the following DOCTYPE declaration:

```
<!DOCTYPE application-client PUBLIC "-//Borland Corporation//DTD J2EE
Application Client
 1.3//EN""http://www.borland.com/devsupport/appserver/dtds/application-
client_1_3-borland.dtd">
```

The contents of the Borland-specific application client DTD are:

```
<!ELEMENT application-client (ejb-ref*, resource-ref*, property*)>
<!ELEMENT ejb-ref (ejb-ref-name, jndi-name)>
<!ELEMENT resource-ref (res-ref-name, jndi-name)>
<!ELEMENT property (prop-name, prop-type, prop-value)>
<!ELEMENT prop-name (#PCDATA)>
<!ELEMENT prop-type (#PCDATA)>
<!ELEMENT prop-value (#PCDATA)>
<!ELEMENT ejb-ref-name (#PCDATA)>
<!ELEMENT jndi-name (#PCDATA)>
<!ELEMENT res-ref-name (#PCDATA)>
```

Here ejb-ref-name and res-ref-names are the names of the corresponding elements in the J2EE XML file, and jndi-name is the absolute JNDI name with which the object is deployed in JNDI.

Example XML using the DTD

As discussed, every application client needs a pair of XML files; a standard file and a vendor-specific file.

Example of a standard file:

```
<?xml version="1.0" encoding="ISO8859_1"?>

<!DOCTYPE application-client PUBLIC '-//Sun Microsystems, Inc.//DTD J2EE
Application Client 1.3//EN' 'http://java.sun.com/j2ee/dtds/application-
client_1_3.dtd'>
<application-client>
  <display-name>SimpleSort</display-name>
  <description>J2EE AppContainer spec compliant Sort client</description>
  <env-entry>
    <description>
      Testing environment entry
    </description>
    <env-entry-name>myStringEnv</env-entry-name>
    <env-entry-type>java.lang.String</env-entry-type>
    <env-entry-value>MyStringEnvEntryValue</env-entry-value>
  </env-entry>
  <ejb-ref>
    <ejb-ref-name>ejb/Sort</ejb-ref-name>
    <ejb-ref-type>Session</ejb-ref-type>
    <home>SortHome</home>
    <remote>Sort</remote>
    <ejb-link>sort</ejb-link>
  </ejb-ref>
  <resource-ref>
    <description>
       reference to a jdbc datasource mentioned down in the DD section
    </description>
    <res-ref-name>jdbc/CheckingDataSource</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
  </resource-ref></application-client>
```

Example of a vendor-specific file:

```
<?xml version="1.0"?>

<!DOCTYPE application-client PUBLIC "-//Borland Corporation//DTD J2EE
   Application Client 1.3//EN"
 "http://www.borland.com/devsupport/appserver/dtds/
   application-client_1_3-borland.dtd">
<application-client>
          <ejb-ref>
            <ejb-ref-name>ejb/Sort</ejb-ref-name>
            <jndi-name>sort</jndi-name>
          </ejb-ref>
          <resource-ref>
            <res-ref-name>jdbc/CheckingDataSource</res-ref-name>
<jndi-name>datasources/OracleDataSource</jndi-name>
          </resource-ref>
</application-client>
```

For more information about environment entries, ejb-refs, or resource-refs, see the relevant sections of Sun Microsystems' EJB 2.0 specifications at www.java.sun.com/j2ee.

Sample code

This example shows the usage of the logical local JNDI naming context. It shows how a client uses the deployment descriptors specified in the preceding section.

```
// get a JNDI context using the Naming service and create a remote object

    javax.naming.Context context = new javax.naming.InitialContext();
    Object ref = context.lookup("java:comp/env/ejb/Sort");
    SortHome home = (SortHome) javax.rmi.PortableRemoteObject.narrow(ref,
     SortHome.class);
    Sort sort = home.create();
    // get the value of an environment entry using JNDI
    Object envValue = context.lookup("java:comp/env/myStringEnv");
    System.out.println("Value of env entry = "+ (java.lang.String) envValue );
    // locate a  UserTransaction object
    javax.transaction.UserTransaction userTransaction =
      (javax.transaction.UserTransaction)  context.lookup("java:comp/
UserTransaction");

userTransaction.begin();
    // locate the datasource using resource-ref name
    Object resRef = context.lookup("java:comp/env/jdbc/CheckingDataSource");
    java.sql.Connection conn = ((javax.sql.DataSource)resRef).getConnection();
    //do some database work.
    userTransaction.commit();
⋮
```

Support of references and links

During application assembly and deployment you must verify that all EJB and resource references have been properly linked. For more information about EJB and resource references, consult Sun Microsystems' EJB 2.0 and J2EE 1.3 specifications.

The Borland AppServer client container supports the use of ejb-links. In the case of a standalone JAR file, the ejb-links have to be resolved before the JAR is deployed. There must be a JNDI name specified for the target bean in the vendor-specific section of the client deployment descriptor.

For a client JAR that is part of an Enterprise Application Archive (EAR), the JNDI name of the target EJB may live in a different ejb-jar. The client verify tool checks that the target EJB with the name specified in the ejb-link tag exists.

During runtime, the container resolves (locates) the target EJB corresponding to the ejb-link name in the EAR and uses the JNDI name of the target EJB. Note that application clients run in their own Java virtual machines. EJB-links are not optimized for application clients like they are for EJBs referring to another EJB located in the same container.

Keep the following rules in mind when working with EJB references and ejb-links in deployment descriptors for application client containers:

1 An ejb-ref that is not an ejb-link must have an entry in a Borland-specific file containing the JNDI name of the referenced (target) EJB.

2 An ejb-ref that has an ejb-link element must follow these rules:

- If the ejb-ref is in a client JAR and is a standalone JAR, rule 1 applies. That is, it should have a Borland-specific file with the JNDI name resolved in the deployment descriptor within the (same) JAR.
- If the ejb-ref is in a client-jar embedded in an application archive (an EAR), the JNDI name of the target EJB is not required to exist in the application-client-borland.xml file. In this case, the name in the ejb-link element is composed of a path name specifying the fully qualified path to the ejb-jar containing the referenced enterprise bean with the ejb-name of the target bean appended and separated from the path name by "#". The path name is relative to the JAR file containing the application client that is referencing the enterprise bean. This allows multiple enterprise beans with the same ejb-name to be uniquely identified.

If the path is not specified, container picks first matching EJB-name that it finds from list of EJB JARs inside EAR and throws an exception if doesn't find a bean with same name in ejb-link element.>

Using the VisiClient Container

The following command line demonstrates the use of the VisiClient Container:

```
Prompt% appclient <client-archive> [-uri <uri>] [client-arg1 client-arg2 ..]
```

The following table describes VisiClient container command line elements and definitions

Element	Definition
<client-archive>	A standalone client JAR or EAR containing client JAR.
-uri	The relative location of the client JAR inside an EAR file. This is required for EAR contained JAR files.
<client-args>	Space separated list of arguments passed to the client's main class.

VisiClient Container usage example

The following command lines demonstrate usage of an application client. In the example, the `appclient` launcher sets the classpath required to launch VisiClient.

This example is also located in the Hello example in the `install_dir/examples/j2ee/hello` directory. When your server (EJB container) is up, to run a client embedded inside an EAR file, the command is:

```
appclient me install_dir\examples\j2ee\build\hello\hello.ear -uri
helloclient.jar
```

To run a client in a standalone JAR file, the command is:

```
appclient me install_dir\examples\j2ee\build\hello\client\helloclient.jar
```

Running a J2EE client application on machines not running AppServer

To run a J2EE application client on a client machine that does not have Borland AppServer installed on it, make sure to copy the following VisiClient files to your client machine and run the following processes.

1 Copy the following JAR files from `<install_dir>/lib` to client machine:

- lm.jar
- xmlrt.jar
- asrt.jar
- vbjorb.jar
- vbsec.jar
- jsse.jar
- jaas.jar
- vbejb.jar

1 Copy the following JAR file from <install_dir>/jms/tibco/clients/java to client machine:

- tibjms.jar

1 Copy <install_dir>/bin/appclient.config to client machine.

2 Copy <install_dir>/BES/bin/appclient.exe to client machine.

To run the J2EE client using the appclient:

1 Set the PATH to appclient.exe and JDK.

2 Edit the appclient.config to change JAVA_HOME, and lib PATH.

3 Run the J2EE client from <client_application_folder>/client.

Embedding VisiClient Container functionality into an existing application

As an alternative to deploying and running a client application in the VisiClient container, it is possible to use a programmatic approach to embed the client container's functionality into an existing application. In this case, the client application can be started in a common Java fashion by running a class implementing the main() method.

To embed the VisiClient container functionality into your application, you need to call the following method:

```
public static void com.borland.appclient.Container.init
  (java.io.InputStream   deploymentDescriptorSun,
  java.io.InputStream deploymentDescriptorBorland)
throws IllegalArgumentException;
```

This method will create and populate the “java:comp/env” naming context based on the information provided in the pair of Sun and Borland deployment descriptors. The deploymentDescriptorSun and deploymentDescriptorBorland parameters must represent text XML data corresponding to the deployment descriptors. An IllegalArgumntException exception is thrown if the data provided is not recognized as a valid deployment descriptor.

Sample code

This example shows usage of this method:

```
public static void main (String[] args) {
   ⋮
     // load deployment descriptor files
     java.io.FileInputStream ddSun = new
   java.io.FileInputStream("META-INF/application-client.xml");
     java.io.FileInputStream ddBorland = new
   java.io.FileInputStream("META-INF/application-client-borland.xml");
     // initialize client container
     com.borland.appclient.Container.init(ddSun, ddBorland);
     // lookup ejb in JNDI using an ejb-ref
   javax.naming.Context context = new javax.naming.InitialContext();
     Object ref = context.lookup ("java:comp/env/ejb/hello");
   ⋮
}
```

Note

Only application client descriptors can be loaded using this method. This means that all ejb-refs must be resolved or located by specifying the jndi-name in the Borland descriptor. This cannot be done using the ejb-link in the Sun descriptor since using ejb-link requires complete knowledge of the whole application including application and EJB JAR deployment descriptors.

Use of Manifest files

VisiClient container relies on the presence of a manifest file to obtain information about launching an application. The manifest file should be saved in the META-INF subdirectory of the client archive. The relevant attributes in the manifest file for the VisiClient container are:

- The main class to be launched by the container on startup. This is an application entry point which must be present in the manifest file.
- The classpath of the dependencies of the main class. If the client-jar is self-contained, or if dependencies are specified using the system `CLASSPATH` during application launch, this attribute can be ignored.

Example of a Manifest file

An example of a Manifest file is shown below.

```
Manifest-Version: 1.0
Main-Class: SortClient
Class-Path:
```

This example shows the container will execute by loading the main method of the class specified in the `Main-Class` attribute of the Manifest file. In this example it is `SortClient`. The container expects to have a method with the following signature in this class:

```
public static void main(String[ ] args) throws Exception {...}
```

The container will report an error and exit if it doesn't find the main method. The client verify utility, which comes with VisiClient, tries to locate a main class and reports an error if it doesn't find one.

Exception handling

Application client code is responsible for taking care of any exceptions that are generated during the program execution. Any unhandled exceptions are caught by the container which will log them and terminate the Java virtual machine process.

Using resource-reference factory types

The client application deployed in a client container can use the VisiTransact JDBC connection pooling and Prepared Statement re-use facilities. Refer to the Deployment, and Transaction chapters of the Borland AppServer *Developer's Guide* for details about configuration and deployment. Client applications in AppServer can use JDBC 2-based datasources.

Note that just like `javax.sql.DataSource` (which is one of the possible res-ref-types) VisiClient allows the application to use URL, JMS, and Mail factories as the resource-ref types.

`java.net.url` and `java_mail.session` factories are deployed in the in-process local JNDI subcontext that resides in the client container virtual machine on startup. Other res-ref-types like JMS and Mail should be configured and deployed using the relevant vendor tools for these products.

Other features

The AppServer includes a number of extra features in the VisiClient container in addition to the requirements for the J2EE specification. These include:

- **User Transaction interface:** This is available in the java:comp/env name space and can be looked up using JNDI. It supports transaction demarcation, and propagation.
- **Client Verify Tool:** This runs on standalone client JARs or client JARs embedded in an EAR file. The verify tool enforces the following rules:
  - The manifest file in the client JAR has the main class specified.
  - The JAR/EAR is valid (it has the correct required manifest entries).
  - ejb-refs are valid (that is, a JNDI name for the target EJB is specified in the Borland-specific file).
  - If ejb-ref is an ejb-link, then the archive should be an EAR file. There must also be an EJB with the same name as the ejb-link value in the EAR file.
  - Resource references are valid.

Using the Client Verify tool

The following command line demonstrates the use of the Client Verify tool:

```
iastool -verify -src <srcjar> -role <DEVELOPER| ASSEMBLER| DEPLOYER>
```

Usage examples of Client Verify tool:

```
iastool -verify -src sort.jar -role DEVELOPER
iastool -verify -src sort.ear clients/sort_client.jar -role DEVELOPER
```

For more information see “verify” on the available options.

12

Caching of Stateful Session Beans

The EJB Container supports stateful session enterprise beans using a high-performance caching architecture based on the Java Session Service (JSS). There are two pools of objects: the ready pool and the passive pool. Enterprise beans transition from the ready pool to the passive pool after a configurable timeout. Transitioning an enterprise bean to the passive pool stores the enterprise bean's state in a database. Passivation of stateful sessions exists for two purposes:

1. Maximize memory resources
2. Implement failover

Configuring Borland's JSS implementation is discussed in "Java Session Service (JSS) configuration." This document explains the use of the properties that control the passivation and persistence of individual session objects.

Passivating Session Beans

At deployment time, the deployer uses the Borland AppServer's (AppServer) tools to set a passivation timeout for the EJB Container in a particular Partition. The container regularly polls active session beans to determine when they are last accessed. If a session bean has not been accessed during the timeout period, its state is sent to persistent storage and the bean instance is removed from memory.

Simple Passivation

Passivation timeouts are set at the container-level. You use the property `ejb.sfsd.passivation_timeout` to configure the length of time a session bean can go un-accessed before its state is persisted and its instance removed from memory. This length of time is specified in seconds. The default value is five seconds. This property can be set in the `partition.xml` properties file for the Partition you are configuring. This file is located in:

```
<install_dir>/var/domains/base/configurations/<configuration_name>
/ mos/<partition_name>/adm/properties
```

Edit this file to set the `ejb.sfsb.passivation_timeout` property.

If you set this property to a non-zero value, you can also set the integer property `ejb.sfsb.instance_max` for each deployed session bean in their deployment descriptors. This property defines the maximum number of instances of a particular stateful session bean that are allowed to exist in the EJB container's memory at the same time. If this

number is reached and a new instance of a stateful session needs to be allocated, the EJB container throws an exception indicating lack of resources. 0 is a special value. It means no maximum set.

If the maximum number of stateful sessions defined by the `ejb.sfsb.instance_max property` is reached, the EJB container blocks a request for an allocation of a new bean for the time defined by the integer property `ejb.sfsb.instance_max_timeout`. The container will then wait for the number to drop below this value before throwing an exception indicating a lack of resources. This property is defined in ms (1/1000th of second). 0 is a special value. It means not to wait and throw an exception indicating lack of resources immediately.

Aggressive Passivation

One of the key advantages in the use of JSS is its ability to fail over. Several containers implementing JSS can be configured to use the same persistent store, allowing them to fail over to each other. Setting up the JSS for failover is discussed in "Java Session Service (JSS) configuration." To facilitate taking advantage of the JSS failover capability, Borland provides the option of using aggressive passivation.

Aggressive passivation is the storage of session state regardless of its timeout. A bean that is set to use aggressive passivation will have its session state persisted every time it is polled, although its instance will not be removed from memory unless it times out. In this way, if a container instance fails in a cluster, a recently-stored version of the bean is available to other containers using identical JSS instances communicating with the same backend. As in simple passivation, if the bean times out, it will still be removed from memory.

Again, aggressive passivation is set Partition-wide using the boolean property `ejb.sfsb.aggressive_passivation`. Setting the property to `true` (the default) stores the session's state regardless of whether it was accessed before the last passivation attempt. Setting the property to false allows the container to use only simple passivation. Again, this property is set in the container's properties file `partition.xml` located in:

```
<install_dir>/var/domains/base/configurations/<configuration_name>
/ mos/<partition_name>/adm/properties
```

Bear in mind that although using aggressive passivation aids in failover, it also results in a performance hit since the container accesses the database more often. If you configure the JSS to use a non-native database (that is, you choose not to use JDataStore), the loss of performance can be even greater. Be aware of the tradeoff between availability and performance before you elect to use aggressive passivation.

Sessions in secondary storage

Most sessions are not kept in persistent storage forever after they timeout. Borland provides a mechanism for removing stored sessions from the database after a discrete period of time known as the *keep alive timeout*. The keep alive timeout specifies the minimum amount of time in seconds to persist a passivated session in stateful storage. The actual amount of time it is kept in the database can vary, since it is not wise from a performance standpoint to constantly poll the database for unused sessions. The actual amount of time a session is persisted is at least the value of the keep alive timeout and not more than twice the value of the keep alive timeout.

Unlike the other passivation properties discussed above, the keep alive timeout can be specified either Partition-wide and/or on the individual session bean. If you set a keep alive timeout for a specific bean, its value will take precedence over any container-wide values. If you do not specify a keep alive timeout for a particular bean, it will use the Partition-wide value.

Setting the keep alive timeout in Containers

The Borland JSS implementation uses the property `ejb.sfsb.keep_alive_timeout` to specify the amount of time (in seconds) to maintain a passivated session in stateful storage. The default value is 86,400 seconds, or twenty-four hours. Like the other properties discussed above, you set the keep alive timeout in the container properties file:

```
<install_dir>/var/domains/base/configurations/<configuration_name>
/ mos/<partition_name>/adm/properties
```

Remember that any value you specify here can be overridden by setting a keep alive timeout for a specific session bean.

Setting the keep alive timeout for a particular session bean

You may wish to have certain session beans hosted in your container have their passivated states stored for greater or lesser periods of time than others. You can use the `<timeout>` element in the `ejb-borland.xml` file to set the keep alive timeout for a particular bean. The DTD element for a session bean provides this element:

```
<!ELEMENT session (ejb-name, bean-home-name?, bean-local-home-name?, timeout?,
ejb-ref*, ejb-local-ref*, resource-ref*, resource-env-ref*, property*)>
```

For example, let's say we have a simple stateful session bean called `personInfo` collecting a bit of personal information for simple message forum. We might be inclined to keep this session highly-available, without aggressive passivation, and have little need to store it in our database for more than a few minutes if it passivates. Since the rest of our session beans need to be kept in stateful storage a bit longer if they passivate, we'll use the Borland-specific deployment descriptor for the bean's JAR to set a shorter keep alive timeout, say 300 seconds (five minutes). In our `ejb-borland.xml` deployment descriptor, we'd have the following:

```
<ejb-jar>
 <enterprise-beans>
  <session>
   <ejb-name>personInfo</ejb-name>
   <timeout>300</timeout>
  </session>
 </enterprise-beans>
</ejb-jar>
```

This value will override any values we entered in the `ejbcontainer.properties` file while allowing other hosted sessions to use the default value found there.

13

Entity Beans and CMP 1.1 in Borland AppServer

Here we'll examine how entity beans are deployed in the Borland AppServer (AppServer) and how persistence of entities can be managed. This is not, however, an introduction to entity beans and should not be treated as such. Rather, this document will explore the implications of using entity beans within Borland Partitions. We'll discuss descriptor information, persistence options, and other container-optimizations. Information on the Borland-specific deployment descriptors and implementations of Container-Managed Persistence (CMP) will be documented in favor of general EJB information that is generally available from the J2EE Specifications from Sun Microsystems.

Entity Beans

Entity beans represent a view of data stored in a database. Entity beans can be fine-grained entities mapping to a single table with a one-to-one correspondence between entity beans and table rows. Or, entity beans can span multiple tables and present data independent of the underlying database schema. Entity beans can have relationships with one another, can be queried for data by clients, and can be shared among different clients.

Deploying your Entity Bean to one of the AppServer Partitions requires that it be packaged as a part of a JAR. The JAR must include two descriptor files: `ejb-jar.xml` and the proprietary `ejb-borland.xml` file. The `ejb-jar.xml` descriptor is fully-documented at the Sun Java Center. The DTD for `ejb-borland.xml` is reproduced in this document and its usage documented here. The Borland proprietary descriptor contains a number of properties that can be set to optimize container performance and manage the persistence of your entity beans.

Container-managed persistence and Relationships

Borland's EJB container provides tools that generate the database access calls at the time that the entity bean is deployed; that is, when the entity bean is installed into a Partition. The tools use the deployment descriptors to determine the instance fields for which they must generate database access calls. Instead of coding the database access directly in the bean, the bean provider of a container-managed entity bean must specify in the deployment descriptor those instance fields for which the container tools

must generate access calls. The container has sophisticated deployment tools capable of mapping the fields of an entity bean to its data source.

Container-managed persistence has many advantages over bean-managed persistence. It is simpler to code because bean provider does not have to code the database access calls. Handling of persistence can also be changed without having to modify and recompile the entity bean code. The Deployer or Application Assembler can do this by modifying the deployment descriptor when deploying the entity bean. Shifting the database access and persistence handling to the container not only reduces the complexity of code in the bean, it also reduces the scope of possible errors. The bean provider can focus on debugging the business logic of the bean rather than the underlying system issues.

The EJB 2.0 specification allows entity beans that use container-managed persistence to also have container-managed relationships among themselves. The container automatically manages bean relationships and maintain the referential integrity of these relationships. This differs from the EJB 1.1 specification, which only allowed you to expose a bean's instance state through its remote interface.

Just as you defined container-managed persistence fields in a bean's deployment descriptor, you can now define container-managed relationship fields in the deployment descriptor. The container supports relationships of various cardinalities, including one-to-one, one-to-many, and many-to-many.

Implementing an entity bean

Implementing an entity bean follows the rules defined in the EJB 1.1 and 2.0 specifications. You must implement a home interface, a remote interface or a local interface (if using the 2.0 container-managed persistence), and the entity bean implementation class. The entity bean class must implement the methods that correspond to those declared in the remote or local and home interfaces.

Packaging Requirements

Like session beans, entity beans can expose their methods with their interfaces. Each Entity Bean must also have corresponding entries in its JAR's deployment descriptors. The standard deployment descriptor, `ejb-jar.xml` contains essentially three different types of deployment information. These are:

1. **General Bean Information:** This corresponds to the `<enterprise-beans>` elements found in the descriptor file and is used for all three types of beans. This information also includes information on the bean's interfaces and class, security information, environmental information, and even query declarations.
2. **Relationships:** This corresponds to the `<relationships>` elements found in the descriptor file and applies to entity beans using CMP only. This is where container-managed relationships are spelled out.
3. **Assembly Information:** This corresponds to the `<assembly-descriptor>` element which explains how the beans interact with the application as a whole. Assembly information is broken down into four categories:
   - Security Roles: simple definitions of security roles used by the application. Any security role references you defined for your beans must also be defined here.
   - Method Permissions: each method of each bean can have certain rules about their execution. These are set here.
   - Container-Transactions: this specifies the transaction attributes as per the EJB 2.0 specification for each method participating in a transaction
   - Exclude List: methods to be uncalled by anyone

All of these can be accessed through the Deployment Descriptor Editor. You should refer to the EJB 2.0 specification for DTD information and the proper use of the descriptor files.

Entity Bean Primary Keys

Each Entity Bean must have a unique primary key that used to identify the bean instance. The primary key can be represented by a Java class that must be a legal value type in RMI-IIOP. Therefore, it extends the `java.io.Serializable` interface. It must also provide an implementation of the `Object.equals(Object other)` and `Object.hashCode()` methods.

Normally, the primary key fields of entity beans must be set in the `ejbCreate()` method. The fields are then used to insert a new record into the database. This can be a difficult procedure, however, bloating the method, and many databases now have built-in mechanisms for providing appropriate primary key values. A more elegant means of generating primary keys is for the user to implement a separate class that generates primary keys. This class can also implement database-specific programming logic for generating primary keys.

Generating primary keys from a user class

With enterprise beans, the primary key is represented by a Java class containing the unique data. This primary key class can be any class as long as that class is a legal value type in RMI-IIOP, meaning it extends the `java.io.Serializable` interface. It must also provide an implementation of the `Object.equals(Object other)` and `Object.hashCode()` methods, two methods which all Java classes inherit by definition.

The primary key class can be specific to an particular entity bean class. That is, each entity bean can define its own primary key class. Or, multiple entity beans can share the same primary key class.

The bank application uses two different entity beans to represent savings and checking accounts. Both types of accounts use the same field to uniquely identify a particular account record. In this case, they both use the same primary key class, AccountPK, to represent the unique identifier for either type of account. The following code shows the definition of the account primary key class:

```
public class AccountPK implements java.io.Serializable {
   public String name;
   public AccountPK() {}
   public AccountPK(String name) {
      this.name = name;
   }
}
```

Generating primary keys from a custom class

To generate primary keys from a custom class, you must write a class that implements the `com.borland.ejb.pm.PrimaryKeyGenerationListener` interface.

Support for composite keys

Primary keys are not restricted to a single column. Sometimes, a primary key is composed of more than one column. In a more realistic example, a course is not identified merely by its name. Instead, the primary key for each course record can be the department in which the course is offered and the course number itself. The department code and the course number are separate columns in the Course table. A select statement that retrieves a particular course, or all courses in which a student is enrolled, must use the entire primary key; that is, it must consider both columns of the primary key.

The Borland CMP engine supports composite primary keys. You can use keys with multiple columns in the where clause of a select statement. You can also select all fields of a compound key in the select clause portion of the statement.

For the where clause, specify multiple field names in the same manner that you specify single field names. Use “and” to separate each field. The format is

```
<column> = :<parameter>[ejb/<entity bean>]
```

Note that the equal (=) sign is one of several possible notations. You could also specify greater than (>), less than (<), greater than or equal (>=), or less than or equal (<=). The colon (:) notation indicates parameter substitution. The parameter field is specified with the bean name first, followed by a dot (.), then the bean attribute.

For example, to find all students taking Art 205, Renaissance Art where classes are identified by the department (Art) and the course number (205), you might have the following select statement defined for the finder method `findByCourse()`:

```
SELECT sname FROM Enrollment WHERE course_department = :c.department[ejb/
Course] AND
   course_number = :c.number[ejb/Course]
```

You can also have the select statement return multiple fields from a compound key. In the select clause of the select statement, list the fields, separated by commas. Note that you use the same dot notation as for parameters; that is, specify the entity bean name, followed by a dot (.), then the attribute name. For example, the finder method `findByStudent()` can have the following select statement:

```
SELECT c.department, c.number FROM Enrollment WHERE student_name = :s
```

Reentrancy

By default, entity beans are not reentrant. When a call within the same transaction context arrives at the entity bean, it causes the exception `java.rmi.RemoteException` to be thrown.

You can declare an entity bean reentrant in the deployment descriptor; however, take special care in this case. The critical issue is that a container can generally not distinguish between a (loopback) call within the same transaction and a concurrent invocation (in the same transaction context) on that same entity bean.

When the entity bean is marked reentrant, it is illegal to allow a concurrent invocation within the same transaction context on the bean instance. It is the programmer's responsibility to ensure this rule.

Container-Managed Persistence in AppServer

The AppServer's EJB Container is fully J2EE 1.3 compliant. The bean provider designs persistence schemas for their entity beans, determined the methods for accessing container-managed fields and relationships, and defines these in the beans' deployment descriptor. The deployer maps this persistence schema to the database and creates any other necessary classes for the beans' maintenance.

Information on J2EE 1.3 entity beans and CMP 2.0 is found in the “Using Borland AppServer Properties for CMP 2.x.”

AppServer CMP engine's CMP 1.1 implementation

While you don't have to be an expert on all aspects of the Borland CMP engine to use it effectively, it is helpful to have some knowledge of certain areas. This section provides information on the areas that users of the CMP engine should understand. In particular, it focuses on the deployment descriptor file and the XML statements contained within the file.

Before continuing, there are some key things to note in the implementation of an entity bean that uses 1.1 container-managed persistence:

- The entity bean has no implementations for finder methods. The EJB Container provides the finder method implementations for entity beans with container-managed persistence. Rather than providing the implementation for finder methods in the bean's class, the deployment descriptor contains information that enables the container to implement these finder methods.
- The entity bean declares all fields `public` that are managed by the container for the bean. The CheckingAccount bean declares `name` and `balance` to be public fields.

- The entity bean class implements the seven methods declared in the `EntityBean` interface: `ejbActivate()`, `ebjPassivate()`, `ejbLoad()`, `ejbStore()`, `ejbRemove()`, `setEntityContext()`, and `unsetEntityContext()`. However, the entity bean is required to provide only skeletal implementations of these methods, though it is free to add application-specific code where appropriate. The CheckingAccount bean saves the context returned by `setEntityContext()` and releases the reference in `unsetEntityContext()`. Otherwise, it adds no additional code to the `EntityBean` interface methods.
- There is an implementation of the `ejbCreate()` method (because this entity bean allows callers of the bean to create new checking accounts), and the implementation initializes the instance's two variables, account name and balance amount, to the argument values. The `ejbCreate()` method returns a `null` value because, with container-managed persistence, the container creates the appropriate reference to return to the client.
- The entity bean provides the minimal implementation for the `ejbPostCreate()` method, though this method could have performed further initialization work if needed. For beans with container-managed persistence, it is sufficient to provide just the minimal implementation for this method because `ejbPostCreate()` serves as a notification callback. Note that the same rule applies to the methods inherited from the `EntityBean` interface as well.

Providing CMP metadata to the Container

According to the EJB Specification, the deployer must provide CMP metadata to the EJB container. The Borland Container captures the CMP-relevant metadata in the XML deployment descriptor. Specifically, the Borland Container uses the vendor-specific portion of the deployment descriptor for the CMP metadata.

This section illustrates some of the information that needs to be provided for container-managed finder methods, particularly if you are constructing container-managed finder methods at the command line level. Because it is not an exhaustive reference, you should refer to the DTD of the deployment descriptor for the detailed syntax. Look for the syntax for the finder methods and Object-Relation (OR) mapping metadata.

Constructing finder methods

When you construct a finder method, you are actually constructing an SQL select statement with a where clause. The select statement includes a clause that states what records or data are to be found and returned. For example, you might want to find and return checking accounts in a bank. The where clause of the select statement sets limits on the selection process; that is, you might want to find only those checking accounts with a balance greater than some specified amount, or accounts with a certain level of activity per month. When the Container uses container-managed persistence, you must specify the terms of the where clause in the deployment descriptor.

For example, suppose you have a finder method called `findAccountsLargerThan(int balance)` and you are using container-managed persistence. This finder method attempts to find all bank accounts with a balance greater than the specified value. When the Container executes this finder method, it actually executes a select statement whose where clause tests the account balances against the `int` value passed as a parameter to the method. Because we're using container-managed persistence, the deployment descriptor needs to specify the conditions of the where clause; otherwise, the Container does not know how to construct the complete select statement.

The value of the where clause for the `findAccountsLargerThan(int balance)` method is "`balance > :balance`". In English, this translates to: "the value of the balance column is greater than the value of the parameter named `balance`." (Note that there is only one argument to the finder method, an `int` value.)

The default container-managed persistence implementation supports this finder method by constructing the complete SQL select statement, as follows:

```
select * from Accounts where ? > balance
```

The CMP engine then substitutes "?" with the `int` parameter. Lastly, the engine converts the result set into either an `Enumeration` or `Collection` of primary keys, as required by the EJB Specification.

It is possible to inspect the various SQL statements that the CMP implementation constructs. To do this, enable the `EJBDebug` flag on the container. When that flag is enabled, it prints the exact statements constructed by the Container.

While other EJB Container products use code generation to support CMP, the Borland Container does not use code generation because it has serious limitations. For example, code generation makes it difficult to support a "tuned update" feature, because of the great number of different update statements to container-managed fields that are required.

Constructing the where clause

The `where` clause is a necessary part of select statements when you want to delimit the extent of the returned records. Because the `where` clause syntax can be fairly complex, you must follow certain rules in the XML deployment descriptor file so that the EJB Container can correctly construct this clause.

To begin with, you are not obligated to use the literal "where" in your `<where-clause>`. You can construct a where clause without this literal and rely on the Container to supply it. However, the Container only does this if the `<where-clause>` is not an empty string; it leaves empty strings empty. For example, you could define a where clause as either:

```
<where-clause> where a = b </where-clause>
```

or:

```
<where-clause> a = b </where-clause>
```

The Container converts `a = b` to the same `where` clause, `where a = b`. However, it leaves unmodified an empty string defined as `<where-clause>` "" `</where-clause>`.

Note

The empty string makes it easy to specify the `findAll()` method. When you specify just an empty string, the Container construes that to mean the following:

```
select [values] from [table];
```

Such a select statement would return all values from a particular table.

Parameter substitution

Parameter substitution is an important part of the `where` clause. The Borland EJB Container does parameter substitution wherever it finds the standard SQL substitution prefix colon (`:`). Each parameter for substitution corresponds to a name of a parameter in the finder specification found in the XML descriptor.

For example, in the XML deployment descriptor, you might define the following finder method which takes a parameter `balance` (note that `balance` is preceded by a colon):

```
<finder>
   <method-signature>findAccountsLargerThan(float balance)</method-signature>
   <where-clause>balance > :balance</where-clause>
</finder>
```

The Container composes a SQL select statement whose `where` clause is:

```
balance > ?
```

Note that the `:balance` parameter in the deployment descriptor becomes a question mark (`?`) in the equivalent SQL statement. When invoked, the Container substitutes the value of the parameter `:balance` for the `?` in the `where` clause.

Compound parameters

The Container also supports compound parameters; that is, the name of a table followed by a column within the table. For this, it uses the standard dot (`.`) syntax,

where the table name is separated from the column name by a dot. These parameters are also preceded by a colon.

For example, the following finder method has the compound parameters `:address.city` and `:address.state`:

```
<finder>
   <method-signature>findByCity(Address address)</method-signature>
   <where-clause>city = :address.city AND state = :address.state</where-clause>
</finder>
```

The `where` clause uses the `city` and `state` fields of the `address` compound object to select particular records. The underlying `Address` object could have Java Beans-style getter methods that correspond to the attributes `city` and `state`. Or, alternatively, it could have public fields that correspond to the attributes.

Entity beans as parameters

An entity bean can also serve as a parameter in a finder method. You can use an entity bean as a compound type. To do so, you must tell the CMP engine which field to use from the entity bean's passed reference to the SQL query. If you do not use the entity bean as a compound type, then the Container substitutes the bean's primary key in the `where` clause.

For example, suppose you have a set of `OrderItems` entity beans associated with an `Order` entity object. You might have the following finder method:

```
java.util.Collection OrderItemHome.findByOrder(Order order);
```

This method returns all `OrderItems` associated with a particular `Order`. The deployment descriptor entry for its `where` clause would be:

```
<finder>
   <method-signature>findByOrder(Order order)</method-signature>
   <where-clause>order_id = :order[ejb/orders]</where-clause>
</finder>
```

To produce this where clause, the Container substitutes the primary key of the `Order` object for the string `:order[ejb/orders]`. The string between the brackets (in this example, `ejb/orders`) must be the `<ejb-ref>` corresponding to the home of the parameter type. In this example, `ejb/orders` corresponds to an `<ejb-ref>` pointing to `OrderHome`.

When you use an EJBObject as a compound type (using the dot notation), you are actually accessing the underlying get method for the field in the `<finder>` definition. For example, the following in the `<finder>` definition:

```
order_id = :order.orderId
```

calls the `getOrderId()` method on the order `EJBObject` and uses the result of the call in the selection criterion.

Specifying relationships between entities

Relational databases (RDBMS) permit records in one table to be associated with records in another table. The RDBMS accomplishes this using foreign keys; that is, a record in one table maintains a field (or column) that is a foreign key or reference to (usually) the primary key of a related record in another table. You can map these same references among entity beans.

For the CMP engine to map references among entity beans, you use an `<ejb-link>` entry in the deployment descriptor. The `<ejb-link>` maps field names to their corresponding entities. The CMP engine uses this information in the deployment descriptor to locate the field's associated entity. (Refer to the pigs example for an illustration of the `<ejb-link>` entry.)

Any container-managed persistence field can correspond to a foreign key field in the corresponding table. When you look at the entity bean code, these foreign key CMP fields appear as object references.

For example, suppose you have two database tables, an `address` table and a `country` table. The `address` table contains a reference to the `country` table. The SQL `create` statements for these tables might look as shown below.

```
create table address (
   addr_id         number(10),
   addr_street1    varchar2(40),
   addr_street2    varchar(40),
   addr_city       varchar(30),
   addr_state       varchar(20),
   addr_zip         varchar(10),
   addr_co_id       number(4)         * foreign key *
);
create table country (
   co_id            number(4),
   co_name          varchar2(50),
   co_exchange      number(8, 2),
   co_currency      varchar2(10)
);
```

Note that the `address` table contains the field `addr_co_id`, which is a foreign key referencing the `country` table's primary key field, `co_id`.

There are two classes that represent the entities which correspond to these tables, the `Address` and `Country` classes. The `Address` class contains a direct pointer, `country`, to the `Country` entity. This direct pointer reference is an EJBObject reference; it is not a direct Java reference to the implementation bean.

Now examine the code for both classes:

```
//Address Class
public class Address extends EntityBean {
   public int id;
   public String street1;
   public String street1;
   public String city;
   public String state;
   public String zip;
   public Country country; // this is a direct pointer to the Country
}
//Country Class
public class Country extends EntityBean {
   public int id;
   public String name;
   public int exchange;
   public String currency;
}
```

In order for the Container to resolve the reference from the `Address` class to the `Country` class, you must specify information about the `Country` class in the deployment descriptor. Using the `<ejb-link>` entry in the deployment descriptor, you instruct the Container to link the reference to the field `Address.country` to the JNDI name for the home object, `CountryHome`. (Look at the pigs example for a more detailed explanation.) The container optimizes this cross-entity reference; because of the optimization, using the cross reference is as fast as storing the value of the foreign key.

However, there are two important differences between using a cross reference and storing the foreign key value:

– When you use a cross reference pointer to another entity, you do not have to call the other entity's home object `findByPrimaryKey()` method to retrieve the corresponding object entity. Using the above example as an illustration, the `Address.country` pointer to the `Country` object lets you retrieve the country object directly. You do not have to call `CountryHome.findByPrimaryKey(address.country)` to get the `Country` object that corresponds to the country `id`.

– When you use a cross reference pointer, the state of the referenced entity is only loaded when you actually use it. It is not automatically loaded when the entity containing the pointer is loaded. That is, merely loading in an `Address` object does not actually load in a `Country` object. You can think of the `Address.country` field as a "lazy" reference, though when the underlying object is actually used does a "lazy" reference load in its corresponding state. (Note that this "lazy" behavior is a part of the EJB model.) This facet of the EJB model results in the decoupling of the life cycle of `Address.country` from the life cycle of the `Address` bean instance itself. According to the model, `Address.country` is a normal entity EJBObject reference; thus, the state of `Address.country` is only loaded when and if it is used. The Container follows the EJB model and controls the state of `AddressBean.country` as it does with any other EJBObject.

Container-managed field names

The Borland Container has changed the container-managed persistent field names so that they are more Java friendly. SQL column names often prepend a shortened form of the table name, followed by an underscore, to each column name. For example, in the `address` table, there is a column for the city called `addr_city`. The full reference to this column is `address.addr_city`. With the Borland Container, this maps to the Java field `Address.city`, rather than the more redundant and more awkward `Address.addr_city`.

You can achieve this Java-friendly column-to-field-name mapping using the deployment descriptor. While this section shows you how to manually edit the deployment descriptor, it is best to use the Deployment Descriptor Editor GUI to accomplish this. See "Using the Deployment Descriptor Editor" in the *Management Console User's Guide* for instructions on using the GUI screens.

Should you choose to manually edit the deployment descriptor, use the `<env-entry-name>`, `<env-entry-type>`, and `<env-entry-value>` subtags within the `<env-entry>` tag. Place the more friendly Java field name in the `<env-entry-name>` tag, noting that it is referencing a JDBC column. Put the type of the field in the `<env-entry-type>` tag. Lastly, place the actual SQL column name in the `<env-entry-value>` tag. The following deployment descriptor code segment illustrates this:

```
<env-entry>
   <env-entry-name>ejb.cmp.jdbc.column:city</env-entry-name>
   <env-entry-type>String</env-entry-type>
   <env-entry-value>addr_city</env-entry-value>
</env-entry>
```

Setting Properties

Most properties for Enterprise JavaBeans can be set in their deployment descriptors. The Borland Deployment Descriptor Editor (DDEditor) also allows you to set properties and edit descriptor files. Use of the Deployment Descriptor Editor is described in the Borland AppServer *Management Console User's Guide*. Use properties in the deployment descriptor to specify information about the entity bean's interfaces, transaction attributes, and so forth, plus information that is unique to an entity bean. In addition to the general descriptor information for entity beans, here are also three sets of properties that can be set to customize CMP implementations, entity properties, table properties, and column properties. Entity properties can be set either by using the EJB Designer or in the XML directly.

Using the Deployment Descriptor Editor

You can use the Deployment Descriptor Editor, which is part of the Borland AppServer to set up all of the container-managed persistence information. You should refer to the *Management Console User's Guide* for complete information on the use of the Deployment Descriptor Editor and other related tools.

J2EE 1.2 Entity Bean using BMP or CMP 1.1

Descriptor Element	Navigation Tree Node/Panel Name	DDEditor Tab
Entity Bean name	Bean	General
Entity Bean class	Bean	General
Home Interface	Bean	General
Remote Interface	Bean	General
Home JNDI Name	Bean	General
Persistence Type (CMP or BMP)	Bean	General
Primary Key Class	Bean	General
Reentrancy	Bean	General
Icons	Bean	General
Environment Entries	Bean	Environment
EJB References to other Beans	Bean	EJB References
EJB Links	Bean	EJB References
Resource References to data objects/connection factories	Bean	Resource References
Resource Reference type	Bean	Resource References
Resource Reference Authentication Type	Bean	Resource References
Security Role References	Bean	Security Role References
Entity Properties	Bean	Properties
Container Transactions	Bean:Container Transactions	Container Transactions
Transactional Method	Bean:Container Transactions	Container Transactions
Transactional Method Interface	Bean:Container Transactions	Container Transactions
Transactional Attribute	Bean:Container Transactions	Container Transactions
Method Permissions	Bean:Method Permissions	Method Permissions
CMP Description	Bean:CMP1.1	CMP 1.1
CMP Tables	Bean:CMP1.1	CMP 1.1
Container-Managed Fields Description	Bean:CMP1.1	CMP 1.1
Finders	Bean:CMP1.1	Finders
Finder Method	Bean:CMP1.1	Finders
Finder WHERE Clause	Bean:CMP1.1	Finders
Finder Load State option	Bean:CMP1.1	Finders

Container-managed data access support

For container-managed persistence, the Borland EJB Container supports all data types supported by the JDBC specification, plus some other types beyond those supported by JDBC.

The following shows the basic and complex types supported by the Borland EJB Container:

- Basic types:
  - `boolean Boolean`
  - `double Double`
  - `long Long`
  - `BigDecimal java.util.Date`
  - `short Short`
  - `byte[]`
  - `char Character`
  - `int Integer`

- byte Byte
- float Float
- String java.sql.Date
- java.sql.Time java.sql.TimeStamp

- Complex types
  - Any class implementing `java.io.Serializable`, such as Vector and Hashtable
  - Other entity bean references

Keep in mind that the Borland Container supports classes implementing the `java.io.Serializable` interface, such as Hashtable and Vector. The container supports other data types, such as Java collections or third party collections, because they also implement `java.io.Serializable`. For classes and data types that implement the `Serializable` interface, the Container merely serializes their state and stores the result into a `BLOB`. The Container does not do any “smart” mapping on these classes or types; it just stores the state in binary format. The Container's CMP engine observes the following rule: the engine serializes as a `BLOB` all types that are not one of the explicitly supported types.

In this context, the Container follows the JDBC specification: a `BLOB` is the type to which `LONGVARBINARY` maps. (For Oracle, this is `LONG RAW`.)

Using SQL keywords

The CMP engine for the Borland Container can handle all SQL keywords that comply with the SQL92 standard. However, you should keep in mind that vendors frequently add their own keywords. For example, Oracle uses the keyword `VARCHAR2`. If you want to ensure that the CMP engine can handle vendor keywords that may differ from the SQL standard, set up an environment property in the deployment descriptor that maps the CMP field name to the column name. By using this sort of environment property, you do not have to modify your code.

For example, suppose you have a CMP field called “select”. You can use the following environment property to map “select” to a column called “SLCT”, as shown below.

```
<cmp-info>
         <database-map>
            <table>Data</table>
            <column-map>
               <field-name>select</field-name>
               <column-name>SLCT</column-name>
               </column-map>
            </database-map>
</cmp-info>
```

Using null values

It is possible that your database values can contain SQL null values. If so, you must map them to fields whose Java data types are permitted to contain Java null values. Typically, you do this by using Java types instead of primitive types. Thus, you use a Java `Integer` type rather than a primitive `int` type, or a Java `Float` type rather than a primitive `float` type.

Establishing a database connection

You must specify a DataSource so that the CMP engine can open a database connection. The DataSource defines the information necessary for establishing a database connection, such as username and password. Define a DataSource and then use a `resource-ref` to refer to the DataSource in the XML deployment descriptor for the bean. The CMP engine can then use the DataSource to access the database via JDBC.

At the point in the vendor-specific XML file where you provide the `jndi` binding for the `resource-ref`, add the element

```
<cmp-resource>True</cmp-resource>
```

For cases where the entity bean declares only one `resource-ref`, you do not need to provide the above XML element. When the entity bean has only one `resource-ref`, the Borland Container knows to automatically choose that one resource as the `cmp-resource`.

Container-created tables

You can instruct the Borland EJB Container to automatically create tables for container-managed entities based on the entity's container-managed fields. Because table creation and data type mappings vary among vendors, you must specify the JDBC database dialect in the deployment descriptor to the Container. For all databases (except for JDataStore) if you specify the dialect, then the Container automatically creates tables for container-managed entities for you. The Container will not create these tables unless you specify the dialect.

However, for the JDataStore database, the Container can detect the dialect from the URL for the JDataStore database. Thus, for JDataStore, the Container will create these tables regardless of whether you explicitly specify the dialect.

The following table shows the names or values for the different dialects (case is ignored for these values):

Database Name	Dialect Value
JDataStore	jdatastore
Oracle	oracle
Sybase	sybase
MSSQLServer	mssqlserver
DB2	db2
Interbase	interbase
Informix	informix
No database	none

Mapping Java types to SQL types

When you develop an enterprise bean for an existing database, you must map the SQL data types specified in the database schema to Java programming language data types.

The Borland EJB Container follows the JDBC rules for mapping Java programming language types to SQL types. JDBC defines a set of generic SQL type identifiers that represent the most commonly used SQL types. You must use these default JDBC mapping rules when you develop an enterprise bean to model an existing database table. (These types are defined in the class `java.sql.Types`.)

The following table shows the default SQL to Java type mapping as defined by the JDBC specification.

Java type	JDBC SQL type
`boolean/Boolean`	`BIT`
`byte/Byte`	`TINYINT`
`char/Character`	`CHAR(1)`
`double/Double`	`DOUBLE`
`float/Float`	`REAL`
`int/Integer`	`INTEGER`
`long/Long`	`BIGINT`
`short/Short`	`SMALLINT`
`String`	`VARCHAR`
`java.math.BigDecimal`	`NUMERIC`
`byte[]`	`VARBINARY`
`java.sql.Date`	`DATE`
`java.sql.Time`	`TIME`
`java.sql.Timestamp`	`TIMESTAMP`
`java.util.Date`	`TIMESTAMP`
`java.io.Serializable`	`VARBINARY`

Automatic table mapping

The Borland EJB container has the capability to automatically map Java types defined in the enterprise bean code to database table types. However, while it may create these tables automatically, it does not necessarily use the most optimal mapping approach. In fact, automatically generating these mappings and tables is more of a convenience for developers.

The Borland-generated tables are not optimized for performance. Often, they overuse database resources. For example, the container maps a Java `String` field to the corresponding SQL `VARCHAR` type. However, the mapping is not sensitive to the actual length of the Java field, and so it maps all string fields to the maximum `VARCHAR` length. Thus, it might map a two-character Java `String` to a `VARCHAR(2000)` column.

In a production situation, it is preferable for database administrators (DBA) to create the tables and do the type mapping. The DBA can override the default mappings and produce a table optimized for performance and use of database resources.

While all relational databases implement SQL types, there may be significant variations in how they implement these types. Even when they support SQL types with the same semantics, they may use different names to identify these types. For example, Oracle implements a Java `boolean` as a`NUMBER(1,0)`, while Sybase implements it as a `BIT` and DB2 implements it as a `SMALLINT`.

When the Borland EJB Container creates the database tables for your enterprise beans, it automatically maps entity bean fields and database table columns. The container must know how to properly specify the SQL types so that it can correctly create the tables in each supported database. As a result, the EJB Container maps some Java types differently, depending on the database in use. The following table shows the mapping for Oracle, Sybase/MSSQL, and DB2:

Java types	Oracle	Sybase/MSSQL	DB2
boolean/Boolean	NUMBER(1,0)	BIT	SMALLINT
byte/Byte	NUMBER(3,0)	TINYINT	SMALLINT
char/Character	CHAR(1)	CHAR(1)	CHAR(1)
double/Double	NUMBER	FLOAT	FLOAT
float/Float	NUMBER	REAL	REAL
int/Integer	NUMBER(10,0)	INT	INTEGER
long/Long	NUMBER(19,0)	NUMERIC(19,0)	BIGINT
short/Short	NUMBER(5,0)	SMALLINT	SMALLINT
String	VARCHAR(2000)	TEXT	VARCHAR(2000)
java.math.BigDecimal	NUMBER(38)	DECIMAL(28,28)	DECIMAL
byte[]	LONG RAW	IMAGE	BLOB
java.sql.Date	DATE	DATETIME	DATE
java.sql.Time	DATE	DATETIME	TIME
java.sql.Timestamp	DATE	DATETIME	TIMESTAMP
java.util.Date	DATE	DATETIME	TIMESTAMP
java.io.Serializable	RAW(2000)	IMAGE	BLOB

The following table shows the Java to SQL type mapping for JDatastore, Informix, and Interbase:

Java types	JDatastore	Informix	Interbase
boolean/Boolean	BOOLEAN	SMALLINT	SMALLINT
byte/Byte	SMALLINT	SMALLINT	SMALLINT
char/Character	CHAR(1)	CHAR(1)	CHAR(1)
double/Double	DOUBLE	FLOAT	DOUBLE PRECISION
float/Float	FLOAT	SMALLFLOAT	FLOAT
int/Integer	INTEGER	INTEGER	INTEGER
long/Long	LONG	DECIMAL(19,0)	NUMBER(15,0)
short/Short	SMALLINT	SMALLINT	SMALLINT
String	VARCHAR	VARCHAR(2000)	VARCHAR(2000)
java.math.BigDecimal	NUMERIC	DECIMAL(32)	NUMBER(15,15)
byte[]	OBJECT	BYTE	BLOB
java.sql.Date	DATE	DATE	DATE
java.sql.Time	TIME	DATE	DATE
java.sql.Timestamp	TIMESTAMP	DATE	DATE
java.util.Date	TIMESTAMP	DATE	DATE
java.io.Serializable	OBJECT	BYTE	BLOB

14

Entity Beans and Table Mapping for CMP 2.x

Here we'll examine how entity beans are deployed in the Borland AppServer (AppServer) and how persistence of entities can be managed. This is not, however, an introduction to entity beans and should not be treated as such. Rather, this document will explore the implications of using entity beans within Borland Partitions. We'll discuss descriptor information, persistence options, and other container-optimizations. Information on the Borland-specific deployment descriptors and implementations of Container-Managed Persistence (CMP) will be documented in favor of general EJB information that is generally available from the J2EE Specifications from Sun Microsystems.

Entity Beans

Entity beans represent a view of data stored in a database. Entity beans can be fine-grained entities mapping to a single table with a one-to-one correspondence between entity beans and table rows. Or, entity beans can span multiple tables and present data independent of the underlying database schema. Entity beans can have relationships with one another, can be queried for data by clients, and can be shared among different clients.

Deploying your Entity Bean to one of the AppServer Partitions requires that it be packaged as a part of a JAR. The JAR must include two descriptor files: `ejb-jar.xml` and the proprietary `ejb-borland.xml` file. The `ejb-jar.xml` descriptor is fully-documented in the J2EE 1.3 Specification. The DTD for `ejb-borland.xml` is reproduced in this document and aspects of its usage documented here. The Borland proprietary descriptor allows for the configuration of a number of properties that can be set to optimize container performance and manage the persistence of your entity beans.

Container-managed persistence and Relationships

Borland's EJB container provides tools that generate the persistence calls at the time that the entity bean is deployed; that is, when the entity bean is installed into a Partition. The tools use the deployment descriptors to determine the instance fields which must be persisted. Instead of coding the database access directly in the bean, the bean provider of a container-managed entity bean must specify in the deployment descriptor those instance fields for which the container tools must generate access

calls. The container has sophisticated deployment tools capable of mapping the fields of an entity bean to its data source.

Container-managed persistence has many advantages over bean-managed persistence. It is simpler to code because the bean provider does not have to code the database access calls. Handling of persistence can also be changed without having to modify and recompile the entity bean code. The Deployer or Application Assembler can do this by modifying the deployment descriptor when deploying the entity bean. Shifting the database access and persistence handling to the container not only reduces the complexity of code in the bean, it also reduces the scope of possible errors. The bean provider can focus on debugging the business logic of the bean rather than the underlying system issues.

Borland's Persistence Manager (PM) not only persists CMP fields but also CMP relationships. The container manages bean relationships and maintains the referential integrity of these relationships. Just as you defined container-managed persistence fields in a bean's deployment descriptor, you can now define container-managed relationship fields in the deployment descriptor. The container supports relationships of various cardinalities, including one-to-one, one-to-many, and many-to-many.

Packaging Requirements

Like session beans, entity beans can expose their methods with a remote interface or with a local interface. The remote interface exposes the bean's methods across the network to other, remote components. The local interface exposes the bean's methods only to local clients; that is, clients located on the same EJB container.

Entity beans that use EJB 2.0 container-managed persistence should use the local model. That is, the entity bean's local interface extends the `EJBLocalObject` interface. The bean's local home interface extends the `EJBLocalHome` interface. You must deploy these interfaces as well as an implementation of your bean's class.

Each Entity Bean must also have corresponding entries in its JAR's deployment descriptors. The standard deployment descriptor, `ejb-jar.xml` contains essentially three different types of deployment information. These are:

1. **General Bean Information:** This corresponds to the `<enterprise-beans>` elements found in the descriptor file and is used for all three types of beans. This information also includes information on the bean's interfaces and class, security information, environmental information, and even query declarations.
2. **Relationships:** This corresponds to the `<relationships>` elements found in the descriptor file and applies to entity beans using CMP only. This is where container-managed relationships are spelled out.
3. **Assembly Information:** This corresponds to the `<assembly-descriptor>` element which explains how the beans interact with the application as a whole. Assembly information is broken down into four categories:
   - **Security Roles:** simple definitions of security roles used by the application. Any security role references you defined for your beans must also be defined here.
   - **Method Permissions:** each method of each bean can have certain rules about their execution. These are set here.
   - **Container-Transactions:** this specifies the transaction attributes as per the EJB 2.0 specification for each method participating in a transaction.
   - **Exclude List:** methods not to be called by anyone.

In addition, each Entity Bean also provides persistence information in the Borland-specific descriptor file, `ejb-borland.xml`. In this descriptor file, you specify information used by the Borland CMP engine and PM to persist entities in a backing store. This information includes:

- **General Bean Information:** Information about deployed Enterprise JavaBeans, including interface locations.
- **Table and Column Properties:** Information about database tables and columns used by entity beans in the JAR.

– **Security Roles:** Authorization information for the deployed Enterprise JavaBeans.

All of these can be accessed from the Deployment Descriptor Editor. You should refer to the EJB 2.0 specification for DTD information and the proper use of the descriptor files.

A note on reentrancy

By default, entity beans are not reentrant. When a call within the same transaction context arrives at the entity bean, it causes the exception `java.rmi.RemoteException` to be thrown.

You can declare an entity bean reentrant in the deployment descriptor; however, take special care in this case. The critical issue is that a Container can generally not distinguish between a (loopback) call within the same transaction and a concurrent invocation (in the same transaction context) on that same entity bean.

When the entity bean is marked reentrant, it is illegal to allow a concurrent invocation within the same transaction context on the bean instance. It is the programmer's responsibility to ensure this rule.

Container-Managed Persistence in AppServer

The AppServer's EJB Container is fully J2EE 1.3 compliant. It implements both container-managed persistence (CMP) for Enterprise JavaBeans implementing either the EJB 1.1 and/or EJB 2.0 specifications. The bean provider designs persistence schemas for their entity beans, determines the methods for accessing container-managed fields and relationships, and defines these in beans' deployment descriptors. The deployer maps this persistence schema to the database and creates any other necessary classes for the beans' maintenance.

The EJB 2.0 Specification from Sun Microsystems details the specifics for the bean and container contracts in Chapters 10 and 11. Creating the persistence schema is not in the scope of this document, but is well discussed in both the Sun specification and in the Borland JBuilder documentation at `http://info.borland.com/techpubs/jbuilder/`, the relevant parts of which are the *Developing Applications with Enterprise JavaBeans* and the *Distributed Application Developer's Guide*.

About the Persistence Manager

The Persistence Manager (PM) provides a data-access layer for reading and writing entity beans. It also provides navigation and maintenance support for relationships between entities and extensions to EJB-QL. Currently, the PM only supports data access to relational database by means of JDBC. The PM uses an optimistic concurrency approach to data access. Conflicts in resource state are resolved before transaction commit or rollback by use of verified SQL update and delete statements.

Although the PM does not manage transactions (this is the Container's responsibility), it is aware of transaction start and completion and can therefore manage entity state. The PM uses the `TxContext` class to represent the root of managed entities during transaction lifecycles. When the container manages a transaction it asks the PM for the associated `TxContext` instance. If none exists, as is the case when a new transaction has started, one is created by the PM. When a transaction is completing, the container calls the method `TxContext.beforeCompletion()` to alert the PM to verify entity state.

The PM has complete responsibility for entity data storage and the maintenance of the state of relationships between entities. Relationship editing is also managed by the PM. This simplifies interactions with the container and allows the PM to optimize its read and write operations. This approach also suppresses duplicate `find` requests by tracking returned primary keys for requested entities. Data from duplicate `find` operations can then be returned from the first load of the entity's data.

Borland CMP engine's CMP 2.x implementation

In CMP 2.x, the details of constructing finder and select methods have been pushed into the EJB 2.0 specification. Users should thoroughly inspect the specification for details on implementing their database SQL. The Borland EJB Container is fully-compliant with the EJB 2.0 specification and supports all of its features.

The implementation class for an entity bean using 2.0 container-managed persistence is different from that of a bean using 1.1 container-managed persistence. The major differences are as follows:

- The class is declared as an abstract class.
- There are no public declarations for the fields that are container-managed fields. Instead, there are abstract `get` and `set` methods for container-managed fields. These methods are abstract because the container provides their implementation. For example, rather than declaring the fields `balance` and `name`, the `CheckingAccount` class might include these get and set methods:

```
public abstract float getBalance();
public abstract void setBalance(float bal);
public abstract String getName();
public abstract void setName(String n);
```

- Container-managed relationship fields are likewise not declared as instance variables. The class instead provides abstract get and set methods for these fields, and the container provides the implementation for these methods.

Table Mapping for CMP 2.x is accomplished using the vendor-specific `ejb-borland.xml` deployment descriptor. The descriptor is a companion to the ejb-jar.xml descriptor described in the EJB 2.0 specification. Borland uses the XML tag `<cmp2-info>` as an enclosure for table mapping data as needed. Then you use the `<table-properties>` and its associated `<column-properties>` elements to specify particular information about the entity bean's implementation. Use the DTD for syntax of the XML grammar.

Optimistic Concurrency Behavior

The container uses optimistic or pessimistic concurrency to control the behavior of multiple transactions accessing the same data. The AppServer has four optimistic concurrency behaviors which are specified as Table Properties. These behaviors are:

- `SelectForUpdate`
- `SelectForUpdateNoWAIT`
- `UpdateAllFields`
- `UpdateModifiedFields`
- `VerifyModifiedFields`
- `VerifyAllFields`

The behavior exhibited by the container corresponds to the value of the `optimisticConcurrencyBehavior` Table Property.

Pessimistic Behavior

In this mode, the container will allow only one transaction at a time to access the data held by the entity bean. Other transactions seeking the same data will block until the first transaction has committed or rolled back. This is achieved by setting the `SelectForUpdate` table property and issuing a tuned SQL statement with the `FOR UPDATE` statement included. You can issue this SQL by overriding the SQL generated by the CMP engine. Other selects on the row are blocked until then. The tuned SQL generated looks like this:

```
SELECT ID, NAME FROM EMP_TABLE WHERE ID=? FOR UPDATE
```

You can also specify the `SelectForUpdateNoWAIT` table property. Doing so instructs the database again to lock the row until the current transaction is committed or rolled back. However, other selects on the row will fail (rather than blocking). The following SQL illustrates a SELECT statement for the above:

```
SELECT ID, NAME FROM EMP_TABLE WHERE ID=? FOR UPDATE NOWAIT
```

These options should be used with caution. Although it does ensure the integrity of the data, your application's performance could suffer considerably. This option will also not function if you are using the Option A cache, since the entity bean remains in memory in this mode and calls to `ejbLoad()` are not made between transactions.

Optimistic Concurrency

This mode permits the container to allow multiple transactions to operate on the same data at the same time. While this mode is superior in performance, there is the possibility that data integrity could be compromised.

The AppServer has four optimistic concurrency behaviors which are specified as Table Properties. These behaviors are:

- `SelectForUpdate`
- `SelectForUpdateNoWAIT`
- `UpdateAllFields`
- `UpdateModifiedFields`
- `VerifyModifiedFields`
- `VerifyAllFields`

SelectForUpdate

Use this option for pessimistic concurrency. With this option specified, the database locks the row until the current transaction is committed or rolled back. Other selects on the row are blocked until then.

SelectForUpdate NoWAIT

Use this option for pessimistic concurrency. With this option specified, the database locks the row until the current transaction is committed or rolled back. Other selects on the row will fail.

UpdateAllFields

With this option specified, the container issues an update on all fields, regardless of whether or not they were modified. For example, consider a CMP entity bean with three fields, KEY, VALUE1, and VALUE2. The following update will be issued at the terminus of every transaction, regardless of whether or not the bean was modified:

```
UPDATE MyTable SET (VALUE1 = value1, VALUE2 = value2) WHERE KEY = key
```

UpdateModifiedFields

This option is the default optimistic concurrency behavior. The container issues an update only on the fields that were modified in the transaction, or suppresses the update altogether if the bean was not modified. Consider the same bean from the previous example, and assume that only VALUE1 was modified in the transaction. Using `UpdateModifiedFields`, the container would issue the following update:

```
UPDATE MyTable SET (VALUE1 = value1) WHERE KEY = key
```

This option can provide a significant performance boost to your application. Very often data access is read-only. In such cases, not sending an update to the database upon every transaction saves quite a bit of processing time. Suppressing these updates also prevents your database implementation from logging them, also enhancing performance. The JDBC driver is also taxed far less, especially in large-scale EJB applications. Even for well-tuned drivers, the less work they have to perform, the better.

VerifyModifiedFields

This option, when enabled, orders the CMP engine to issue a tuned update while verifying that the updated fields are consistent with their previous values. If the value has changed in between the time the transaction originally read it and the time the transaction is ready to update, the transaction will roll back. (You will need to handle these rollbacks appropriately.) Otherwise, the transaction commits. Again using the same table, the CMP engine generates the following SQL using the `VerifyModifiedFields` behavior if only VALUE1 was updated:

```
UPDATE MyTable SET (VALUE1 = value1) WHERE KEY = key AND VALUE1 = old-VALUE1
```

VerifyAllFields

This option is very similar to `VerifyModifiedFields`, except that all fields are verified. Again using the same table, the CMP engine generates the following SQL using this option:

```
UPDATE MyTable SET (VALUE1 = value1) WHERE KEY = key AND VALUE1 = old-VALUE1
AND VALUE2 =  old-VALUE2
```

Note

The two verify settings can be used to replicate the `SERIALIZABLE` isolation level in the Container. Often your applications require serializable isolation semantics. However, asking the database to implement this can have a significant performance impact. Using the verify settings allows the CMP engine to implement optimistic concurrency using field-level locking. The smaller the granularity of the locking, the better the concurrency.

Persistence Schema

The Borland CMP 2.x engine can create the underlying database schema based on the structure of your entity beans and the information provided in the entity bean deployment descriptors. You don't need to provide any CMP mapping information in such cases. Simply follow the instructions for “Specifying tables and datasources,” below. Or, the CMP engine can adapt to an existing underlying database schema. Doing so, however, requires you to provide information to the CMP engine about your database schema. In such cases, see “Basic Mapping of CMP fields to columns” as well as CASE 2 in “Specifying tables and datasources.”

Specifying tables and datasources

The minimum information required `in ejb-borland.xml` is an entity bean name and an associated datasource. A datasource is used to obtain connections to a database. Information on datasource configuration is given in “Connecting to Resources with Borland AppServer: using the Definitions Archive (DAR).” There are two means of providing this information.

CASE 1: A development environment without existing database tables using either JDataStore or Cloudscape databases.

In this case, the Borland CMP engine creates tables automatically, assuming that the entity bean name is the same as the desired table name. You need only provide the bean's name and its associated datasource as a property:

```
<entity>
 <ejb-name>CustomerEJB</ejb-name>
 <property>
  <prop-name>ejb.datasource</property>
  <prop-value>serial://ds/myDatasource</prop-value>
 </property>
</entity>
```

The Borland CMP engine will automatically create tables in this datasource based on the bean's name and fields.

CASE 2: A deployment environment with (or without) existing database tables using supported databases.

In this case, you need to supply information on the tables to which the entities map. You'll provide a table name in the `<entity>` portion of the descriptor, and some properties in the `<table-properties>` portion:

```
<entity>
 <ejb-name>CustomerEJB</ejb-name>
 <cmp2-info>
```

```
  <table-name>CUSTOMER</table-name>
 </cmp2-info>
</entity>
⋮
<table-properties>
 <table-name>CUSTOMER</table-name>
  <property>
   <prop-name>datasource</prop-name>
   <prop-value>serial://ds/myDatasource</prop-value>
  </property>
</table-properties>
```

Note that the datasource property is called `datasource` when specified in the `<table-properties>` element and `ejb.datasource` when in the `<entity>` element. If you are using a database other than JDataStore or Cloudscape and would like to have the Borland CMP engine automatically create this table, add the following XML to the `<table-properties>` element:

```
⋮
<table-properties>
 <table-name>CUSTOMER</table-name>
  <property>
   <prop-name>create-tables</prop-name>
   <prop-value>True</prop-value>
  </property>
</table-properties>
```

Basic Mapping of CMP fields to columns

Basic field mapping is accomplished using the `<cmp-field>` element in the ejb-borland.xml deployment descriptor. In this element, you specify a field name and a corresponding column to which it maps. Consider the following XML for an entity bean called `LineItem`, which maps two fields, `orderNumber` and `line`, to two columns, `ORDER_NUMBER` and `LINE`:

```
<entity>
 <ejb-name>LineItem</ejb-name>
 <cmp2-info>
  <cmp-field>
   <field-name>orderNumber</field-name>
   <column-name>ORDER_NUMBER</column-name>
  </cmp-field>
  <cmp-field>
   <field-name>line</field-name>
   <column-name>LINE</column-name>
  </cmp-field>
 </cmp2-info>
</entity>
```

Mapping one field to multiple columns

Many users may employ coarse-grained entity beans that implement a Java class to represent more fine-grained data. For example, an entity bean might use an `Address` class as a field, but may need to map elements of the class (like `AddressLine1`, `AddressCity`, and so forth) to an underlying database. To do this, you use the `<cmp-field-map>` element, which defines a field map between your fine-grained class and its underlying database representation. Note that such classes must implement `java.io.Serializable` and all their data members must be public.

Consider an entity bean called `Customer` that uses the class `Address` to represent a customer's address. The `Address` class has fields for `AddressLine`, `AddressCity`, `AddressState`, and `AddressZip`. Using the following XML, we can map the class to its representation in a database with corresponding columns:

```
<entity>
 <ejb-name>Customer</ejb-name>
```

```
 ⋮
 <cmp2-info>
  <cmp-field>
   <field-name>Address</field-name>
   <cmp-field-map>
    <field-name>Address.AddressLine</field-name>
    <column-name>STREET</column-name>
   </cmp-field-map>
   <cmp-field-map>
    <field-name>Address.AddressCity</field-name>
    <column-name>CITY</column-name>
   </cmp-field-map>
   <cmp-field-map>
    <field-name>Address.AddressState</field-name>
    <column-name>STATE</column-name>
   </cmp-field-map>
   <cmp-field-map>
    <field-name>Address.AddressZip</field-name>
    <column-name>ZIP</column-name>
   </cmp-field-map>
  </cmp-field>
 </cmp2-info>
 ⋮
</entity>
```

Note that we use one `<cmp-field-map>` element per database column.

Mapping CMP fields to multiple tables

You may have an entity that contains information persisted in multiple tables. These tables must be linked by at least one column representing a foreign key in the linked table. For example, you might have a `LineItem` entity bean mapping to a table `LINE_ITEM` with a primary key `LINE` that is a foreign key in a table called `QUANTITY`. The `LineItem` entity also contains some fields from the `QUANTITY` table that correspond to `LINE` entries in `LINE_ITEM`. Here's what our `LINE_ITEM` table might look like:

LINE	ORDER_NO	ITEM	QUANTITY	COLOR	SIZE
001	XXXXXXX01	Kitty Sweater	2	red	XL

QUANTITY, COLOR, and SIZE are all values that are also stored in the QUANTITY table, shown here. Note the identical values for some of the fields. This is because the LINE_ITEM table itself stores information in the QUANTITY table, using the LineItem entity to provide composite information.

LINE	QUANTITY	COLOR	SIZE
001	2	red	XL

Again, we can describe these relationships using a combination of `<cmp-field>` elements and a `<table-ref>` element. The `<cmp-field>` elements define the fields found in `LineItem`. Since there are some fields that require information from `QUANTITY`, we'll specify that generically by using a `TABLE_NAME.COLUMN_NAME` syntax. For instance, we'd define `LINE_ITEM`'s `COLOR` column as `QUANTITY.COLOR`. Finally, we'll specify the linking column, `LINE`, that makes up our primary key/foreign key relationship. We'll do this using the `<table-ref>` element.

Now let's look at the XML. First we define the CMP fields for the `LineItem` entity bean:

```
<entity>
 <ejb-name>LineItem</ejb-name>
 ⋮
 <cmp2-info>
  <cmp-field>
   <field-name>orderNumber</field-name>
   <column-name>ORDER_NO</column-name>
  </cmp-field>
```

```
<cmp-field>
 <field-name>line</field-name>
 <column-name>LINE</column-name>
</cmp-field>
<cmp-field>
 <field-name>item</field-name>
 <column-name>ITEM</column-name>
</cmp-field>
<cmp-field>
 <field-name>quantity</field-name>
 <column-name>QUANTITY.QUANTITY</column-name>
</cmp-field>
<cmp-field>
 <field-name>color</field-name>
 <column-name>QUANTITY.COLOR</column-name>
</cmp-field>
<cmp-field>
 <field-name>size</field-name>
 <column-name>QUANTITY.SIZE</column-name>
</cmp-field>
```

Next, we specify the linking column between `LINE_ITEM` and `QUANTITY` by using a `<table-ref>` element.

```
<table-ref>
 <left-table>
  <table-name>LINE_ITEM</table-name>
  <column-list>
   <column-name>LINE</column-name>
  </column-list>
 </left-table>
 <right-table>
  <table-name>QUANTITY</table-name>
  <column-list>
   <column-name>LINE</column-name>
  </column-list>
 </right-table>
</table-ref>
</cmp2-info>
</entity>
```

Specifying relationships between tables

To specify relationships between tables, you use the `<relationships>` element in `ejb-borland.xml`. Within the `<relationships>` element, you define an `<ejb-relationship-role>` containing the role's source (an entity bean) and a `<cmr-field>` element containing the relationship. The descriptor then uses `<table-ref>` elements to specify relationships between two tables, a `<left-table>` and a `<right-table>`. You must observe the following cardinalities:

- One `<ejb-relationship-role>` must be defined per direction; if you have a bi-directional relationship, you must define an `<ejb-relationship-role>` for each bean with each referencing the other.
- Only one `<table-ref>` element is permitted per relationship.

Within the `<left-table>` and `<right-table>` elements, you specify a column list that contains the column names to be linked together. The column list corresponds to the `<column-list>` element in the descriptor. The XML is:

```
<!ELEMENT column-list (column-name+)>
```

Let's look at some relationships to see how this XML is put into practice:

CASE 1: a unidirectional one-to-one relationship.

Here, we have a `Customer` entity bean with a primary key, `CUSTOMER_NO`, that is also used as a primary key for an entity called `SpecialInfo`, which contains special customer information stored in a separate table. We need to specify a relationship between these two entities. The `Customer` entity uses a field called `specialInformation` to map to the `SpecialInfo` bean. We specify two relationship roles, one for each bean and assign either to left- and/or right-table. Then we specify the name of their related column for both.

```
<relationships>
<ejb-relation>
 <ejb-relationship-role>
  <relationship-role-source>
   <ejb-name>Customer</ejb-name>
  </relationship-role-source>
  <cmr-field>
   <cmr-field-name>specialInformation</cmr-field-name>
   <table-ref>
    <left-table>
     <table-name>CUSTOMER</table-name>
     <column-list>CUSTOMER_NO</column-list>
    </left-table>
    <right-table>
     <table-name>SPECIAL_INFO</table-name>
     <column-list>CUSTOMER_NO</column-list>
    </right-table>
   </table-ref>
  </cmr-field>
 </ejb-relationship-role>
```

Next, we finish the `<ejb-relation>` entry by providing its other half, the `SpecialInfo` bean. Since this is a mono-directional relationship, we don't need to specify any table elements. We only need add the following, defining the other half of the relationship and its source:

```
  <ejb-relationship-role>
   <relationship-role-source>
    <ejb-name>SpecialInfo</ejb-name>
   </relationship-role-source>
  </ejb-relationship-role>
 </ejb-relation>
</relationships>
```

CASE 2: a bidirectional one-to-many relationship.

Here, we have a `Customer` entity bean with a primary key, `CUSTOMER_NO`, that is also a foreign key in an `Order` entity bean. We want the Borland EJB Container to manage this relationship. The `Customer` bean uses a field called “orders” that links a customer to his orders. The `Order` bean uses a field called “customers” for linking in the reverse direction. First, we define the relationship and its source for the first direction: setting up the mapping for a `Customer`'s orders.

```
<relationships>
 <ejb-relation>
  <ejb-relationship-role>
   <relationship-role-source>
    <ejb-name>Customer</ejb-name>
   </relationship-role-source>
   <cmr-field>
    <cmr-field-name>orders</cmr-field-name>
```

Then, we add the table references to specify the relationship between the tables. We're basing this relationship on the `CUSTOMER_NO` column, which is a primary key for `Customer` and a foreign key for `Orders`:

```
    <table-ref>
     <left-table>
      <table-name>CUSTOMER</table-name>
      <column-list>
       <column-name>CUSTOMER_NO</column-name>
      </column-list>
     </left-table>
     <right-table>
      <table-name>ORDER</table-name>
      <column-list>
       <column-name>CUSTOMER_NO</column-name>
      </column-list>
     </right-table>
     </table-ref>
    </cmr-field>
   </ejb-relationship-role>
```

We're not quite done with our relationship, though. Now, we need to complete it by specifying the relationship role for the other direction:

```
   <ejb-relationship-role>
    <relationship-role-source>
     <ejb-name>Customer</ejb-name>
    </relationship-role-source>
    <cmr-field>
     <cmr-field-name>customers</cmr-field-name>
     <table-ref>
      <left-table>
       <table-name>ORDER</table-name>
       <column-list>
        <column-name>CUSTOMER_NO</column-name>
       </column-list>
      </left-table>
      <right-table>
       <table-name>CUSTOMER</table-name>
       <column-list>
        <column-name>CUSTOMER_NO</column-name>
       </column-list>
      </right-table>
     </table-ref>
    </cmr-field>
   </ejb-relationship-role>
  </ejb-relation>
  ⋮
 </relationships>
```

CASE 3: a many-to-many relationship.

If you define a many-to-many relationship, you must also have the CMP engine create a cross-table which models a relationship between the left table and the right table. Do this using the `<cross-table>` element, whose XML is:

```
<!ELEMENT cross-table (table-name, column-list, column-list)>
```

You may name this cross-table whatever you like using the `<table-name>` element. The two `<column-list>` elements correspond to columns in the left and right tables whose relationship you wish to model. For example, consider two tables, `EMPLOYEE` and `PROJECT`, which have a many-to-many relationship. An employee can be a part of multiple projects, and projects have multiple employees. The `EMPLOYEE` table has three elements, an employee number (`EMP_NO`), a last name (`LAST_NAME`), and a project ID number (`PROJ_ID`). The `PROJECT` table contains columns for the project ID number (`PROJ_ID`), the project name (`PROJ_NAME`), and assigned employees by number (`EMP_NO`).

To model the relationship between these two tables, a cross-table must be created.. For example, to create a cross-table that shows employee names and the names of the projects on which they are working, the <table-ref> element would look like the following:

```
<table-ref>
    <left-table>
       <table-name>EMPLOYEE</table-name>
          <column-list>
                <column-name>EMP_NO</column-name>
         <column-name>LAST_NAME</column-name>
             <column-name>PROJ_ID</column-name>
          </column-list>
    </left-table>
    <cross-table>
        <table-name>EMPLOYEE_PROJECTS</table-name>
           <column-list>
                <column-name>EMP_NAME</column-name>
          <column-name>PROJ_ID</column-name>
           </column-list>
           <column-list>
            <column-name>PROJ_ID</column-name>
               <column-name>PROJ_NAME</column-name>
           </column-list>
    </cross-table>
    <right-table>
        <table-name>PROJECT</table-name>
           <column-list>
               <column-name>PROJ_ID</column-name>
            <column-name>PROJ_NAME</column-name>
            <column-name>EMP_NO</column-name>
           </column-list>
    </right-table>
</table-ref>
```

Since these are “secondary tables” and therefore have no primary keys, the `PROJ_ID` column appears in both column lists. This could also be the common column `EMP_NO`, depending upon how you wish to model the data.

Using cascade delete and database cascade delete

Use `<cascade-delete>` when you want to remove entity bean objects. When cascade delete is specified for an object, the container automatically deletes all of that object's dependent objects. For example you may have a Customer bean which has a one-to-many, uni-directional relationship to an Address bean. Because an address instance must be associated to a customer, the container automatically deletes all addresses related to the customer when you delete the customer.

To specify cascade delete, use the <cascade-delete> element in the `ejb-jar.xml` file as follows:

```
<ejb-relation>
 <ejb-relation-name>Customer-Account</ejb-relation-name>
 <ejb-relationship-role>
  <ejb-relationship-role-name>Account-Has-Customer
  </ejb-relationship-role-name>
  <multiplicity>one</multiplicity>
  <cascade-delete/>
 </ejb-relationship-role>
</ejb-relation>
```

Database cascade delete support

AppServer supports the database cascade delete feature, which allows an application to take advantage of a database's built in cascade delete functionality. This reduces the number of SQL operations sent to the database by the container, therefore improving performance.

To use database cascade delete, the tables corresponding to the entity beans have to be created with the appropriate table constraints on the respective database. For example, if you are using cascade delete in EJB 2.0 entity beans on `Order` and `LineItem` entity beans, the tables have to be created as follows:

```
create table ORDER_TABLE (ORDER_NUMBER integer, LAST_NAME varchar(20),
FIRST_NAME varchar(20), ADDRESS varchar(48));
create table LINE_ITEM_TABLE (LINE integer, ITEM varchar(100), QUANTITY
numeric, ORDER_NUMBER integer CONSTRAINT fk_order_number REFERENCES
ORDER_TABLE(ORDER_NUMBER) ON DELETE CASCADE);
```

The `<cascade-delete-db>` element in the `ejb-borland.xml` file specifies that a cascade delete operation will use the cascade delete functionality of the database. By default this feature is turned off.

Note

If you specify the `<cascade-delete-db>` element in the `ejb-borland.xml` file, you must specify `<cascade-delete>` in `ejb-jar.xml`.

The XML for `<cascade-delete-db>` in the `ejb-borland.xml` is shown in the following relationship:

```
<relationships>
        <!--
        ONE-TO-MANY: Order LineItem
        -->
        <ejb-relation>
            <ejb-relationship-role>
                <relationship-role-source>
                    <ejb-name>OrderEJB</ejb-name>
                </relationship-role-source>
                <cmr-field>
                    <cmr-field-name>lineItems</cmr-field-name>
                    <table-ref>
                        <left-table>
                            <table-name>ORDER_TABLE</table-name>
                            <column-list>
                                <column-name>ORDER_NUMBER</column-name>
                            </column-list>
                        </left-table>
                        <right-table>
                            <table-name>LINE_ITEM_TABLE</table-name>
                            <column-list>
                                <column-name>ORDER_NUMBER</column-name>
                            </column-list>
                        </right-table>
                    </table-ref>
                </cmr-field>
            </ejb-relationship-role>
            <ejb-relationship-role>
                <relationship-role-source>
                    <ejb-name>LineItemEJB</ejb-name>
                </relationship-role-source>
                <cmr-field>
                    <cmr-field-name>order</cmr-field-name>
                    <table-ref>
                        <left-table>
                            <table-name>LINE_ITEM_TABLE</table-name>
                            <column-list>
```

```
                                <column-name>ORDER_NUMBER</column-name>
                            </column-list>
                        </left-table>
                        <right-table>
                            <table-name>ORDER_TABLE</table-name>
                            <column-list>
                                <column-name>ORDER_NUMBER</column-name>
                            </column-list>
                        <right-table>
                    </table-ref>
                </cmr-field>
            </ejb-relationship-role>
<cascade-delete-db />
        </ejb-relation>
    </relationships>
```

15

Using Borland AppServer Properties for CMP 2.x

Setting Properties

Most properties for Enterprise JavaBeans can be set in their deployment descriptors. The Borland Deployment Descriptor Editor (DDEditor) also allows you to set properties and edit descriptor files. Use of the Deployment Descriptor Editor is described in the Borland Management Console *User's Guide*. For more information, see “Using the Deployment Descriptor Editor”. Use properties in the deployment descriptor to specify information about the entity bean's interfaces, transaction attributes, and so forth, plus information that is unique to an entity bean. In addition to the general descriptor information for entity beans, here are also three sets of properties that can be set to customize CMP implementations, entity properties, table properties, and column properties. Entity properties can be set either by using the Deployment Descriptor Editor or in the XML directly.

Using the Deployment Descriptor Editor

You can use the Deployment Descriptor Editor, which is part of the Borland AppServer (AppServer), to set up all of the container-managed persistence information. The following table shows descriptor information and where in the Deployment Descriptor Editor that information can be entered.

For complete information on the use of the Deployment Descriptor Editor and other related tools, see “Using the Deployment Descriptor Editor” in *Management Console User's Guide*.

The EJB Designer

CMP 2.x properties are set using the EJB Designer. For more information about the EJB Designer, see “The EJB Designer” in the *Borland Management Console User's Guide*.

J2EE 1.3 and 1.4 Entity Bean

Descriptor Element	Navigation Tree Node/Panel Name	DDEditor Tab
Entity Bean name	Bean	General
Entity Bean class	Bean	General
Home Interface	Bean	General
Remote Interface	Bean	General
Local Home Interface	Bean	General
Local Interface	Bean	General
Home JNDI Name	Bean	General
Local Home JNDI Name	Bean	General
Persistence Type (CMP or BMP)	Bean	General
CMP Version	Bean	General
Primary Key Class	Bean	General
Reentrancy	Bean	General
Icons	Bean	General
Environment Entries	Bean	Environment
EJB References to other Beans	Bean	EJB References
EJB Links	Bean	EJB References
Resource References to data objects/connection factories	Bean	Resource References
Resource Reference type	Bean	Resource References
Resource Reference Authentication Type	Bean	Resource References
Security Role References	Bean	Security Role References
Entity Properties	Bean	Properties
Security Identity	Bean	Security Identity
EJB Local References to beans in the name JAR	Bean	EJB Local References
EJB Local Links	Bean	EJB Local References
Resource Environmental References for JMS	Bean	Resource Env Refs
Container Transactions	Bean:Container Transactions	Container Transactions
Transactional Method	Bean:Container Transactions	Container Transactions
Transactional Method Interface	Bean:Container Transactions	Container Transactions
Transactional Attribute	Bean:Container Transactions	Container Transactions
Method Permissions	Bean:Method Permissions	Method Permissions
Entity, Table, and Column Properties	JAR	EJB Designer (see below)

Setting CMP 2.x Properties

The AppServer uses the EJB Designer, a component of the Deployment Descriptor Editor, to set CMP 2.x properties. The EJB Designer is fully-documented in “The EJB Designer” in the *Borland Management Console User's Guide*..

Editing Entity properties

To edit Entity properties using the EJB Designer:

1. Start the DDEditor and open the deployment descriptor for the JAR containing your entity beans.
2. Select the top-level object in the DDEditor's Navigation Pane. In the Properties Pane you will see three tabs—General, XML, and EJB Designer.
3. Choose the EJB Designer Tab and left-click on any of the bean representations that appear. Click the Properties button. The Entity Beans Properties window appears.
4. Edit the properties you desire and click OK. The properties themselves are discussed below.

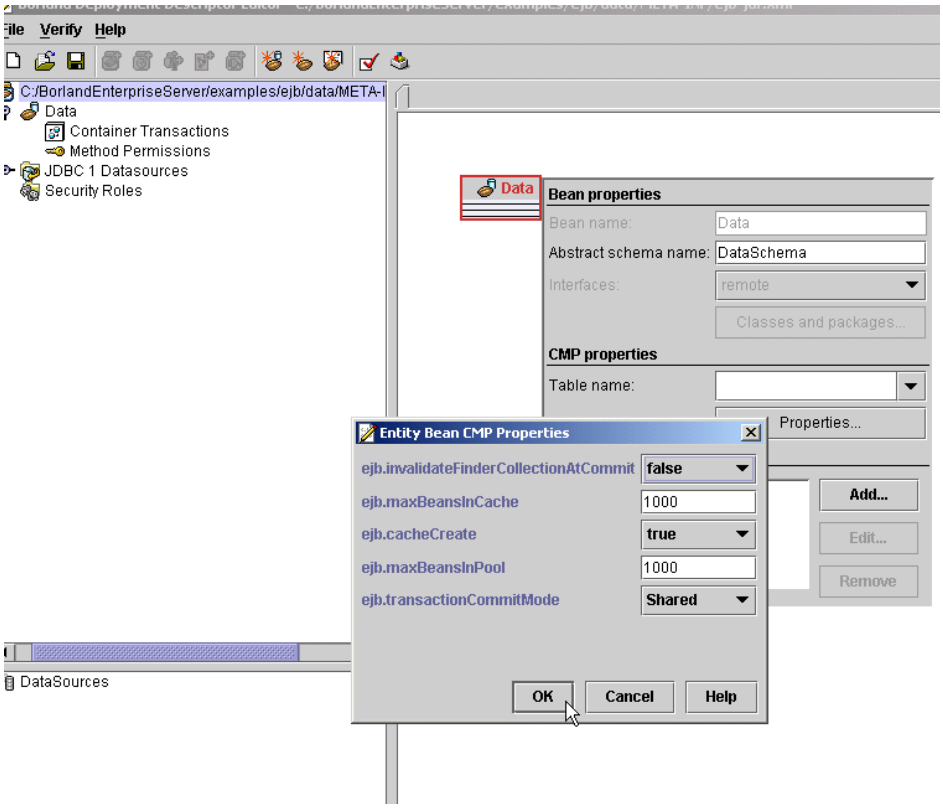

Editing Table and Column properties

Table and Column properties can only be set by editing the `ejb-borland.xml` descriptor file from the DDEditor's Vendor XML Tab, or by using the EJB Designer. To edit or add Table and Column properties:

1. Start the DDEditor and open the deployment descriptor for the JAR containing your entity beans.
2. Select the top-level object in the DDEditor's Navigation Pane. In the Properties Pane you will see three tabs: General, XML, and EJB Designer.
3. Select the XML Tab. Two additional Tabs are now available in the Properties Pane; Standard and Vendor. Choose Vendor.

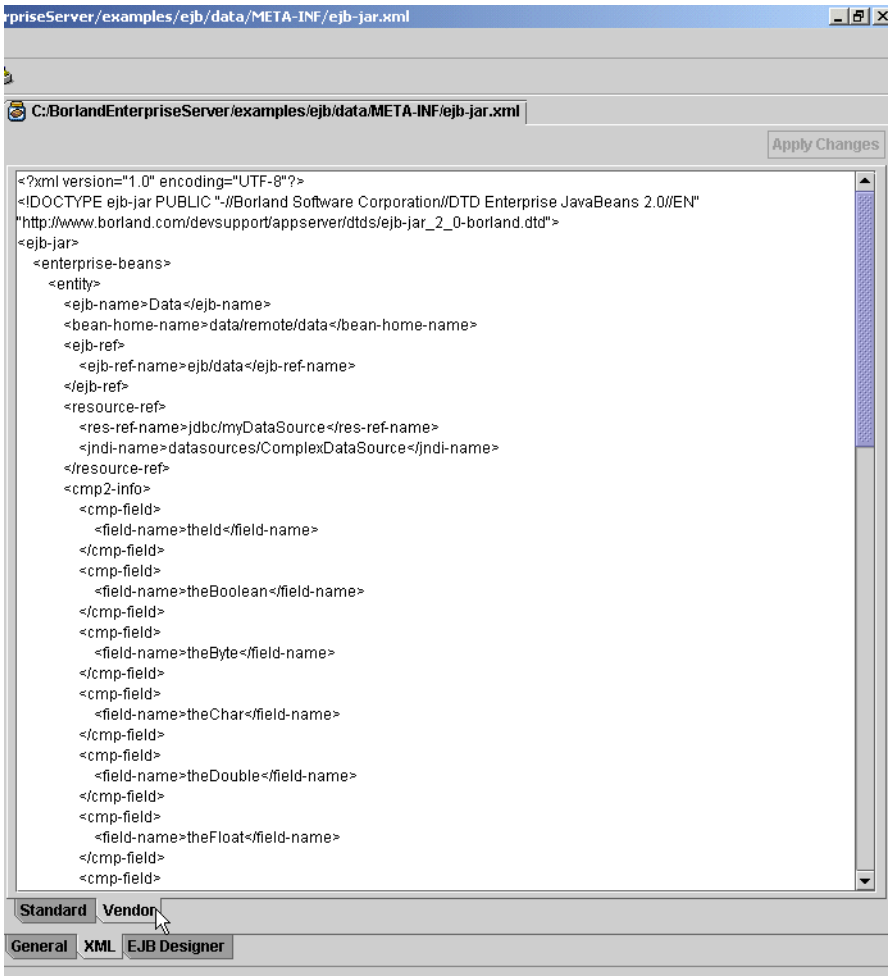

4. Locate or add either the `<column-properties>` or `<table-properties>` elements and add property definitions in accordance with the borland-specific DTD (see the ejb-borland.xml). Germane entries are in bold. Descriptions of the entity, table, and column properties follow, including their data type, default values, and a property description.

Entity Properties

These properties are for CMP 1.1 and above implementations:

Property	Type	Default	Description
ejb.maxBeansInCache	lava.lang.Integer	1000	This option specifies the maximum number of beans in the cache that holds on to beans associated with primary keys, but not transactions. This is relevant for Option “A” and “B” (see ejb.transactionCommitMode below). If the cache exceeds this limit, entities will be moved to the ready pool by calling ejbPassivate.
ejb.maxBeansInPool	java.lang.Integer	1000	The maximum number of beans in the ready pool. If the ready pool exceeds this limit, entities will be removed from the container by calling unsetEntityContext().
ejb.maxBeansInTransactions	lava.lang.Integer	500* (see Description)	A transaction can access any/large number of entities. This property sets an upper limit on the number of physical bean instances that EJB container will create. Irrespective of the number of database entities/rows accessed, the container will manage to complete the transaction with a smaller number of entity objects (dispatchers). The default for this is calculated as ejb.maxBeansInCache/2. If the ejb.maxBeansInCache property is not set, this translates to 500.
ejb.TransactionCommitMode	Enumerated	Shared	Indicates the disposition of an entity bean with respect to a transaction. Acceptable values are: ■ **Exclusive**: This entity has exclusive access to the particular table in the database. The state of the bean at the end of the last committed transaction can be assumed to be the state of the bean at the beginning of the next transaction. ■ **Shared**: This entity shares access to the particular table in the database. However, for performance reasons, a particular bean remains associated with a particular primary key between transactions to avoid extraneous calls to ejbActivate() and ejbPassivate() between transactions. The bean stays in the active pool. ■ **None**: This entity shares access to the particular table in the database. A particular bean does not remain associated with a particular primary key between transactions, but goes back to the ready pool after every transaction.

These properties are for CMP 2.x implementations only:

Property	Type	Default	Description
ejb.invalidateFinderCollectionAtCommit	java.lang.Boolean	False	Whether or not to optimize transaction commit by invalidating finder collections. CMP 2.x only.
ejb.cacheCreate	java.lang.Boolean	True	Whether or not to attempt to cache the insert of the entity bean until the ejbPostCreate is processed.
ejb.datasource	java.lang.String	N/A	Default JDBC datasource to use in case no table-properties have been set. CMP 2.x only.

Property	Type	Default	Description
ejb.truncateTableName	java.lang.Boolean	False	If no table name is specified, CMP2.x engine will use the EJB name as the table name. EJB names can be more than 30 characters in length. Moreover, certain databases have a restriction on the table length to be 30 characters or less. This property is used to force the table name to be truncated to be 30 characters or less. CMP 2.x only.
ejb.eagerLoad	java.lang.Boolean	False	eager-loads the entire row and keeps the data in the transactional cache. After loading, all database resources are released. Subsequent getters could get data in cache and not having to require any more database resources. CMP 2.x only.

Table Properties

The following properties apply to CMP 2.x only. If you are migrating from CMP 1.1 to CMP 2.x, you must update your CMP properties. CMP 1.1 properties were formerly of the format ejb.<property-name>, and were all specified in the <entity> portion of the deployment descriptor. With CMP 2.x, the AppServer adds Table and Column Properties, which manage persistence. Refer to these properties below to see where migration issues may appear.

Property	Type	Default	Description
datasource	java.lang.String	None	JNDI datasource name of the database for this table.
optimisticConcurrencyBehavior	java.lang.String	UpdateModifiedFields	The container uses optimistic or pessimistic concurrency to control multiple transactions (updates) that access shared tables. Acceptable values are: ■ SelectForUpdate: database locks the row until the current transaction is committed or rolled back. Other selects on the row are blocked (wait) until then. ■ SelectForUpdateNoWAIT: database locks the row until the current transaction is committed or rolled back. Other selects on the row will fail. ■ UpdateAllFields: perform an update on all of an entity's fields, regardless if they were modified or not. ■ UpdateModifiedFields: perform an update only on fields known to have been modified prior to the update being issued. ■ VerifyModifiedFields: verify the entity's modified fields against the database prior to update. ■ VerifyAllFields: verify all the entity's fields against the database prior to update regardless if they were modified or not. Pessimistic concurrency specifies the container to allow only one transaction at a time to access the entity bean. Other transactions that try to access the same data will block (wait) until the first transaction completes. This is achieved by issuing a tuned SQL with FOR UPDATE when the entity bean is loaded. To achieve pessimistic concurrency set SelectForUpdate or SelectForUpdateNoWAIT.

Property	Type	Default	Description
useGetGeneratedKeys	java.lang.Boolean	False	Whether to use the JDBC3 java.sql.Statement.getGeneratedKeys() method to populate the primary key from autoincrement/sequence SQL fields. Currently, only Borland JDataStore supports this statement.
primaryKeyGenerationListener	java.lang.String	None	Specifies a class, written by the user, that implements com.borland.ejb.pm.PrimaryKeyGenerationListener interface and generates primary keys..
dbcAccesserFactory	java.lang.String	None	A factory class that can provide accessor class implementations to get values from a java.sql.ResultSet, and set values for a java.sql.PreparedStatement.
getPrimaryKeyBeforeInsertSql	java.lang.String	None	SQL statement to execute before inserting a row to provide primary key column names.
getPrimaryKeyAfterInsertSql	java.lang.String	None	SQL statement to execute after inserting a row to provide primary key column names.
useAlterTable	java.lang.Boolean	false	Whether or not to use the SQL ALTER statement to alter an entity's table to add columns for fields that do not have a matching column.
createTableSql	java.lang.String	None	SQL statement used to create the table if it needs to be created automatically.
create-tables	java.lang.Boolean	false	The Borland CMP engine automatically creates tables for Cloudscape and JDataStore databases—that is, in the development environment. To enable automatic table creation in other databases, you must set this flag to true.

Column Properties

Property	Type	Default	Description
ignoreOnInsert	java.lang.String	false	Specifies the column that must not be set during the execution of an INSERT statement. This property is used in conjunction with the getPrimaryKeyAfterInsertSql property.
createColumnSql	java.lang.String	None	Use this property to override the standard data-type lookup and specify the data type manually, use this property. ■ Local transactions support the javax.ejb.EJBContext methods setRollbackOnly() and getRollbackOnly(). ■ Local transactions support time-outs for database connections and transactions. ■ Local transactions are lightweight from a performance standpoint.

Property	Type	Default	Description
columnJavaType	java.lang.String	None	Java type used to create this column if the table needs to be created automatically. The acceptable values are: ■ java.lang.Boolean ■ java.lang.Byte ■ java.lang.Character ■ java.lang.Short ■ java.lang.Integer ■ java.lang.Long ■ java.lang.Float ■ java.math.BigDecimal ■ java.lang.String ■ java.sql.Time ■ java.sql.Date ■ java.sql.TimeStamp ■ java.io.Serializable This property is ignored if createColumnSql is set.

Security Properties

These security properties are specified in the <entity> portion of the deployment descriptor.

Property	Type	Default	Description
ejb.security.transportType	Enumerated	SECURE_ONLY	This property configures the Quality of Protection of a particular EJB. If set to CLEAR_ONLY, only non-secure connections are accepted from the client to this EJB. This is the default setting, if the EJB does not have any method permissions. If set to SECURE_ONLY, only secure connections are accepted form the client to this EJB. This is the default setting, if the EJB has at least one method permission set. If set to ALL, both secure and non-secure connections are accepted from the client. Setting this property controls a transport value of the ServerQoPConfig policy. See the “Security API” chapter from the Programmer's Reference for details.
ejb.security.trustInClient	java.lang.Boolean	False	This property configures the Quality of Protection of a particular EJB. If set to true, the EJB container requires the client to provide an authenticated identity. By default, the property is set to false, if there is at least one method with no method permissions set. Otherwise, it is set to true. Setting this property controls a transport value of the ServerQoPConfig policy. See the “Security API” chapter from the Programmer's Reference for details.

16

EJB-QL and Data Access Support

EJB-QL allows you to specify queries in an object oriented query language, EJB-QL. The Borland CMP engine translates these queries into SQL queries. The Borland AppServer (AppServer) provides some extensions to the EJB-QL functionality described in the Sun Microsystems EJB 2.x Specification.

Selecting a CMP Field or Collection of CMP Fields

When only one cmp-field of an otherwise large EJB is required, you can use EJB-QL to select a single instance of collection of that cmp-field. Using EJB-QL in this way improves application performance by eliminating the need to load an entire EJB. For example, this query method selects only the balance field from the Account table:

```
<query>
   <query-method>
     <method-name>ejbSelectBalanceOfAccountLineItem</method-name>
      <method-params>
       <method-param>java.lang.Long</method-param>
   </method-params>
   </query-method>
   <result-type-mapping>Local</result-type-mapping>
   <ejb-ql>SELECT l.balance FROM Account a, IN (a.accountLineItem)
     l WHERE l.lineItemId=?1</ejb-ql>
</query>
```

The return types of the EJB-QL query method are:

- If the Java type of the cmp-field is an object type, and the query method is a single-object query method, the return type is an instance of that object type.
- If the Java type of the cmp-field is an object type and the query method returns multiple objects, a collection of instances of the object type is returned.
- If the Java type of the cmp-field is a primitive Java type, and the SELECT method is a single-object method, the return type is that primitive type.
- If the Java type of the cmp-field is a primitive Java type, and the SELECT method is for multiple objects, a collection of the wrappered Java type is returned.

Selecting a ResultSet

When more than one cmp-field is to be returned by a single query method, the return type must be of type ResultSet. This allows you to select multiple cmp-fields from the same or multiple EJBs in the same query method. You then write code to extract the desired data from the ResultSet. This feature is a Borland extension of the CMP 2.x specification.

Aggregate Functions in EJB-QL

Aggregate functions are MIN, MAX, SUM, AVG, and COUNT. For the aggregate functions MIN, MAX, SUM, and AVG, the path expression that forms the argument for the function must terminate in a cmp-field. Also, database queries for MAX, MIN, SUM, and AVG will return a null value if there are no rows corresponding to the argument to the aggregate function. If the return type is an object-type, then null is returned. If the return type is a primitive type, then the container will throw a `ObjectNotFoundException` (a sub-class of `FinderException`) if there is no value in the query result.

The path expression to the COUNT functions may terminate in either a cmp-field or cmr-field, or may be an identification variable.

For example, the following EJB-QL aggregate function terminates in a CMP field:

```
<query>
   <query-method>
     <method-name>ejbSelectMaxLineItemId</method-name>
      <method-params>
       <method-param>java.lang.String</method-param>
      </method-params>
   </query-method>
   <result-type-mapping>Local</result-type-mapping>
   <ejb-ql>SELECT MAX(l.lineItemId) FROM Account AS a, IN (a.accountLineItem) l
WHERE l.accountId=?1</ejb-ql>
</query>
```

The following restrictions must be observed for aggregate functions:

- Arguments to the SUM and AVG functions must be numeric (Integer, Byte, Long, Short, Double, Float, and BigDecimal).
- Arguments to the MAX and MIN functions must correspond to orderable cmp-field types (numeric, string, character, and dates).
- The path expression that forms the argument for the COUNT function can terminate in either a cmp-field or a cmr-field. Application performance is greatly enhanced when the COUNT function is used to determine the size of a collection of cmr-fields.

Data Type Returns for Aggregate Functions

The following table shows the data types that can be arguments for the various aggregate functions in EJB-QL selecting a single object, and what data types will be returned.

An aggregate function that selects multiple objects returns a collection of the wrappered Java data type that is returned.

Aggregate Function	Argument data type	Expected return type
MIN, MAX, SUM	java.lang.Integer	java.lang.Integer
AVG	java.lang.Integer	java.lang.Double
COUNT	java.lang.Integer	java.lang.Long
MIN, MAX, SUM	java.lang.Integer	java.lang.Integer
AVG	java.lang.Integer	java.lang.Double
COUNT	java.lang.Integer	java.lang.Long

Aggregate Function	Argument data type	Expected return type
MIN, MAX, SUM	java.lang.Byte	java.lang.Byte
AVG	java.lang.Byte	java.lang.Double
COUNT	java.lang.Byte	java.lang.Long
MIN, MAX, SUM	java.lang.Byte	java.lang.Byte
AVG	java.lang.Byte	java.lang.Double
COUNT	java.lang.Byte	java.lang.Long
MIN, MAX, SUM	java.lang.Long	java.lang.Long
AVG	java.lang.Long	java.lang.Double
COUNT	java.lang.Long	java.lang.Long
MIN, MAX, SUM	java.lang.Long	lonjava.lang.Long
AVG	java.lang.Long	java.lang.Double
COUNT	java.lang.Long	java.lang.Long
MIN, MAX, SUM	java.lang.Short	java.lang.Short
AVG	java.lang.Short	java.lang.Double
COUNT	java.lang.Short	java.lang.Long
MIN, MAX, SUM	java.lang.Short	java.lang.Short
AVG	java.lang.Short	java.lang.Double
COUNT	java.lang.Short	java.lang.Long
MIN, MAX, SUM	java.lang.Double	java.lang.Double
AVG	java.lang.Double	java.lang.Double
COUNT	java.lang.Double	java.lang.Long
MIN, MAX, SUM	java.lang.Double	java.lang.Double
AVG	java.lang.Double	java.lang.Double
COUNT	java.lang.Double	java.lang.Long
MIN, MAX, SUM	java.lang.Float	java.lang.Float
AVG	java.lang.Float	java.lang.Double
COUNT	java.lang.Float	java.lang.Long
MIN, MAX, SUM	java.lang.Float	java.lang.Float
AVG	java.lang.Float	java.lang.Double
COUNT	java.lang.Float	java.lang.Long
MIN, MAX, SUM	java.math.BigDecimal	java.math.BigDecimal
AVG	java.math.BigDecimal	java.lang.Double
COUNT	java.math.BigDecimal	java.lang.Long
MIN, MAX	java.lang.String	java.lang.String
COUNT	java.lang.String	java.lang.Long
MIN, MAX	java.util.Date	java.util.Date
COUNT	java.util.Date	java.lang.Long
MIN, MAX	java.sql.Date	java.sql.Date
COUNT	java.sql.Date	java.lang.Long
MIN, MAX	java.sql.Time	java.sql.Time
COUNT	java.sql.Time	java.lang.Long
MIN, MAX	java.sql.Timestamp	java.sql.Timestamp
COUNT	java.sql.Timestamp	java.lang.Long

Support for ORDER BY

The EJB 2.0 Specification supports three SQL clauses in EJB-QL: `SELECT`, `FROM`, and `WHERE`.

The Borland CMP engine also supports the SQL clause `ORDER BY` in the same EJB-QL statement, provided it is placed **after** the `WHERE` clause. This is done in the standard `ejb-`

jar.xml deployment descriptor in the <ejb-ql> entity. For example, the following EJB-QL statement selects distinct objects from a Customer Bean and orders them by the LNAME field:

```
<query>
 <description></description>
 <query-method>
  <method-name>findCustomerByNumber</method-name>
  <method-params />
 <ejb-ql>SELECT Distinct Object(c) from CustomerBean c WHERE c.no >
     1000 ORDER BY c.LNAME</eql-ql>
 </query-method>
<query>
```

You can specify either ASC (ascending) or (DESC) descending in your EJB-QL as well. If you do not specify either, the results will be ordered ascending by default.

For example, consider the following table:

NAME	DEPARTMENT	SALARY	HIRE DATE
Timmy Twitfuller	Mail Room	1000	1/1/01
Sam Mackey	The Closet with the Light Out	800	1/2/02
Ralph Ossum	Coffee Room	900	1/4/01

The query:

```
SELECT OBJECT(e) FROM EMPLOYEE e ORDER BY e.HIRE_DATE
```

will produce the following result:

NAME	DEPARTMENT	SALARY	HIRE DATE
Timmy Twitfuller	Mail Room	1000	1/1/01
Ralph Ossum	Coffee Room	900	1/4/01
Sam Mackey	The Closet with the Light Out	800	1/2/02

Support for GROUP BY

The GROUP BY clause is used to group rows in the result table prior to the SELECT operation being performed. Consider the following table:

NAME	DEPARTMENT	SALARY	HIRE DATE
Mike Miller	Mail Room	1200	11/18/99
Timmy Twitfuller	Mail Room	1000	1/1/01
Buddy	Coffee Room	1000	4/13/97
Sam Mackey	The Closet with the Light Out	800	1/2/02
Todd Whitmore	The Closet with the Light Out	900	4/12/01
Ralph Ossum	Coffee Room	900	1/4/01

We can get the average salary of each department using a single query method:

```
SELECT e.DEPARTMENT, AVG(e.SALARY) FROM EMPLOYEE e GROUP BY e.DEPARTMENT
```

The results are:

DEPARTMENT	AVG(SALARY)
Coffee Room	950
Mail Room	1100
The Closet with the Light Out	850

Sub-Queries

Sub-queries are permitted as deep as the database implementation being queried allows. For example, you could use the following sub-query (in bold) specified in ejb-jar.xml. Note that the sub-query includes ORDER BY as well, and the results are to be returned in descending (DESC) order.

```
<query>
   <query-method>
     <method-name>findApStatisticsWithGreaterThanAverageValue</method-name>
     <method-params />
   </query-method>
     <ejb-ql>SELECT Object(s1) FROM ApStatistics s1 WHERE s1.averageValue >
SELECT AVG(s2.averageValue) FROM ApStatistics s2 ORDER BY s1.averageValue
DESC</ejb-ql>
</query>
```

See your database implementation documentation for details on the appropriate use of sub-queries.

Dynamic Queries

There are situations where you may need to search dynamically for data, based on variable criteria. Unfortunately EJB-QL queries do not support this scenario. Since EJB-QL queries are specified in the deployment descriptor, any changes to the queries require re-deployment of the bean. The AppServer offers a Dynamic Query feature which allows you to construct and execute EJB-QL queries dynamically and programmatically in the bean code.

Dynamic queries offer these benefits:

- allow you to create and execute new queries without having to update and deploy an EJB.
- reduce the size of the EJB's deployment descriptor file because finder queries can be dynamically created instead of statically defined in the deployment descriptors.

Dynamic queries don't need to be added to the deployment descriptor. They are declared in the bean class for dynamic `ejbSelects`, or in the local or remote home interfaces for dynamic finders.

A finder method for a dynamic query is:

```
public java.util.Collection findDynamic(java.lang.String ejbql,
   Class[] types, Object[] args)
  throws javax.ejb.FinderException

public java.util.Collection findDynamic(java.lang.String ejbql,
   Class[] types, Object[] args, java.lang.String sql)
  throws javax.ejb.FinderException
```

The `ejbSelect`s for dynamic queries are:

```
  public java.util.Collection selectDynamicLocal(java.lang.String ejbql,
Class[] types, Object[] params)
      throws javax.ejb.FinderException

public java.util.Collection selectDynamicLocal(java.lang.String ejbql, Class[]
types, Object[] params, java.lang.String sql)
      throws javax.ejb.FinderException

  public java.util.Collection selectDynamicRemote(java.lang.String ejbql,
Class[] types, Object[] params)
      throws javax.ejb.FinderException

public java.util.Collection selectDynamicRemote(java.lang.String ejbql, Class[]
types, Object[] params, java.lang.String sql)
      throws javax.ejb.FinderException

  public java.sql.ResultSet selectDynamicResultSet(java.lang.String ejbql,
Class[] types, Object[] params)
      throws javax.ejb.FinderException

public java.sql.ResultSet selectDynamicResultSet(java.lang.String ejbql,
Class[] types, Object[] params, java.lang.String sql)
      throws javax.ejb.FinderException
```

where the following applies:

- **`java.lang.String ejbql`**: this represents the actual EJB-QL syntax.
- **`Class[] types`**: this array gives the class types of the parameters to the select or finder method (it can be an empty array if there are no parameters).

- `Object[] params`: this array gives the actual values of the parameters. This is the same as the parameters argument of the regular select or finder method.

  The return type of a dynamic select or finder is always `java.util.Collection`, with the exception of the `selectDynamicResultSet`. If there is a single instance of the object or value type returned from the query, it is the first member of the collection. Dynamic queries follow the same rules as regular queries.

- `java.lang.String sql`: User specified sql. If specified, this will override the sql generated by EJB-QL.

Note

There should not be any trace of the eight methods associated with dynamic queries in your deployment descriptor.

Overriding SQL generated from EJB-QL by the CMP engine

Important

This feature is for advanced users only!

The Borland CMP engine generates SQL calls to your database based on the EJB-QL you enter in your deployment descriptors. Depending on your database implementation, the generated SQL may be less than optimal. You can capture the generated SQL using tools supplied by your backing-store implementation or another development tool. If the generated SQL is not optimal, you can replace it with your own. However, we offer no validation on the user SQL.

Note

A problem with your SQL may generate an exception which can potentially crash the system.

You specify your own optimized SQL in the Borland proprietary deployment descriptor, `ejb-borland.xml`. The XML grammar is identical to that found in `ejb-jar.xml`, except that the `<ejb-ql>` element is replaced with a `<user-sql>` element. This proprietary element contains a SQL-92 statement (**not** an EJB-QL statement) that is used to access the database instead of the CMP engine-generated SQL.

Important

The `SELECT` clause for this statement must be identical to the `SELECT` clause generated by the Borland CMP engine.

Subsequent clauses are user-optimized. The ordering of the fields in the `SELECT` clause is proprietary to the CMP engine and therefore must be preserved.

For example:

```
<entity>
 <ejb-name>EmployeeBean</ejb-name>
 ⋮
 <query>
  <query-method>
   <method-name>findWealthyEmployees</method-name>
   <method-params />
  </query-method>
  <user-sql>SELECT E.DEPT_NO, E.EMP_NO, E.FIRST_NAME, E.FULL_NAME,
          E.HIRE_DATE, E.JOB_CODE, E.JOB_COUNTRY,
          E.JOB_GRADE, E.LAST_NAME, E.PHONE_EXT, E.SALARY
       FROM EMPLOYEE E WHERE E.SALARY > 200000
  </user-sql>
 </query>
 ⋮
</entity>
```

Note

The extensive `SELECT` statement reflects the type of SQL generated by the CMP engine.

When the CMP engine encounters an EJB-QL statement in the `ejb-jar.xml` deployment descriptor, it checks `ejb-borland.xml` to see if there is any user SQL provided in the same bean's descriptor.

If none is present, the CMP engine generates its own SQL and executes it.

If the `ejb-borland.xml` descriptor does contain a query element, it uses the SQL within the `<user-sql>` tags instead.

Important

The `<query>` element in `ejb-borland.xml` **does not** replace the `<query>` element in the standard `ejb-jar.xml` deployment descriptor. If you want to override the CMP engine's SQL, you must provide the elements in **both** descriptors.

Container-managed data access support

For CMP, the Borland EJB Container supports all data types supported by the JDBC specification, including types beyond those supported by JDBC.

The following shows the basic and complex types supported by the Borland EJB Container:

- Basic types:
  - `boolean Boolean`
  - `double Double`
  - `long Long`
  - `BigDecimal java.util.Date`
  - `byte Byte`
  - `float Float`
  - `short Short`
  - `byte[]`
  - `char Character`
  - `int Integer`
  - `String java.sql.Date`
  - `java.sql.Time java.sql.TimeStamp`
- Complex types
  - Any class implementing `java.io.Serializable`, such as Vector and Hashtable
  - Other entity bean references

Note

The Borland CMP engine now supports using the Long value type for dates, as well as `java.sql.Date` for `java.util.Date`.

Keep in mind that the Borland Container supports classes implementing the `java.io.Serializable` interface, such as Hashtable and Vector. The container supports other data types, such as Java collections or third party collections, because they also implement `java.io.Serializable`. For classes and data types that implement the `Serializable` interface, the Container merely serializes their state and stores the result into a `BLOB`. The Container does not do any “smart” mapping on these classes or types; it just stores the state in binary format. The Container's CMP engine observes the following rule: the engine serializes as a `BLOB` all types that are not one of the explicitly supported types.

Depending on your database implementation, the following data types require fetching based on column index:

Database	Data Types
Oracle	■ LONG RAW

Database	Data Types
Sybase	■ NTEXT ■ IMAGE
MS SQL	■ NTEXT ■ IMAGE

Note

If you use either of the two data types `BINARY` (MS SQL) or `RAW` (Oracle) as primary keys, you must explicitly specify their size.

Support for Oracle Large Objects (LOBs)

There are two types of Large Objects (LOBs), Binary Large Objects (BLOBs) and Character Large Objects (CLOBs).

BLOBs are mapped to CMP fields with the following data types:

- byte[]
- java.io.Serializable
- java.io.InputStream

CLOBs, by virtue of being *Character* Large Objects, can only be mapped to cmp-fields with the java.lang.String data type.

By default, the Borland CMP engine does not automatically map cmp-field to LOBs. If you intend to use LOB data types, you must inform the CMP engine explicitly in the `ejb-borland.xml` deployment descriptor. You do this by setting the Column Property `createColumnSql`. For example:

```
<column-properties>
 <column-name>CLOB-column</column-name>
 <property>
  <prop-name>createColumnSql</prop-name>
  <prop-type>String></prop-type>
  <prop-value>CLOB</prop-value>
 </property>
</column-properties>

<column-properties>
 <column-name>BLOB-column</column-name>
 <property>
  <prop-name>createColumnSql</prop-name>
  <prop-type>String></prop-type>
  <prop-value>BLOB</prop-value>
 </property>
</column-properties>
```

Note

The default BLOB size limit for adding and finding BLOB data from AppServer using CMP EJBs is 10,000 bytes. The default can be changed by setting the system property below:

```
-DEJBCmpMaxBlobSize=xxxxxxxx
```

This limit does not actually control the size of the BLOB. The database and its driver can also limit this size. For example, Oracle treats BLOBs that are greater than 4GB size differently than BLOB less than 4GB.

Container-created tables

You can instruct the Borland EJB Container to automatically create tables for container-managed entities based on the entity's container-managed fields by enabling the `create-tables` property. Because table creation and data type mappings vary among vendors, you must specify the JDBC database dialect in the deployment descriptor to the Container. For all databases (except for JDataStore) if you specify the dialect, then the Container automatically creates tables for container-managed entities for you if the `create-tables` property is set to `true`. The Container will not create these tables unless you specify the dialect.

The following table shows the names or values for the different dialects (case is ignored for these values):

Database Name	Dialect Value
JDataStore	jdatastore
Oracle	oracle
Sybase	sybase
MSSQLServer	mssqlserver
DB2	db2
Interbase	interbase
Informix	informix

17

Generating Entity Bean Primary Keys

Each entity bean must have a unique primary key that is used to identify the bean instance. The primary key can be represented by a Java class, which must be a legal value type in RMI-IIOP. Therefore, it extends the `java.io.Serializable` interface. It must also provide an implementation of the `Object.equals(Object other)` and `Object.hashCode()` methods.

Normally, the primary key fields of entity beans must be set in the `ejbCreate()` method. The fields are then used to insert a new record into the database. This can be a difficult procedure, however, bloating the method, and many databases now have built-in mechanisms for providing appropriate primary key values. A more elegant means of generating primary keys is for the user to implement a separate class that generates primary keys. This class can also implement database-specific programming logic for generating primary keys.

You may either generate primary keys by hand, use a custom class, or allow the container to use the database tools to perform this for you. If you use a custom class, implement the `com.borland.ejb.pm.PrimaryKeyGenerationListener` interface, discussed below. To use the database tools, you can set properties for the CMP engine to generate primary keys depending upon the database vendor.

Generating primary keys from a user class

With enterprise beans, the primary key is represented by a Java class containing the unique data. This primary key class can be any class as long as that class is a legal value type in RMI-IIOP, meaning it extends the `java.io.Serializable` interface. It must also provide an implementation of the `Object.equals(Object other)` and `Object.hashCode()` methods, two methods which all Java classes inherit by definition.

Generating primary keys from a custom class

To generate primary keys from a custom class, you must write a class that implements the `com.borland.ejb.pm.PrimaryKeyGenerationListener` interface.

Note

This is a new interface for generating primary keys. In previous versions of Borland AppServer, this class was `com.inprise.ejb.cmp.PrimaryKeyGenerator`. This interface is still supported, but Borland recommends using the newer interface when possible.

Next, you must inform the container of your intention to use your custom class to generate primary keys for your entity beans. To do this, you set a table property `primaryKeyGenerationListener` to the class name of your primary key generator.

Implementing primary key generation by the CMP engine

Primary key generation can also be implemented by the CMP engine. Borland provides four properties to support primary key generation using database specific features. These properties are:

- `getPrimaryKeyBeforeInsertSql`
- `getPrimaryKeyAfterInsertSql`
- `ignoreOnInsert`
- `useGetGeneratedKeys`

All of these properties are table properties except `ignoreOnInsert`, which is a column property.

Oracle Sequences: using getPrimaryKeyBeforeInsertSql

The property `getPrimaryKeyBeforeInsertSql` is typically used in conjunction with Oracle Sequences. The value of this property is a SQL statement used to select a primary key generated from a sequence. For example, the property could be set to:

```
SELECT MySequence.NEXTVAL FROM DUAL
```

The CMP engine would execute this SQL and then extract the appropriate value from the ResultSet. This value will then be used as the primary key when performing the subsequent `INSERT`. The extraction from the ResultSet is based on the primary key's type

SQL Server: using getPrimaryKeyAfterInsertSql and ignoreOnInsert

Two properties need to be specified for cases involving SQL Server. The `getPrimaryKeyAfterInsertSql` property specified the SQL to execute after the `INSERT` has been performed. As above, the CMP engine extracts the primary key from the ResultSet based on the primary key's type. The property `ignoreOnInsert` must also be set to the name of the identity column. The CMP engine will then know not to set that column in the `INSERT`.

JDataStore JDBC3: using useGetGeneratedKeys

Borland's JDataStore supports the new JDBC3 method `java.sql.Statement.getGeneratedKeys()`. This method is used to obtain primary key values from newly inserted rows. No additional coding is necessary, but note that this method is unsupported in other databases and is recommended for use only with Borland JDataStore. To use this method, set the boolean property `useGetGeneratedKeys` to `True`.

Automatic primary key generation using named sequence tables

A named sequence table is used to support auto primary key generation when the underlying database (such as Oracle `SEQUENCE`) and the JDBC driver (`AUTOINCREMENT` in JDBC 3.0) do not support key generation . The named sequence table allows you to specify a table that holds a key to use for primary key generation. The container uses this table to generate the keys.

The table must contain a single row with a single column

To use the name sequence table your table must have a single row with a single column that is an integer (for the sequence values). You must create a table with one column named "`SEQUENCE`" with any initial value. For example:

```
CREATE TABLE TAB_A_SEQ (SEQUENCE int);
INSERT into TAB_A_SEQ values (10);
```

In this example key generation starts from value 10.

To enable this feature, set it in `<column-properties>` in `ejb-borland.xml`:

```
<table-properties>
        <table-name>TABLE_A</table-name>
        <column-properties>
       <column-name>ID</column-name>
            <property>
                <prop-name>autoPkGenerator</prop-name>
                <prop-type>java.lang.String</prop-type>
                <prop-value>NAMEDSEQUENCETABLE</prop-value>
            </property>
       <property>
                <prop-name>namedSequenceTableName</prop-name>
                <prop-type>java.lang.String</prop-type>
                <prop-value>TAB_A_SEQ</prop-value>
            </property>
            <property>
                <prop-name>keyCacheSize</prop-name>
                <prop-type>java.lang.Integer</prop-type>
                <prop-value>2</prop-value>
            </property>
            </column-properties>
⋮
    </table-properties>
```

Note that "ID" is the primary key column, which is marked for auto Pk Generation using `NAMEDSEQUENCETABLE`. The table used is `TAB_A_SEQ`.

Note

Set the `ejb.CacheCreate` property to `false` while using `getPrimaryKeyAfterInsert` or `useGetGeneratedKeys`. The container needs to know the primary key to dispatch calls to the bean instance. Therefore, it needs to know the primary key at the same time the `Create` method returns.

Key cache size

When generating the primary key, the container fetches the key from the table in the database. You can improve performance by reducing trips to the database by specifying a key cache size. To use this feature, in the `ejb-borland.xml` file, you set the `<key-cache-size>` element to specify how many primary key values the database will fetch. The container will cache the number of keys used for primary key generation when the value of the cache size is > 1.

The default value for key cache size, if not specified, is 1. Although key cache size is optional, it is recommended you specify a value > 1 to utilize performance optimization.

Note

There may be gaps in the keys generated if the container is rebooted or used in a clustered mode.

18

Transaction management

This chapter describes how to handle transactions.

Understanding transactions

Application programmers benefit from developing their applications on platforms such as Java 2 Enterprise Edition (J2EE) that support transactions. A transaction-based system simplifies application development because it frees the developer from the complex issues of failure recovery and multi-user programming. Transactions are not limited to single databases or single sites. Distributed transactions can simultaneously update multiple databases across multiple sites.

A programmer typically divides the total work of an application into a series of units. Each unit of work is a separate transaction. As the application progresses, the underlying system ensures that each unit of work, each transaction, fully completes without interference from other processes. If not, it rolls back the transaction and completely undoes whatever work the transaction had performed.

Characteristics of transactions

Typically, transactions refer to operations that access a shared resource like a database. All access to a database is performed in the context of a transaction. All transactions share the following characteristics:

- Atomicity
- Consistency
- Isolation
- Durability

These characteristics are denoted by the acronym ACID.

A transaction often consists of more than a single operation. Atomicity requires that either all or none of the operations of a transaction are performed for the transaction to be considered complete. If any of a transaction's operations cannot be performed, then none of them can be performed.

Consistency refers to resource consistency. A transaction must transition the database from one consistent state to another. The transaction must preserve the database's semantic and physical integrity.

Isolation requires that each transaction appear to be the only transaction currently manipulating the database. Other transactions can run concurrently. However, a

transaction must not see the intermediate data manipulations of other transactions until and unless they successfully complete and commit their work. Because of interdependencies among updates, a transaction can get an inconsistent view of the database were it to see just a subset of another transaction's updates. Isolation protects a transaction from this sort of data inconsistency.

Transaction isolation is qualified by varying levels of concurrency permitted by the database. The higher the isolation level, the more limited the concurrency extent. The highest level of isolation occurs when all transactions can be serialized. That is, the database contents look as if each transaction ran by itself to completion before the next transaction started. However, some applications can tolerate a reduced level of isolation for a higher degree of concurrency. Typically, these applications run a greater number of concurrent transactions even if transactions are reading data that may be partially updated and perhaps inconsistent.

Lastly, durability means that updates made by committed transactions persist in the database regardless of failure conditions. Durability guarantees that committed updates remain in the database despite failures that occur after the commit operation and that databases can be recovered after a system or media failure.

Transaction support

The Borland AppServer (AppServer) supports flat transactions, but not nested transactions. Transactions are implicitly propagated. This means that the user does not have to explicitly pass the transaction context as a parameter, because the J2EE container transparently handles this for the client.

Transaction management can be performed programmatically by calling the standard JTS or JTA APIs. An alternative, and more recommended approach, when writing J2EE components such as Enterprise JavaBeans (EJBs) is to use declarative transactions where the J2EE Container transparently starts and stops transactions.

Transaction manager services

There are two transaction managers, or engines, available in AppServer:

- Transaction Manager (formerly known as Partition Transaction Service)
- OTS (formerly known as 2PC Transaction Service)

A Transaction Manager exists in each AppServer Partition. It is a Java implementation of the CORBA Transaction Service Specification. The Transaction Manager supports transaction timeouts, one-phase commit protocol and can be used in a two-phase commit protocol under special circumstances.

Use the Transaction Manager under the following conditions:

- When using one-phase commit protocol.
- When you need faster performance. Currently, only the Transaction Manager can be configured to be in-process. The transaction management APIs and other transaction components are in-process JVM calls, so it is much faster than the OTS engine.
- When using a two-phase commit protocol but do not care about transaction recovery. For example, when checking business logic during development of an Enterprise JavaBean there is no need for transaction recovery. If you use the Transaction Manager for two-phase commit, you must set the “Allow unrecoverable completion” property to true in “Properties” for the Transaction Manager as displayed under the Partition in the AppServer Management Console. Alternatively, you can set system property `EJBAllowUnrecoverableCompletion` for the partition.

The OTS engine exists in a separate address space. It provides a complete solution for distributed transactional CORBA applications. Implemented on top of the VisiBroker ORB, the OTS engine simplifies the complexity of distributed transactions by providing an essential set of services—including a transaction service, recovery and logging, integration with databases, and administration facilities—within one, integrated architecture.

Distributed transactions and two-phase commit

The Borland EJB Container supports distributed transactions. Distributed transactions are those transactions that cross systems, platforms, and Java Virtual Machines (JVMs).

Transactions that manipulate data across multiple resources use a two-phase commit process. This process ensures that the transaction correctly updates all resources involved in the transaction. If it cannot update all resources, then it updates none of the resources.

Note

Although support is provided by AppServer for two-phase commit transactions, they are inherently expensive due to number of remote procedure calls (RPCs) and should be used only when needed. See “When to use two-phase commit transactions”.

There are two steps to a two-phase commit. The first step is the preparation phase. In this phase the transaction service requests that each resource involved in the transaction readies its updates and signal to the transaction service whether it can commit the updates. The second step is the commit phase. The transaction service initiates the actual resource updates only when all resources have signaled that they can complete the update process. Should any resource signal they cannot perform updates, the transaction service instructs all other resources to rollback all updates involved in the transaction.

The Transaction Manager and OTS engine support both heterogeneous distributed (two-phase commit) transactions and two-phase commit for homogeneous resources.

By default, the Transaction Manager does not allow multiple resources to participate in a global transaction, but it can be configured to allow multiple resource participation through its support for unrecoverable transaction completion. This can be enabled on the Transaction Manager by setting either “Allow unrecoverable completion” option from the Management Console (right-click the Transaction Manager and select “Properties”), or the Partition system property `EJBAllowUnrecoverableCompletion`. When unrecoverable transaction completion is enabled, the container makes a one-phase commit call on each participating resource during the transaction commit process. Care must taken when enabling unrecoverable transaction completion; as the name suggests, no recovery is available when a failure occurs prior to transaction completion, which may lead to inconsistent states in participating resources.

To support heterogeneous two-phase commit transactions, the OTS engine must integrate with XA support in the underlying resources. With availability of XA-enabled JDBC drivers from DBMS vendors and JMS support provided by message service providers, the EJB container and OTS engine allow multiple resources to participate in a single transaction.

Two-phase commit for homogeneous databases requires some configuration of the DBMS servers. While the container controls the commit to the first database, the DBMS server controls the commits to the subsequent databases using the DBMS's built-in transaction coordinator. For more information, see your vendor's manual for the DBMS server.

When to use two-phase commit transactions

One of the basic principals of building high performance distributed applications is to limit the number of remote procedure calls (RPCs). The following explains typical situations; when and when not to use two-phase commit transactions. Avoiding a two-phase commit transaction when it is not needed, therefore avoiding unnecessary RPCs involving JTA XAResource objects and the OTS engine, greatly improves your application's performance.

Using multiple JDBC connections for access to multiple database resources from a single vendor in the same transaction

In scenarios involving multiple databases from a single vendor, it is often possible to avoid using two-phase commit. You can access one database and use it to access the second database by tunneling access through the connection to the first database. Oracle and other DBMSs provide this capability. In this case the AppServer Partition

can be configured with only one JDBC connection to the “fronting” database. Access to the “backing” database is tunneled through the first JDBC connection.

Using multiple JDBC connections to the same database resource in the same transaction

When multiple JDBC connections to the same database are obtained and used by distributed participants within a single transaction, a two-phase commit can be avoided. The JDBC connections, as expected, need to be obtained from a XA datasource. But, rather than performing a two-phase commit, a one-phase commit can be used to complete the transaction since only a single resource is involved. This is achieved by using the Transaction Manager rather than the OTS engine. An alternative is to collocate all EJBs involved in the transaction, rather than having them deployed in distributed Partitions. In this case, a non-XA datasource is used and no two-phase commit is required.

Using multiple disparate resources in a single transaction

In this case there is a need for a two-phase commit transaction. This situation arises when, for example, you are running a single transaction against both Oracle and Sybase, or if you have a transaction that includes access to an Oracle database and a JMS provider, such as MQSeries. In the latter case, the transaction is coordinated using JTA XAResource object, obtained from Oracle via JDBC and MQSeries via JMS, and enables both resources to participate in the two-phase commit transaction completion. It is worth noting that two-phase commit capabilities (provided by the OTS engine), are only needed when a single transaction involves access to multiple incompatible resources.

Note

In order to utilize the OTS engine as the default transaction service, the Transaction Manager must be stopped first.

EJBs and 2PC transactions

With the introduction of messaging in the J2EE platform, a number of common scenarios now exist involving access to multiple resources from EJBs in a single transaction . As we know, when more than one resource is involved in a transaction, the OTS engine is needed to reliably complete the transaction using the two-phase commit protocol. Sample scenarios include:

- A session bean accesses two types of entity beans in a transaction where each are persisted in a different database.
- A session bean accesses an entity bean and in the same transaction does some messaging work, such as sending a message to a JMS queue.
- In the `onMessage` method of a message-driven bean, access entity beans on message delivery.

In each of the above examples, two heterogeneous resources need to be accessed from within a session bean or a message-driven bean as part of a single transaction. These EJBs have the `REQUIRED` transaction attribute defined and need access to the OTS engine. However, if the OTS engine is running, then all modules deployed to that Partition are able to discover it and can attempt to use it. The OTS engine will perform a one-phase commit when only one resource is registered in a transaction, but suffers the extra RMI overhead since it is an external process. Ideally, the in-process Transaction Manager should be used for EJBs not involved in a two-phase commit transaction. To better utilize the transaction services available in AppServer, a bean-level property, `ejb.transactionManagerInstanceName` may be specified for EJBs that require 2PC transaction completion. This property provides the name of the OTS engine to be used by the EJB container doing transaction demarcation on any of the methods for the relevant bean. Both the Transaction Manager and the OTS engine may be available for all EJBs but only those that do not have `ejb.transactionManagerInstanceName` specified will discover the Transaction Manager.

This property can be commonly used for session or message-driven beans since transactions are usually demarcated in a session bean facade or the `onMessage` method of a message-driven bean.

To set the `ejb.transactionManagerInstanceName` property use the Management Console. Navigate to your deployed EJB module, right-click on it and select “DDEditor”. In the DDEditor select the required bean from the Navigation Pane. Select the “Properties” tab and add the `ejb.transactionManagerInstanceName` property. Define the property as a `String` and specify a unique name value such as “MyTwoPhaseEngine”.

Next, you must modify the OTS engine factory name with the `ejb.transactionManagerInstanceName` value. In the Management Console, select the OTS engine from the “corbaSample” configuration, identified as the “OTS engine” managed object type. Right-click and select “Properties” from the drop-down menu. In the Properties dialog choose the Settings tab and modify the value for “Factory Name”. Click OK, and restart the service. The OTS engine may also be started from the command line, independent of an AppServer server. The factory name can be provided using property `vbroker.ots.name` as follows:

```
prompt> ots -Dvbroker.ots.name=<MyTwoPhaseEngine>
```

The EJB will now use the OTS engine named “MyTwoPhaseEngine”. As mentioned, the Partition may be hosting several J2EE modules, but only those beans that have `ejb.transactionManagerInstanceName` set go to the (non-default) OTS engine. Other beans in the Partition that require method invocation in a transaction, but do not require 2PC, always find the Transaction Manager due to local service affinity.

Following is a deployment configuration usage example. Displayed below is an extract from deployment descriptor `ejb-borland.xml` packaged with the deployed EJB module and viewable in the DDEditor. The `ejb.transactionManagerInstanceName` property is set for Session bean “`OrderSesEJB`” where **`OrderSesEJB`** takes orders from customers, creates an order in the database and sends messages to the manufacturers for making parts.

```
<ejb-jar>
    <enterprise-beans>
        <session>
            <ejb-name>OrderSesEJB</ejb-name>
            <bean-home-name>OrderSes</bean-home-name>
            <bean-local-home-name />
            <ejb-local-ref>
                <ejb-ref-name>ejb/OrderEntLocal</ejb-ref-name>
                <jndi-name>OrderEntLocal</jndi-name>
            </ejb-local-ref>
            <ejb-local-ref>
                <ejb-ref-name>ejb/ItemEntLocal</ejb-ref-name>
            </ejb-local-ref>
            <resource-ref>
                <res-ref-name>jms/QueueConnectionFactory</res-ref-name>
                <jndi-name>QueueConnectionFactory</jndi-name>
            </resource-ref>
            <resource-env-ref>
                <resource-env-ref-name>jms/OrderQueue</resource-env-ref-name>
                <jndi-name>OrderQueue</jndi-name>
            </resource-env-ref>
            <property>
                <prop-name>ejb.transactionManagerInstanceName</prop-name>
                <prop-type>String</prop-type>
                <prop-value>TwoPhaseEngine</prop-value>
            </property>
        </session>

<ejb-jar>
```

Example runtime scenarios

The following diagrams show configurations where the standard Transaction Manager and the OTS engine co-exist. The deployment configuration is done in a manner in

which the beans participating in 2PC transactions have their transaction management done by the OTS engine, named "TwoPhaseEngine", and those that don't need 2PC transactions use the default in- process Transaction Manager.

The example archive used is complex.ear, in an AppServer Partition. It has three beans:

- **OrderSesEJB:** takes orders from customers, creates an order in the database, and sends messages to the manufacturers for making parts.
- **UserSesEJB:** creates new users in the company database. Only accesses a single database, therefore only needs to access a 1PC engine (Transaction Manager).
- **OrderCompletionMDB:** receives a notification from the manufacturer about the part delivery, and also updates the database using entity beans.

To configure this example deployment scenario:

1. Using the DDEditor, add the ejb.transactionManagerInstance property to the beans OrderSesEJB and OrderCompletionMDB. Refer to the above XML extract for this example.
2. Next, using the Management Console, start the OTS engine with factory name set as "TwoPhaseEngine".
3. Keep the local Transaction Manager enabled.

The following diagrams show example interactions between the client, the AppServer Partition, and how the AppServer Partition locates the right transaction service based on the above configuration. All of the beans are assumed to have container-managed transactions.

Example 1PC usage

1. The client calls a method of UserSesEJB. This is an implementation of the method that creates users in the database.
2. Before the call is actually invoked, as shown below, the Partition uses its in-process Transaction Manager to begin the transaction.
3. The session bean does some database work.
4. When the call is over, the Partition issues commit.
5. The Transaction Manager calls commit_one_phase() on the database resource.

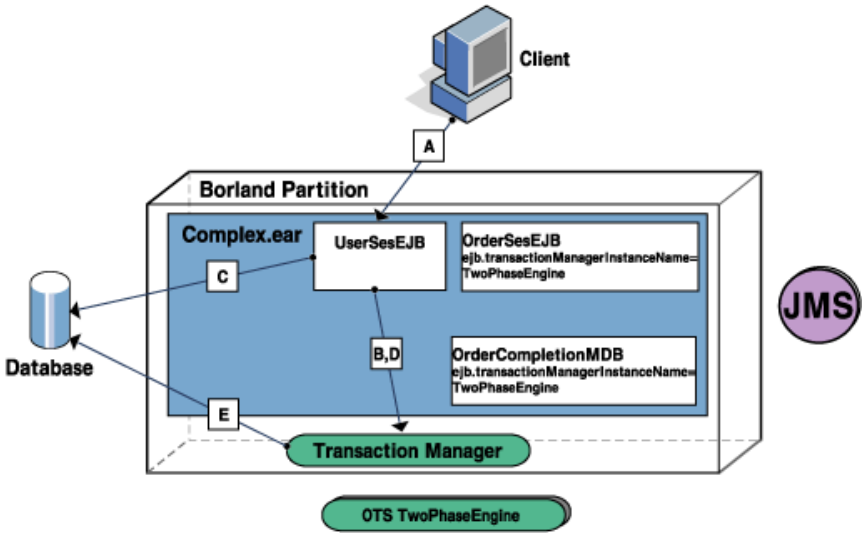

Example 2PC usage

1. The client calls `OrderSesEJB.create()` method to create a new order.
2. Since the bean is configured to use the OTS engine named **TwoPhaseEngine**, the container locates the right transaction service named TwoPhaseEngine, and uses it for beginning the transaction.
3. The session bean does some database work, and sends a message to a JMS queue.
4. When the call is over, the Partition issues commit.
5. The OTS engine coordinates the transaction completion with the database and the JMS resources.

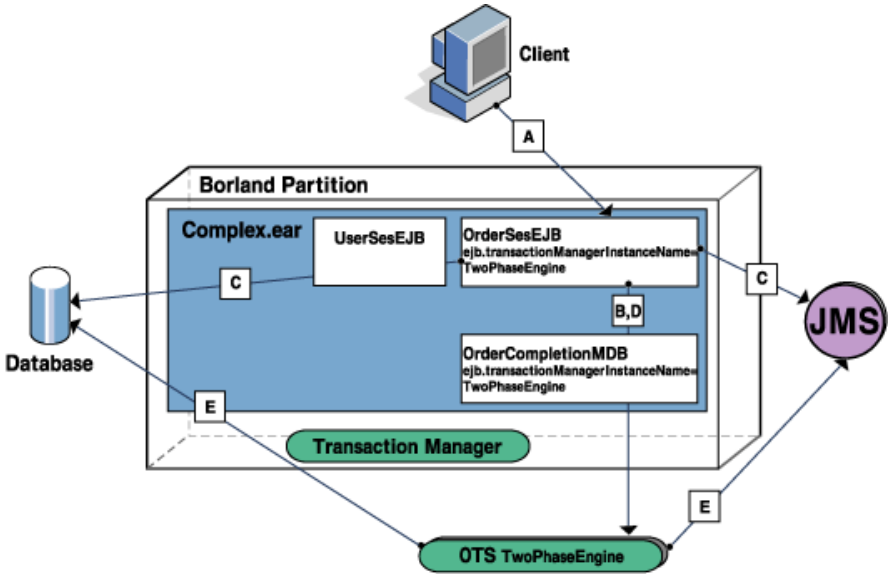

Example 2PC usage with MDBs

At some point in time, an asynchronous message is delivered to `OrderCompletionMDB` by invoking its `onMessage()` method, which has a `REQUIRED` transaction attribute. The container starts a transaction using ITS and then invokes the `onMessage()` method. In the body of the method, the bean updates the database to indicate order delivery. It is important to note that there are 2 resources involved. The first one is the JMS resource, which is associated with the MDB instances that got the message, and the second is the database that the MDB instance updated. This scenario is similar to the example diagram above.

Note

`ejb.transactionManagerInstanceName` is also supported for MDBs. See "MDBs and Transactions" for more information.

Declarative transaction management in Enterprise JavaBeans

Transaction management for Enterprise JavaBeans (EJBs) is handled by the EJB Container and the EJBs. Enterprise JavaBeans make it possible for applications to update data in multiple databases within a single transaction.

EJBs utilize a declarative style of transaction management that differs from the traditional transaction management style. With declarative management, the EJB declares its transaction attributes at deployment time. The transaction attributes indicate whether the EJB container manages the bean's transactions or whether the

bean itself manages its own transactions, and, if so, to what extent it does its own transaction management.

Traditionally, the application was responsible for managing all aspects of a transaction. This entailed such operations as:

- Creating the transaction object.
- Explicitly starting the transaction.
- Registering resources involved in the transaction.
- Keeping track of the transaction context.
- Committing the transaction when all updates completed.

It requires a developer with extensive transaction processing expertise to write an application that is responsible for managing a transaction from start to finish. The code for such an application is more complex and difficult to write, and it is easy for “pilot error” to occur.

With declarative transaction management, the EJB container manages most if not all aspects of the transaction for you. The EJB container handles starting and ending the transaction, plus maintains its context throughout the life of the transaction object. This greatly simplifies an application developer's responsibilities and tasks, especially for transactions in distributed environments.

Understanding bean-managed and container-managed transactions

When an EJB programmatically performs its own transaction demarcation as part of its business methods, that bean is considered to be using bean-managed transaction. On the other hand, when the bean defers all transaction demarcation to its EJB container, and the container performs the transaction demarcation based on the Application Assembler's deployment instructions, then the bean is referred to as using container-managed transaction.

EJB session beans, both stateful and stateless varieties, can use either container- or bean-managed transactions. However, a bean cannot use both types of transaction management at the same time. EJB entity beans can only use container-managed transaction. It is the bean provider who decides the type of transaction which an EJB can use.

An EJB can manage its own transaction if it wishes to start a transaction as part of one operation and then finish the transaction as part of another operation. However, such a design might be problematic if one operation calls the transaction starting method, but no operation calls the transaction ending method.

Whenever possible, enterprise beans should use container-managed transactions as opposed to bean-managed transactions. Container-managed transactions require less programming work and are less prone to programming error. In addition, a container-managed transaction bean is easier to customize and compose with other beans.

Local and Global transactions

A transaction involves an atomic unit of work performed against data maintained by one or more resource managers. Examples of resource managers are database managements systems and JMS message providers. A local transaction involves work performed against a single resource manager independent of an external transaction manager. For instance, a JDBC connection obtained from a database can have SQL operations performed on it to update the database and then have the work committed in a local transaction using a `commit()` operation, assuming autoCommit mode for the connection is turned off, otherwise each operation is performed within a local transaction. A global transaction is coordinated by a transaction manager, such as the partition Transaction Manager or OTS engine, and it can involve work performed for one or more distributed resource managers. Transaction management for EJBs, both container-managed and bean-managed, implies use of global transactions. When a single resource manager participates in a global transaction, all work may be performed within a local transaction on behalf of the global transaction (refer to EJB

Specification Version 2.0, Section 17.6.4 Local transaction optimization for more details).

Methods of an EJB defined with bean-managed transactions must obtain an implementation handle to JTA interface `javax.transaction.UserTransaction` and invoke operations on it to explicitly participate in a global transaction.

With container-managed transactions, the EJB Container interposes each EJB method call and follows certain rules to determine whether or not work should be processed as part of a global transaction. The decision taken by the Container depends on the transaction attribute value set for the method, by the Application Assembler in the components deployment descriptor, and whether a global transaction context exists upon invocation of the method (refer to Table 14 in EJB Specification Version 2.0, Section 17.6.2.7 Transaction attribute summary). Should the method be processed without the presence of a global transaction context, work performed against an external resource manager from within the method is completed using local transaction(s). The following are specific examples of when local transactions are used for EJB methods of an EJB with container-managed transaction demarcation:

- If the transaction attribute is set to `NotSupported` and the container detects that resources were accessed.
- If the transaction attribute is set to `Supports` and the container detects that a) the method was not invoked from within a global transaction, and b) resources were accessed.
- If the transaction attribute is set to `Never` and the container detects that resources were accessed.

Transaction attributes

EJBs that use bean-managed transaction have transaction attributes associated with each method of the bean. The attribute value tells the container how it must manage the transactions that involve this bean. There are six different transaction attributes that can be associated with each method of a bean. This association is done at deployment time by the Application Assembler or Deployer.

These attributes are:

- `Required`: This attribute guarantees that the work performed by the associated method is within a global transaction context. If the caller already has a transaction context, then the container uses the same context. If not, the container begins a new transaction automatically. This attribute permits easy composition of multiple beans and co-ordination of the work of all the beans using the same global transaction.
- `RequiresNew`: This attribute is used when the method does not want to be associated with an existing transaction. It ensures that the container begins a new transaction.
- `Supports`: This attribute permits the method to avoid using a global transaction. This must only be used when a bean's method only accesses one transaction resource, or no transaction resources, and does not invoke another enterprise bean. It is used solely for optimization, because it avoids the cost associated with global transactions. When this attribute is set and there is already a global transaction, the EJB Container invokes the method and have it join the existing global transaction. However, if this attribute is set, but there is no existing global transaction, the Container starts a local transaction for the method, and that local transaction completes at the end of the method.
- `NotSupported`: This attribute also permits the bean to avoid using a global transaction. When this attribute is set, the method must not be in a global transaction. Instead, the EJB Container suspends any existing global transaction and starts a local transaction for the method, and the local transaction completes at the conclusion of the method.
- `Mandatory`: It is recommended that this attribute not be used. Its behavior is similar to Requires, but the caller must already have an associated transaction. If not, the container throws a `javax.transaction.TransactionRequiredException`. This attribute

makes the bean less flexible for composition because it makes assumptions about the caller's transaction.

– `Never`: It is recommended that this attribute not be used. However, if used, the EJB Container starts a local transaction for the method. The local transaction completes at the conclusion of the method.

Under normal circumstances only two attributes, `Required` and `RequiresNew`, must be used. The attributes `Supports` and `NotSupported` are strictly for optimization. The use of `Never` and `Mandatory` are not recommended because they affect the composibility of the bean. In addition, if a bean is concerned about transaction synchronization and implements the `javax.ejb.SessionSynchronization` interface, then the Assembler/ Deployer can specify only the attributes `Required`, `RequiresNew`, or `Mandatory`. These attributes ensure that the container invokes the bean only within a global transaction, because transaction synchronization can only occur within a global transaction.

Note

When a client calls an EJB which in turn calls another EJB, and both EJBs access the same database, only one JDBC connection will be used if the Transaction attribute of the methods involved is set to Required. The reason is that work done in each of the beans becomes part of the single transaction.

Programmatic transaction management using JTA APIs

All transactions use the Java Transaction API (JTA). When transactions are container managed, the platform handles the demarcation of transaction boundaries and the container uses the JTA API; you do not need to use this API in your bean code.

A bean that manages its own transactions (bean-managed transaction), however, must use the JTA `javax.transaction.UserTransaction` interface. This interface allows a client or component to demarcate transaction boundaries. Enterprise JavaBeans that use bean-managed transactions use the method `EJBContext.getUserTransaction()`.

In addition, all transactional clients use JNDI to look up the `UserTransaction` interface. This simply involves constructing a JNDI `InitialContext` using the JNDI naming service, as shown in the following line of code:

```
javax.naming.Context context = new javax.naming.InitialContext();
```

Once the bean has obtained the `InitialContext` object, it can then use the JNDI `lookup()` operation to obtain the `UserTransaction` interface, as shown in the following code sample.

```
javax.transaction.UserTransaction utx = (javax.transaction.UserTransaction)
      context.lookup("java:comp/UserTransaction");
```

Note that an EJB can obtain a reference to the `UserTransaction` interface from the `EJBContext` object. This is because an enterprise bean by default inherits a reference to the `EJBContext` object. Thus, the bean can simply use the `EJBContext.getUserTransaction()` method rather than having to obtain an `InitialContext` object and then using the JNDI `lookup()` method. However, a transactional client that is not an enterprise bean must use the JNDI lookup approach.

When the bean or client has the reference to the `UserTransaction` interface, it can then initiate its own transactions and manage these transactions. That is, you can use the `UserTransaction` interface methods to begin and commit (or rollback) transactions. You use the `begin()` method to start the transaction, then the `commit()` method to commit the changes to the database. Or, you use the `rollback()` method to abort all changes made within the transaction and restore the database to the state it was in prior to the start of the transaction. Between the `begin()` and `commit()` methods, you include code to carry out the transaction's business.

JDBC API Modifications

The standard Java Database Connectivity (JDBC) API is used by the AppServer to access databases that support JDBC through vendor provided drivers. Requests for access to a database is centralized through the AppServer JDBC Connection Pool.

This section describes modifications the AppServer JDBC pool makes to JDBC behavior for transactions.

The JDBC pool is a pseudo JDBC driver that allows a transactional application to obtain a JDBC connection to a database. The JDBC pool associates JDBC connections with the Transaction Manager's transactions, and delegates connection requests to JDBC drivers that factory the JDBC connections. Once a connection is obtained using the JDBC pool, the transaction is coordinated automatically by the transaction service.

The JDBC pool and its associated resources provide complete transactional access to the DBMS. The JDBC pool registers resources transparently with the transaction coordinator. Because of limitations of the 1.x version of the JDBC API, the JDBC pool can only provide one-phase commit. Version 2.0 of the JDBC API supports full two-phase commit.

Modifications to the behavior of the JDBC API

To enable JDBC access for transactional applications written in Java, you use the JDBC API. The JDBC API is fully documented at the following web site:

```
http://www.javasoft.com/products/jdk/1.2/docs/guide/jdbc/spec/
jdbc-spec.frame.html
```

However, the behavior of some JDBC methods is overridden by the partition's transaction service when they are invoked within the context of a transaction managed by the partition. The following methods are affected:

- `Java.sql.Connection.commit()`
- `Java.sql.Connection.rollback()`
- `Java.sql.Connection.close()`
- `Java.sql.setAutoCommit(boolean)`

The rest of this section explains the changes to the semantics of these methods for partition-managed transactions.

Note

If a thread is not associated with a transaction, all of these methods will use the standard JDBC transaction semantics.

Overridden JDBC methods

Java.sql.Connection.commit()

As defined in the JDBC API, this method commits all work that was performed on a JDBC connection since the previous `commit()` or `rollback()`, and releases all database locks.

If a global transaction is associated with the current thread of execution do not use this method. If the global transaction is not a container-managed transaction, that is the application manages its own transactions, and a commit is required use the JTA API to perform the commit rather than invoking `commit()` directly on the JDBC connection.

Java.sql.Connection.rollback()

As defined in the JDBC API, this method rolls back all work that was performed on a JDBC connection since the previous `commit()` or `rollback()`, and releases all database locks.

If a global transaction is associated with the current thread of execution do not use this method. If the global transaction is not a container-managed transaction, that is the application manages its own transactions, and a rollback is required use the JTA API to perform the rollback rather than invoking `rollback()` directly on the JDBC connection.

Java.sql.Connection.close()

As defined in the JDBC API, this method closes the database connection and all JDBC resources associated with the connection.

If the thread is associated with a transaction this call simply notifies the JDBC pool that work on the connection is complete. The JDBC pool releases the connection back to the connection pool once the transaction has completed. JDBC connections opened by the JDBC pool cannot be closed explicitly by an application.

Java.sql.Connection.setAutoCommit(boolean)

As defined in the JDBC API, this method is used to set the auto commit mode of a transaction. The `setAutoCommit()` method allows Java applications to either:

- Execute and commit all SQL statements as individual transactions (when set to true). This is the default mode, or
- Explicitly invoke `commit()` or `rollback()` on the connection (when set to false).

If the thread is associated with a transaction, the JDBC pool turns off the auto-commit mode for all connections factoried in the scope of a partition's transaction service transaction. This is because the transaction service must control transaction completion. If an application is involved with a transaction, and it attempts to set the auto commit mode to true, the `java.sql.SQLException()` will be raised.

Handling of EJB exceptions

Enterprise JavaBeans can throw application and/or system level exceptions if they encounter errors while handling transactions. Application-level exceptions pertain to errors in the business logic and are intended to be handled by the calling application. System-level exceptions, such as runtime errors, transcend the application itself and can be handled by the application, the bean, or the bean container.

The EJB declares application-level exceptions and system-level exceptions in the `throws` clauses of its `Home` and `Remote` interfaces. You must check for checked exceptions in your program `try`/`catch` block when calling bean methods.

System-level exceptions

An EJB throws a system-level exception, which is a `java.ejb.EJBException` (but may also be a `java.rmi.RemoteException`), to indicate an unexpected system-level failure. For example, it throws this exception if it cannot open a database connection. The `java.ejb.EJBException` is a runtime exception and does not have to be listed in the `throws` clause of the bean's business methods.

System-level exceptions usually require the transaction to be rolled back. Often, the container managing the bean does the rollback. Other times, especially with bean-managed transactions, the client must rollback the transaction.

Application-level exceptions

An EJB throws an application-level exception to indicate application-specific error conditions, that is, business logic errors and not system problems. These application-level exceptions are exceptions other than `java.ejb.EJBException`. Application-level exceptions are checked exceptions, which means you must check for them when you call a method that potentially can throw this exception.

The EJB's business methods use application exceptions to report abnormal application conditions, such as unacceptable input values or amounts beyond acceptable limits. For example, a bean method that debits an account balance can throw an application exception to report that the account balance is not sufficient to permit a particular debit operation. A client can often recover from these application-level errors without having to rollback the entire transaction.

The application or calling program gets back the same exception that was thrown and this allows the calling program to know the precise nature of the problem. When an

application-level exception occurs, the EJB instance does not automatically rollback the client's transaction. The client now has the knowledge and the opportunity to evaluate the error message, take the necessary steps to correct the situation, and recover the transaction. Otherwise, the client can abort the transaction.

Handling application exceptions

Because application-level exceptions report business logic errors, the client is expected to handle these exceptions. While these exceptions can require transaction rollback, they do not automatically mark the transaction for rollback. You often have the option to retry the transaction, though there are times when you must abort and rollback the transaction.

The bean Provider is responsible for ensuring that the state of the bean is such that, if the client continues with the transaction, there is no loss of data integrity. If the Provider cannot ensure this degree of integrity, then the bean marks the transaction for rollback.

Transaction rollback

When your client program gets an application exception, you must first check if the current transaction has been marked for “rollback only”. For example, a client can receive a `javax.transaction.TransactionRolledbackException`. This exception indicates that the helper enterprise bean failed and the transaction has been aborted or marked “rollback only”. In general, the client does not know the transaction context within which the called enterprise bean operated. The called bean may have operated in its own transaction context separate from the calling program's transaction context, or it may have operated in the calling program's context.

If the EJB operated in the same transaction context as the calling program, then the bean itself (or its container) may have already marked the transaction for rollback. When the EJB container has marked a transaction for rollback, the client should stop all work on the transaction. Normally, a client using declarative transactions will get an appropriate exception, such as `javax.transaction.TransactionRolledbackException`. Note that declarative transactions are those transactions where the container manages the transaction details.

A client that is itself an EJB calls the `javax.ejb.EJBContext.getRollbackOnly` method to determine if its own transaction has been marked for rollback or not.

For bean-managed transactions—those transactions managed explicitly by the client—the client should rollback the transaction by calling the `rollback` method from the `java.transaction.UserTransaction` interface.

Options for continuing a transaction

When a transaction is not marked for rollback, then the client has three options:

- Rollback the transaction.
- Pass the responsibility by throwing a checked exception or re-throwing the original exception.
- Retry and continue the transaction. This can entail retrying portions of the transaction.

When a client receives a checked exception for a transaction not marked for rollback, its safest course is to rollback the transaction. The client does this by either marking the transaction as “rollback only” or, if the client has actually started the transaction, calling the `rollback` method to actually rollback the transaction.

The client can also throw its own checked exception or re-throw the original exception. By throwing an exception, the client lets other programs further up the transaction chain decide whether or not to abort the transaction. However, in general it is preferable for the code or program closest to the occurrence of the problem to make the decision about saving the transaction.

Lastly, the client can continue with the transaction. The client can evaluate the exception message and decide if invoking the method again with different parameters is likely to succeed. However, you need to keep in mind that retrying a transaction is

potentially dangerous. You have no knowledge of nor guarantee that the enterprise bean properly cleaned up its state.

Clients that are calling stateless session beans, on the other hand, can retry the transaction with more confidence if they can determine the problem from the thrown exception. Because the called bean is stateless, the client does not have the problem of not knowing the state in which the bean left the transaction.

19

Message-Driven Beans and JMS

JMS and EJB

According to the specification, there are no limitations on a bean acting as a JMS message producer or synchronous consumer. It can use the regular JMS APIs to send a message to a queue or publish to a topic. As long as you perform synchronous style consumption of messages (that is, not based on `javax.jms.MessageListener`), then there are no problems on the consumption side either. The complexity lies in the need for a JMS message send or receive request to participate in a transaction context shared by other work in an application. We already know how to solve this problem using JMS and JTA in a non-EJB application. The EJBs demand no special treatment.

Since EJB method invocations are synchronous, some calls will have to wait until the bean has completed its processing. This may include calling other beans, databases, and so forth. This RMI behavior can be undesirable in many situations. For example, you may just want to call the method and have it return before doing any heavy processing, allowing the caller to proceed with other tasks in the meantime. Threading in the client is an obvious way to achieve this, but it suffers from two problems:

- the client's programming model is not a true asynchronous style
- if the client is an EJB, threading is prohibited in its method implementations

The most desirable scenario is for an appclient, servlet, EJB, or other component to have the capability to fire a message and then have an EJB be driven asynchronously by that message. In turn, that EJB can send a message to another EJB or perform direct data access or other business logic. The caller does not wait beyond the time the message is successfully queued. On the other side, the EJB can process the message at its convenience. This EJB's processing typically involves a unit of work made up of three operations:

1. dequeueing the message,
2. activating an instance and performing whatever work the business logic demands, and
3. optionally queuing a reply message back

Enterprise systems require that it be possible to have transactional and other container-managed guarantees for this unit of work.

EJB 2.0 Message-Driven Bean (MDB)

The EJB 2.0 specification formalizes the integration between JMS and asynchronous invocation of enterprise beans by pushing these responsibilities to the EJB Container. This eases the burden on the developer, who now simply provides a class that is a JMS listener and also an EJB. This is done by implementing `javax.jms.MessageListener` and `javax.ejb.MessageDrivenBean` in the class. This and an XML descriptor containing all the deployment settings is all that the application programmer needs to provide.

From a client's perspective, this EJB is nonexistent. The client simply publishes messages to the queue or topic. The EJB container associates the MDB with the published queue/topic and handles all lifecycle, pooling, concurrency, reentrance, security, transaction, message handling, and exception handling issues.

EJB 2.1 MDB

With integration of the J2EE Connector Architecture 1.5 (JCA 1.5) in EJB 2.1 the MDB can now process messages from non-JMS messaging servers in addition to JMS based providers. A JCA 1.5 compliant resource adapter implementation can be developed for any type of messaging server and deployed to an application server. When configured to pass inbound messages from the messaging server to the application server, the resource adapter can be selected as the source for messages driving 2.1 MDBs.

JCA 1.5 defines a Message Inflow contract which is a messaging contract between the EJB container and an asynchronous connector, so that MDBs automatically process incoming messages from an EIS or some other type of messaging provider. EJB 2.1 MDBs must implement the standard javax.ejb.MessageDrivenBean interface as well as a specific messaging interface defined by the connector. If the connector is a JMS-based provider, the MDB must implement javax.jms.MessageListener, but for non-JMS providers it must implement some other type of interface that is specific to the provider.

In Borland Application Server 6.7, EJB 2.1 MDBs can be configured to process messages from JMS providers either indirectly through a JCA resource adapter or directly without the need for a pre-deployed JCA resource adapter.

Client View of an MDB

Clients do not bind to an MDB like they do for session beans and entity beans. The client only needs to send a message to the destination to which the MDB is configured to listen. Typically clients also use the `<resource-ref>` and `<resource-env-ref>`(in EJB 2.0) or `<message-destination-ref>`(in EJB 2.1) for JMS destination specification in their deployment descriptor and then point to the same JNDI names as configured in the MDB deployment descriptor. See “Obtaining JMS Connection Factories and Destinations in J2EE Application Components” for information on how to configure your client deployment descriptors to communicate with the JMS provider.

This being the case, there is no EJB metadata or handle of which the client needs to be aware. This is because there is no RMI client view of a Message Driven Bean.

MDB Configuration

Since MDBs do not expose EJB interfaces, they do not have JNDI names in the normal sense like `EJBHome` objects do. When an MDB is deployed, it communicates with a message provider in preparation for processing incoming messages.

EJB 2.0 MDBs are associated with two JMS resource objects that must pre-exist in JNDI before the MDB is deployed. These are:

- a JMS connection factory to use for connecting to the JMS provider and
- a JMS queue/topic on that provider to listen to for incoming messages

The JNDI names for these objects are specified in the MDB's `ejb-borland.xml` deployment descriptor. The `<connection-factory-name>` captures the resource connection factory used to connect to the JMS service provider. The `<message-driven-`

destination-name> element captures the actual topic/queue on which the MDB is to listen. Once these elements are specified, the MDB has all the information it needs to connect to the JMS service provider, receive messages, and send replies.

EJB 2.1 MDBs can be configured in one of two ways. If the EJB 2.1 MDB implements javax.jms.MessageListener, indicating a JMS based MDB, it can be configured to communicate directly with the JMS provider rather than use a JCA 1.5 connector. In this case, JNDI names for JMS resource objects can be specified in the MDB's ejb-borland.xml deployment descriptor under <jms-provider-ref> element. An EJB 2.1 MDB can alternatively be configured to receive messages from a JCA 1.5 connector using element <resource-adapter-ref> in the Borland-specific deployment descriptor file ejb-borland.xml.

Connecting to a JMS Server from EJB 2.0 MDBs

EJB 2.0 MDBs provide a special case for connecting to the JMS server, the source for inbound messages. In the standard deployment descriptor file, ejb-jar.xml, the type of JMS destination from which the inbound messages are received, is defined using the <message-driven-destination> element in the MDB's declaration. For example:

```
<message-driven>
 <ejb-name>MyMDBTopic</ejb-name>
⋮
<message-driven-destination>
  <destination-type>javax.jms.Topic</destination-type>
  <subscription-durability>Durable</subscription-durability>
 </message-driven-destination>
⋮
</message-driven>
```

Consult the J2EE 1.3 Specification for the proper use of this element. The Borland-specific XML file, ejb-borland.xml binds the logical name of the JMS destination with the JNDI name using equivalent element, <message-driven-destination>. The JNDI name for JMS connection factory required to connect to JMS server must also be defined using <connection-factory-name>. For example,

```
<message-driven>
<ejb-name>MyMDBTopic</ejb-name>
⋮
<message-driven-destination>jms/resources/Topic</message-driven-destination>
<connection-factory-name>jms/resources/tcf</connection-factory-name>
⋮
</message-driven>
```

See the Configuring JMS Connection Factories and Destinations section in the *Using JMS* section for detailed information on configuring these JMS resource objects bound under JNDI.

Note

You must use an XA connection factory when the MDB is deployed with the REQUIRED transaction attribute. The whole idea of this deployment is to enable the consumption of the message that drives the MDB to share the same transaction as any other work that is done from within the MDB.onMessage() method. To achieve this the container performs XA coordination with the JMS service provider and any other resources enlisted in the transaction.

Connecting to message source from EJB 2.1 MDBs

As a result of EJB 2.1 and JCA 1.5, there have been changes to both the standard deployment descriptor, ejb-jar.xml, and the Borland proprietary deployment descriptor, ejb-borland.xml for J2EE 1.4.

Changes to ejb-jar.xml

Each EJB 2.1 MDB is connected to its message source based on the information in the deployment descriptor. The standard deployment descriptor, `ejb-jar.xml`, in EJB 2.1 has changed to accommodate the connector-based MDB.

EJB 2.1 adds new elements, `<messaging-type>`, `<message-destination-type>`, and `<activation-config>`, in the `ejb-jar.xml` file:

The `<messaging-type>` element indicates which message type will be used, it does this by stating the fully-qualified interface name that the MDB implements. If no interface name is given, the container defaults to JMS message type, `javax.jms.MessageListener`.

The optional `<message-destination-type>` element designates a fully-qualified interface name that represents the type of destination from which the bean will get messages. For MDBS that represent JMS message type, `javax.jms.MessageListener`, the allowed values are `javax.jms.Topic` or `javax.jms.Queue`.

Since the connector-based MDBs no longer rely exclusively on JMS, the EJB 2.0 elements `<message-driven-destination><message-selector>` and `<acknowledge-mode>` elements are eliminated in EJB 2.1. Configuration properties required by EJB 2.1 MDB activation can be defined in a generic set of name-value pairs under the `<activation-config>` element. The property names and values used to describe the messaging service vary depending on the type of service used. These `<activation-config>` properties are examined when the message-driven bean is deployed. Each eliminated JMS-related element from EJB 2.0 can be represented by an `<activation-config-property>` element when `<messaging-type>` element specifies a JMS messaging type (`javax.jms.MessageListener`)

Here is an example of how a JMS-based MDB can be defined in EJB 2.1 `ejb-jar.xml` file:

```
<enterprise-beans>
 <message-driven>
   <ejb-name>EJB_SEC_MDB_TOPIC_CMT</ejb-name>
   <ejb-class>com.sun.ts.tests.ejb.ee.sec.mdb.MsgBean</ejb-class>
   <messaging-type>javax.jms.MessageListener</messaging-type>
   <transaction-type>Container<transaction-type>
   <message-destination-type>javax.jms.Topic</message-destination-type>
   <message-destination-link>StockTopic</message-destination-link>
   <activation-config>
    <activation-config-property>
     <activation-config-property-name>acknowledgeMode
        </activation-config-property-name>
     <activation-config-property-value>Auto-acknowledge
        <activation-config-property-value>
    </activation-config-property>
    <activation-config-property>
     <activation-config-property-name>destinationType
        </activation-config-property-name>
     <activation-config-property-value>javax.jms.Topic
        <activation-config-property-value>
    </activation-config-property>
    <activation-config-property>
     <activation-config-property-name>subscriptionDurability
        </activation-config-property-name>
     <activation-config-property-value>DURABLE
        <activation-config-property-value>
    </activation-config-property>
   <activation-config>
 </message-driven>
 ⋮
</enterprise-beans>
```

The property names and values used in the <activation-config> to describe the messaging service vary depending on the type of message service used, but EJB 2.1 defines the following four fixed properties for JMS-based MDBs:

<activation-config-property-name>	Description	<activation-config-property-value>
acknowledgeMode	Allows an MDB container to notify the JMS provider that the MDB has received the message.	Auto-acknowledge (default) or Dups-ok-acknowledge
messageSelector	Allows an MDB to be selective about which messages it receives. It allows an MDB to set properties based on which messages will be received. These properties can be expressions or Boolean logic.	String selector
destinationType	Indicates the type of destination from which the MDB receives messages	javax.jms.Queue or javax.jms.Topic
subscriptionDurablity	Determines whether the JMS provider must store any messages that an MDB receives when the MDB container is disconnected from the provider.	NonDurable (default) or Durable

The Message Inflow contract in the JCA 1.5 specification is a contract between the messaging service provider and the application server for the delivery of messages to an MDB. As a part of this contract, the messaging provider implements a JavaBean component called ActivationSpec. The ActivationSpec defines the properties that are required by the messaging provider in order to deliver messages. The administrator may define default values for these properties, but when the application containing the MDB is deployed, these will be overridden by any `<activation-config-property>` elements defined in the MDB's deployment descriptor. A JMS provider that confirms to the Sun specification will therefore have the above mentioned properties on its ActivationSpec. You can leave a property out of the MDB's deployment descriptor and define it administratively instead. Conversely, there may be other, provider-specific properties that you previously had to define administratively that you can now include in the MDB's deployment descriptor.

Standard descriptor element `<message-destination-link>` is used to define a logical name for the message destination. It is used with a `<message-destination>` element to illustrate message flow within an application. For MDBs that specify a JMS provider message source, the JMS destination object is resolved from the target `<message-destination>` of `<message-destination-link>`, if present in the MDBs deployment descriptor.

In EJB 2.1 MDBs, standard descriptor element `<message-destination-ref>` can be used instead of `<resource-env-ref>` for definition of JMS destinations used within application logic of the MDB.

Changes to ejb-borland.xml

The Borland proprietary deployment descriptor has been modified to accommodate the new connector-based MDB. It now includes a new element, `<message-source>`. This element allows an application assembler to specify activation of the MDB through a JCA 1.5 resource adapter or in the case of a JMS messaging type MDB directly to a JMS provider. If you are using a JMS provider, you must use the `<jms-provider-ref>` element as follows:

```
<enterprise-beans>
   <message-driven>
      <ejb-name>EJB_SEC_MDB_TOPIC_CMT</ejb-name>
      <message-source>
         <jms-provider-ref>
            <message-driven-destination-name>
              Jms/MyTopic
            </message-driven-destination-name>
            <connection-factory-name>jms/myTCF</connection-factory-name>
            <pool>
```

```
                <max-size>120</max-size>
                <init-size>100</init-size>
                <wait-timeout>600</wait-timeout>
             </pool>
          </jms-provider-ref>
       </message-source>
           ⋮
    </message-driven>
 </enterprise-beans>
```

If you are using a Connector-based non-JMS messaging provider, use the following <message-source>:

```
 <enterprise-beans>
    <message-driven>
       <ejb-name>EJB_SEC_MDB_TOPIC_CMT</ejb-name>
       <message-source>
          <resource-adapter-ref>
             <instance-name>
               MyResourceApadter
             </instance-name>
          </resource-adapter-ref>
       </message-source>
    </message-driven>
 </enterprise-beans>
```

The resource adapter may provide the Java class name and the interface type of an optional set of JavaBean classes representing various administered objects. Administered objects are specific to a messaging style or message provider and can be referenced using <resource-env-ref> from the application logic of an MDB. For example, some messaging styles may need applications to use special administered objects for sending and synchronously receiving messages via connection objects using messaging-style specific APIs. Borland deployment descriptor element <resource-env-ref> is extended to override property values of an administered object.

For example:

```
⋮
  <message-driven>
     <message-source>
        <resource-adapter-ref>
           <instance-name>ResourceAdapter1</instance-name>
        </resource-adapter-ref>
     </message-source>
     ⋮
     <resource-env-ref>
        <resource-env-ref-name>mdbRequiredConnFactory</resource-env-ref-name>
        <admin-object>
           <property>
               <prop-name>serverUrl</prop-name>
               <prop-type>String<prop-type>
               <prop-value>localhost:7222</prop-value>
           </property>
        </admin-object>
     </resource-env-ref>
     ⋮
  </message-driven>
  ⋮
```

Clustering of MDBs

The clustering of MDBs differs from the clustering of other enterprise beans. With MDBs, producers put messages into a destination. The messages will reside in the destination until a consumer takes the messages off the destination (or, if the messages are non-durable, when the server hosting the destination crashes). This is a *pull* model since the message will just reside on the destination until a consumer asks for it. The containers contend to get the next available message on the destination. MDBs provide an ideal load-balancing paradigm, one that is smoother than other enterprise bean implementations for distributing a load . The server that is the least burdened can ask for and obtain the message. The tradeoff for this optimal load-balancing is that messaging has extra container overhead by virtue of the destination's position between the producer and the consumer.

There is not, however, the same concept of failover with a messaging service as exists in VisiBroker. If the consumer disappears, the queue fills up with messages. As soon as the consumer is brought back online, the messages resume being consumed. Of course, the JMS server itself should be fault-tolerant. The client should never notice any “failure” with the exception of response delays if such messages are expected. This kind of fault tolerance demands only a way of detecting failed consumers and activating them after failure.

That said, it is possible to deploy MDBs in more than one Partition with the Messaging Server pushing messages to only one, switching to the other in case of failure. Most JMS products allow queues to behave in load-balancing or fault-tolerant modes. That is, MDB replicas can register to the same queue and the messages are distributed to them using a load-balancing algorithm. Alternately, messages may all go to one consumer until it fails, at which point delivery shifts to another. The connection established to the JMS service provider from the MDB can also provide a load-balancing and/or fault-tolerant node. JMS service providers may provide fault-tolerance features. For specific information on clustering and fault-tolerance features, see “JMS provider pluggability.”

Keep in mind that only one MDB instance in a container that subscribes to a topic will consume any given message. This means that, for all parallel instances of an MDB to concurrently process messages, only one of the instances will actually receive any particular message. This frees up the other instances to process other messages that have been sent to the topic. Note that each container that binds to a particular topic will consume a message sent to that topic. The JMS subsystem will treat each message-driven bean in separate containers as a separate subscriber to that message! This means that if the same MDB is deployed to many containers in a cluster, then each

deployment of the bean will consume a message from the topic to which it subscribes. If this is not the behavior you desire, and you require exactly one consumption of a message, then you should consider deploying a queue rather than a topic.

Error Recovery

The following section deals with JMS server connection failures and setting connection rebind attempt properties. It also covers the redelivery of messages when an MDB fails to consume a message.

Rebinding EJB 2.0 and EJB 2.1 MDBs configured with a JMS provider message source

A connection failure usually occurs after you deploy your bean, causing a need for rebind attempt. You also receive an error if you are trying to deploy your bean and a connection to the JMS server was never established. Whether a failure occurs post deployment or no connection was found during deployment, the container will transparently attempt to rebind the JMS service provider connection when you set the rebind attempt properties. This ensures even greater fault-tolerance from an MDB instance.

The two bean-level properties that control the number of rebind attempts made and the time interval between attempts are:

– `ejb.mdb.rebindAttemptCount`: this is the number of times the EJB Container tries to re-establish a failed JMS connection for this MDB. The default value is `5` (five).

  To make the container attempt to rebind infinitely you need to explicitly specify `ejb.mdb.rebindAttemptCount=0`.

– `ejb.mdb.rebindAttemptInterval`: the time in seconds between successive retry attempts. The default value is `60`.

Redelivered messages for EJB 2.0 and EJB 2.1 MDBs configured with a JMS provider message source

Should the MDB fail to consume a message for any reason, the message will be re-delivered by the JMS service. The message will only be re-delivered five times. After five attempts, the message will be delivered to a dead queue (if one is configured). There is one bean-level property that controls the re-deliver attempt count:

– `ejb.mdb.maxRedeliverAttemptCount`: the max number of times a message will be re-delivered by the JMS service provider if an MDB is unable to consume it. The default value is `5`.

There are two bean-level properties for delivering a message to a dead queue:

– `ejb.mdb.unDeliverableQueueConnectionFactory`: looks up JNDI name for the connection factory to create connection to the JMS service.

– `ejb.mdb.unDeliverableQueue`: looks up the JNDI name of the queue.

  The XML example for `unDeliverableQueueConnectionFactory` and `unDeliverableQueue` is shown here:

```
<ejb-jar>
   <enterprise-beans>
       <message-driven>
           <ejb-name>MyMDB</ejb-name>
          <message-driven-destination-name>serial://jms/q
            </message-driven-destination-name>
           <connection-factory-name>serial://jms/xaqcf
            </connection-factory-name>
           <pool>
```

```
            <max-size>20</max-size>
            <init-size>0</init-size>
        </pool>
        <resource-ref>
            <res-ref-name>jms/QueueConnectionFactory</res-ref-name>
            <jndi-name>jms/xaqcf</jndi-name>
        </resource-ref>
        <property>
            <prop-name>ejb.mdb.maxRedeliverAttemptCount</prop-name>
            <prop-type>String</prop-type>
            <prop-value>3</prop-value>
        </property>
        <property>
            <prop-name>ejb.mdb.unDeliverableQueueConnectionFactory
               </prop-name>
            <prop-type>String</prop-type>
            <prop-value>serial://jms/qcf</prop-value>
        </property>
        <property>
            <prop-name>ejb.mdb.unDeliverableQueue</prop-name>
            <prop-type>String</prop-type>
            <prop-value>serial://jms/q2</prop-value>
        </property>
        <property>
            <prop-name>ejb-designer-id</prop-name>
            <prop-type>String</prop-type>
            <prop-value>MyMDB</prop-value>
        </property>
    </message-driven>
  </enterprise-beans>
  <assembly-descriptor />
</ejb-jar>
```

You can set these properties with the DDEditor. From the Console, navigate the tree on the left until you find the module containing your MDBs. Right-click the module and select DDEditor. When the DDEditor appears, select the bean node in the Navigation Pane to open the editor's panels for that bean. Select the “Properties” tab from the Content pane, and Add properties.

MDBs and Transactions

For information about using JMS within transactions, see “JMS and Transactions”. This section deals exclusively with using MDBs in transactions.

A common scenario for using the MDBs involves transactions requiring two-phase commit (2PC). Such an MDB has the `REQUIRED` transaction attribute. The MDB application method could be written to access and possibly update an external resource. Completion of the container managed transaction for the MDB method must include receipt of the message that triggered the method, and any work performed against the external resource from within the method. To achieve this, the transaction must be coordinated by a 2PC transaction service, such as the OTS engine. See “EJBs and 2PC transactions” for more details on optimal ways to use the OTS engine with MDBs.

20

Registering a Java Bean type object into JNDI

This chapter describes how you can register a Java Bean type object into JNDI and look it up from JNDI.

Adding a Java Bean Type Object into JNDI

To put a java bean type object into JNDI:

1. Add a new entry called `jndi-object` into the jndi-definitions.xml file.
2. Set the `jndi-name` as the name of the object. This is used to bind the object into JNDI and then for the client to lookup the object.
3. Set the `class-name` as the class of the object. You must place the class under partition's classpath since it is deployed as a library.
4. Set the list of properties. Each property entry consists of a property name, property type, and property value.

Example

```
<jndi-object>
   <jndi-name>TestObject</jndi-name>
   <class-name>examples.j2ee.jndi.Foo</class-name>
   <property>
       <prop-name>street</prop-name>
       <prop-type>String</prop-type>
       <prop-value>Park west</prop-value>
   </property>
   <property>
       <prop-name>postal</prop-name>
       <prop-type>String</prop-type>
       <prop-value>1262445</prop-value>
   </property>
</jndi-object>
```

21

Connecting to Resources with Borland AppServer: using the Definitions Archive (DAR)

J2EE specifies a uniform mechanism for establishing connections to resources using Java standard interfaces. A resource related object containing resource manager location details and connection attributes is bound under a JNDI service provider, and can be retrieved by your application as a resource connection factory in a JNDI lookup. Sample resource connection factories include JDBC datasources and JMS connection factories. Once a resource connection factory is obtained from JNDI, a connection to the desired resource manager can then be established. A connection to a relational database is obtained through a JDBC datasource, a connection to a message broker is obtained through a JMS connection factory, and general Enterprise Information Systems (EIS) connections are obtained through JCA resource adapters.

Use the Borland Management Console and Borland Deployment Descriptor Editor (DDEditor) to create, edit, and deploy resource connection factories and other resource related JNDI objects, such as JMS destinations. An XML descriptor file (`jndi-definitions.xml`), generally called the JNDI Definitions module, captures the properties representing resource related objects. This file is packaged in a Data ARchive (DAR) module.

In the Borland AppServer (AppServer), a partition hosted Naming Service represents the default JNDI service provider—its an implementation of a CosNaming service provider. Resource related objects are bound in the Naming Service of an AppServer partition through deployment of DAR or RAR module using standard AppServer deployment procedures. At that time, only properties required to create an instance of a resource connection factory, or JMS destination, are stored in the JNDI bound object. During JNDI lookup for a resource related object, an instance of the desired resource object is created using the stored property values from the object retrieved. The newly created instance is then passed back to the caller of JNDI `lookup()` method. In this way, DARs can be successfully deployed to an AppServer partition without having to load classes for vendor specific resource objects. Class libraries for resource vendors are only required by application processes that actually perform JNDI lookup of resource related objects.

Note

In prior versions of AppServer, a file-system service provider called the Serial Provider was the default JNDI service provider for deployment of DARs and JNDI Definitions modules. Resource related objects bound to this provider involved creation of the

resource objects during deployment, and hence required vendor class libraries to be deployed in advance. In addition, JNDI names for resource related objects required a serial URL prefix, that is “serial://”. With the Naming Service being the default service provider, this prefix is no longer required in JNDI name specification. Deployment of existing DARs/JNDI Definitions modules that have JNDI names with this prefix are now automatically bound to the Naming Service.

A resource related object is obtained in J2EE through a resource reference. You can reference resource connection factories or JMS destinations from EJBs, servlets and other J2EE application components using resource-reference elements in the component's deployment descriptors. See “Using JDBC” for more specific information on how to define JDBC datasource resource references, and “Using JMS” for resource reference definition examples of JMS connection factories and destinations.

Each AppServer partition has a predeployed DAR module named `default-resources.dar` containing example definitions for JDBC datasources, JMS connection factories and JMS destinations. This module can be examined, updated and redeployed in the following steps:

1. Navigate to **default-resources.dar** under the Deployed Modules node of a partition in the left pane of the Borland Management console.
2. Right-click on **default-resources.dar** and select **Edit deployment descriptor** from the context menu. The Borland Deployment Descriptor Editor window will open. The available datasources and connection factories should be visible in the left pane.
3. Right-click on the root node in the navigation pane of the Borland Deployment Descriptor Editor and select the appropriate option to create a New object you want to add.

When a J2EE component attempts a JNDI lookup for a resource reference, vendor classes associated with the resource object must be available in the runtime environment. If the J2EE component is deployed to an AppServer partition, the vendor class libraries must be deployed to the partition as a library archive. Exceptions to this rule include JNDI lookups for resource objects whose dependent class libraries are bundled with AppServer. For example, a JDataStore datasource or any JMS resource object of the JMS message server installed with AppServer.

JNDI Definitions Module

Resource related objects are bound under the Naming Service through deployment of a DAR file containing the JNDI Definitions Module. A DAR has a special .dar file extension. It must be deployed to an AppServer partition either individually or packaged with other J2EE modules in an EAR file.

Note

A DAR is *not* a part of the J2EE specification. It is a Borland-specific implementation designed to simplify deployment and management of resource connection factories and JMS destinations. You do *not* package connection factory or JMS destination vendor classes in this archive type. Those classes must be deployed as a library to individual Partitions.

The only contents of the DAR that you must provide is an XML descriptor file called `jndi-definitions.xml`. It contains definitions for resource related objects, each with a JNDI name identifying its location in JNDI. Like other descriptors, this is placed within the `META-INF` directory of the DAR. The contents of the DAR hence is as follows:

```
META-INF/jndi-definitions.xml
```

You deploy the DAR containing the descriptor file just as you would any other J2EE module using either the Console or command-line utilities, or as part of an EAR. You can deploy any number of distinctly named DARs in the same Partition or to an AppServer cluster. Should two or more deployed DARs have resource object definitions with identical JNDI names, the last deployed module overwrites any existing object binding on the same node.

Once deployed, resource objects defined in the DAR can be examined in the Naming Service namespace using the JNDI Browser.

Migrating to DARs from previous versions of Borland AppServer

Previous product versions, including IAS 4.1 and BAS 4.5, did not have a DAR module to contain the `jndi-definitions.xml` descriptor. If you have a customized `jndi-definitions.xml` file that needs to be transferred to AppServer, follow these migration steps:

1. If you want the entire contents of the default resources overridden, make a temporary directory called `META-INF` and place your existing `jndi-definitions.xml` file within it.
2. Open a command window and use the following `jar` command:

```
prompt>jar uvMf default-resources.dar META-INF/jndi-definitions.xml
```

3. Now deploy this module following the usual procedures.

If you have performed only a few customizations on your old `jndi-definitions.xml` file then it may be easier to simply move the appropriate XML stanzas from the old file into the one contained within the pre-deployed DAR.

Creating and Deploying a DAR

The DDEditor walks you through creating a new JNDI Definitions Module. Open the DDEditor and select “File|New...” The Object Gallery appears.

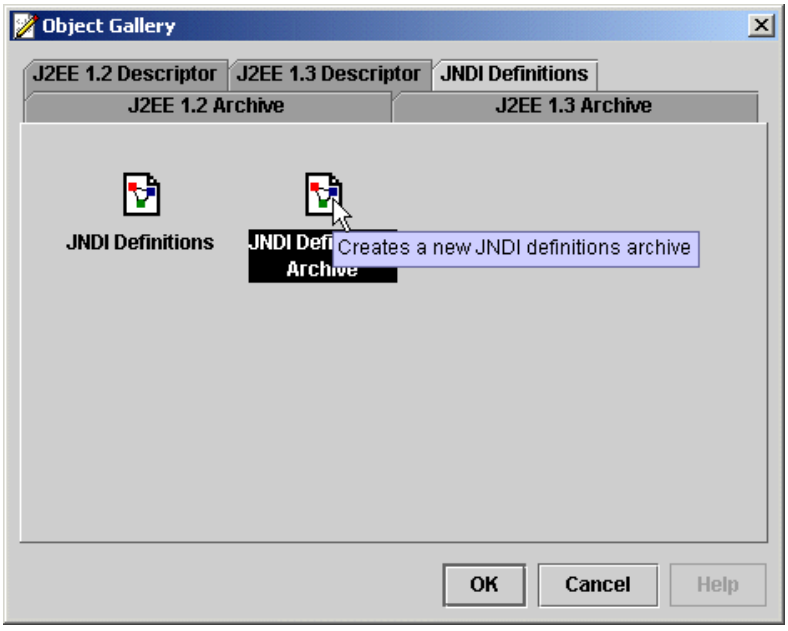

Select the JNDI Definitions tab and select JNDI Definitions Archive to create a new DAR. Click OK. You may now add JDBC datasources or add JMS resources. Or, you may do this later. When you are finished save the module by choosing “File|Save...” .

After saving the archive, use the J2EE Deployment Wizard to deploy the module. The wizard reads resource definitions from the DAR and binds them to the Naming Service of the target partition(s). To start the wizard, open the Console and select “Wizards|Deployment Wizard.” Follow the on-screen instructions.

Disabling and Enabling a Deployed DAR

Once a DAR module is deployed to a partition, it is enabled. This means that its resource object definitions are bound to the Naming Service of the partition for as long as the Naming Service remains active. Enabling a DAR module rebinds its resource object definitions to the Naming Service and in the process overwrites any pre-existing content for JNDI names specified. Disabling a DAR module has no immediate impact on the contents of an active Naming Service. Upon a subsequent restart of the partition, disabled DARs are not deployed to the partition, that is, resource object definitions are not bound to the Naming Service. By default, the Naming Service stores object bindings in memory. Each time the host partition is restarted, resource object bindings from previously deployed DARs are destroyed. If the Naming Service is configured with a JDBC backing store, resource object bindings for all DARs are retained, even those that were once deployed but now marked disabled. Use the JNDI Browser to locate and permanently remove these bindings.

To manipulate a deployed DAR module, use the Console to select it from the set of Deployed Modules for a partition, right-click and choose the appropriate action.

Packaging DAR Modules in an Application EAR

Sometimes it is useful to package all archives that make up a complete application into a single deployable unit. The common scenario is that you have some EJBs in an EJB Archive, some servlets and JSPs in a Web Archive and they both depend on some datasources or JMS administrative objects defined in a DAR. Using the Archive Tools in the Console it is easy to package a set of individual archives into a single EAR Module.

Note

Because DARs are not a part of the J2EE specification, you must include at least one other valid J2EE module along with your DAR within the EAR. An EAR containing only a DAR file is not a valid J2EE archive.

22

Using JDBC

Resource related objects such as JDBC datasources are obtained in a portable J2EE mandated way through JNDI. A JDBC datasource is resolved by performing a JNDI lookup of a J2EE Resource Reference defined in the deployment descriptors of an application component. Resource Reference definitions involve both standard J2EE and Borland's proprietary deployment descriptors. In the standard deployment descriptor, a Resource Reference specifies a logical name relative to the application's JNDI environment naming context, `java:comp/env/`. Borland's deployment descriptor complements the standard descriptor by associating the Resource Reference logical name with the actual JNDI location of the JDBC resource definition. For example, in an EJB Jar component, the standard J2EE deployment descriptor, `ejb-jar.xml`, specifies Resource References for an EJB using a <resource-ref> element for a JDBC datasource. In Borland AppServer (AppServer), a JNDI lookup of a Resource Reference involves retrieval of the JDBC datasource definition from which the desired datasource object is created and returned to caller of lookup. The property values present in the JDBC datasource definition determine the type and characteristics of datasource object created.

Before a Resource Reference lookup can be attempted, the required datasource definition must first be bound to its physical JNDI location. In AppServer, JDBC datasource definitions are bound into a JNDI service provider during deployment of a Definitions ARchive (DAR) module. By default, these objects are bound to the partitions Naming Service, the JNDI CosNaming service provider in AppServer. This chapter describes how to define JDBC datasource definitions in a DAR module and how to obtain a reference to a JDBC datasource from your J2EE application.

Configuring JDBC Datasources

Using the Console, navigate to the “Deployed Modules” list in the Partition whose datasources you need to configure. By default, every Partition has a predeployed JNDI Definitions Module (DAR) called `default-resources.dar`. Right-click on the module and choose “Edit deployment descriptor” from the context menu. The Deployment Descriptor Editor (DDEditor) appears.

In the Navigation Pane of the DDEditor is a list of datasources preconfigured in the product. If needed they can be individually edited to suit the user's requirements.

To create a new JDBC datasource, Right-click on the node at the top of the tree in the navigation pane and select “New JDBC Datasource...” from the Context Menu.

A dialog box prompts for a JNDI name for the newly-created datasource. Once given, a representation of this datasource appears in the tree in the Navigation Pane. Click its representation to open its configuration panel.

The DDEditor has knowledge of some common JDBC drivers and can autofill the class names and essential properties for the appropriate JDBC datasource. If the JDBC datasource type you want appears in the Datasource Type list then choose it, otherwise select “Other(JDBC2)”.

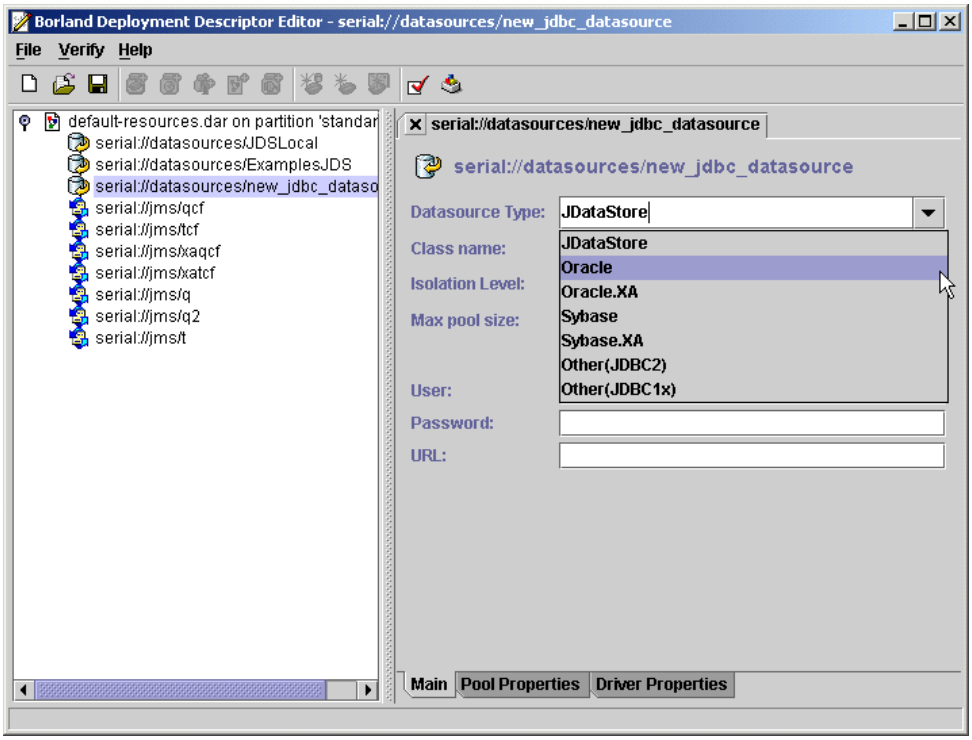

The Main tab in the content pane captures the essential properties needed to define the chosen database. If the database is known to the DDEditor, it will automatically complete these properties.

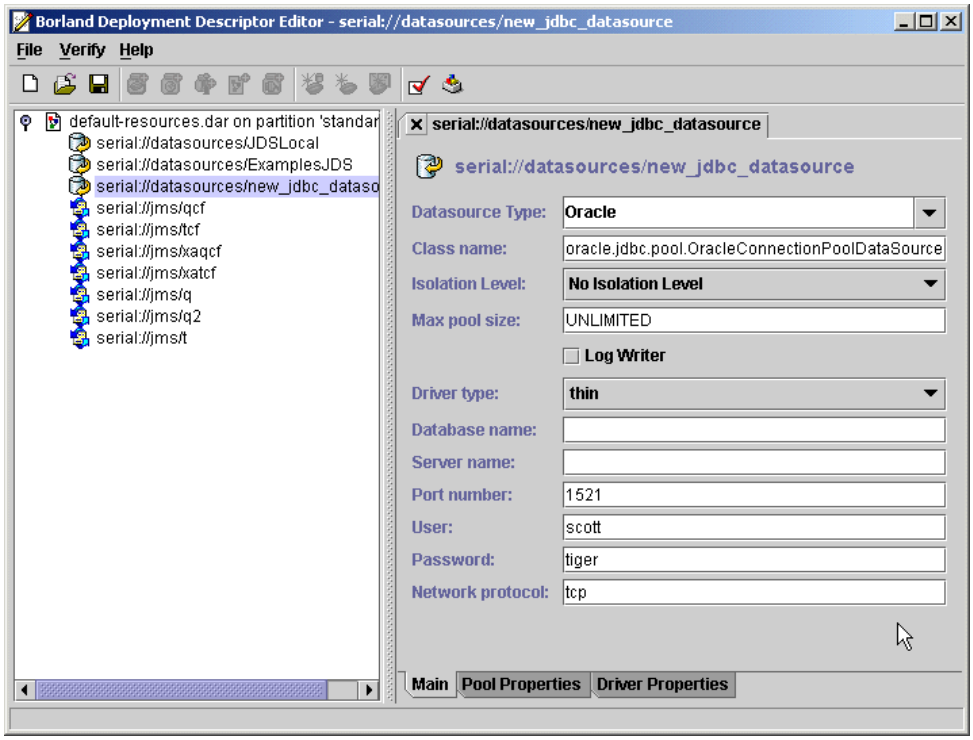

The Driver Properties and Pool Properties tabs capture some of the information from the Main tab, but also allow you to set any less common properties that do not appear in the Main tab.

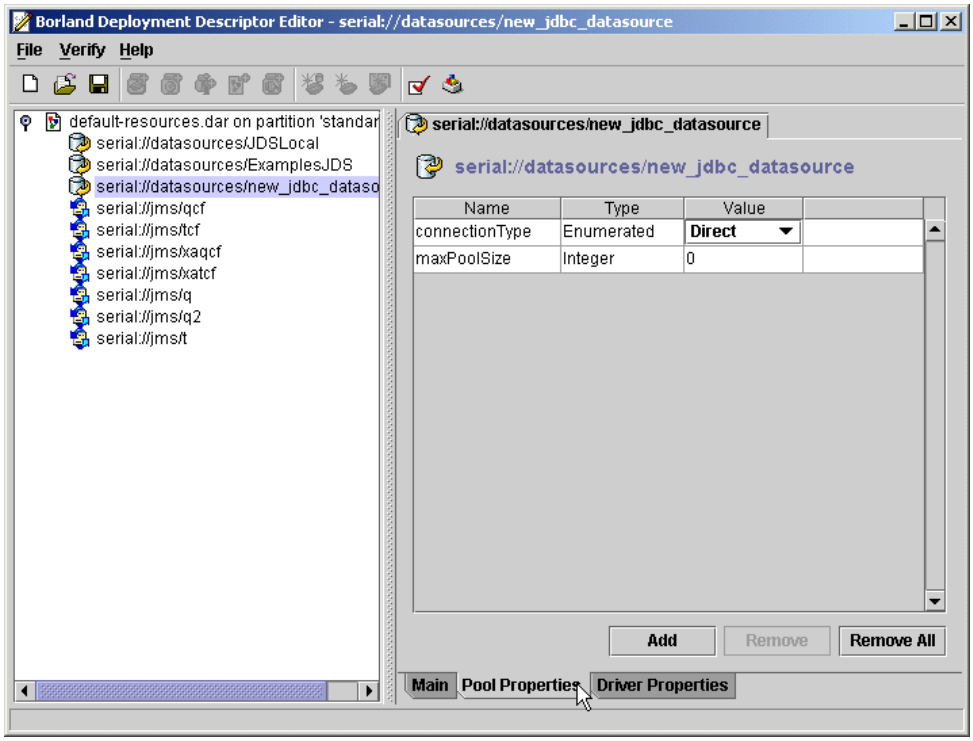

To add pool properties, click the Add button and select the property you want to add from the drop-down list under “Name.” Pool Properties are described in “Defining the

Connection Pool Properties for a JDBC Datasource". The same procedure is used for adding Driver Properties.

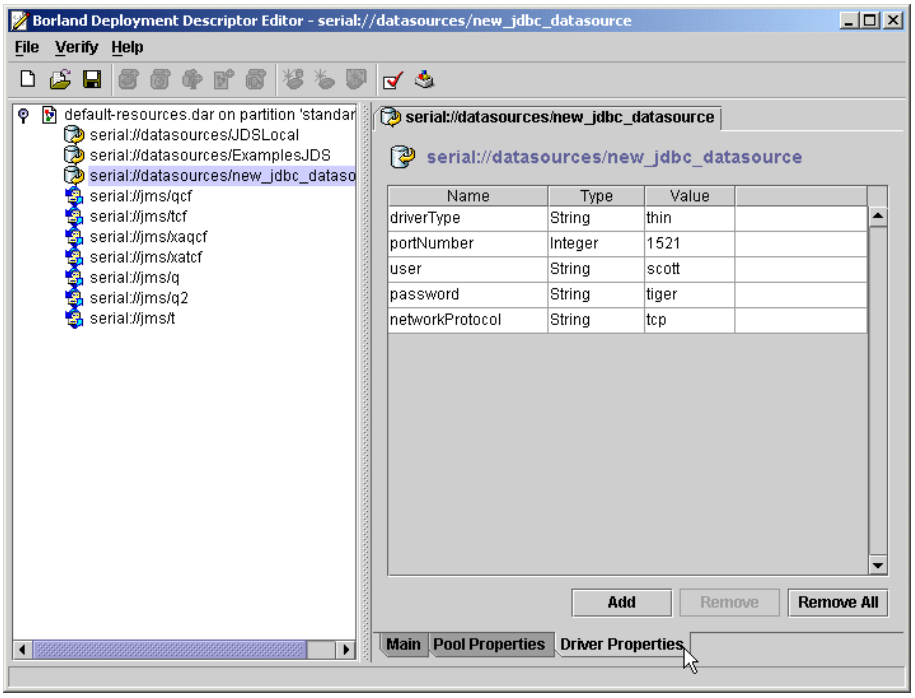

Consult your database documentation for any properties you need to define.

Once you're finished, save the module and dismiss the final modal window. The JNDI Definitions Module is automatically re-deployed to the Partition.

Deploying Driver Libraries

If a deployed application component contains a JNDI lookup of third party JDBC datasource, vendor libraries are required and must be deployed to the target Partition as a library archive before the lookup is performed. Note that these steps are not necessary if you are using the native all-Java database, JDataStore. When trying to connect to another database like Oracle or Sybase the respective JDBC driver must first be deployed to the target Partition. To deploy a library to multiple partitions do the following:

1. From the Console, select Wizard->Deployment Wizard. This will open the Deployment Wizard.
2. Click on the Add button and navigate to the library file in the resulting window and click OK. The library name should now appear in the selection box in the Deployment Wizard.
3. Click on the Next button. The names of the partitions will appear in the Deployment Wizard window.
4. Select the partitions to which you want to deploy the library and click on the Finish button. The deployment status will appear in a separate window.
5. Click the Close button to close this window. You can verify that the libraries are deployed correctly by checking the partition's Deployed Modules node in the Management Console navigation pane. The name of the library should appear under the Deployed Modules node.
6. Stop and restart the partition for the deployment to take effect.

To deploy the library to a single partition:

1. Right-click on the partition's name in the navigation pane in the Management Console and select Deploy Modules from the context menu. The Deployment Wizard will open.
2. Click on the Add button and navigate to the library file in the resulting window and click OK. The library name should now appear in the selection box in the Deployment Wizard.
3. Click on the Next button. The partition name will appear in the Deployment Wizard window.
4. Select the partition to which you want to deploy the library and click on the Finish button. The deployment status will appear in a separate window.
5. Click the Close button to close this window. You can verify that the libraries are deployed correctly by checking the partition's Deployed Modules node in the Management Console navigation pane. The name of the library should appear under the Deployed Modules node.
6. Stop and restart the partition for the deployment to take effect.

Defining the Connection Pool Properties for a JDBC Datasource

At runtime each JDBC datasource corresponds to an instance of a connection pool. Connection pools provide for the reuse of connections and optimization of database connectivity. Some datasources may require different treatment as connection pools than others. A number of configuration options exist for these connection pools. Control of pool sizes, statement execution behavior, and transaction parameters are specified as properties in the `<visitransact-datasource>` element in the DAR descriptor file. You specify properties using the `<property>` element, which includes the `<prop-name>`, `<prop-type>`, and `<prop-value>` elements. The complete list of properties, allowed values, defaults, and descriptions appear in the following table:

Name	Allowed Values	Description	Default Value
`FinalizeNoTxBusy Connections`	This property differs from the rest of the Connection Pool properties because it must be set in the BAS Partition by editing the `partition_server.config` to add the following line: `vmpram -DFinalizeNoTxBusyConnections` When set, it will impact all BAS JDBC Connection Pools active in that partition.	When a JDBC Connection enters a state of NoTxBusy, the application must close the connection, otherwise the JDBC Connection Pool will hold onto a reference of it indefinitely and the underlying database connection is never released. This property, when set, configures BES connection pools to hold a weak reference for JDBC connections, allowing out of scope NoTxBusy connections, not closed by applications, to be freed when JVM garbage collection occurs.	`connectionType`
`Enumerated`: ■ `Direct` ■ `XA`	Indicates type transaction association of all connections retrieved from the connection pool, whether “Direct” or “XA”	Not Applicable. Property specification is mandatory	`optimizeXA`

Name	Allowed Values	Description	Default Value
Boolean	By default, XAResource API calls are kept to a minimum for optimization of AppServer JDBC connection pool performance. Setting optimizeXA to a value of false disables this optimization. Under certain conditions, datasources must have optimizeXA property set to false. For instance, conflicts can arise between default XAResource optimizations in AppServer JDBC connection pool and certain vendor resource managers, such as Oracle, during two-phase commit protocol. When setting optimizeXA to false, applications using a JDBC connection in the scope of a distributed transaction must issue a connection close() before completion of the transaction, otherwise unexpected transaction completion conditions will arise.	True	maxPoolSize
Integer	Specifies the maximum number of database connections that can be obtained from this datasource connection pool.	0 (zero), implying unbounded size	waitTimeout
Integer	The number of seconds to wait for a free connection when maxPoolSize connections are already opened. When using the maxPoolSize property and the pool is at its max and can't serve any more connections, the threads looking for JDBC connections end up waiting for the connection(s) to become available for a long time if the wait time is unbounded (set to 0 seconds). You can set the waitTimeout period to suit your needs.	30	busyTimeout
Integer	The number of seconds to wait before a busy connection is released	600 (ten minutes)	idleTimeout
Integer	A pooled connection remaining in an idle state for a period of time longer than this timeout value should be closed. All idle connections are checked for idleTimeout expiration every 60 seconds. The value of the idleTimeout is given in seconds.	600 (ten minutes)	queryTimeout
Integer	Specifies in seconds the time limit for executing database queries by this datasource.	0 (zero), implying indefinite period	dialect
Enumerated: ■ oracle ■ sybase ■ interbase ■ jdatastore	Specifies the database vendor as a hint for automatic table creation performed during Container Managed Persistence	This property is optional. There is no default value.	isolationLevel
Enumerated: ■ TRANSACTION_NONE ■ TRANSACTION_READ_COMMITTED ■ TRANSACTION_READ_UNCOMMITTED ■ TRANSACTION_REPEATABLE_READ ■ TRANSACTION_SERIALIZEABLE	Indicates database isolation level associated with all connections opened by this datasource's connection pool. See the J2EE 1.3 specification for details on these isolation levels.	Defaults to whatever level is provided by the JDBC driver vendor.	reuseStatements

Name	Allowed Values	Description	Default Value
`Boolean`	Optimization directive requesting prepared SQL statements to be cached for reuse. Applies to all connections obtained from the connection pool.	`True`	`initSQL`
`String`	Specifies a list of ";" separated SQL statements to be executed each time a connection is obtained for a fresh transaction. The SQL is performed before any application work is performed on the connection.	This property is optional. There is no default value.	`refreshFrequency`
`Integer`	Using `dbPingSQL`, this property specifies a timeout, in seconds, for each connection in an idle state. Once the timeout expires, the connection is examined to determine if it is still a valid connection. All idle connections are checked for `refreshFrequency` at sixty second intervals.	`300` (five minutes)	`dbPingSQL`
`String`	Specifies an SQL statement used to validate open connections present in the connection pool and to refresh connections during a `refreshFrequency` timeout.	Not defined. When no SQL is specified, the container uses `java.sql.ConnectionisClosed()` method to validate the connection.	`resSharingScope`
`Enumerated`: ■ `Shareable` ■ `Unshareable`	Indicates whether connection statements and result sets are cached for reuse. If set to `Shareable`, connection statements and results sets are cached, thereby optimizing throughput of connections. If set to `Unshareable`, connections are closed once the application closes the connection.	`Shareable`	`maxPreparedStatement CacheSize`

Name	Allowed Values	Description	Default Value
Integer	Each connection within an AppServer JDBC pool caches java.sql.PreparedStatement objects for reuse. Each PreparedStatement cache is organized by SQL literal strings representing unique SQL statement requests. This property limits the number of PreparedStatements cached per pooled JDBC connection. It specifies the maximum size of the cache. If a cache reaches the limit, any subsequent javax.sql.Connection.prepareStatement() calls result in non-cached instances of PreparedStatement objects being created and returned to the caller. The lifecycle of the cache is the same as the JDBC connection lifecycle. For example, if an idle connection times out, both the connection and its PreparedStatement cache are discarded. Unresolved parameterized SQL statements are cached, for example, the statement SELECT NAME FROM CUSTOMER WHERE AGE=20 is cached as SELECT NAME FROM CUSTOMER WHERE :age='?'. Note that this property is only effective when the reuseStatements property of the datasource is set to true (default). The default value is 40, which is usually sufficient for applications.	40	maxPreparedStatementsPerQuery
Integer	Under certain conditions such as high concurrency or when CMP 2.0 entity beans are processed, more than one PreparedStatement can be processed concurrently for the same SQL query on the same pooled connection. For example, a SQL query SELECT name FROM table1 WHERE id=? can return distinct result sets when different values are used for ?. Although the PreparedStatement cache has a single entry for each SQL query, two or more PreparedStatements can exist in the cache for the query. This property specifies the maximum number of cached PreparedStatements for a single query. If the limit is exceeded for a particular query, subsequent javax.sql.Connection.prepareStatement() calls result in non-cached instances of PreparedStatement objects created and returned to the caller. Like maxPreparedStatementCacheSize, this property is only effective when the reuseStatements property of the datasource is set to true (default).	20	

Getting debug output

A number of system properties can be set to log activity at datasource, connection pool, connection and statement levels during application processing. It is not necessary to configure these properties during normal application runtime execution but should a situation arise where details of JDBC flow of control is needed these options are useful. Runtime output generated with these properties set can be provided to Borland Technical Support to help resolve issues involving JDBC datasource and connections. Setting these properties for a partition results in log message generation during JDBC activity. Note that additional `log4j` configuration is required to ensure that the messages are actually written to the partition log. Locate the partition's `log4j` configuration file, called `logConfiguration.xml`, and add the following <logger> element:

```
⋮
<log4j:configuration>
 ⋮
    <logger name="com.inprise.visitransact.jdbc2" additivity="true">
       <level value="DEBUG" />
    </logger>
    ⋮
</log4j:configuration>
```

Note

BAS logging is based on the Log4j infrastructure. Some user applications which use Log4j may cause the Partition to hang. User applications should use the per-partition log4j Configuration file, rather than deploying a Configuration file in the archive. By default, this file is located under the Partition's Managed Object footprint at: `<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/mos/<partition_name>/adm/properties/logConfiguration.xml`. Alternately, you can uncomment the following line in the `<install_dir>/bin/partition.config` file:

```
vmprop borland.enterprise.server.partition.disableSystemRedirect=true
```

System Property Name	Type	Description	Default
`DataSourceDebug`	`Boolean`	Reports activity at datasource level for all datasources	`False`
`ConnectionPoolDebug`	`Boolean`	Reports activity at connection pool level for all datasources	`False`
`ConnetionPoolStateDebug`	`Boolean`	Reports transitions of connections in connection pool	`False`
`JDBCProxyDebug`	`Boolean`	Reports activity at connection level for all connections	`False`
`PreparedStatementCacheDebug`	`Boolean`	Reports activity at prepared statement level for all statements	`False`

Descriptions of AppServer's Pooled Connection States

When the EJB container's statistic gathering option is enabled, the Partition event log contains useful statistics about the JDBC connections pool. The log lists the number of connections in the various lifecycle states of a pooled JDBC2 connection. Following is a description of each state:

- **Free:** a cached/pooled connection that is available for use by an application
- **TxBusy:** a cached connection that is in use in a transaction
- **NoTxBusy:** a cached connection that is in use by an application with no transaction context
- **Committed:** a connection that was associated with a transaction received a `commit()` call from the transaction service
- **RolledBack:** a connection that was associated with a transaction received a `rollback()` call from the transaction service
- **Prepared:** a connection that was associated with a transaction received a `prepare()` call from the transaction service
- **Forgot:** a connection that was associated with a transaction received a `forget()` call from the transaction service
- **TxBusyXaStart:** a pooled connection that is associated with a transaction branch.
- **TxBusyXaEnd:** a pooled connection that has finished its association with a transaction branch
- **BusyTimedOut:** a cached connection that was removed from the pool after it stayed with the transaction longer than the `busyTimeout` pool property
- **IdleTimedOut:** a connection that was removed from the pool due to being idle for longer than the pool's `idleTimeout` property
- **JdbcHalfCompleted:** a transitionary state where the connection is participating in a background housekeeping activity related to pool management (being refreshed, for example) and therefore, unavailable until the activity completes
- **Closed:** the underlying JDBC connection was closed
- **Discarded:** A cached connection got discarded (due to timeout errors, for example)
- **JdbcFinalized:** an unreferenced connection was garbage collected

Support for older JDBC 1.x drivers

JDBC 1x drivers do not provide a datasource object. Under the J2EE specification, however, database connections are always fetched using the `javax.sql.DataSource` interface. To allow users to still use JDBC 1x drivers, AppServer provides an implementation of a JDBC 1x datasource to allow writing portable J2EE code. This implementation is a facade provided on top of the DriverManager connection mechanism of the JDBC 1x specification.

If you want to define a datasource on top of such a driver then in the DD Editor choose the Datasource Type field as “Other(JDBC1x)”. Then in the Main panel you can input the Driver Manager classname and connection URL for your particular database and driver.

The class name `com.inprise.visitransact.jdbc1w2.InpriseConnectionPoolDataSource` is not the `DriverManager` class of the JDBC driver; it is a wrapper class. The vendor's class should be specified in the Driver Class Name text box of the editor panel.

Advanced Topics for Defining JDBC Datasources

Whether you choose to use the server's graphical tools or not, defining a datasource means providing some information to the container in XML format. Let's look at what it takes to define a JDBC datasource and bind it to JNDI. Let's start by examining the DTD of the `jndi-definitions.xml` file. The elements in bold are the main elements specific to JDBC datasources.

```
<!ELEMENT jndi-definitions (visitransact-datasource*, driver-datasource*, jndi-
object*)>
<!ELEMENT visitransact-datasource (jndi-name, driver-datasource-jndiname,
property*)>
<!ELEMENT driver-datasource (jndi-name, datasource-class-name,
log-writer?, property* )>
<!ELEMENT jndi-object (jndi-name, class-name, property* )>
<!ELEMENT property (prop-name, prop-type, prop-value)>
  <!ELEMENT prop-name (#PCDATA)>
  <!ELEMENT prop-type (#PCDATA)>
  <!ELEMENT prop-value (#PCDATA)>
  <!ELEMENT jndi-name (#PCDATA)>
  <!ELEMENT driver-datasource-jndiname (#PCDATA)>
  <!ELEMENT datasource-class-name (#PCDATA)>
  <!ELEMENT log-writer (#PCDATA)>
  <!ELEMENT class-name (#PCDATA)>
```

Defining a JDBC datasource involves two XML elements. The first is the `<visitransact-datasource>` element. This is where you define the datasource your application code will look up. You include the following information:

- **jndi-name:** this is the name of the datasource as it will be referenced by JNDI. It is also the name found in the resource references of your enterprise beans.
- **driver-datasource-jndiname:** this is the JNDI name of the driver class supplied by the database or JMS vendor that you deployed as a library to your Partitions. It is also the name that will be referenced by the `<driver-datasource>` element discussed next.
- **properties:** these are the properties for the datasource's role in its connection pool. We'll discuss these properties in a little more detail in the Defining the Connection Pool Properties for a JDBC Datasource section.

So, let's look at an example of this portion of the datasource definition in the XML. In the following example, we'll look at an example using Oracle:

```
<jndi-definitions>
        <visitransact-datasource>
            <jndi-name>datasources/Oracle</jndi-name>
            <driver-datasource-jndiname>datasources/OracleDriver
              </driver-datasource-jndiname>
            <property>
               <prop-name>connectionType</prop-name>
               <prop-type>Enumerated</prop-type>
               <prop-value>Direct</prop-value>
            </property>
            ⋮
            <!-- other properties as needed -->
            ⋮
        </visitransact-datasource>
        ⋮
    </jndi-definitions>
```

We're not done. Now we must perform the other half of the datasource definition by providing information on the driver. We do this in the `<driver-datasource>` element, which includes the following information:

– **jndi-name:** This is the JNDI name of the driver class, and its value *must* be identical to the `<driver-datasource-jndiname>` value from the `<visitransact-datasource>` element.

– **datasource-class-name:** Here is where you provide the name of the connection factory class supplied from the resource vendor. It must be the same class you deployed to the Partition as a library.

– **log-writer:** This is a boolean element that activates verbose modes for some vendor connection factory classes. Consult your resource's documentation for the use of this property.

– **properties:** These are properties specific to the JDBC resource, such as usernames, passwords, and so forth. These properties are passed to the driver class for processing. Consult your JDBC resource documentation for property information. Specifying the properties in XML is shown below.

Armed with this information, let's complete our datasource definition for the Oracle datasource we started above. In order to be thorough, let's first reproduce the XML we started above:

```
<jndi-definitions>
 <visitransact-datasource>
       <jndi-name>datasources/Oracle</jndi-name>
       <driver-datasource-jndiname>datasources/OracleDriver
          </driver-datasource-jndiname>
      <log-writer>False</log-writer>
      <property>
           <prop-name>connectionType</prop-name>
           <prop-type>Enumerated</prop-type>
           <prop-value>Direct</prop-value>
      </property>
 </visitransact-datasource>
    ⋮
```

Note the driver datasource JNDI name in bold. Now we'll add the following:

```
   <driver-datasource>
  <jndi-name>datasources/OracleDriver</jndi-name>
  <datasource-class-name>oracle.jdbc.pool.OracleConnectionPoolDataSource</
datasource-class-name>
  <property>
    <prop-name>user</prop-name>
    <prop-type>String</prop-type>
    <prop-value>MisterKittles</prop-value>
  </property>
  <property>
    <prop-name>password</prop-name>
    <prop-type>String</prop-type>
    <prop-value>Mittens</prop-value>
  </property>
   ⋮
   // other properties as needed
   ⋮
    </driver-datasource>
</jndi-definitions>
```

Now the JDBC datasource is fully defined. Once you've packaged the XML file as a DAR, you can deploy it to a Partition. Doing so registers the datasource with the Naming Service and makes it available for lookup.

Connecting to JDBC Resources from J2EE Application Components

In Borland proprietary deployment descriptors, such as `ejb-borland.xml` for EJB components, the `<resource-ref>` element is used to map a datasource logical name to actual JNDI location of a JDBC datasource definition. Mapping of the logical name to its location occurs when a JNDI lookup is performed for a desired datasource in the application component. You use the element within your individual component definitions. For example, a `<resource-ref>` for an entity bean must be found within the `<entity>` tags. Let's look at the DTD representation of the `<resource-ref>` element of Borland deployment descriptors:

```
<!ELEMENT resource-ref (res-ref-name, jndi-name, cmp-resource?)>
```

In this element you specify the following:

- **res-ref-name:** this is the logical name for the resource, the same logical name you use in the `<resource-ref>` element of the standard `ejb-jar.xml` descriptor file. This is the name your application components use to look up the datasource.
- **jndi-name:** this is the JNDI name of the datasource that will be bound to its logical name. It must match the value of the corresponding `<jndi-name>` element of the `<visitransact-datasource>` element deployed with the DAR.
- **cmp-resource:** this is an optional boolean element that is relevant to entity beans only. If set to `True`, the container's CMP engine will monitor this datasource.

Let's look at an example entity bean that uses the Oracle datasource we defined above:

```
<entity>
 <ejb-name>entity_bean</ejb-name>
 ⋮
    <resource-ref>
  <res-ref-name>jdbc/MyDataSource</res-ref-name>
  <jndi-name>datasources/Oracle</jndi-name>
  <cmp-resource>True</cmp-resource>
    </resource-ref>
 ⋮
</entity>
```

As you can see, we used the identical JNDI name from the `<visitransact-datasource>` element from the datasource definition. Now let's see how we obtain a datasource object reference. To do so, the application performs a lookup of the `<res-ref-name>` value of the deployed components and the object references are retrieved from the remote CosNaming provider. For example:

```
javax.sql.DataSource ds1;

try {
 javax.naming.Context ctx = (javax.naming.Context)
    new javax.naming.InitialContext();
 ds1 = (DataSource)ctx.lookup("java:comp/env/jdbc/MyDataSource");
}
catch (javax.naming.NamingException exp) {
 exp.printStackTrace();
}
```

A database `java.sql.Connection` can now be obtained from `ds1`.

23

Using JMS

Resource related objects such as JMS connection factories and JMS Queue/Topic destinations are obtained in a portable J2EE mandated way through JNDI. A JMS resource object is resolved by performing a JNDI lookup of a J2EE Resource Reference defined in the deployment descriptors of an application component. Resource Reference definitions involve both standard J2EE and Borland's proprietary deployment descriptors. In the standard deployment descriptor, a Resource Reference specifies a logical name relative to the application's JNDI environment naming context, `java:comp/env/`. Borland's deployment descriptor complements the standard descriptor by associating the Resource Reference logical name with the actual JNDI location of the JMS resource definition. For example, in an EJB Jar component, the standard J2EE deployment descriptor, `ejb-jar.xml`, specifies Resource References for an EJB using a `<resource-ref>` element for a JMS connection factory and `<resource-env-ref>` elements for JMS Topics and Queues. In Borland AppServer(AppServer), a JNDI lookup of a Resource Reference involves retrieval of the JMS resource definition from which the desired JMS object is created and returned to caller of lookup. The property values present in the JMS resource definition determine the type and characteristics of resource object created.

Before a Resource Reference look up can be attempted, the required resource definition must first be bound to its physical JNDI location. In the AppServer, JMS resource definitions are bound into a JNDI service provider during deployment of a Definitions ARchive (DAR) module. By default, these objects are bound to the partitions' Naming Service, the JNDI CosNaming service provider in AppServer. This chapter describes how to define JMS resource object definitions in a DAR module and delves into the details of how to get a handle to a JMS resource object from a J2EE application. A discussion of JMS activity and how it relates transactions is also provided.

A DAR contains the JNDI definitions module (`jndi-definitions.xml` file) which contains properties for each resource related object that you want to bind to a JNDI provider (Naming Service). When an application EAR is deployed, the contents of the DAR file get deployed in the Naming Service of a partition. The properties of the resource related object defined in the `jndi-definitions.xml` file are stored in a JNDI bound object in the partition-hosted Naming Service.

When an application client or an EJB component does a JNDI lookup for a resource related object, it calls a `lookup()` method which communicates with the JNDI provider:

1 The Application Client refers to the `<resource-ref>` element in the standard deployment descriptor (in the case of EJBs it is `ejb-jar.xml`) to get the logical name of the resource. (It does a lookup in the component's local namespace, `java:comp/env`, to obtain the logical name of the object.) This logical name is specified in its `<resource-ref-name>` sub element. For example, in `ejb-jar.xml`:

```
⋮
 <description>This example demonstrates JMS XA and JDBC XA in a two-phase
commit transaction.</description>
 <enterprise-beans>
  <session>
   ⋮
     <resource-ref>
        <description />
        <res-ref-name>jms/insurance/ConnectionFactory</res-ref-name>
        <res-type>javax.jms.ConnectionFactory</res-type>
        <res-auth>Container</res-auth>
     </resource-ref>
     ⋮
  </session>
```

2 Using this logical name, the Container obtains the actual JNDI location of the JMS resource definition (a JNDI bound object) from the Borland proprietary deployment descriptor, `ejb-borland.xml`:

```
⋮
 <enterprise-beans>
  <session>
  ⋮
     <resource-ref>
        <res-ref-name>jms/insurance/ConnectionFactory</res-ref-name>
        <jndi-name>jms/xacf</jndi-name>
     </resource-ref>
  </session>
```

3 The Container then creates an instance of the resource object by using the stored property values in the bound object. The following properties were stored in the ConnectionFactory object from the `jndi-definitions.xml` file when the related DAR was deployed:

```
⋮
<jndi-definitions>
  <jndi-object>
     <jndi-name>jms/xacf</jndi-name>
     <classname>com.tibco.tibjms.TibjmsXAConnectionFactory</class-name>
     <property>
        <prop-name>serverUrl</prop-name>
        <prop-type>String</prop-type>
        <prop-value>localhost:7222</prop-value>
     </property>
  </jndi-object>
```

4 This instance is then wrapped in the Borland proprietary API (in the JMS proxy layer) by the container and passed back to the caller (could be an application client or could be another J2EE component) of the `lookup()`.

JMS 1.1 Common APIs

With JMS 1.1, JMS client applications have the option to use domain-independent unified APIs. A client can obtain a handle to a generic JMS ConnectionFactory, and from it a generic Session object that can be used with either a Queue or Topic for message processing. The common APIs along with their domain specific APIs are listed in the table below:

JMS Common APIs (in JMS 1.1)	Point-to-Point Domain APIs	Publish/Subscribe Domain
ConnectionFactory	QueueConnectionFactory	TopicConnectionFactory
Connection	QueueConnection	TopicConnection

JMS Common APIs (in JMS 1.1)	Point-to-Point Domain APIs	Publish/Subscribe Domain
Destination	Queue	Topic
Session	QueueSession	TopicSession
MessageProducer	QueueSender	TopicPublisher
MessageConsumer	QueueReceiver	TopicSubscriber
XAConnectionFactory	XAQueueConnectionFactory	XATopicConnectionFactory
XAConnection	XAQueueConnection	XATopicConnection
XASession	XAQueueSession	XATopicSession

The major change when using common interfaces is that one or more Queue and/or Topic destinations can now be simultaneously accessed through the same session all within the same transaction. With this change, either all the messages in a single transaction (messages to/from the queue(s) and the topic(s)) get sent and the transaction is considered successful or the whole transaction is aborted and none of the messages are delivered.

Borland AppServer supports the domain-independent APIs of JMS 1.1, and the constraints imposed by J2EE 1.4 in use of all JMS 1.1 APIs.

Configuring JMS Connection Factories and Destinations

Using the Management Console, navigate to the “Deployed Modules” list in the Partition whose JMS resource objects you need to configure. By default, every Partition has a predeployed JNDI Definitions Module (DAR) called `default-resources.dar`. Right-click on the module and choose “Edit deployment descriptor” from the context menu. The Deployment Descriptor Editor (DDEditor) opens.

In the Navigation Pane of the DDEditor is a list of JMS connection factories, and queues/topics preconfigured in the product. Click on the connection factory name. The right pane will display the properties for it. For each connection factory, you can choose Tibco, Sonic, WMQ or another (“Other”) JMS provider. The DDEditor has knowledge of Tibco and auto fills the class names for each. You can also choose the object resource type from the JMS Object type drop-down list. For OpenJMS, you must edit the `openjms.xml` file to configure connection factories and destination. See “Configuring JNDI objects for OpenJMS” for details on how to access this file.

If you selected “Other” from the JMS Provider list, look up the JMS vendor's documentation to ascertain the correct name of its connection factory, topic, or queue implementation class. In addition, the Main panel will not suggest any properties to fill in and you will need to use the Properties tab to set any appropriate properties.

To create a new JMS object, right-click the root node in the navigation pane and select “New JMS Object” from the Context Menu.

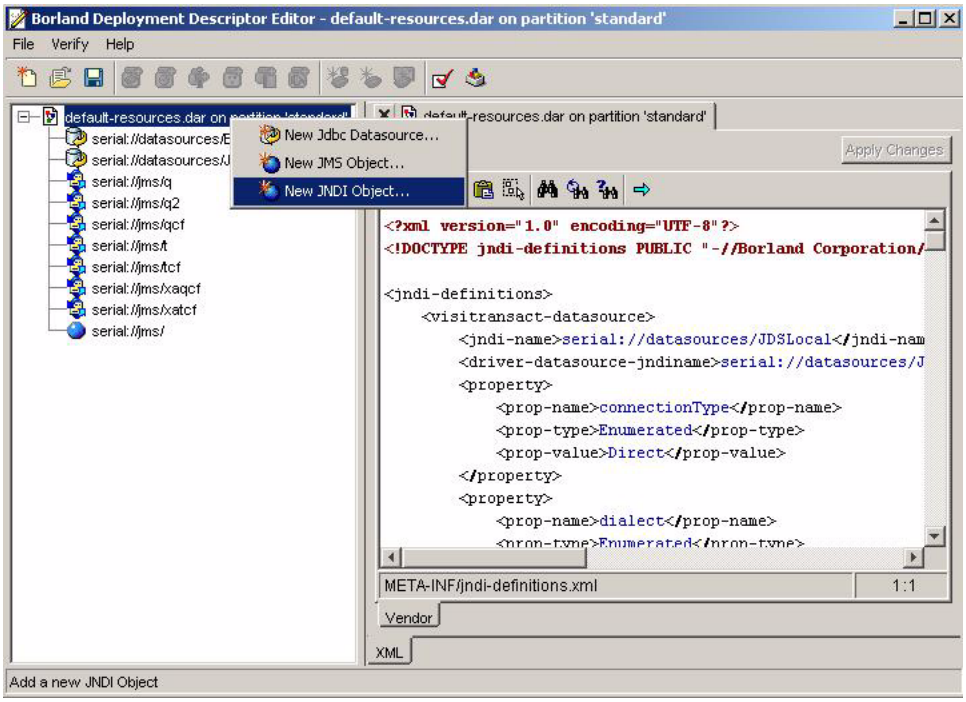

A dialog box prompts for JNDI name of JMS object to be created. By default, the name specified corresponds to a location in the Naming Service. If a JNDI name is specified with a “serial//” prefix, the remaining name following the prefix corresponds to a location in the Naming Service. Once given, a representation of this JMS object appears in the tree in the Navigation Pane. Click the representation to open its configuration panel.

The DDEditor has knowledge of Tibco, Sonic, WMQ and can auto fill the class names for it.

Note

The Main panel will not suggest any JMS objects other than Tibco, Sonic and WMQ. You need to use the Properties tab to set any appropriate properties.

When you're finished, choose “File|Save...” and the module will be saved back to the Partition and redeployed.

Defining Connection Pool Properties for JMS Connection Factories

Each JMS connection factory defined in AppServer has an associated connection pool. You can specify connection pool properties for each JMS connection factory defined in the `jndi-definitions.xml` file of a DAR module. The AppServer partition system properties can be specified to dictate the default behavior for all JMS connection pools established in a partition's Java virtual machine. However, properties defined for individual JMS connection factories in the `jndi-definitions.xml` file override system property values.

The AppServer partition system properties used for default configuration of JMS connection pools are listed in the table below :

Property Name	Description	Type	Values
JMSConnectionMaxPoolSize	Maximum number of JMS connections allowed for a JMS connection pool	Integer	0<n, where n is the maximum number of connections allowed for a JMS connection pool. The default is 0, indicating an unlimited number
JMSConnectionWaitTimeout	Time allowed to wait for a free connection in the JMS connection pool	Integer	Default is 30 seconds
JMSConnectionIdleTimeout	Time allowed for connection to remain idle in the JMS connection pool before being discarded	Integer	Default is 60 seconds
JMSConnectionPoolDebug	Enable display of debug messages associated with AppServer JMS connection pooling	Boolean	Default value is true true – Enable Debug false – Disable debug
JMSConnectionPoolDisable	Controls use of AppServer connection pooling for JMS Connection factories.	Boolean	Default value is false false – Enable JMS connection true – Disable JMS connection
JMSConnectionPoolMonitorLevel	Enable JMS pool monitoring for AppServer JMS connection and session pools. **Important:** Each JMS connection can have a set of JMS sessions created, which are maintained in a JMS session pool. Consider the impact on the performance of the connection when you set a value for this property, because each level corresponds to an increased degree of maintenance and collection of values for sets of counters and states. If you choose “maximum”, it will have the greatest performance impact on the connection.	String	■ “none” (Default) ■ “minimum” ■ “medium” ■ “maximum”
JMSSessionMaxPoolSize	Maximum number of JMS sessions for each JMS session pool of a connection factory	Integer	0<n, where n is the maximum number of JMS sessions allowed for a JMS session pool associated with a JMS connection. The default is 0, indicating an unlimited number
JMSSessionWaitTimeout	Time allowed to wait for a free session in a JMS session pool	Integer	Default is 30 seconds
JMSSessionPoolDisable	Controls use of AppServer session pooling per JMS connection. When a JMS connection factory is looked up under JNDI, if the value of this property is true the vendor JMS connection factory is returned. It is not wrapped by any AppServer proxy class.	Boolean	Default value set to false false – Enable JMS session true – Disable JMS session
JMSSessionPoolDebug	Enable display of debug messages associated with AppServer JMS session pooling	Boolean	Default value is false false – Disable debug true – Enable debug

Defining Individual JMS Connection Factory Properties

You can define JMS pool properties for individual connection factories in `jndi-definitions.xml` file. These properties override partition system properties. Use the `<property>` element to add a pool property. For example:

```
<jndi-definitions>
   <!-- *************************** -->
   <!-- * JMS Connection Factories * -->
   <!-- *************************** -->
       <jndi-object>
           <jndi-name>jms/cf</jndi-name>
           <class-name>com.tibco.tibjms.TibjmsConnectionFactory</class-name>
           <property>
               <prop-name>serverUrl</prop-name>
               <prop-type>String</prop-type>
               <prop-value>localhost:7222</prop-value>
           </property>
           <property>
               <prop-name>besConnectionPoolMaxPoolSize</prop-name>
               <prop-type>Integer</prop-type>
               <prop-value>11</prop-value>
           </property>
           <property>
               <prop-name>besConnectionPoolDebug</prop-name>
               <prop-type>Boolean</prop-type>
               <prop-value>true</prop-value>
           </property>
           <property>
               <prop-name>besSessionPoolDisable</prop-name>
               <prop-type>Boolean</prop-type>
               <prop-value>true</prop-value>
           </property>
       </jndi-object>
       ⋮
   </jndi-definitions>
```

The full set of JMS Connection factory pool properties are listed below together with the corresponding system property that each overrides:

Individual Pool Property	Associated System Property
besConnectionPoolMaxPoolSize	JMSConnectionMaxPoolSize
besConnectionPoolWaitTimeout	JMSConnectionWaitTimeout
besConnectionPoolIdleTimeout	JMSConnectionIdleTimeout
besConnectionPoolMonitorLevel	JMSConnectionPoolMonitorLevel
besConnectionPoolDisable	JMSConnectionPoolDisable
besConnectionPoolDebug	JMSConnectionPoolDebug
besSessionPoolMaxPoolSize	JMSSessionMaxPoolSize
besSessionPoolWaitTimeout	JMSSessionWaitTimeout
besSessionPoolDisable	JMSSessionPoolDisable
besSessionPoolDebug	JMSSessionPoolDebug

Obtaining JMS Connection Factories and Destinations in J2EE Application Components

A JMS connection factory object is obtained in much the same way as a JDBC datasource object. The factory object is declared in a `<resource-ref>` element of both

the standard J2EE and Borland-specific deployment descriptors. However, extra configuration is required if an application needs to interact with destinations of a JMS provider. A <resource-env-ref> element must be specified in both descriptors with definition of at least one JMS destination, that is a target queue or topic on which messages can be produced/consumed. While the standard J2EE deployment descriptor provides the logical name and type of a JMS connection factory and a JMS destination, the Borland specific deployment descriptor maps the logical name to a reference of the actual target object, resolved through JNDI lookup.

J2EE 1.2 and J2EE 1.3

Details of <resource-ref> and <resource-env-ref> elements for standard J2EE deployment descriptors are described in J2EE 1.3 specifications. These elements apply to all application components, for instance EJBs, Servlets and application clients, that wish to use JMS APIs. Similarly, corresponding <resource-ref> and <resource-env-ref> elements exist in accompanying Borland-specific deployment descriptors. Let's look at deployment descriptors for an EJB session bean that uses JMS. First, from standard EJB descriptor, ejb-jar.xml:

```
⋮
<session>
 <ejb-name>session_bean</ejb-name>
 ⋮
 <resource-ref>
     <res-ref-name>jms/MyJMSQueueConnectionFactory</res-ref-name>
  <res-type>javax.jms.QueueConnectionFactory</res-type>
  <res-auth>Container</res-auth>
  <res-sharing-scope>Shareable</res-sharing-scope>
 </resource-ref>
 <resource-env-ref>
  <res-env-ref-name>jms/MyJMSQueue</res-env-ref-name>
  <res-env-ref-type>javax.jms.Queue</res-env-ref-type>
 </resource-env-ref>
 ⋮
</session>
```

The portable descriptor above defines logical names for a JMS connection factory and a JMS Queue through <resource-ref> and <resource-env-ref> respectively. In Borland's proprietary deployment descriptors, such as ejb-borland.xml for EJB components, the <resource-ref> element is used to resolve the logical name to actual JNDI location of a JMS connection factory definition when a JNDI lookup is performed for the desired connection factory in the application component. This element is used within descriptor definitions of individual components. For example, a <resource-ref> for an entity bean must be found within the <entity> tags. Let's examine the DTD representation of the <resource-ref> element for J2EE 1.2 and 1.3 Borland deployment descriptors:

```
<!ELEMENT resource-ref (res-ref-name, jndi-name, cmp-resource?)>
```

In this element you specify the following:

- **res-ref-name**: this is the logical name for the resource object, the same logical name you use in the <resource-ref> element of the standard ejb-jar.xml descriptor file. This is the name your application components use to look up the JMS connection factory.
- **jndi-name**: this is the JNDI name of the connection factory that will be bound to the logical name. It must match the value of the corresponding <jndi-name> element of the <jndi-object> element in the deployed DAR where the connection factory is defined.

Just as the <resource-ref> element is used to map logical names for JMS connection factories to actual JNDI location of desired connection factory definition, the <resource-env-ref> element maps the logical name for JMS destinations, such as Queues and Topics, to actual JNDI location of destination definition. The DTD representation of this element for Borland deployment descriptors is as follows:

```
<!ELEMENT resource-env-ref (resource-env-ref-name, jndi-name)>
```

Two elements are specified:

- **resource-env-ref-name**: this is the logical name of the Topic or Queue, and its value must be identical to the value of the <res-env-ref-name> of J2EE standard descriptor.
- **jndi-name**: this is the JNDI name of the topic or queue that resolves the logical name.

The final contents for Borland descriptor ejb-borland.xml accompanying the ejb-jar.xml defined above is:

```
<session>
 <ejb-name>session_bean</ejb-name>
 ⋮
 <resource-ref>
  <res-ref-name>jms/MyJMSQueueConnectionFactory</res-ref-name>
  <jndi-name>resources/qcf</jndi-name>
 </resource-ref>
 <resource-env-ref>
  <resource-env-ref-name>jms/MyJMSQueue</resource-env-ref-name>
  <jndi-name>resources/q</jndi-name>
 </resource-env-ref>
 ⋮
</session>
```

Keep in mind that <resource-ref> and <resource-env-ref> elements can be used for all J2EE components that require JMS related resource objects. For instance, application clients using JMS APIs should obtain connection factories and destinations the same way as EJBs, through JNDI lookup or Resource References in application code and specification of <resource-ref> and <resource-env-ref> elements in the clients deployment descriptors. For example, in the J2EE standard descriptor, application-client.xml:

```
<application-client>
⋮
<resource-ref>
  <res-ref-name>jms/MyJMSTopicConnectionFactory</res-ref-name>
  <res-type>javax.jms.TopicConnectionFactory</res-type>
  <res-auth>Application</res-auth>
</resource-ref>
<resource-env-ref>
  <res-env-ref-name>jms/MyJMSTopic</res-env-ref-name>
  <res-env-ref-type>javax.jms.Topic</res-env-ref-type>
</resource-env-ref>
⋮
  </application-client>
```

and its accompanying Borland descriptor application-client-borland.xml:

```
  <application-client>
 ⋮
```

```
<resource-ref>
 <res-ref-name>jms/MyJMSTopicConnectionFactory</res-ref-name>
 <jndi-name>resources/tcf</jndi-name>
</resource-ref>
<resource-env-ref>
 <resource-env-ref-name>jms/MyJMSTopic</resource-env-ref-name>
 <jndi-name>resources/t</jndi-name>
</resource-env-ref>
⋮
 </application-client>
```

Now let's see how we obtain object references to a JMS connection factory and a destination in the application logic. In retrieval of the connection factory, the application performs a JNDI lookup of the `<res-ref-name>` value from `<resource-ref>` element in the J2EE deployment descriptor. To retrieve the destination object, a JNDI lookup is performed against the `<res-env-ref-name>` value of `<resource-env-ref>` element in the J2EE deployment descriptor. The names specified for `<jndi-name>` are identical to JNDI names in `<jndi-object>` elements of the JMS resource definitions in a deployed DAR module. When a lookup succeeds a JMS resource object is obtained, that is, for JMS connection factory identified through logical name `jms/MyJMSTopicConnectionFactory` a deployed JMS definition object is retrieved from the Naming Service under `resources/tcf` and from it a connection factory object is created and returned to the application.

For example, the application client code associated with client descriptors provided above resolves JMS resource objects as follows:

```
javax.jms.TopicConnectionFactory myTCF;
javax.jms.Topic myTopic;
try {
 javax.naming.Context ctx = (javax.naming.Context) new
javax.naming.InitialContext();
  myTCF = (TopicConnectionFactory) ctx.lookup("java:comp/env/jms/
MyJMSTopicConnectionFactory");
  // Now ready to obtain a connection from myTCF
  myTopic = (Topic) ctx.lookup("java:comp/env/jms/MyJMSTopic");
  ⋮
}
catch (javax.naming.NamingException exp) {
 exp.printStackTrace();
}
```

J2EE 1.4

In earlier versions of J2EE, each application component had to declare an `<resource-env-ref>` in the standard deployment descriptor for look up of a JMS destination from its own local namespace. If separate application components have references to the same destination, there was no way for a deployer to know that these `<resource-env-ref>`s should be bound to the same destination.

Here is an example which uses the `<resource-env-ref>` to define the same JMS destination from two separate application components, Session beans, in this case:

```
⋮
<ejb-jar ... >
   <enterprise-beans>
         <session>
            <ejb-name>SenderEJB</ejb-name>
            ⋮
            <resource-ref>
               <res-ref-name>jms/ConnectionFactory</res-ref-name>
               <res-type>javax.jms.ConnectionFactory</res-type>
               <res-auth>Container</res-auth>
            </resource-ref>
            <resource-env-ref>
```

```
            <resource-env-ref-name>jms/LogicalNameA</resource-env-ref-name>
            <resource-env-ref-type>javax.jms.Queue</resource-env-ref-type>
         </resource-env-ref>
      </session>
      <session>
         <ejb-name>ReceiverEJB</ejb-name>
         ⋮
         <resource-ref>
            <res-ref-name>jms/ConnectionFactory</res-ref-name>
            <res-type>javax.jms.ConnectionFactory</res-type>
            <res-auth>Container</res-auth>
         </resource-ref>
         <resource-env-ref>
            <resource-env-ref-name>jms/LogicalNameB</resource-env-ref-name>
            <resource-env-ref-type>javax.jms.Queue</resource-env-ref-type>
         </resource-env-ref>
      </session>
```

In J2EE 1.4, although you can continue to use `<resource-env-ref>`, a new element, `<message-destination-ref>` has been introduced for specification of JMS destinations. The same example above can be rewritten using the `<message-destination-ref>` element instead of `<resource-env-ref>` in the standard deployment descriptor, `ejb-jar.xml`, as follows:

```
⋮
<ejb-jar ... >
   <enterprise-beans>
      <session>
         <ejb-name>SenderEJB</ejb-name>
         ⋮
         <resource-ref>
            <res-ref-name>jms/ConnectionFactory</res-ref-name>
            <res-type>javax.jms.ConnectionFactory</res-type>
            <res-auth>Container</res-auth>
         </resource-ref>
         <message-destination-ref>
            <message-destination-ref-name>jms/LogicalNameA
              </message-destination-ref-name>
           <message-destination-type>javax.jms.Queue
             </message-destination-type>
            <message-destination-usage>Produces</message-destination-usage>
             <message-destination-link>MsgQueue1</message-destination-link>
         </message-destination-ref>
      </session>
      <session>
         <ejb-name>ReceiverEJB</ejb-name>
         ⋮
         <resource-ref>
            <res-ref-name>jms/ConnectionFactory</res-ref-name>
            <res-type>javax.jms.ConnectionFactory</res-type>
            <res-auth>Container</res-auth>
         </resource-ref>
         <message-destination-ref>
            <message-destination-ref-name>jms/LogicalNameB
              </message-destination-ref-name>
           <message-destination-type>javax.jms.Queue
             </message-destination-type>
            <message-destination-usage>Consumes</message-destination-usage>
           <message-destination-link>MsgQueue1</message-destination-link>
         </message-destination-ref>
      </session>
```

```
    </enterprise-beans>
    <assembly-descriptor>
    <message-destination>
          <message-destination-name>MsgQueue1</message-destination-name>
       </message-destination>
    </assembly-descriptor>
</ejb-jar>
```

The above example shows two `<message-destination-ref>` elements that refer to the same destination, the Queue MsgQueue1. Each of these `<message-destination-ref>`s have a `<message-destination-link>` element whose value is MsgQueue1. The link elements maps the `<message-destination-ref>`s to a `<message-destination>` element in the `<assembly-descriptor>` resolving the two `<message-destination-ref>`s to the same queue. The ejb-borland has a corresponding `<message-destination>` element within its `<assembly-descriptor>` element. At runtime, the `<message-destination>` element from `ejb-jar.xml` is resolved against the `<jndi-name>` of the `<message-destination>` element specified in the `ejb-borland.xml` descriptor:

```
⋮
   <ejb-jar ... >
      <enterprise-beans>
         ⋮
      </enterprise-beans>
      <assembly-descriptor>
      <message-destination>
            <message-destination-name>MsgQueue1</message-destination-name>
            <jndi-name>jms/TibcoQueue1</jndi-name>
         </message-destination>
      </assembly-descriptor>
   </ejb-jar>
```

To show JMS message flow in an application where more than one `<message-destination-ref>`s resolves to the same underlying destination, each one of them should declare a `<message-destination-link>` with a value corresponding to a `<message-destination>` element in the `<assembly-descriptor>`. The value in the `<message-destination-link>` must match the value of the `<message-destination-name>` in the `<message-destination>` element. At runtime, the `<message-destination-ref>`s will resolve to the same destination.

You can also link to a `<message-destination>` defined in a different J2EE module within the same application. For example, `<message-destination-link> ../other/other.jar#destination </message-destination-link>` would link to a <message-destination> with the name destination in the JAR file at the relative path `../other/other.jar`.

You must also specify a `<message-destination>` element in the `ejb-jar.xml` for every destination you use.

Important

The JNDI name of a `<message-destination>` in Borland deployment descriptor takes precedence over a specified JNDI name of a `<message-destination-ref>` that has a link to the `<message-destination>` element.

JMS and Transactions

The rules for using JMS APIs in EJB bean code with transactions are discussed in the EJB 2.0 specification section 17.3.5.

Following is an extract:

> 17.3.5 Use of JMS APIs in transactions
>
> The Bean Provider must not make use of the JMS request/reply paradigm (sending of a JMS message, followed by the synchronous receipt of a reply to that message) within a single transaction.

> Because a JMS message is not delivered to its final destination until the transaction commits, the receipt of the reply within the same transaction will never take place. Because the container manages the transactional enlistment of JMS sessions on behalf of a bean, the parameters of the createSession(boolean transacted,int acknowledgeMode), createQueueSession(boolean transacted,int acknowledgeMode) and createTopicSession(boolean transacted, int acknowledgeMode) methods are ignored. It is recommended that the Bean Provider specify that a session is transacted, but provide 0 for the value of the acknowledgment mode.
>
> The Bean Provider should not use the JMS acknowledge() method either within a transaction or within an unspecified transaction context. Message acknowledgment in an unspecified transaction context is handled by the container. Section 17.6.5 describes some of the techniques that the container can use for the implementation of a method invocation with an unspecified transaction context.

Avoiding use of the JMS request/reply paradigm and JMS `acknowledge()` method is equally relevant for other J2EE components such as application clients, as it is to EJB bean code. In addition to rules described above, application code should not use any JMS XA APIs. The program should look exactly as if the code is written in a non-transactional JMS program. It is the Container's responsibility to handle any XA handshakes required when a global transaction is active. The only configuration required is that deployment descriptor element `<resource-ref>`, with reference to the JMS Connection factory JNDI object, be set up to use the XA variant. If it is non-XA, the program still runs, but there are no atomicity guarantees, in other words, it is a local transaction. Also note that for AppServer to automatically handle the transaction handshakes it is necessary to have the application run in a Container, either EJB, Web or appclient. For example, a java client with no JMS XA API calls will not have its JMS activity participate in a global transaction, one has to write it as a J2EE application client instead. Also make sure that all connection factories are looked up through deployment descriptor element `<resource-ref>`. This allows the Container to trap the JMS API calls and insert appropriate hooks.

Let us examine in more detail the following sentences extracted from the EJB 2.1 specification:

> Because the container manages the transactional enlistment of JMS sessions on behalf of a bean, the parameters of the createSession(boolean transacted,int acknowledgeMode), createQueueSession(boolean transacted,int acknowledgeMode) and createTopicSession(boolean transacted, int acknowledgeMode) methods are ignored. It is recommended that the Bean Provider specify that a session is transacted, but provide 0 for the value of the acknowledgment mode.

The assumption here is that messages produced/consumed by JMS sessions should be included as part of the unit of work maintained by a global transaction, should a global transaction be active. In order for transactional enlistment to occur, the parent connection factory of connections on which `createSession()`, `createQueueSession()` or `createTopicSession()` are invoked must be defined as a `javax.jms.XAConnectionFactory`, `javax.jms.XAQueueConnectionFactory` or `javax.jms.XATopicConnectionFactory`, respectively. That is, the value for `<res-type>` of J2EE deployment descriptor element `<resource-ref>`, with definition of JMS connection factory to be used for the J2EE component, must be either `javax.jms.XAConnectionFactory`, `javax.jms.XAQueueConnectionFactory` or `javax.jms.XATopicConnectionFactory`. If the connection factory has a non XA connection factory `<res-type>`, the program still runs but work performed on JMS sessions will not be included in the global transaction; in this case the `transacted` and `acknowledgeMode` parameters will influence the behavior of message production/consumption. For instance:

```
import javax.jms.*;

QueueConnectionFactory nonXAQCF;
Queue                  myQueue;

try {
 javax.naming.Context ctx = (javax.naming.Context) new
javax.naming.InitialContext();
```

```
  nonXAQCF = (QueueConnectionFactory) ctx.lookup("java:comp/env/jms/
MyJMSQueueConnectionFactory");
  myQueue = (Queue) ctx.lookup("java:comp/env/jms/MyJMSQueue");
 }
 catch (javax.naming.NamingException exp) {
  exp.printStackTrace();
 }

 // Note: A global transaction context is currently active when the Session
         is being created

 QueueSession qSession = conn.createQueueSession(false,
Session.AUTO_ACKNOWLEDGE);
 QueueSender = qSession.createSender(myQueue);
 TextMessage msg = qSession.createTextMessage("A Message ");
 sender.send(msg);
```

Here, the TextMessage `msg` is queued regardless of the outcome of the active global transaction. This is in line with test cases in the J2EE Compatibility Test Suite. It seems useful to have this capability in a global transaction, whereby a log message needs to be sent irrespective of the enclosing global transaction's completion result.

Multiple resource access within a single global transaction is supported in AppServer. This provides the capability to do a unit of work which is composed of sending/receiving JMS messages along with some other type of resource manager access. That is, it is desirable to write code (in an EJB for example) which does some work against a non JMS resource such as a database and also send a message to a queue with the Container providing transactional completion for all work performed. Upon completion of the transaction, either the work performed against the database is committed AND the message is queued, or should something fail during the transaction database work is rolled back AND the message is not delivered to the queue.

In application code, an EJB method such as `doSomeWork()` shown below is supported in AppServer:

```
// Business method in a session bean, the EJB container marks the transaction
     void doSomeWork()
     {
       // Establish a database connection
          java.sql.Connection dbConn = datasource.getConnection();

        // Execute SQL
          ⋮

       // Call a remote EJB in the same transaction

          ejbRemote.doWork();

       // Send a JMS message to a queue
          jmsSender.send(msg);

      }
```

Enabling the JMS services security

See "JMS provider pluggability" for vendor-specific information on JMS services such as security.

Advanced Concepts for Configuring JMS Connection Factories and Destinations

JDBC datasource resource objects and JMS Connection Factories and Destinations for JMS providers Tibco, SonicMQ and WMQ, are defined in a DAR module using a `jndi-definitions.xml` descriptor. Module `tibco-resources.dar`, deployed to Welcome Partition in sample BAS configuration j2eeSample contains some default Tibco Connection Factories, Topics and Queues, defined for the AppServer install with Tibco JMS server option. You can edit these existing definitions to suit your environment or create new definitions using the DDEditor. JMS connection factories, similar to JDBC datasources, are classes that wrap the connection factory classes provided by JMS vendors. If you want to use a JMS vendor not bundled or certified to work with AppServer, you need to deploy that vendor's connection factory classes to your Partition.

See "JMS provider pluggability" for vendor-specific information on JMS queues.

24

Using Hibernate

This section provides information on the integration of Hibernate with Borland AppServer, how to create a Hibernate application, how to package Hibernate applications, how to enable and disable the Hibernate service, and how to enable and disable Hibernate statistics logging.

Overview

Hibernate is a powerful, high performance object/relational persistence and query service. A hibernate application is a standard J2EE application that uses Hibernate entities for Object Relational Mapping (ORM) and persistence. For more information on Hibernate, see www.hibernate.org.

Borland AppServer integrates tightly with Hibernate and simplifies the development and deployment of hibernate applications.

The Borland AppServer support for Hibernate provides the following capabilities:

- Integrating Hibernate with Borland AppServer facilitates the use of Hibernate entities that leverage Borland AppServer's infrastructural capabilities, such as Transaction, Naming Service, Connection Pooling, Container-Managed Data Sources, and so on.
- You no longer need to package Hibernate libraries with the Hibernate application. Hibernate 3.1.3 libraries are available in the Partition's class path.
- During deployment, Borland AppServer detects hibernate applications, and automatically creates hibernate artifacts (SessionFactory) and publishes them to the JNDI.
- Borland AppServer enables Hibernate applications to work with AppServer Data Source and Transaction Manager (JTS and OTS) using JTA and CMT.
- You can use the ORM capabilities of Hibernate and all its features within applications deployed in the Borland AppServer.
- Borland AppServer offers the flexibility of using existing CMP 2.0 entity beans in conjunction with Hibernate entities.
- The development of your model (persistence layer) will be easier since persistent entities can be POJOs and they do not need to implement any interface like entity beans.

Creating a Hibernate Application

Creating a Connection Data Source

The Hibernate session factory requires a connection data source to be specified in the Hibernate configuration files. You can create a data source in the Borland Application Server using a DAR (JNDI definitions module) or using the jndi-definitions.xml file. The default-resources.dar file includes some default data sources.

If you want to use a particular database, you need to ensure that a data source is created for that database and you need to use the jndi-name of the data source in the `connection.datasource` property in the hibernate session factory configuration file (hibernate.cfg.xml).

Example

```
<visitransact-datasource>
   <jndi-name>datasources/JDSLocal</jndi-name> …
```

Use `datasource/JDSLocal` in the hibernate.cfg.xml for the `connection.datasource` property.

Working with Configuration Files and Session Factories

Borland AppServer detects hibernate applications and automatically creates session factories and publishes them to the JNDI.

Working with the Hibernate Configuration File

Note

This section provides information on how to configure a Hibernate SessionFactory for use within Borland AppServer. For more information on configuring a Hibernate SessionFactory, see the Hibernate documentation at www.hibernate.org.

You can configure the SessionFactory using the hibernate configuration file (hibernate.cfg.xml). You must specify a JNDI name in the `name` attribute of the `session-factory` element within the hibernate configuration file.

```
<session-factory name="hibernate/SF">
…
```

Once you configure and package the hibernate configuration file, Borland AppServer will automatically create the session factories and publish them to the JNDI.

In the Hibernate application code, obtain the SessionFactory from the JNDI by using the following:

```
Context ctx = new InitialContext();
ctx.lookup(“java:comp/env/hibernate/SF”);
```

Note

Using this code avoids the manual step of creating a session factory programmatically.

Configuring the Data Source

You must use the same connection data source that you created in the “Creating a Connection Data Source” section. To specify the data source, use the following property:

```
<property name="connection.datasource">datasources/JDSLocal</property>
```

Configuring Transactions

You can configure Application Server Transaction Support using the following properties:

- `transaction.factory_class`

  To configure this class, use the following:

  ```
  <property name="transaction.factory_class">
    org.hibernate.transaction.JTATransactionFactory
  </property>
  ```

  Note

  If you use JTA directly (BMT), you should use `org.hibernate.transaction.JTATransactionFactory`. In a CMT session bean, you should use `org.hibernate.transaction.CMTTransactionFactory`.

- `transaction.manager_lookup_class`

  To configure this class, use the following:

  ```
  <property name="transaction.manager_lookup_class">
     org.hibernate.transaction.BESTransactionManagerLookup
  </property>
  ```

Setting the Property for Automatic Session Flushing

If you are using JTA Transaction boundaries (BMT or CMT) for a session, you need to set the auto session flush property to flush a session before the transaction completes.

```
<property name="hibernate.transaction.flush_before_completion">true</property>
```

Setting the Property for Automatic Session Closing

If you are using Container Managed Transactions, when you want to close a Session after the Transaction completes, you need to set the following property:

```
<property name="hibernate.transaction.auto_close_session">true</property>
```

Setting the Property for Collecting Statistics Information

The Hibernate service displays statistics information for each SessionFactory. To enable Hibernate to collect statistics information, you need to set the following property on the SessionFactory:

```
<property name="hibernate.generate_statistics">true</property>
```

For more information on Hibernate Statistics, see “Working with Hibernate Statistics”.

Setting Entity Mapping

You can include all the Entity Mapping files for a SessionFactory. To do this, you need to set the following property:

```
<mapping resource="com/borland/examples/ejb/insurance/Claim.hbm.xml"/>
```

Note

The path of the mapping XML file is relative to the Hibernate configuration file.

Setting the Dialect

You can specify dialects using the `dialect` property. For example, for using a JdataStore data source, you can specify the dialect as follows:

```
<property name="dialect">
  org.hibernate.dialect.JdataStoreDialect
</property>
```

For information on other supported dialects, see the Hibernate documentation at www.hibernate.org.

Setting the Property to Drop and Recreate the Database Schema

To drop and recreate database schema on startup, you need to set the `hbm2ddl.auto` property to `create`.

```
<property name="hbm2ddl.auto">create</property>
```

Setting the property to Log Executed SQL

To echo all executed SQL to stdout, you need to set the `show_sql` property to `true`.

```
<property name="show_sql">true</property>
```

Packaging

EJB Module: EJB-JAR

For an EJB-JAR, you must include the `hibernate.cfg.xml` file under the `Archive` root.

For example, if the EJB-JAR file is `Hibernate-EJB.jar`, then the packaging should be as shown in the image below.

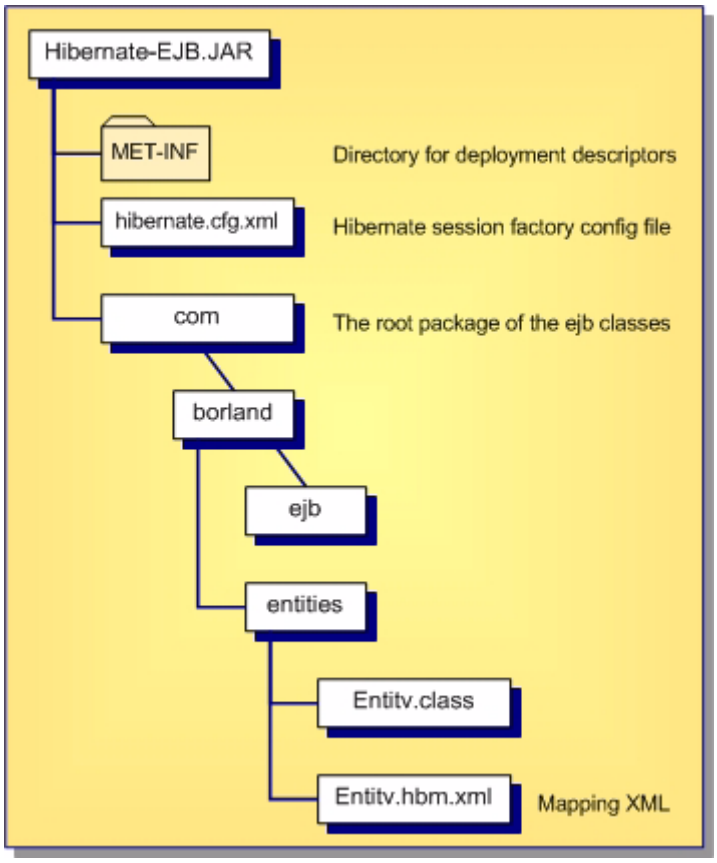

WEB Module: WAR

For a WAR, you must include the `hibernate.cfg.xml` file under the `WEB-INF/classes` folder.

For example, if the WAR file is `Hibernate-web.war`, then the packaging should be as shown in the image below.

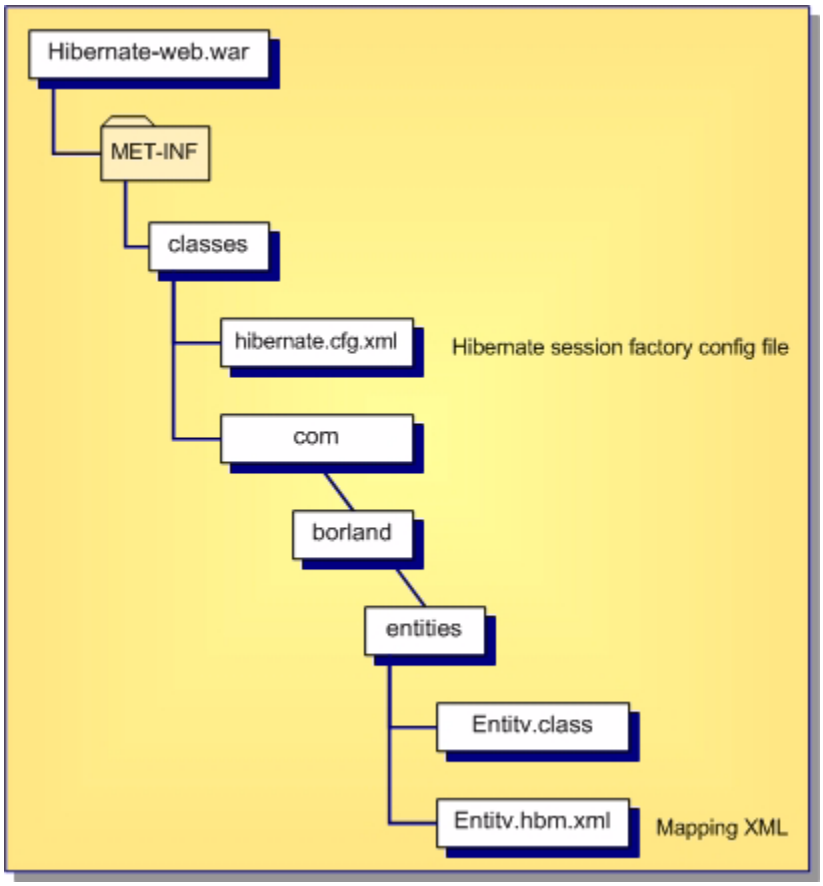

Packaging the Mapping XMLs

We recommend that you put the Mapping XML files and the entities at the same location. It is also a good practice to name the Mapping XML files as `<Entity-Class-Name>.hbm.xml`. The diagram given in the "EJB Module: EJB-JAR" and "WEB Module: WAR" sections indicate where you should put the Mapping XML files. You can refer to these Mapping XML files in the `hibernate.cfg.xml` relative to the configuration file location.

Packaging Multiple Session Factory configuration Files

If the application needs to use multiple databases (data sources) you might need to configure multiple hibernate session factories. By default, the Borland AppServer deployment only looks for `hibernate.cfg.xml`. To configure multiple session factories, do the following:

1. Include a file named `hibernate-borland.properties` in the same location where the hibernate configuration file (`hibernate.cfg.xml`) is located.
2. Specify the list of hibernate session factory configuration files in the `hibernate-borland.properties` file. Ensure that you separate each file with a semi-colon (;).

Example

```
hibernate.cfg.files=hibernate.cfg.xml;hibernate.cfg1.xml …
```

Note

Once you specify the list of files in the property file, the Hibernate deployment configures only these SessionFactories.

Working with EAR Modules

You must include the hibernate configuration files and the entities either within an EJB-JAR or a WAR. If an EJB-JAR uses Hibernate entities and session factories, you must include all the configuration files, Mapping XMLs, entity classes, and other files that the EJB-JAR needs within the module itself (EJB-JAR or WAR) and in the appropriate locations. These EJB-JARs and WARs can then be packaged within an EAR module.

Note

You cannot have hibernate entities and configuration files existing right under the EAR as a library.

Enabling and Disabling the Hibernate Service

The Borland Management Console allows you to enable and disable the Hibernate service.

To enable or disable the Hibernate service:

1. Open the Borland Management Console.
2. Open Configurations.
3. To enable the Hibernate service, right-click on Hibernate Service and select Enable.

   To disable the Hibernate service, right-click on Hibernate Service and select Disable.

Working with Hibernate Statistics

The Borland Management Console allows you to enable and disable Hibernate statistics logging.

To enable or disable Hibernate statistics logging:

1. Open the Borland Management Console.
2. Open Configurations.
3. Right-click on Hibernate Service and select Properties.
4. To enable the Hibernate statistics logging, select the hibernate.statistic.enable check box.

   To disable the Hibernate statistics logging, deselect the hibernate.statistic.enable check box.
5. In the hibernate.statistic.period field, enter the interval time, in milliseconds, with which the Hibernate statistics should be logged.

25

JMS provider pluggability

The Borland AppServer (AppServer) is designed to support any arbitrary JMS provider as long as certain requirements are met. There are three aspects of JMS pluggability: runtime pluggability, configuration of JMS admin objects (connection factories and queues/topics), and service management. You will achieve the best results if all three are met, but just having the runtime level pluggability, as well as vendor-specific ways to achieve the other levels, may be sufficient in many situations.

Borland AppServer 6.7 bundles the Tibco EMS 4.2.0 V12 and OpenJMS 0.7.6.1 JMS providers. OpenJMS is bundled as a partition level service.

Runtime pluggability

Runtime pluggability is determined by compliance to the J2EE specification. A CTS compliant JMS product that additionally implements the JMS specification optional APIs can seamlessly plug into the AppServer runtime. All features like transactions and MDB support are retained.

JMS products must possess the capability to perform transactional messaging to support MDBs and J2EE container intercepted messaging. That is, a JMS queue or topic must be a transactional resource. AppServer requires that JMS products implement the JTA XAResource interface and support JMS XA APIs.

In addition, the JMS product should support the `javax.jms.ConnectionConsumer` interface. The latter is vital since a central idea of MDBs is the concurrent consumption of messages. The ConnectionConsumer interface achieves this. The mechanism also works in conjunction with some optimal methods of the `javax.jms.Session` objects, namely `Session.run()` and `Session.setMessageListener()`.

Configuring JMS administered objects (connection factories, queues and topics)

If the JMS providers' admin objects, like connection factories and destinations follow the JavaBeans specification (as encouraged in the JMS specification), the Borland Deployment Descriptor Editor tool can define, edit and deploy these objects into the AppServer JNDI tree without needing a JMS product-specific mechanism.

For specific information on using other JMS service providers with AppServer and requirements for admin objects (queues, topics, and connection factories), see “Other JMS providers”.

Setting Admin Objects Using Borland Deployment Descriptor

You can set the admin object properties for Tibco from the Borland Deployment Descriptor Editor. To do so:

1. Launch the Borland Deployment Descriptor Editor from within the Management Console or standalone from the Start menu.
2. Select File|New and click on the JNDI Definitions tab to bring it forward.
3. Select the JNDI Definitions Archive and click OK to create a new JMS object.
4. Right-click on the Untitled JMS object in the left pane and select New JMS Object.
5. Give the JMS object a name in the New JMS Object window and click OK. Your JMS object will appear under your archive.
6. Click on your JMS object, and select the Main tab.
7. Configure your object by selecting various fields from the drop-down menus and entering information in the properties' fields.
8. To add additional properties for the JMS object, select the Properties tab and click “Add” to add properties (name, type, value).

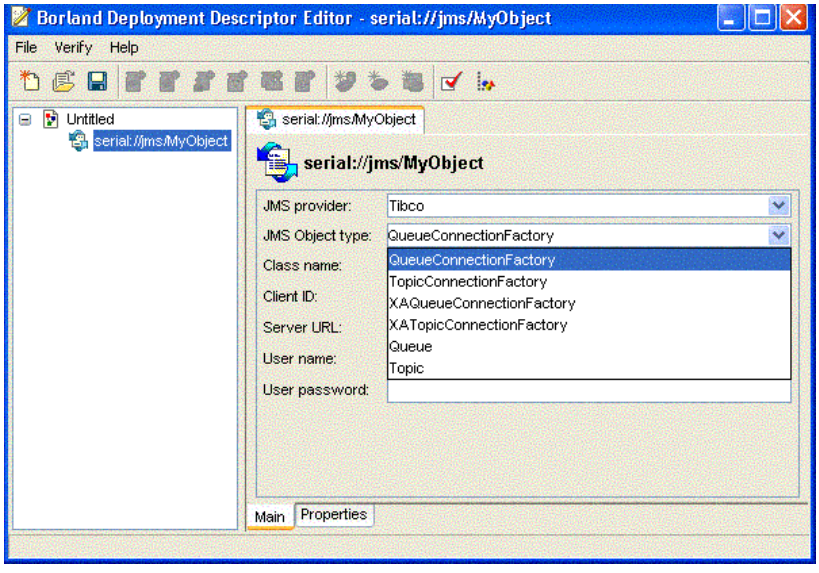

Service Management for JMS Providers

The AppServer service control infrastructure can manage the JMS service process (either a JVM or native process, whatever form it takes in the JMS provider) as a first

class managed object. Operations like starting, stopping and server configuration is provided for supported (Tibco, and OpenJMS) providers out of the box.

Tibco EMS 4.2

Tibco has achieved the runtime level of pluggability that is determined by the compliance to the J2EE specification. Tibco 4.2 is JMS 1.1 compliant and supports unified JMS APIs.

Added value for Tibco

Tibco provides this added value:

- Transparent installation
- Tibco Admin Console is available from AppServer Management Console *Tools* menu.

Configuring Admin Objects for Tibco

Tibco's admin object properties are defined in AppServer and can be configured graphically using the Borland Deployment Descriptor Editor.

See “Setting Admin Objects Using Borland Deployment Descriptor”.

Auto Queue Creation Feature in Tibco

Tibco has an auto queue creation feature by which if a specified queue does not exist in the server, the Tibco server will create the queue as necessary.

Tibco Admin Console

Note

You can launch the Tibco Admin Console from within AppServer only for the Windows platform. For all other platforms, run the executable from `<tibco_home>` directory to launch the console.

AppServer includes the Tibco Admin Console for additional configuration. To launch the Tibco Admin Console, select it from the Tools menu in the AppServer Management Console.

Configuring clients for fault tolerant Tibco connections

To connect to a backup server in the event of failure of a primary server, a client application must specify multiple server URLs in the `jndi-object` XML for the connection factories as below:

```
<jndi-object>
   <jndi-name>jms/XAConnectionFactory</jndi-name>
   <class-name>com.tibco.tibjms.TibjmsXAConnectionFactory</class-name>
   <property>
      <prop-name>serverUrl</prop-name>
      <prop-type>String</prop-type>
      <prop-value>localhost:7222,anotherhost:7222</prop-value>
   </property>
</jndi-object>
<jndi-object>
   <jndi-name>jms/ConnectionFactory</jndi-name>
   <class-name>com.tibco.tibjms.TibjmsConnectionFactory</class-name>
   <property>
      <prop-name>serverUrl</prop-name>
      <prop-type>String</prop-type>
      <prop-value>localhost:7222,anotherhost:7222</prop-value>
   </property>
</jndi-object>
<jndi-object>
 <jndi-name>jms/XAQueueConnectionFactory</jndi-name>
 <class-name>com.tibco.tibjms.TibjmsXAQueueConnectionFactory</class-name>
 <property>
  <prop-name>serverUrl</prop-name>
  <prop-type>String</prop-type>
  <prop-value>localhost:7222,anotherhost:7222</prop-value>
 </property>
</jndi-object>
<jndi-object>
 <jndi-name>jms/QueueConnectionFactory</jndi-name>
 <class-name>com.tibco.tibjms.TibjmsQueueConnectionFactory</class-name>
 <property>
  <prop-name>serverUrl</prop-name>
  <prop-type>String</prop-type>
  <prop-value>localhost:7222,anotherhost:7222</prop-value>
 </property>
</jndi-object>
<jndi-object>
 <jndi-name>jms/XATopicConnectionFactory</jndi-name>
 <class-name>com.tibco.tibjms.TibjmsXATopicConnectionFactory</class-name>
 <property>
  <prop-name>serverUrl</prop-name>
  <prop-type>String</prop-type>
  <prop-value>localhost:7222,anotherhost:7222</prop-value>
 </property>
</jndi-object>
 <jndi-name>jms/TopicConnectionFactory</jndi-name>
 <class-name>com.tibco.tibjms.TibjmsTopicConnectionFactory</class-name>
 <property>
  <prop-name>serverUrl</prop-name>
  <prop-type>String</prop-type>
  <prop-value>localhost:7222,anotherhost:7222</prop-value>
 </property>
</jndi-object>
```

Enabling Security for Tibco

Note

For information on SSL, please refer to the Tibco documentation. Tibco documentation is located in `<install_dir>\jms\tibco\doc\html`.

To enable security for Tibco, you can either modify the `tibemsd.conf` file located in `/<install_dir>/jms/tibco/bin`, or you can set it using the Tibco Admin tool.

Note

Make sure that the Tibco service is active before following the steps below.

1. From the Tools menu in Borland Management Console launch the Tibco Admin Console tool.
2. Type `connect`.
3. Enter Login name and Password.
4. Once connected, type `set server authorization=enabled`.
5. Security is now enabled. For client authentication, users should be created and added to an authorization group. For instance, create a user using the command:

   ```
   create user <name> [<description>] [password=<password>]
   ```

6. Add a member, type `add member <group-name> <user-name> [,<user-name2>,...]`.

Disabling security for Tibco

Follow the steps for Enabling Security for Tibco described above, but in step 4 rather than enabling security set the server authorization to disabled:

```
set server authorization=disabled
```

OpenJMS

OpenJMS is tied to the lifecycle of an AppServer partition. The AppServer contains a complete footprint of OpenJMS.

Note

OpenJMS 7.6 is JMS 1.0 compliant and does not support the unified APIs.

OpenJMS provides the following added value:

- Transparent Installation
- Support for automatic table creation
- Out-of-box integration with AppServer's Naming Service (JNDI), Transaction Service and Datasource
- Provides partition-level service management
- Support for RMI connector using VisiBroker
- Lifecycle management using Borland Management Console

When you install AppServer with OpenJMS in Borland AppServer version 6.7, OpenJMS comes packaged as a partition level template. This means that partitions created from this template get OpenJMS as an in-process service.

The following properties in the partition template are set to `true` by default when you install AppServer with OpenJMS:

- `ejb.mdb.use_jms_threads=true`

  This property is needed to allow the transaction that was started by OpenJMS to propagate into AppServer

- `ejb.mdb.local_transaction_optimization=true`

  This property is needed to allow the use of non-XA JMS connection factories to be used in transactional scenarios. If not set, the MDBs that have transaction `onMessage` method will fail to deploy

- `jts.allow_unrecoverable_completion=true`

  By default, AppServer uses the JDataStore database for message persistence. If your application database is different from your message persistence database, you will need to set this property to true in order to achieve the two phase commit.

See "EJB, JSS, and JTS Properties" for more information on each of the properties.

Even though OpenJMS can be used in a standalone mode where it is the only service in the partition, there are some advantages to using OpenJMS as an in-process service:

- It helps avoid the two-phase commit (2PC), and hence the performance cost and deployment complexity associated with it. This involves JDBC connection sharing among different components in the partition that access the database. It can be achieved by making OpenJMS persist messages in the same database as the one in which application data is stored. This configuration where 2PC can be avoided is possible only when OpenJMS is accessed through its embedded or RMI connector inside the AppServer partition. See "Configuring Datasource to Achieve 2PC Optimization" for details.
- Since all the components are centralized in a single virtual machine, you can avoid the cost of TCP/IP. The JMS client library calls an in-process JMS service using regular Java call and vice versa.
- Since VBJ provides local call optimizations, the application won't need to have two types of connectors. It can use only the RMI connector irrespective of client location with respect to the JMS server.

You can find the OpenJMS product documentation in the `<appserver_install>/jms/openjms/docs` directory.

Configuring JNDI objects for OpenJMS

Since each partition in AppServer can host an instance of OpenJMS, there is a dedicated configuration file, `openjms.xml`, for each partition. The `openjms.xml` file contains information about various OpenJMS connectors and JNDI objects that the instance hosts.

Note

AppServer does not support the administration GUI in OpenJMS. You can create and delete queues for OpenJMS by editing the `openjms.xml` file.

To add new Queues, Topics or Connection Factories for OpenJMS, you must modify its configuration file, `openjms.xml`. To access this file:

1. Select the relevant partition in the left pane of the Borland Management Console.
2. Right-click on the OpenJMS service in the left pane.
3. Select **Properties** from the drop-down menu.
4. Click on the openjms.xml tab in the properties pane to bring it forward.

5 Edit the file to add the JNDI objects.

See the OpenJMS documentation in the <appserver_install>/jms/openjms/docs directory for more details on this file.

The code sample below shows you how to add a Queue and a Topic connection factory.

- You can add any number of factories as necessary for your application.
- Make sure that each object you add has a unique name so as to avoid being overwritten in JNDI. The name has to be unique among multiple instances.
- You must have at least one TopicConnectionFactory and one QueueConnectionFactory in the embedded scheme.
- If you do not specify a port for the connector, the default port will be used. Refer to the OpenJMS documentation for details at <appserver_install>/jms/openjms/docs.

```
<Configuration>

  <ServerConfiguration host="localhost" embeddedJNDI="false" />

  <JndiConfiguration>
    <property name="java.naming.factory.initial"
              value="com.inprise.j2ee.jndi.CtxFactory" />
    <property name="java.naming.provider.url"
              value="serial://" />
  </JndiConfiguration>

  <Connectors>
    <Connector scheme="embedded">
      <ConnectionFactories>
        <QueueConnectionFactory name="jms/EmbeddedQueueConnectionFactory" />
        <TopicConnectionFactory name="jms/EmbeddedTopicConnectionFactory" />
      </ConnectionFactories>
    </Connector>

   <Connector scheme="tcp">
      <ConnectionFactories>
        <QueueConnectionFactory name="jms/TcpQueueConnectionFactory" />
        <TopicConnectionFactory name="jms/TcpTopicConnectionFactory" />
        <QueueConnectionFactory name="jms/qcf" />
        <QueueConnectionFactory name="jms/QueueConnectionFactory" />
        <QueueConnectionFactory name="jms/xaqcf" />
        <TopicConnectionFactory name="jms/tcf" />
        <TopicConnectionFactory name="jms/TopicConnectionFactory" />
        <TopicConnectionFactory name="jms/xatcf" />
      </ConnectionFactories>
    </Connector>

    <Connector scheme="rmi">
      <ConnectionFactories>
        <QueueConnectionFactory name="jms/qcf" />
        <QueueConnectionFactory name="jms/QueueConnectionFactory" />
        <QueueConnectionFactory name="jms/xaqcf" />
        <TopicConnectionFactory name="jms/tcf" />
        <TopicConnectionFactory name="jms/TopicConnectionFactory" />
        <TopicConnectionFactory name="jms/xatcf" />
      </ConnectionFactories>
    </Connector>

  </Connectors>
```

Note

When configuring JMS resource objects in Borland deployment descriptors, in preparation for an application JNDI lookup, be sure you add a serial:// prefix to the value of jndi-name elements. For example, serial://jms/q. OpenJMS resource objects are deployed independent of DAR files. They are bound directly under JNDI with a serial:// prefixed name upon BAS partition start. Applications that perform a JNDI lookup of an OpenJMS resource object must use the serial:// prefix to resolve the object.

Connection Modes in OpenJMS

OpenJMS supports multiple ways for the client to access it—using Embedded, TCP, and RMI connectors.

Use the Embedded connector when OpenJMS is installed as an in-process service. Specify all the connection factories needed locally under the embedded connector section in the `openjms.xml` file. You can take advantage of the 2PC optimization only if you use OpenJMS as a partition level (in-process) service using the embedded or RMI connector. In embedded mode, the JMS clients access the JMS server using local Java calls and use embedded queue/connection factories. These factories provide optimal ways to avoid the cost of TCP/IP.

If you use OpenJMS as an out-of-process service, you must use the RMI or TCP connectors. The RMI connector in AppServer is configured to use RMI-over-IIOP and hence can carry transaction context from client to JMS server. It has the capability to optimize local calls when the clients are co-located with the OpenJMS server and so makes it more efficient. The TCP connector, which is based on custom protocol doesn't carry transaction context.

Important

You can disable the TCP or RMI connectors when not using them. Do not disable the Embedded connector even if you are not using it. The Embedded connector is used internally to control service management (start, stop etc.) of OpenJMS as a part of the AppServer partition level service.

Changing the Datasource for OpenJMS

By default, when the OpenJMS service starts, it looks in the `partition.xml` file for the location of the datasource where the OpenJMS messages will be persisted. This datasource entry must exist in a DAR file. In the event that the entry is not found in the DAR file, OpenJMS service will default to the datasource mentioned in the `openjms.xml` file. If you want OpenJMS to use the datasource configured in the `partition.xml` file only, you can comment the <DatabaseConfiguration> entry out in `openjms.xml` file. This way, if the datasource is not found, you will see an error message. You can change the datasource and make it point to the same datasource which your J2EE application uses. To change the datasource:

1. Right-click on the OpenJMS service in the left pane of the Borland Management Console.
2. Select **Properties** from the drop-down menu.
3. Enter the path to the datasource in the **Name** text box in the **General** tab.
4. (Optional) Uncheck the “Clean messages on startup” checkbox if you do not want the previously stored (undelivered) messages to be deleted when you restart the partition. The messages that get delivered, are automatically deleted from the database. The messages that could not be delivered for any reason remain in the database. It is these messages that get cleaned up when you check this box. This checkbox is checked by default.

You must also specify the right database driver in the `openjms.xml` configuration file. To access this file, right-click on the OpenJMS service and select Properties from the menu. Click on the `openjms.xml` tab in the properties pane.

Creating Tables for OpenJMS

If you choose to use a database other than JDataStore, you must create appropriate tables in the database before using it. In case of JDataStore, the tables are pre-created. Use the scripts provided by OpenJMS to create tables in other databases. You can find these scripts in `<bas_install>\jms\openjms\config\db` directory. See the OpenJMS User's Guide for details. This guide is available in `<appserver_install>/jms/openjms/docs` directory.

Configuring Datasource to Achieve 2PC Optimization

By using OpenJMS as a partition-level service you can achieve the two-phase commit optimization. OpenJMS can be configured to persist data in any relational database. By default, AppServer uses the partition's JDataStore database for persisting JMS messages. You can change the default datasource and make it point to the same datasource which your J2EE application uses. (See "Changing the Datasource for OpenJMS" for details on how to change the datasource.) This way, OpenJMS and your application will use a single transaction resource and you will be able to avoid the two-phase commit.

Important

When there are multiple messaging applications in the partition, and each of them uses a distinct datasource for application data, the two-phase commit optimization is not possible for each of those applications. The two-phase commit optimization will work only with the one that has the same datasource as OpenJMS. OpenJMS can persist data in only one datasource for all applications in a partition, regardless of the number of applications in that partition. So, if a partition has multiple applications and each application stores its data in its own separate database, you can make OpenJMS datasource to point to only one of those databases. The application that stores its data in this database will achieve the 2PC optimization.

Configuring Security with OpenJMS

The authentication and security features that come with OpenJMS version 0.7.6. work with the following AppServer configurations:

1. OpenJMS authentication with TCP Connector
2. OpenJMS authentication with VBJ-based RMI Connector

Note

HTTPS and TCPS connections are not supported in AppServer 6.7 release.

The following XML code in `openjms.xml` file shows an example of how to turn on security for configurations 1 only. It provides a list of authenticated users:

```
<SecurityConfiguration securityEnabled="true"/>
  <Users>
    <User name="admin" password="admin"/>
    <User name="j2ee" password="j2ee"/>
  </Users>
```

3. OpenJMS authentication with VBJ-based RMI connector using AppServer security

   Refer to the `<appserver_install>/examples/security/Readme.html` document on how to configure security for this configuration.

Refer to the OpenJMS documentation in `<appserver_install>/jms/openjms/docs` directory for details on how to use security in OpenJMS.

Specifying Partition Level Properties for OpenJMS

To integrate OpenJMS as a partition level service in AppServer, it is introduced as a new service in the partition's configuration. The following properties are for OpenJMS in the `partition.xml` file. This file is located in `<appserver_install>/var/domains/base/configurations/<my_config>/mos/<openjms_partition>/adm/properties` directory.

The following code in the `partition.xml` file creates OpenJMS as a partition level service:

```
- <service name="jms"
        runas.propstorage="management_runas.properties"
        version="6.7" description="JMS Service based on OpenJMS(tm)
           version 0.7.6.1"
        vendor="Borland Software Corporation"
        class="com.borland.enterprise.server.services.PartitionService"
        startup.synchronization="service_ready"
        startup.service_ready.max_wait="0"
        shutdown.synchronization=""
        shutdown.phase="1">
    <properties lifecycle.class="com.borland.jms.JmsPartitionService"
        openjms.configfile="adm/openjms/conf/openjms.xml"
        openjms.home="../../../../../../../../jms/openjms"
        openjms.clean_messages_on_startup="true"
        openjms.datasource="serial://datasources/OpenJmsDataSource"
        openjms.sql_file="adm/openjms/conf/openjms.sql"
        openjms.datasource_lookup_interval="1"
        openjms.max_datasource_lookup_retries="1" />
    </service>
```

The properties are described in the table below:

Property Name	Description	Default Value
`lifecycle.class`	Used to add an inprocess JMS service. In case of OpenJMS(tm) this property has to have a value of com.borland.jms.JmsPartitionService. The reflection based code of the partition launcher dynamically detects if the class specified in this property is in the Java CLASSPATH. If it finds it, it tries to load and start the service.	`com.borland.jms.JmsPartitionService`
`openjms.configfile`	This property specifies the location of the configuration file. The location is relative to the current working directory of your partition. This file is the central place where configuration of OpenJMS(tm) is stored. This file is needed for the embedded OpenJMS(tm) service to come up with a AppServer partition. This file also contains the list of JNDI objects (Queues, and Topics) that need to be created as part of the OpenJMS(tm) service startup.	`adm/openjms/conf/openjms.xml`
`openjms.home`	This property specifies where OpenJMS(tm) is installed. OpenJMS uses the value specified here to locate various resources.	`<AppServerInstallRoot>/jms/openjms`

Property Name	Description	Default Value
openjms.clean_messages_on_startup	This property indicates whether to clean up the database tables containing JMS messages across partition restart. Currently, this property is only available for JDataStore. For other databases, you must delete messages manually.	true
openjms.datasource	This property specifies the JNDI name of the datasource that is used to persist messages in OpenJMS(tm). If this datasource is the same as the one that your application uses, then the JDBC connection pool will be shared among them and a single JDBC connection will be used both to persist messages and provide data access to your application, thereby avoiding a need for 2PC. If the specified datasource is not available in the JNDI namespace at the startup time, the startup code will use the properties openjms.datasource_lookup_interval, and openjms.max_datasource_lookup_retries described below, to make multiple attempts to access the target datasource. Despite that, if the lookup fails, the initialization code will internally construct a datasource from the information specified in the configuration (openjms.xml) file. **Note:** The startup code will use the information from openjms.xml file only if the user specified datasource is not available. If the datasource is pre-deployed or available in JNDI , RDBMS configuration in the openjms.xml file will be ignored.	serial://datasources/JDSLocal
openjms.sql_file	This property is used to specify the file that contains the SQL statements to drop, and create database tables. These tables are used by OpenJMS(tm) for message persistence.	adm/openjms/conf/openjms.sql
openjms.datasource_lookup_interval	This property specifies the duration between successive datasource lookup attempts. See the property openjms.datasource property above.	1 second
openjms.max_datasource_lookup_retries	This property specifies the number of attempts to lookup for the datasource before attempting to use the default datasource.See the property openjms.datasource property above.	5 seconds
openjms.recreate_database_on_startup	This property when set will cause re-creation of database across each startup of the service. This is useful for cases where previous messages are not needed in the subsequent runs (for example, while testing).	false
openjms.database.softcommit	This property is only applicable when JDataStore is used to persist JMS messages. This property provides improved performance during commit process, but with lack of recoverability in some rare failure scenarios. See JDataStore documentation for more details.	true

Property Name	Description	Default Value
openjms.database	This property only applies to JDS and is used to specify the name of the JDS database.	openjms.jds
openjms.use_bes_transactions	OpenJMS starts a transaction before it dispatches messages. It uses the transaction service that is part of the partition which contains OpenJMS. If no transaction service is available in the partition, one from the Smart Agent domain is selected. Use this property when you use OpenJMS during transactions. This property does not affect the messaging applications that don't involve transactions. However, to avoid an extra cost of transaction startup and propagation turn this property off.	true

OpenJMS Topologies

Important

If you have two OpenJMS services with both of them using the TCP connector, make sure you have different port numbers specified for them in the `openjms.xml` file. To open this file, right-click on the OpenJMS service in the Borland Management Console and select Properties from the drop down menu. Click on the openjms.xml tab in the properties pane.

OpenJMS can be configured to run in the following two topologies:

- **Server shared mode**—where the OpenJMS service is hosted in a dedicated partition with other services in that partition disabled. It is available as a shared service to all partitions that are in the same osagent domain as the configuration in which OpenJMS partition resides. The remote clients can access OpenJMS via an RMI or TCP connector. Since OpenJMS needs a naming service to bind the JNDI objects that are specified in its configuration file, the Transaction and Naming services must be enabled in the partition that hosts OpenJMS or they should be available in the SmartAgent domain.
- **Embedded Service mode**—where OpenJMS is run as an embedded service in each of the partitions in the configuration. Each partition uses the embedded (intra-virtual machine) connector of OpenJMS instead of the TCP or RMI connector. The JMS clients use the embedded queue/topic connection factories. These factories provide optimal ways to avoid the cost of TCP/IP. Even though JMS clients can use the RMI connector in this mode, to achieve maximum performance it must use the local (embedded) connector.

Note

If multiple OpenJMS service instances are running in a single SmartAgent domain of AppServer, there will be no database sharability or automatic failover to a running instance. This is because there is no support for clustering for OpenJMS.

Using Message Driven Beans (MDB) with OpenJMS

For an AppServer partition to support MDBs, the MDB must be able to access a JMS server. To make the MDB access the OpenJMS server, make sure that:

1. OpenJMS is installed and enabled as an in-process service in your partition or available in the domain. Right-click on the OpenJMS service and select Start from the menu to enable the service.

2 The resource references are properly configured in the `ejb-jar.xml` file to point to the right type of connection factory.

Important

If your MDB needs transactional access, you must use Embedded or RMI connection factories with your MDB so as to support transaction propagation.

Other JMS providers

Borland AppServer supports the SonicMQ 6.0/6.1 and WebSphereMQ 5.3/6.0 JMS providers. For information on how to integrate SonicMQ with AppServer see “Integrating SonicMQ into Borland AppServer.” For information on how to integrate WebSphereMQ with AppServer see “Integrating WebSphereMQ into Borland AppServer (BAS).”

26

Integrating SonicMQ into Borland AppServer

This document provides the steps to enable Borland AppServer (AppServer) to work with an independent installation of a SonicMQ 6.0/6.1 JMS provider. Both SonicMQ versions 6.0 and 6.1 are JMS 1.1 compliant.

Note

You must purchase SonicMQ separately. It is not bundled with AppServer 6.7.

Installing SonicMQ

Install SonicMQ to a location independent of the AppServer installation. Make sure that the Management features are installed so that you can manage the SonicMQ services through the Sonic Management Console.

Configuring SonicMQ Administered Objects in AppServer

You must define the JMS administered objects accessed through JNDI in the Borland proprietary DAR modules. The Borland Deployment Descriptor Editor(DDEditor) tool in AppServer allows you to create the administered objects in a DAR module. See “Setting Admin Objects Using Borland Deployment Descriptor”.

See the *SonicMQ V6.1 Configuration and Management Guide* for information on all the properties that can be configured for administered objects using the Sonic JMS Administered Objects tool.

See “Using JMS” for a description of AppServer related properties that can be applied to JMS connection factory object definitions in DAR modules.

Resolving SonicMQ library modules in the AppServer environment

SonicMQ 6.0/6.1 client libraries, `sonic_Client.jar` and `sonic_XA.jar`, and their dependent libraries must be loaded by AppServer for deployment of J2EE application(s) that wish to access a SonicMQ server.

The suggested approach for enabling SonicMQ client libraries in AppServer is to apply the following updates to JMS related configuration files located under `<AppServer>/bin`:

– Edit the `sonic.config` file and set the value of `jms.home` to the root directory of the external SonicMQ installation. For example:

```
set jms.home=C:/SonicMQ/V61
```

– Edit the `jms.config` file. Uncomment the statement to include `sonic.config`. Make sure that the include statements for other JMS providers are commented out:

```
#include $var(installRoot)/bin/tibco.config
#include $var(installRoot)/bin/openjms.config
include $var(installRoot)/bin/sonic.config
```

This allows SonicMQ client libraries to be resolved by all AppServer partitions and by J2EE client applications run by AppServer appclient tool.

Configuring Automatic Queue Creation for SonicMQ Queues deployed to AppServer

When a DAR module containing definition of SonicMQ JMS Queues is deployed to a partition, AppServer can be configured to automatically create the JMS Queues in the target SonicMQ server. Certain SonicMQ management libraries need to be available to AppServer for automatic JMS Queue creation to occur. These libraries must be loaded from the partition's classpath. This can be achieved by updating AppServer configuration files `sonic.config` and `jms.config` as described above. Additionally, the following steps must be performed:

– Make sure that the naming service definition in the partition's configuration file, `partition.xml`, has `jns.auto-create-queues` property set to true as follows:

```
<service name="visinaming"
        runas.propstorage="management_runas.properties" version="6.7"
      description="Naming Service" vendor="Borland Software Corporation"
     class="com.borland.enterprise.server.services.naming.NamingService"
         startup.synchronization="service_ready"
startup.service_ready.max_wait="0"
         shutdown.synchronization="" shutdown.phase="1">
         <properties jns.name="namingservice"
         jns.auto-create-queues="true">
         </properties>
</service>
```

– Update `partition-server.config` file of the partition to ensure target SonicMQ server can be located:

  a Open the Management Console.

  b Switch to the Installations view by clicking on the Installations icon on the extreme left side of the console.

  c In the left pane, navigate to the partition for which you want to make the change. The General Properties page for the partition will open in the right pane.

  d Click on the Files tab at the bottom of the right pane.

e Select partition-server.config in the lower left pane.

f Scroll to the end of the file and enter the following for only the properties you want to change:

```
vmprop <property_name>=<value>
```

You can do this for the following 5 properties:

Property	Default Value
sonicmq.domainName	domain1
sonicmq.brokerURL	tcp://localhost:2506
sonicmq.user	Administrator
sonicmq.pwd	Administrator
sonicmq.brokerName	/Brokers/Broker1

g Save edits and restart the partition.

Note

SonicMQ Server must be active prior to deployment of a DAR module with SonicMQ JMS Queues in order to successfully auto-create the queues.

27

Integrating WebSphereMQ into Borland AppServer (BAS)

This document provides the steps to enable Borland AppServer (AppServer) to work with an independent installation of a WebSphereMQ 5.3/6.0 JMS provider.

Note

You must purchase WebSphereMQ separately. It is not bundled with AppServer 6.7.

Supported Versions

Both WebSphereMQ 5.3 and 6.0 are certified to work with the product.

WebSphereMQ Configuration

To configure WebSphereMQ:

WebSphereMQ 5.3

Out of the box installation of WMQ 5.3 does not support JMS 1.1 APIs. To take advantage of the JMS1.1 features, fix pack 06 (CSD06) or above should be installed on top of WMQ 5.3 installation.

The “standard” WebSphereMQ Client supports only the local (i.e. one-phase commit) transactions, managed by the queue manager to which the client application is connected. To support distributed transactions(2PC), you must install WebSphereMQ Extended Transactional Client.

The WebSphereMQ Extended Transactional Client is a fee-based feature of WebSphereMQ version 5.3. It extends the WebSphereMQ capabilities by allowing WebSphereMQ client applications to participate in globally coordinated transactions. In other words, it offers two-phase commit (XA compliant) processing support to WebSphereMQ client applications so that they can participate in global transactions managed by some external transaction managers.

WebSphereMQ 6.0

The default installation of WebSphereMQ 6.0 has support for JMS 1.1 APIs.

WebSphereMQ 6.0 has built in support for distributed transactions(2PC) and MQ Extended Transactional Client installation is not required.

Configuring Admin Objects with WebSphereMQ

WebSphereMQ's admin object properties are defined in BES and can be configured graphically using the Borland Deployment Descriptor Editor. See "Setting Admin Objects Using Borland Deployment Descriptor".

For a complete list of JNDI properties and other configuration options available with WebSphereMQ 5.3, refer to WebSphereMQ's *Using Java* document published at `http://publibfp.boulder.ibm.com/epubs/pdf/csqzaw12.pdf`.

For a complete list of JNDI properties and other configuration options available with WebSphereMQ 6.0, refer the WebSphereMQ *Using Java* document published at `http://publibfp.boulder.ibm.com/epubs/pdf/csqzaw13.pdf`.

Locating WebSphereMQ Library modules at runtime

WMQ 5.3 Client libraries need to be loaded in BAS partition for deployment of J2EE application(s) that wish to access a WMQ5.3 server.Following are the full set of libraries required by a BAS partition

- `com.ibm.mq.jar`
- `com.ibm.mqjms.jar`
- `com.ibm.mqbind.jar`
- `com.ibm.mqetclient.jar` (this jar is a part of WMQ Extended Transactional Client installation)

One approach in making these available to BAS would be to deploy them to the BAS partition hosting the J2EE application. However, a better approach is to update JMS related configuration files located under `<BAS_install>/bin`:

- Edit the `wmq53.config` and set the value of `jms.home` to the root directory of the external WMQ5.3 installation.
- Edit the `jms.config` file. Uncomment the include statement to include `wmq53.config`. Make sure that the include statements for other JMS providers are commented out:

```
#include $var(installRoot)/bin/tibco.config
#include $var(installRoot)/bin/openjms.config
#include $var(installRoot)/bin/sonic.config
include $var(installRoot)/bin/wmq53.config
```

This allows WMQ5.3 client libraries to be resolved by all BAS partitions and by J2EE client applications run by BAS tool appclient.

WebSphereMQ 6.0

WebSphereMQ 6.0 Client libraries need to be loaded in BAS partition for deployment of J2EE application(s) that will be accessing a WebSphereMQ 6.0 server. A full set of libraries required by a BAS partition are:

- `com.ibm.mq.jar`
- `com.ibm.mqjms.jar`
- `dhbcore.jar`
- `com.ibm.mqetclient.jar` (Extended transactional client)

One approach in making these available to BAS would be to deploy them to the BAS partition hosting the J2EE application. However, a better approach is to update JMS related configuration files located under `<BAS_install>/bin`:

- Edit the `wmq60.config` file and set the value of `jms.home` to the root directory of the external WebSphereMQ 6.0 installation.
- Edit the `jms.config` file. Uncomment the include statement to include `wmq53.config`. Make sure that the include statements for other JMS providers are commented out:

```
#include $var(installRoot)/bin/tibco.config
#include $var(installRoot)/bin/openjms.config
#include $var(installRoot)/bin/sonic.config
include $var(installRoot)/bin/wmq60.config
```

This allows WebSphereMQ 6.0 client libraries to be resolved by all BAS partitions and by J2EE client applications run by BAS tool appclient.

28

Using JACC

The Java Authorization Contract for Containers (JACC) specification defines a contract between a J2EE application server and an authorization policy provider. All J2EE application containers, web containers, and enterprise bean containers are required to support this contract. The contract defined by this specification is divided into three subcontracts. Taken together, these subcontracts describe the installation and configuration of authorization providers such that they will be used by containers in performing their access decisions.

JACC Contracts

The three subcontracts are the Provider Configuration Subcontract, the Policy Configuration Subcontract, and the Policy Decision and Enforcement Subcontract.

Provider Configuration Subcontract

The Provider Configuration Subcontract defines the requirements placed on providers and containers such that Policy providers may be integrated with containers.

Policy Configuration Subcontract

The Policy Configuration Subcontract defines the interactions between container deployment tools and providers to support the translation of declarative J2EE authorization policy into policy statements within a J2SE Policy provider.

Policy Decision and Enforcement Subcontract

The Policy Decision and Enforcement Subcontract defines the interactions between container policy enforcement points and policy decisions required by J2EE containers.

How the JACC-based authorization works

JACC allows the EJB and Web containers in an application server to interact with third party authorization providers to make authorization decisions when a J2EE resource is accessed. The Web and EJB containers in a J2EE application server use JACC-

compliant authorization providers to restrict client access to the resources and services. The providers do this based on the policy information propagated to them by the deploy tool during application deployment. The provider stores this information in its repository for use when making the authorization decisions. Authorization decisions are made by the provider based on whether the principal (user) belongs to a role that has the necessary privileges to access a particular resource.

When an application is being deployed, the AppServer does the following:

1. Create a unique contextID that uniquely identifies the module that is being deployed.
2. Build the PolicyConfiguration with the set of Permissions that will be required to access each resource of the module.
3. Propagate the security policy information to the provider through the JACC APIs.

When a client/user makes a request to access an EJB method or a servlet or URL:

1. The EJB container or the Web container creates an appropriate permission object and the ProtectionDomain object containing the principals of the caller.
2. The container then calls the Policy.implies method of the java.security.Policy object implemented by the provider and passes the two objects to the provider.
3. The provider makes the decision based on the policy information it has stored (by checking its principal-to-role map) and returns a boolean value to the container.
4. If the role to which the principal belongs has access permissions to the resource, the implies method returns true and the user is allowed to access the resource by the container. Otherwise, it returns a false and the user is denied access to the resource.

Configuring JACC provider in Borland AppServer

The JACC provider in AppServer implements the standard java.security.Policy object specified in the Provider Configuration Subcontract section, which it uses to make the access decisions. The JACC provider also implements the PolicyConfigurationFactory class and the PolicyConfiguration interface, which enables deployment tools to propagate all security elements to the provider during application deployment.

The following properties control the installation of the AppServer JACC provider:

Property Name	Description	Default Value
`javax.security.jacc.policy.provider`	Specifies the policy implementation class that will be used by the application server for policy replacement.	`com.borland.security.jacc.provider.BESJACCPolicy`
`javax.security.jacc.PolicyConfigurationFactory.rovider`	Specifies the providers PolicyConfigurationFactory implementation class.	`com.borland.security.jacc.provider.BESPolicyConfigurationFactory`

Configuring a JACC provider using AppServer Management Console

You can configure the JACC provider using the AppServer Management Console or you can configure the JACC provider properties in the `partition_server.config` file.

To configure the properties using the AppServer Management Console:

1. Select the Partition name in the left pane of the console.
2. Right-click on the partition name and select Properties from the resulting menu.
3. The Partition Properties page will open.
4. Click on the Security tab.
5. Configure the two properties in the JACC Properties box.

Configuring a JACC provider through the configuration file

To configure the JACC provider properties in the `partition_server.config` file:

1 Go to the following directory:

```
<install_dir>\var\domains\base\configurations\j2eeSample\mos\
<partition_name>\adm\properties
```

2 Open the `partition_server.config` file.

3 Locate the following lines:

```
#JACC provider configuration
vmprop javax.security.jacc.policy.provider=com.borland.security.jacc.
   provider.BESJACCPolicy
vmprop javax.security.jacc.PolicyConfigurationFactory.provider=com.borland.
   security.jacc.provider.BESPolicyConfigurationFactory
```

4 Configure the properties as desired.

Note

If you leave these properties blank, the JACC provider will not be enabled and the system will fall back to security framework as existed in the previous AppServer releases.

Enabling/Disabling the JACC provider

You have the option of using one of the following:

- Configure AppServer security as a JACC provider (this is the default setting)
- JACC disabled in AppServer security—the underlying security mechanism is the same as it existed in the previous releases of AppServer
- Configure AppServer to use external JACC providers

By default, when you install the AppServer you will receive Borland VisiSecure as the JACC provider. The JACC provider shipped with AppServer is compliant with all JACC APIs and implements the Provider Configuration subcontracts specified in the JACC specification.

The default settings for security properties in the Management Console of the AppServer are set such that you can use AppServer security with the JACC APIs. If you choose to not use the JACC with the AppServer security provider, you must clear the security properties in the management console.

Alternatively, you can extend your security infrastructure by plugging in a third party JACC-based security provider into AppServer. If you choose to use an external provider, you must enter the appropriate values for the properties in the JACC Properties box in the Partition Properties dialog box. Also, make sure that the external JACC provider related jar files are deployed to the partition as library modules.

Configuring external JACC providers

Any JACC-compliant external provider can be plugged into the AppServer. The provider implementation and configuration should follow the guidelines as mentioned below:

- The provider should provide an implementation for `java.security.Policy` and the configuration has to happen correctly through the admin console or the configuration file as discussed in the earlier sections.

- The provider should provide an implementation for the `PolicyConfigurationFactory` and the configuration has to happen correctly through the admin console or the configuration file.
- All the provider dependent jar files should be deployed to the partition as library modules.

An example which demonstrates how a provider should be implemented and configured with BES is shipped with the product. Please refer to `<install dir>/examples/security/jacc` for details.

You can configure an external JACC provider using the Borland Management Console or you can configure the security properties in `partition_server.config` file.

29

Using ADLoginModule in BAS

Active Directory is Microsoft's implementation of directory service for the Windows platform. It provides the means to manage the identities , resources, and the relationships between them, all of which make up the network environment. ADLoginModule is a new LoginModule bundled with BAS which inherits from the LDAPLoginModule and specifically works with Active Directory as backend user store.

How ADLoginModule works

User Principal Name

Different from the LDAPLoginModule, by default, ADLoginModule uses user principal name (UPN) to bind to Active Directory Server, thus performs the authentication. UPN is formed by combining object name with the Fully Qualified Domain Name (FQDN)—*objectname@QFDN*. For example, for `user1` in domain `abc.def.net`, the user principal name `user1@abc.def.net` will be used as the security principal (instead of the DN as in LDAPLoginModule).

Authentication

Authentication process includes two steps:

1. Validate the username/password pair against the user backend store
2. Populate user attributes, which will be used for authorization at latter stage

During the first step, ADLoginModule forms the User Principal Name with the provided username and the domain name option. With the password provided by the user, ADLoginModule binds to the Active Directory. A successful binding operation means the user is authenticated by the Active Directory server.

After the successful authentication, ADLoginModule gets the Distinguished Name (DN) for the user entry from the Active Directory, and populates the designated set of attributes (from options specified in JAAS configuration). In doing so, ADLoginModule searches from SEARCHBASE context and looks for entry satisfying the filter “userPrincipalName=*UPN*”.

With the DN information in hand, ADLoginModule populates the required attributes of that entry based on options specified in JAAS configuration.

Configuring ADLoginModule

A new option DOMAINNAME is added specifically for ADLoginModule, which indicates the domain to which this entity is authenticated against. A sample configuration looks as follows:

```
adrealm {
    com.borland.security.provider.authn.ADLoginModule required
    INITIALCONTEXTFACTORY=com.sun.jndi.ldap.LdapCtxFactory
    PROVIDERURL="ldap://testing.net"
    DOMAINNAME=abc.def.net
    SEARCHBASE="cn=users,dc=abc,dc=def,dc=net"
  };
```

with this configuration, the user will be authenticated against Active Directory Server at host `testing.net`, and for the domain `adc.def.net`. The user entry will be searched from SEARCHBASE “cn=users,dc=abc,dc=def,dc=net”.

Detailed Configuration Options

Similar to LDAPLoginModule, ADLoginModule can be configured with following entry inside JAAS configuration file:

```
<realm-name> {
     com.borland.security.provider.authn.ADLoginModule
        authentication-requirements-flag
     INITIALCONTEXTFACTORY=connection-factory-name
     PROVIDERURL=backend-url
     DOMAINNAME=[domain name as in DNS-mapped format, for example, abc.def.net]
     SEARCHBASE=search-start-point
     USERATTRIBUTES=attribute1, attribute2, ...
     USERNAMEATTRIBUTE=attribute
     QUERY=dynamic-query
  };
```

The detailed description for the options are summarized below:

Property Name	Description
INITIALCONTEXTFACTORY	The InitialContextFactory class that is used by JNDI to bind to LDAP.
PROVIDERURL	The URL to the directory server of the form ldap://<servername>:<port>. This attribute is mandatory.
DOMAINNAME	Anew attribute for Active Directory, indicates the domain name for the user. This is the recommend way to perform login with AD though is not mandatory. For login using DN, USERNAMEATTRIBUTE must be set to “DN”.
SEARCHBASE	Explicitly set the search base for the directory to lookup. This attribute is optional, if this is not specified, the search will be performed from the root context of domain.
USERATTRIBUTES	Comma-separated list of attributes that will be retrieved and stored for an authenticated user. This attribute is optional, if this is not specified all the attributes for the entry will be populated. Refer to LDAPLoginModule in the Security User Guide for more information.
USERNAMEATTRIBUTE	When user is being authenticated to the system—either through CallbackHandler or IdentityWallet, a name-password pair is required. This attribute defines the meaning of “name”—a username within a domain or a DN. This attribute is optional—if this is not specified (which is the default case), DOMAINNAME option must be specified, and user's input is treated as the username within the domain—a UPN is formed as <username>@<domainname>. On the other hand, if DN is used for login, this option must be set to “DN”. In this case, the input from user will be treated as the DN directly.
QUERY	Provides a mechanism to dynamically query the directory server for other information and represent the results as attributes. Refer to LDAPLoginModule in the Security User Guide for more information. This attribute is optional.

30

Using JAXR

This document describes the Java API for XML Registries (JAXR). JAXR is part of J2EE 1.4 specification. It gives the J2EE developer a common standard API to access various XML registries particularly used in web services. The JAXR specification from Sun is available at `http://java.sun.com/xml/jaxr/index.jsp`.

The Borland AppServer (BAS) integrates Apache jUDDI and Apache scout to provide a UDDI registry and JAXR compliance. Apache jUDDI is an open source Java implementation of the Universal Description, Discovery, and Integration (UDDI) specification for Web Services.

JAXR specification defines two types of providers each with a different Capability Level. Each provider offers a different level of support for interacting with the two popular registry specifications, UDDI and ebXML. A type 0 provider offers support for accessing UDDI registries and type 1 provider supports access to both UDDI AND ebXML registries.

Apache scout, which is integrated with BAS, is a type 0 jUDDI JAXR provider. It adapts the jUDDI client to standard JAXR API.

Using JAXR in BAS

Before you use the JAXR APIs, you must set the classpath and system properties settings for the running JVM. You must deploy the `juddi.ear` to a BAS partition. The `juddi.ear` file is located at a BAS repository, `<BAS_home>/var/repository/archives/ears`.

You must include the following libraries that are required by the BAS partition to host the `juddi.ear`:

- `<BAS_home>/lib/scout.jar`
- `<BAS_home>/lib/juddi.jar`
- `<BAS_home>/lib/axis/axis.jar`
- `<BAS_home>/lib/axis/commons-discovery-0.2.jar`

You can include the jar files as library in to your J2EE application (ear, jar or war files) or you can deploy the jar files as a static library into the BAS partition.

If you are running JAXR in a Java client application, all the above mentioned libraries and the libraries below must be included in classpath:

- `<BAS_home>/lib/axis/commons-logging.jar`
- `<BAS_home>/lib/axis/asrt.jar`

System Property

To use the JAXR Provider for UDDI, the name of the ConnectionFactory implementation class must first be specified by setting the System Property `javax.xml.registry.ConnectionFactoryClass` to `org.apache.ws.scout.registry.ConnectionFactoryImpl`. By default, BAS partition has this property automatically set to its JVM. If you are an application user you do not need to set this property. If you are running JAXR in a standalone java application, this system property must be set to point to the JVM. Failure to specify this will result in the value defaulting to `com.sun.xml.registry.common.ConnectionFactoryImpl`, which will not be found. This will result in a `JAXRException` when the `ConnectionFactory.newInstance()` method is called. The BAS JAXR Provider for UDDI does not support lookup of the ConnectionFactory via JNDI.

JAXR Connection Properties

Connection specific properties must be set to ConnectionFactory before getting Connection from the factory. See the JAXR specification for a detailed list of the properties and their descriptions. The following is a subset of properties that are required to get a connection:

Property	Description
javax.xml.registry.queryManagerURL	The URL of the jUDDI registry's inquiry API for UDDI. This url will be of the form: http://<hostname>:<port>/juddi/inquiry. This property is required.
javax.xml.registry.lifeCycleManagerURL	The URL of the UDDI registry's publish API for UDDI. This url will be of the form: http://<hostname>:<port>/juddi/publish.
javax.xml.registry.authenticationMethod	The method of authentication to use when authenticating with the registry. This may take one of two values, `UDDI_GET_AUTHTOKEN` or `HTTP_BASIC`. The default value is `UDDI_GET_AUTHTOKEN` if none is specified.

BAS JAXR Example code

The following example shows you how to create a connection using JAXR API:

```
import javax.xml.registry.Connection;
import javax.xml.registry.ConnectionFactory;
import java.util.Properties;

public class TestConnection
{
   public static void main(String[] args)
   {
      Properties prop = new Properties();
      try
      {
         String queryurl = "http://localhost:8080/juddi/inquiry";
         prop.setProperty("javax.xml.registry.queryManagerURL", queryurl);
         prop.setProperty("javax.xml.registry.lifeCycleManagerURL", queryurl);
         ConnectionFactory factory = ConnectionFactory.newInstance();
         factory.setProperties(prop);
         Connection con = factory.createConnection();
         if(con == null)
            System.out.println("No Connection");
      }
      catch(Exception e)
      {
         e.printStackTrace();
      }
   }
}
```

31

Using the Scheduler Service

Borland AppServer 6.7 (AppServer) supports J2EE 1.4 compliant EJB timer service. In AppServer this service is known as the Scheduler Service. The AppServer Scheduler Service is based on Quartz. Refer to the EJB 2.1 specification for generic information on EJB Timer Service. To obtain Quartz-related documentation, go to `http://www.opensymphony.com/quartz/documentation.action`

The Scheduler Service is a partition-level service, which means that each time you create a partition, it will automatically be included as one of the partition services. The Scheduler Service can be used even when the EJB container goes down.

Configuring the Scheduler Service

You can configure some of the commonly used scheduler service properties in the AppServer Management Console. To do so:

1. Open the AppServer Management Console.
2. Double-click on the partition name whose scheduler service you want to configure to expand the node.
3. Right-click on Scheduler Service node under the partition.
4. Select Properties... from the resulting menu. The Properties dialog box will open.
5. Configure the following Scheduler Service Settings in the General tab:

   Transaction Timeout—specifies the time within which a transaction should be successful. If a transaction is not successfully completed within the time set in this field, the transaction will be marked for rollback.

   Max Redelivery Count—specifies the number of attempts that the Scheduler Service will make to redeliver a message to an application whose transaction (of which the scheduler event was a part) has been rolled back.

   Clean events on startup—If this checkbox is checked, all the jobs and triggers will be deleted from the database when the partition is (re)started. It is applicable only if you specify JobStoreCMT to persist the scheduler events. This option is currently supported only for JDataStore.

   Soft Commit—Check this box if you want to enable soft commit. With soft commit enabled, the operating system cache can buffer file writes from committed

transactions. Soft commit improves performance, but cannot guarantee the durability of the most recently committed transactions.

6 Click on the Quartz tab to bring it forward.

7 Configure the following properties:

Maximum number of threads—Specifies the maximum number of threads in a thread pool

Job Store Type—The default choice in the drop-down menu is Memory. This allows you to store scheduler events in-memory. Select JDBC(CMT) from the menu if you want to persist events in a database.

If you select JDBC(CMT) as your Job Store Type, you must configure the following in the Settings for Job Store box:

Database—Select a database from the drop-down menu

Container Managed DataSource—Specifies the URL for the container managed datasource. Please see the Quartz documentation for details on Container Managed DataSource.

Non Container Managed DataSource—Specifies the URL for non-container managed datasource

8 To set more properties, click on the Advanced... button. The Scheduler (Quartz) Properties page will open. You can configure additional properties here.

Using JDataStore to persist scheduler events

The AppServer Scheduler Service can be configured to persist data in any relational database. By default, AppServer uses the JDataStore for persistence. If you do not specify a database in which to store the scheduler events, AppServer defaults to storing these events in the JDataStore database.

Configuring other databases to persist scheduler events

The partition's JDataStore database is used by default to persist scheduler data. However, you can configure a different database if you want to use the database used for the application data to persist the scheduler data too. To use a database other than JDataStore, you must do the following:

- Create appropriate tables in the database using the scripts provided by Quartz for that database. These scripts are available in the Quartz footprint.
- Choose the right database driver in Quartz's configuration file located at `<partition_working_directory>/adm/scheduler/bes.properties`

Setting up for 2PC Optimization

If the timer is tied to a transaction in an application, if for any reason the transaction is rolled back, then the creation or deletion of the timer will also be rolled back along with the transaction. Similarly, if the scheduler event is delivered to an EJB as part of a transaction which is eventually rolled back, the scheduler service will attempt to redeliver the event. As per EJB 2.1 specification, there must be at least one re-delivery attempt. You can configure the number of redelivery attempts that the scheduler service makes. The default is 1. This means that when a transaction is rolledback, the Scheduler Service in AppServer will try to redeliver the message once. See “Configuring the Scheduler Service” for details on how to configure the maximum redelivery count.

To achieve 2PC optimization, you must use a common datasource to persist scheduler events and store application data which the J2EE application uses. If there are multiple applications in the partition and each of them uses a distinct datasource than 2PC optimization is not possible for each of those applications, but would work only with the one that has the same datasource as the Scheduler Service.

In some deployments, it will be necessary to use a 2PC-enabled (XA) datasource. This means the datasource JNDI name that you specify for transaction use in the `bes.properties` file will need to point to an XA-datasource in the DAR file.

Note

The transactional behavior, for example the rollback operation, is only applicable if you set the persistent store to CMT.

Partition Service properties for Scheduler Service

Quartz is introduced as a new service in the partition's configuration file, `partition.xml`. The table below lists the partition service properties that are specific to Quartz integration.

Property Name	Description	Default Value
lifecycle.class	BES partition makes it possible to dynamically add new services which can follow the life cycle of the partition process.	com.borland.jms.SchedulerPartitionService
properties.location	Specifies the location of the configuration file.	<appserverInstallRoot>\var\domains\base\configurations\<configName>\mos\<partitionName>\adm\scheduler\bes.properties
sql.location	Specifies the location of the sql scripts which you can use to create tables in the database	<partition_dir>\adm\scheduler\tables_jdatastore.sql
scheduler.clean_persistent_data_on_startup	Indicates whether to clean up the database tables containing scheduling data across partition restart.	false
scheduler.database_softcommit	This property is only relevant when JDataStore is used as the backing store for persistence. This property provides improved performance during commit process, but with lack of recoverability in some rare failure scenarios. See the JDataStore documentation for more details. This property is also used in AppServer JSS.	true
scheduler.transaction_timeout	Transaction timeout	No timeout. You can override the default by specifying the time in seconds before timeout occurs.
scheduler.auto_create_tables	Auto create Quartz tables if they do not already exist	true
scheduler.max_redelivery_count	Number of times scheduler service will try to redeliver the event in case transaction rolls back	1
scheduler.use_default_datasource	Whether or not to use our default datasource which has the JNDI URL `jdbc/quartz` and points to JDataStore database at `adm/scheduler/database/scheduler.jds`	yes

Quartz properties used in AppServer

The table below lists the properties in Quartz that are used by AppServer Scheduler Service. These properties are listed in the `<appserver-install>\var\domains\base\configurations\<configuration_name>\mos\<partition_name>\adm\scheduler\bes.properties` file. For a detailed description of these properties, see the Quartz documentation.

Property Name	Description	Default Value
org.quartz.scheduler.instanceName	Specifies the name of the scheduler	TestScheduler
org.quartz.scheduler.instanceId	Specifies the id of the scheduler	AUTO
org.quartz.scheduler.wrapJobExecutionInUserTransaction	Set this property to true to start a UserTransaction before calling execute on the job. The transaction will commit after the job's execute method completes, and the JobDataMap is updated	True
org.quartz.scheduler.userTransactionURL	Specifies the JNDI URL of Application Server's UserTransaction manager. This is only used together with JobStoreCMT	java:comp/UserTransaction
org.quartz.threadPool.class	Specifies the threadpool class	org.quartz.simpl.SimpleThreadPool
org.quartz.threadPool.threadCount	Specifies the number of threads that are available for concurrent execution of jobs. The practical value is from 1-100	30
org.quartz.threadPool.threadPriority	Specifies the thread priority. The value is between Thread.MIN_PRIORITY (1) and Thread.MAX_PRIORITY(10)	5
org.quartz.threadPool.makeThreadsDaemons	Set this property to true to make the threads in the pool created as daemon threads	True
org.quartz.jobStore.class	Specifies the JobStore class. Set this property to RAMJobStore for non-persistent and to JobStoreTx or JobStoreCMT for persistent jobs and triggers. JobStoreTx is for standalone Scheduler Service; JobStoreCMT is used if datasources are to managed by the appserver.	RAMJobStore (Memory). Currently, the AppServer Scheduler Service only supports RAMJobStore and JobStoreCMT
org.quartz.jobStore.driverDelegateClass	org.quartz.impl.jdbcjobstore.oracle.OracleDelegate for Oracle database and org.quartz.impl.jdbcjobstore.HSQLDBDelegate for JDataStore	org.quartz.impl.jdbcjobstore.HSQLDBDelegate. Currently AppServer Scheduler Service only supports JdataStore and Oracle database
org.quartz.jobStore.dataSource	Specifies the name of the container managed transaction(CMT) datasource. JobStoreCMT requires one CMT and one non-CMT datasource.	myDS
org.quartz.jobStore.nonManagedTXDataSource	Specifies the name of the non-container managed transaction datasource	myDSNoTx
org.quartz.dataSource.NAME_CMT.jndiURL	Specifies the JNDI URL of the CMT data source. NAME_CMT is the name of the CMT datasource	jdbc/Quartz
org.quartz.dataSource.NAME_NOT_CMT.jndiURL	Specifies the JNDI URL of the non-CMT data source. NAME_NOT_CMT is the name of the non-CMT datasource	jdbc/Quartz

Clustering support

The Borland AppServer provides clustering support for the Scheduler Service. For example, if you have two identical partitions with Scheduler Service enabled on both of them. If you deploy the same application on them and register a timer in one of the applications, if that partition goes down, assuming that both applications are pointing to the same database, the replica could continue to get the timer events. The AppServer Scheduler Service supports failover.

32

VisiConnect overview

J2EE Connector Architecture

In the information technology environment, enterprise applications generally access functions and data associated with Enterprise Information Systems (EIS). This traditionally has been performed using non-standard, vendor-specific architectures. When multiple vendors are involved, the number of architectures involved exponentiate the complexity of the enterprise application environment. With the introduction of the Java 2 Enterprise Edition (J2EE) 1.4 Platform and the J2EE Connector Architecture (Connectors) 1.5 standards, this task has been greatly simplified.

VisiConnect, the Borland implementation of the Connectors 1.5 standard, provides a simplified environment for integrating various EISs with the Borland AppServer (AppServer). The Connectors provides a solution for integrating J2EE-platform application servers and EISs, leveraging the strengths of the J2EE platform—connection, transaction and security infrastructure—to address the challenges of EIS integration. With the Connectors, EIS vendors need not customize integration to their platforms for each application server. Through VisiConnect's strict conformance to the Connectors, the AppServer itself requires no customization in order to support integration with a new EIS.

Connectors enables EIS vendors to provide standard Resource Adapters for their EISs. These Resource Adapters are deployed to the AppServer, each providing the integration implementation between the EIS and the AppServer. With VisiConnect, the AppServer ensures access to heterogeneous EISs. In turn, the EIS vendors need provide only one standard Connectors-compliant resource adapter. By default, this resource adapter has the capability to deploy to the AppServer.

Components

The Connectors environment consists of two major components—the implementation of the Connectors in the application server, and the EIS-specific Resource Adapter.

In the J2EE 1.4 Architecture, the Connectors is an extension of the J2EE Container, otherwise known as the application server. In compliance with the J2EE 1.4 Platform and Connectors 1.5 specifications, VisiConnect is an extension of the AppServer, and not a service in and of itself. The following diagram illustrates VisiConnect within the AppServer Architecture:

Figure 32.1 VisiConnect within the AppServer

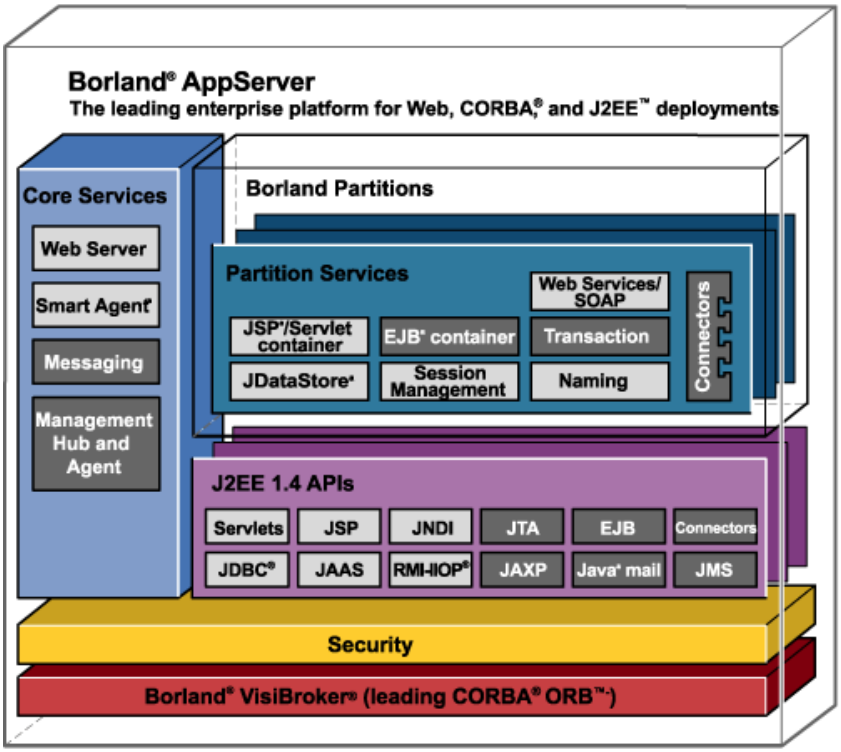

(VisiConnect is represented above by the module titled "**Connectors**.")

A Resource Adapter is a system-level driver specific to an EIS, which provides access to that EIS. To put it simply, a Resource Adapter is analogous to a JDBC driver. The interface between a Resource Adapter and the EIS is specific to the EIS. It can be either a Java interface or a native interface.

The Connectors consists of three main components:

- **System Contracts** that provide the integration between the Resource Adapter and the application server (AppServer).
- **Common Client Interface** that provides a standard client API for Java applications, frameworks, and development tools to interact with the Resource Adapter.
- **Packaging and Deployment** that provides the capacity for various Resource Adapters to plug into J2EE applications in a modular manner.

The following diagram illustrates the Connectors architecture:

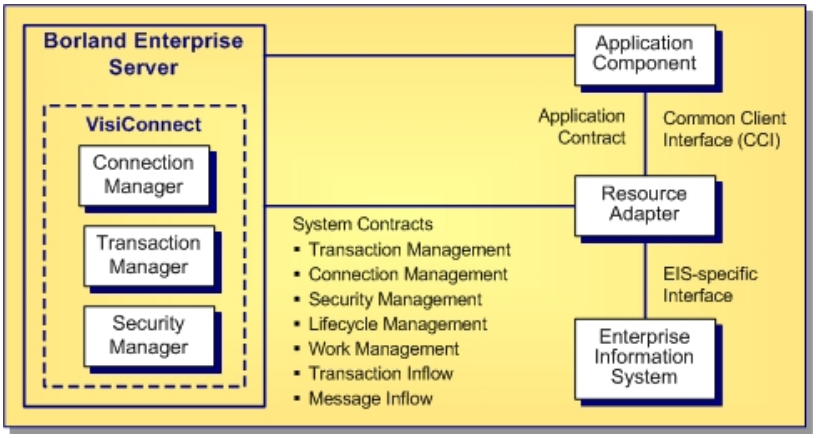

A Resource Adapter and its collateral serve as the Connector. VisiConnect supports Resource Adapters developed by EIS vendors and third-party application developers written to the Connectors 1.5 standard. Resource Adapters contain the components—Java, and if necessary, native code—required to interact with the specific EIS.

System Contracts

The Connectors specification defines a set of system level contracts between the application server and an EIS-specific Resource Adapter. This collaboration keeps all system-level mechanism transparent from the application components. Thus, the application component provider focuses on the development of business and presentation logic, and need not delve into the system-level issues related to EIS integration. This promotes the development of application components with greater ease and maintainability.

VisiConnect, in compliance with the Connectors specification, has implemented the standard set of defined contracts for:

- **Connection Management**, that allows an application server to pool connections to underlying EISs, providing application components with connection services to EISs. This leads to a highly scalable application environment that supports a large number of clients requiring access to heterogeneous EISs.
- **Transaction Management**, the contract between the application server transaction manager and an EIS supporting transactional access to EIS resource managers, that enables the application server to manage transactions across multiple resource managers.
- **Security Management**, that enables secure access to underlying EISs. This provides support for a secure application environment, which reduces security threats to the EIS and protects valuable information resources managed by the EIS.
- **Lifecycle Management** allows an application server to manage the lifecycle of a resource adapter. This contract provides a mechanism for the application server to bootstrap a resource adapter instance during its deployment or application server startup, and to notify the resource adapter instance during its undeployment or during an orderly shutdown of the application server.

- **Work Management** allows a resource adapter to do work (monitor network endpoints, call application components, etc.) by submitting `Work` instances to an application server for execution. The application server dispatches threads to execute submitted `Work` instances. This allows a resource adapter to avoid creating or managing threads directly, and allows an application server to efficiently pool threads and have more control over its runtime environment. The resource adapter can control the security context and transaction context with which `Work` instances are executed.
- **Transaction Inflow** allows a resource adapter to propagate an imported transaction to an application server. This contract also allows a resource adapter to transmit transaction completion and crash recovery calls initiated by an EIS, and ensures that the ACID properties of the imported transaction are preserved.
- **Message Inflow** allows a resource adapter to asynchronously deliver messages to message endpoints residing in the application server independent of the specific messaging style, messaging semantics, and messaging infrastructure used to deliver messages. This contract also serves as the standard message provider pluggability contract that allows a wide range of message providers (Java Message Service (JMS), Java API for XML Messaging (JAXM), etc.) to be plugged into any J2EE compatible application server via a resource adapter.

Connection Management

Connections to an EIS are expensive resources to create and destroy. To support scalable applications, the application server needs to be able to pool connections to the underlying EISs. To simplify application component development, this connection pooling mechanism needs to be transparent to the components accessing the underlying EISs.

The Connectors specification supports connection pooling and management, optimizing application component performance and scalability. The connection management contract, defined between the application server and the Resource Adapter, provides:

- A consistent application development model for connection acquisition for both managed (n-tier) and non-managed (two-tier) applications.
- A framework for the Resource Adapter to provide a standard connection factory and connection interface based on the Common Client Interface (CCI), opaque to the implementation for the underlying EIS.
- A generic mechanism for providing different quality of services (QoS) advanced connection pooling, transaction management, security management, error tracing and logging—for a configured set of Resource Adapters.
- Support for the application server to implement its connection pooling facility.

VisiConnect uses connection management to:

- Create new connections to an EIS
- Configure connection factories in the Java Naming and Directory Interface (JNDI) namespace.
- Find the right connection to an EIS from an existing set of pooled connections, and reuse that connection.
- Hook in AppServer's transaction and security services.

The AppServer establishes, configures, caches and reuses connections to the EIS automatically through VisiConnect.

The application component performs a lookup of a Resource Adapter connection factory in the JNDI namespace, using the connection factory to get a connection to the underlying EIS. The connection factory delegates the connection creation request to the VisiConnect connection manager instance. On receiving this request, the connection manager performs a lookup in the connection pool. If there is no connection in the pool that can satisfy the connection request, VisiConnect uses the `ManagedConnectionFactory` implemented by the Resource Adapter to create a new physical connection to the underlying EIS. If VisiConnect finds a matching connection in the pool, it then uses the matching `ManagedConnection` instance to satisfy the connection request. If a new `ManagedConnection` instance is created, the server adds the new `ManagedConnection` instance to the connection pool.

VisiConnect registers a `ConnectionEventListener` with the `ManagedConnection` instance. This listener enables VisiConnect to receive event notifications related to the state of the ManagedConnection instance. VisiConnect uses these notifications to manage connection pooling, transactions, connection cleanup and handle error conditions.

VisiConnect uses the `ManagedConnection` instance to provide a Connection instance that acts as an application-level handle to the underlying physical connection, to the application component. The component in turn uses this handle—and not the underlying physical connection directly—to access EIS resources.

Transaction Management

Transactional access to multiple EISs is an important and often critical requirement for enterprise applications. The Connectors supports transaction access to multiple, heterogeneous EISs—where a number of interactions must be committed together, or not at all, in order to maintain data consistency and integrity.

VisiConnect utilizes the AppServer's transaction manager and supports Resource Adapters conforming to the following transaction support levels.

- **No Transaction support**: if a Resource Adapter supports neither Local Transactions nor XA Transactions, it is non-transactional. If an application component uses a non-transactional Resource Adapter, the application component must not involve any connections to the respective EIS in a transaction. If the application component is required to involve EIS connections in a transaction, the application component must use a Resource Adapter which support Local or XA Transactions.
- **Local Transaction support**: the application server manages resources directly, which are local to the Resource Adapter. Unlike XA Transactions, local transactions can neither participate in the two-phase commit (2PC) protocol, nor participate as a distributed transaction (whereas the transaction context is simply propagated); instead, local transactions solely target one-phase commit (1PC) optimization. A Resource Adapter defines the type of transaction support in its Sun standard deployment descriptor. When an application component requests an EIS connection as part of a transaction, AppServer starts a local transaction based on the current transaction context. When the application closes the connection, AppServer commits the local transaction, and cleans up the EIS connection once the transaction is completed.
- **XA Transaction support**: a transaction is managed by a transaction manager external to the Resource Adapter and the EIS. A Resource Adapter defines the type of transaction support in its Sun-standard deployment descriptor. When an application component demarcates an EIS connection request as part of a transaction, the AppServer is responsible for enlisting the XA resource with the transaction manager. When the application component closes that connection, the application server unlists the XA resource from the transaction manager, and cleans up the EIS connection once the transaction is completed.

  In compliance with the Connectors 1.5 specification, VisiConnect provides full support for all three specified transaction levels.

One-Phase Commit Optimization

In many cases, a transaction is limited in scope to a single EIS, and the EIS resource manager performs its own transaction management—this is the Local Transaction. An XA Transaction can span multiple resource managers, thus requiring transaction coordination to be performed by an external transaction manager, typically one packaged with an application server. This external transaction manager can either use the 2PC protocol, or propagate the transaction context as a distributed transaction, to manage a transaction that spans multiple EISs. If only one resource manager is participating in an XA Transaction, it uses the 1PC protocol. In an environment where a singleton resource manager is handling its own transaction management, 1PC optimization can be performed, as this involves a less expensive resource than a 1PC XA Transaction.

Security Management

In compliance with the Connectors 1.5 specification, VisiConnect supports both container-managed and component-managed sign-on. At runtime, VisiConnect determines the selected sign-on mechanism based on information specified in deployment descriptor of the invoking component. If VisiConnect is unable to determine the sign-on mechanism requested by the component (most often due to an improper JNDI lookup of the Resource Adapter connection factory), VisiConnect will attempt container-managed sign-on. If the component has specified explicit security information, this will be presented in the call to obtain the connection, even in the case of container-managed sign-on.

Component-Managed Sign-on

When employing component-managed sign-on, the component provides all the required security information—most commonly a username and a password—when requesting to obtain a connection to an EIS. The application server provides no additional security processing other than to pass the security information along on the request for the connection. The Resource Adapter uses the component-provided security information to perform EIS sign-on in an implementation-specific manner.

Container-Managed Sign-on

When employing container-managed sign-on, the component does not present any security information, and the container must determine the necessary sign-on information, providing this information to the Resource Adapter in the request to obtain a connection. The container must determine an appropriate resource principal and provide this resource principal information to the Resource Adapter in the form of a Java Authentication and Authorization Service (JAAS) Subject object.

EIS-Managed Sign-on

When employing EIS-managed sign-on, the Resource Adapter internally obtains all of its EIS connections with a pre-configured, hard-coded set of security information. In this scenario the Resource Adapter does not depend upon the security information passed to it in the invoking component's requests for new connections.

Authentication Mechanisms

The AppServer user must be authenticated whenever they request access to a protected AppServer resource. For this reason, each user is required to provide a credential (a username/password pair or a digital certificate) to AppServer. The following types of authentication mechanisms are supported by AppServer:

- Password authentication a user ID and password are requested from the user and sent to AppServer in clear text. Borland Enterprise Server checks the information and if it is trustworthy, grants access to the protected resource.

- The SSL (or HTTPS) protocol can be used to provide an additional level of security to password authentication. Because the SSL protocol encrypts the data transferred between the client and AppServer, the user ID and password of the user do not flow in the clear. Therefore, AppServer can authenticate the user without compromising the confidentiality of the user's ID and password.
- Certificate authentication: when an SSL or HTTPS client request is initiated, AppServer responds by presenting its digital certificate to the client. The client then verifies the digital certificate and an SSL connection is established. The CertAuthenticator class then extracts data from the client's digital certificate to determine which AppServer User owns the certificate and then retrieves the authenticated User from the AppServer security realm.
- You can also use mutual authentication. In this case, Borland Enterprise Server not only authenticates itself, it also requires authentication from the requesting client. Clients are required to submit digital certificates issued by a trusted certificate authority. Mutual authentication is useful when you must restrict access to trusted clients only. For example, you might restrict access by accepting only clients with digital certificates provided by you.

For more information, see “Getting Started with Security” in the Developer's Guide.

Security Map

In Section 8.5 of the Connectors 1.5 specification, a number of possible options are identified for defining a Resource Principal on the behalf of whom sign-on is being performed. VisiConnect implements the Principal Mapping option identified in the specification.

Under this option, a resource principal is determined by mapping from the identity of the initiating caller principal for the invoking component. The resulting resource principal does not inherit the identity of security attributes of the principal that is it mapped from. Instead, the resource principal derives its identity and security attributes based on the defined mapping. Thus, to enable and use container-managed sign-on, VisiConnect provides the Security Map to specify the initiating principal association with a resourceprincipal. Expanding upon this model, VisiConnect provides a mechanism to map initiating caller roles to resource roles.

If container-managed sign-on is requested by the component and no Security Map is configured for the deployed Resource Adapter, an attempt is made to obtain the connection using a null JAAS Subject object. This is supported based upon the Resource Adapter implementation.

While the defined connection management system contracts define how security information is exchanged between the AppServer and the Resource Adapter, the determination to use container-managed sign-on or component-managed sign-on is based on deployment information defined for the component requesting a connection.

The Security Map is specified with the security-map element in the ra-borland.xml deployment descriptor. This element defines the initiating role association with a resource role. Each security-map element provides a mechanism to define appropriate resource role values for the Resource Adapter and EIS sign-on processing. The security-map elements provide the means to specify a defined set of initiating roles and the corresponding resource role to be used when allocating managed connections and connection handles.

A default resource role can be defined for the connection factory in the security-map element. To do this, specify a user-role value of “*” and a corresponding resource-role value. The defined resource-role is then utilized whenever the current identity if not matched elsewhere in the Security Map.

This is an optional element. However, it must be specified in some form when container-managed sign-on is supported by the Resource Adapter and any component uses it. Additionally, the deployment-time population of the connection pool is attempted using the defined default resource role, given that one is specified.

Security Policy Processing

The Connectors 1.5 specification defines default security policies for any Resource Adapters running in an application server. It also defines a way for a Resource Adapter to provide its own specific security policies overriding the default.

In compliance with this specification, AppServer dynamically modifies the runtime environment for Resource Adapters. If the Resource Adapter has not defined specific security policies, AppServer overrides the runtime environment for the Resource Adapter with the default security policies specified in the Connectors 1.5 specification. If the Resource Adapter has defined specific security policies, Borland Enterprise Server first overrides the runtime environment for the Resource Adapter first with a combination of the default security policies for Resource Adapters and the specific policies defined for the Resource Adapter. Resource Adapters define specific security policies using the security-permission-spec element in the ra.xml deployment descriptor file.

For more information on security policy processing requirements, see Section 18.2, “Security Permissions”, in the Connectors 1.5 specification (`http://java.sun.com/j2ee/download.html#connectorspec`).

Common Client Interface (CCI)

The Common Client Interface (CCI) defines a standard client API for application components. The CCI enables application components, Enterprise Application Integration (EAI) frameworks, and development tools to drive interactions across heterogeneous EISs using a common client API.

The CCI is targeted for use by EAI and enterprise tool vendors. The Connectors 1.5 specification recommends that the CCI be the basis for richer functionality provided by the tool vendors, rather than being an application-level programming interface used by most application developers. Application components themselves may also write to the API. As the CCI is a low-level interface, this use is generally reserved for the migration of legacy modules to the J2EE 1.4 Platform. Through the CCI, legacy EIS clients can integrate directly with the AppServer; this provides for a smoother, less costly migration path to J2EE 1.4.

The CCI defines a remote function call interface that focuses on executing functions on an EIS and retrieving the results. The CCI is independent of a specific EIS; in other words, it is not bound to the data types, invocation hooks, and signatures of a particular EIS. The CCI is capable of being driven by EIS-specific metadata from a repository.

The CCI enables the AppServer to create and manage connections to an EIS, execute an interaction, and manage data records as input, output, or return values. The CCI is designed to leverage the Java Beans architecture and Java Collection framework.

The Connectors 1.5 specification recommends that a Resource Adapter support CCI as its client API, while it requires the Resource Adapter to implement the system contracts. A developer may choose to write the Resource Adapter to provide a client API different from the CCI, such as:

- the Java Database Connectivity (JDBC) API (an example of a general EIS-type interface), or
- for example, the client API based on the IBM CICS Java Gateway (an example of a EIS-specific interface)

The CCI (which form the application contract) consists of the following:

- **ConnectionFactory** A ConnectionFactory implementation creates a connection and interaction object as a means of interacting with an EIS. Its getConnection method gets a connection to an EIS instance.
- **Connection** A Connection implementation represents an application level handle to an EIS instance. The actual connection is represented by a ManagedConnection. An application gets a Connection object by using the getConnection method of a ConnectionFactory object.

- **Interaction** An Interaction implementation is what drives a particular interaction. It is created using the ConnectionFactory. The following three arguments are needed to carry out an interaction via the Interaction implementation: InteractionSpec, which identifies the characteristics of the concrete interaction, and Input and Output, which both carry the exchanged data.
- **InteractionSpec** An InteractionSpec implementation defines all interaction-relevant properties of a connector (for example, the name of the program to call, the interaction mode, and so forth). The InteractionSpec is passed as an argument to an Interaction implementation when a particular interaction has to be carried out.
- **Input and output** The input and output are records.

A **record** is a logical collection of application data elements that combines the actual record bytes together with its type. Examples are COBOL and C data structures. Record implementation in CCI uses streams. In the javax.resource.cci.Streamable interface, reading and writing from streams is handled by read and write methods. In the javax.resource.cci.Record interface, getRecordName() and getRecordShortDescription(), and setRecordName() and setRecordShortDescription() get and set the record data.

You must create records for all of the data structures that are externalized by the EIS functions you want to reuse. You then use the records as input and output objects that pass data via a Resource Adapter to and from an EIS. You will want to consider the following options when creating a record:

- **Having direct access to nested, or hierarchical, records** A direct, or 'flattened', set of accessor methods may be more convenient, or seem more natural, to some users. For example, programmers accustomed to COBOL may expect to be able to refer directly to the field of a sub-record if the field name is unique within the record. This is similar to the way COBOL field names are scoped. There is no need to qualify field names if the field name is unique.
- **Custom and Dynamic Records** You can generally create two types of records: custom and dynamic. The main difference between these is the way fields are accessed. For dynamic records, the fields are found by taking the field name, looking up the offset and the marshalling of the information, and then accessing it. For custom records, the offset and the marshalling of the information is in the code, resulting in faster access. Generating custom records results in more efficient code, but there are restrictions on their use.
- **Records with or without notification** If a record is created with notification, then the properties of the record are bound.

  Note

  If bound properties are not required, then it is more efficient to create a record without notification.

Packaging and Deployment

The Connectors provides packaging and deployment interfaces so that various Resource Adapters can be deployed to J2EE 1.4 Platform compliant application servers, such as the AppServer.

Figure 32.2 Packaging and Deployment in the AppServer and VisiConnect

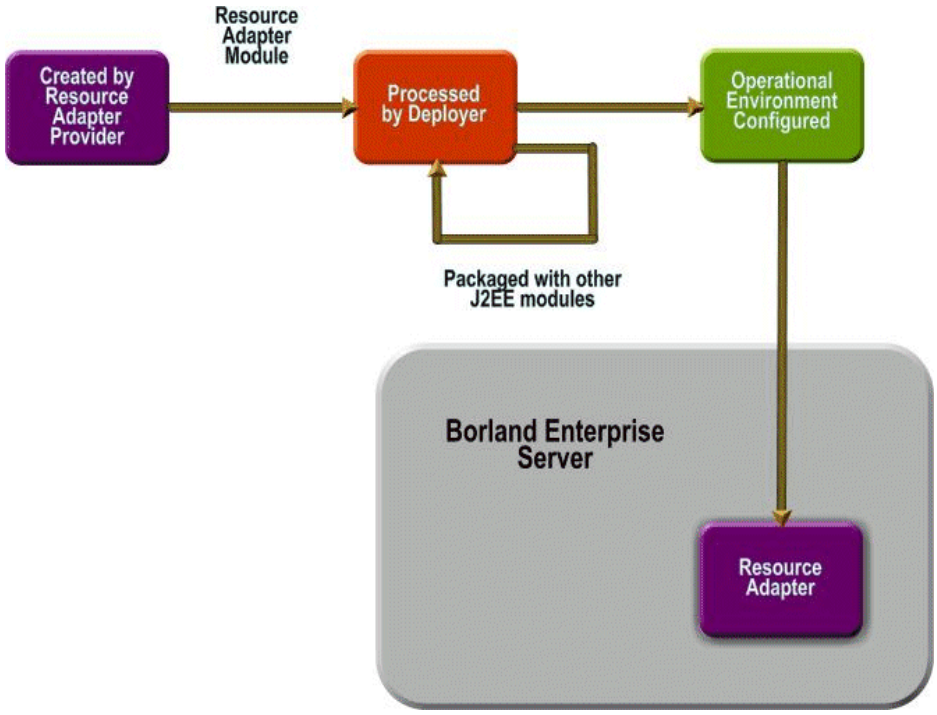

A Resource Adapter packages a set of Java interfaces and classes, which implement the Connectors-specified system contracts and EIS-specific functionality to be provided by the Resource Adapter. The Resource Adapter can also require the use of native libraries specific to the underlying EIS, and other collateral, for example:

- Documentation
- Help files
- A code generator for EJBs
- A tool that directly provides configuration utilities so you can configure the EIS directly
- A tool that provides additional deployment facilities for remote Resource Adapter components
- For example, with IBM CICS, a set of JCL scripts that you may need to run on the mainframe

The Java interfaces and classes are packaged together, with required collateral and deployment descriptors, to create a Resource Adapter module. The deployment descriptors define the deployment contract between a Resource Adapter and the application server

A Resource Adapter can be deployed as a shared, standalone module, or packaged as part of a J2EE application. During deployment, the Resource Adapter module is installed on the AppServer and configured for the target operational environment. The configuration of a Resource Adapter is based on the properties defined in the deployment descriptors.

VisiConnect Features

Among the value-added features provided by VisiConnect as enhancements to the Connectors standard are the following:

- VisiConnect Partition Service
- Additional Classloading Support
- Secure Password Credential Storage
- Connection Leak Detection
- Security Policy Processing of ra.xml Specifications

VisiConnect Partition Service

The Borland Partition with the VisiConnect service enabled is designed to support development and deployment of J2EE applications which bundle Resource Adapters, or standalone Resource Adapter components. The AppServer Partition provides integrated VisiConnect services. Tools include a Deployment Descriptor Editor (DDE) and a set of task wizards for packaging and deploying Resource Adapters and their related descriptor files.

This provides a highly modular environment for running VisiConnect. The AppServer provides a default VisiConnect Service in Partitions for deployment.

Additional Classloading Support

VisiConnect supports the loading of properties or classes that are specified in ClassPath entry of the Resource Adapter's Manifest.mf file. The following is a description of how you configure properties and classes that are in and used by a Resource Adapter.

The Resource Adapter (RAR) archive file and the application component using it (for example, an EJB jar) are contained in an Enterprise Application (EAR) archive. The RAR requires resources such as Java properties that are stored in a JAR file, and that JAR file is contained within the EAR file (not in the RAR itself).

You specify a reference to the RAR Java classes by adding a ClassPath= entry in the RAR Manifest.mf file. You can also store the EJB Java classes in the same JAR file that is contained within the EAR. This scenario provides a “support” JAR file that contains Java classes for the components in the EAR that require them.

Secure Password Credential Storage

VisiConnect provides a standard method for Resource Adapter deployers to plug in their specified authorization/authentication mechanism through secure password credential storage.

This storage mechanism is used to map user roles (AppServer roles, which may be associated with AppServer username and password combinations or credentials) to resource roles (EIS roles, which may be associated with EIS user name and password combinations or credentials).

Connection Leak Detection

VisiConnect provides two mechanisms for preventing connection leaks:

- Leveraging a garbage collector
- Providing an idle timer for tracking the usage of connection objects

Security Policy Processing of ra.xml Specifications

VisiConnect provides a set of security permissions for execution of a Resource Adapter in a managed runtime environment. The AppServer also grants a Resource Adapter explicit permissions to access system resources.

Resource Adapters

Source code for several Resource Adapters are provided with VisiConnect as examples. Some of these Resource Adapters are wrappers for JDBC 2.0 calls, some using the CCI and some not. Deployment descriptors supporting the three transaction levels are provided for each Resource Adapter.

Simplified application examples for these JDBC Resource Adapters are provided with VisiConnect. An EJB is used to model the data in the EIS, and a J2EE client and a Servlet are used to query the Resource Adapter and display the output. The example uses any RDBMS which is supported by a JDBC 2.0 compliant driver. By default, the examples are configured to use JDataStore as the EIS, but it is a straightforward task to configure them to use any JDBC 2.0 RDBMS. The components are packaged as a J2EE Application. For more information, refer to the VisiConnect example `README` provided with the AppServer.

Other sample resource adapters provided with the product include an open-source generic JMS resource adapter with instructions for integrating with JMS providers such as Tibco and OpenJMS, and a mail resource adapter which allows you to use an email server as an EIS. These samples demonstrate the use of message inflow to allow for inbound communication from the EIS to the application server, as well as outbound connection capability.

33

Implementing Partition Interceptors

Implementing Partition Interceptors requires the following steps:

1. Defining your interceptor using the `module-borland.xml` descriptor file.
2. Creating the interceptor class.
3. JARing the class and the descriptor file.
4. Deploy the JAR to the Partition of interest.

Defining the Interceptor

You define the interceptor by creating a `module-borland.xml` file. This file uses the following DTD:

```
<!ELEMENT module (Partition-interceptor?)>
<!ELEMENT Partition-interceptor (class-name, argument?, priority?)>
<!ELEMENT class-name (#PCDATA)>
<!ELEMENT argument (key, value)>
<!ELEMENT key (#PCDATA)>
<!ELEMENT value (#PCDATA)>
<!ELEMENT priority (#PCDATA)>
```

The `<class-name>` element must contain the full-path class name of the implementation contained within the JAR.

The `<priority>` element is an optional field that controls the order in which a set of interceptors for a particular Partition are fired. This value must be between 0 and 9. Priority 0 ranks before priority 9. Interceptors are fired in order during load time and in reverse order during shutdown. If two or more interceptors share the same priority, there is no way to determine or enforce which of that set will be fired relative to the other.

The `<argument>` is an optional element which contains a pair of elements, `<key>` and `<value>`. These are passed into your class implementation as a `java.util.HashMap`. Your code must extract the appropriate values from this type. The limit on arguments is imposed by the JVM implementation.

For example, the following XML defines an interceptor called `InterceptorImpl`:

```
<module>
<Partition-interceptor>
    <class-name>com.borland.enterprise.examples.InterceptorImpl</class-name>
```

```
    <argument>
        <key>key1</key>
        <value>value1</value>
    </argument>
    <argument>
        <key>key2</key>
        <value>value2</value>
    </argument>
    <argument>
        <key>key3</key>
        <value>value3</value>
    </argument>
    <priority>1</priority>
</Partition-interceptor>
</module>
```

Creating the Interceptor Class

Your class must implement:

```
com.borland.enterprise.server.Partition.service.PartitionInterceptor
```

The following methods are available:

- `public void initialize(java.util.HashMap args);`

  This method is called before any Partition services like the Tomcat container are created and initialized. This method is not subject to the <priority> parameter, since it is invoked as each interceptor is loaded.

- `public void startupPreLoad();`

  This method is called after Partition services are started and before the Partition services load modules.

- `public void startupPostLoad();`

  This method is invoked after all Partition services have loaded their respective modules.

- `public void shutdownPreUnload();`

  This method is called before the Partition services unload their respective modules. The <priority> parameter now reverses its meaning; priority 9 interceptors are called first, then priority 8, and so forth.

- `public void shutdownPostUnload();`

  This method is called after the services have unloaded their modules.

- `public void PartitionTerminating();`

  This method is called after the services have been shut down, just before the Partition shuts down.

The following code sample shows the class `InterceptorImpl` defined in the `module-borland.xml` descriptor above:

```
package com.borland.enterprise.examples;

// This interface is contained in xmlrt.jar
import com.borland.enterprise.server.Partition.service.PartitionInterceptor;

public class InterceptorImpl implements PartitionInterceptor {
    static final String _className = "InterceptorImpl";

    public void initialize(java.util.HashMap args) {
    // Writing to System.out and System.err will
    // cause the output to be logged.
    // There is no requirement to log.
     System.out.println(_className + ": initialize");
     System.out.println("key1 has value " + args.get("key1").toString());
     System.out.println("key2 has value " + args.get("key2").toString());
     System.out.println("key3 has value " + args.get("key2").toString());
    }
    public void startupPreLoad() {
    // Writing to System.out and System.err will
    // cause the output to be logged.
    // There is no requirement to log.
    System.out.println(_className + ": startupPreLoad");
    }
    public void startupPostLoad() {
    // Writing to System.out and System.err will
    // cause the output to be logged.
    // There is no requirement to log.
      System.out.println(_className + ": startupPostLoad");
    }
    public void shutdownPreUnload() {
    // Writing to System.out and System.err will
    // cause the output to be logged.
    // There is no requirement to log.
     System.out.println(_className + ": shutdownPreUnload");
    }
    public void shutdownPostUnload() {
    // Writing to System.out and System.err will
    // cause the output to be logged.
    // There is no requirement to log.
      System.out.println(_className + ": shutdownPostUnload");
    }
    public void PartitionTerminating() {
    // Writing to System.out and System.err will
    // cause the output to be logged.
    // There is no requirement to log.
    System.out.println(_className + ": PartitionTerminating");
    }
}
```

Creating the JAR file

Use Java's JAR utility to create a JAR file of the class and its descriptor file.

Deploying the Interceptor

Use the Deployment Wizard to deploy the interceptor to the Partition. **Do not** check either the “Verify deployment descriptors” or the “Generate stubs” checkboxes.

Important

You must restart the Partition after deploying your interceptor.

You can also simply copy your JAR file into one of these two directories, making sure you restart the Partition manually afterward:

- `<install_dir>/var/servers/<server_name>/Partitions/<Partition_name>/lib`
- `<install_dir>/var/servers/<server_name>/Partitions/<Partition_name>/lib/system`

34

Using VisiConnect

The Java 2 Enterprise Edition (J2EE) Connector Architecture enables EIS vendors and third-party application developers to develop Resource Adapters that can be deployed to any application server supporting the J2EE 1.4 Platform Specification. The Resource Adapter provides platform-specific integration between the J2EE component and the EIS. When a Resource Adapter is deployed to the Borland AppServer (AppServer), it enables the development of robust J2EE applications which can access a wide variety of heterogeneous EISs. Resource Adapters encapsulate the Java components, and if necessary, the native components required to interact with the EIS.

Before using VisiConnect, Borland recommends that you read the Connectors 1.5 specification.

VisiConnect service

Resource adapters are hosted by Partitions with the VisiConnect Partition Service enabled. Multiple Resource Adapters can be deployed in the same Partition. VisiConnect is responsible for making the connection factories of its deployed Resource Adapters available to the client through JNDI. Thus, the client can look up the connection factory for a specific Resource Adapter using JNDI.

Service overview

The VisiConnect Service is a complete implementation of the Connectors 1.5 specification, including all optional functionality.

Every Resource Adapter object in the deployed Connector is simultaneously both a Resource Adapter object and a CORBA object.

Unlike other Connectors implementations, VisiConnect has no restrictions on partitioning. Any number of Resource Adapters can go into any number of Partitions running on any number of machines. Plus, support for distributed transactions protocol allows Resource Adapters to be partitioned arbitrarily. Partitioning enables you to configure the application during deployment to optimize its overall performance.

Connection management

The `ra.xml` deployment descriptor file contains a config-property element to declare a single configuration setting for a `ManagedConnectionFactory` instance. The resource

adapter provider typically sets these configuration properties. However, if a configuration property is not set, the resource adapter deployer is responsible for providing a value for the property.

Borland provides its own deployment descriptor for defining connectors and their connection factory properties: `ra-borland.xml`. See Borland DTDs for more information on using the `ra-borland.xml` descriptor.

Configuring connection properties

The following connection pool properties can be set:

Property	Value type	Description	Default
`wait-timeout`	`Integer`	The number of seconds to wait for a free connection when `maximum-capacity` connections are already opened. When using the `maximum-capacity` property and the pool is at its max and can't serve any more connections, the threads looking for connections end up waiting for the connection(s) to become available for a long time if the wait time is unbounded (set to 0 seconds). You can set the `wait-timeout` period to suit your needs.	`30`
`busy-timeout`	`Integer`	The number of seconds to wait before a busy connection is released. If a connection is busy for a long time, the application using it may have hung and be unable to release the connection. This timeout will ensure that connections will be timed out when they have been busy for much longer than necessary.	`600` (ten minutes)
`idle-timeout`	`Integer`	A pooled connection remaining in an idle state for a period of time longer than this timeout value should be closed to conserve resources. All idle connections are checked for `idle-timeout` expiration every 60 seconds. The value of the `idle-timeout` is given in seconds. A value of 0 (zero) indicates that connection cleanup is disabled.	`600` (ten minutes)
`maximum-capacity`	`Integer`	Identifies the maximum number of managed connections which VisiConnect will allow. Throws `ResourceAllocationException` when requests for newly allocated managed connections go beyond this limit.	10

The following properties have been deprecated and are now ignored by VisiConnect. They have been replaced by the pool properties `busy-timeout`, `idle-timeout`, and `wait-timeout`, listed in the table above. You do not have to delete the old-style properties from `ra-borland.xml`.

Unused Pool properties

Property	Default	Description
initial-capacity	1	Identifies the initial number of managed connections which VisiConnect will attempt to obtain during deployment.
capacity-delta	1	Identifies the number of additional managed connections which the VisiConnect will attempt to obtain during resizing of the maintained connection pool.

Property	Default	Description
cleanup-enabled	true	Indicates whether or not the Connection Pool should have unused Managed Connections reclaimed as a means to control system resources.
cleanup-delta	1	Identifies the amount of time the Connection Pool Management will wait between attempts to reclaim unused Managed Connections.

Security management with the Security Map

The Security Map enables the definition of user roles that can be

1 Used directly with the EIS for container-managed sign-on (use-caller-identity).

2 Mapped to an appropriate resource role for container-managed sign-on (run-as).

In the first case, when the user role identified at run time is found in the mapping, the user role itself is used to provide security information for interacting with an EIS. In the second case, when the user role identified at run time is found in the mapping, the associated resource role is used to provide security information for interacting with an EIS.

The use-caller-identity option is used when user identities in the user role identified at run time are available to the EIS as well. For example, a user identity, “borland”/”borland”, belonging to role “Borland”, is available to the AppServer, and the available EIS, a JDataStore database, has an identity of “borland”/”borland” available to it. When a Resource Adapter serving JDataStore is deployed with a Security Map specifying:

```
<security-map>
 <user-role>Borland</user-role>
   <use-caller-identity></use-caller-identity>
</security-map>
```

Applications on this server instance which use this JDataStore database can use use-caller-identity to access it.

Note

Due to a limitation currently in VisiSecure, you must define the caller identity in the resource vault as well as the user vault.

The run-as option is used when it makes sense to map user identities in the user role identified at run time to identities in the EIS. For example, a user identity, “demo”/”demo”, belonging to role “Demo”, is available to the AppServer, and the available EIS, an Oracle database, has an identity of “scott”/”tiger”, which is ideal for a demo user. When a Resource Adapter serving Oracle is deployed with a Security Map specifying:

```
<security-map>
 <user-role>Demo</user-role>
 <run-as>
 <role-name>oracle_demo</role-name>
  <role-description>Oracle demo role</role-description>
 </run-as>
</security-map>
```

The role `oracle_demo` is defined in the resource vault (see below), applications on this server instance which use this Oracle database can use run-as to access it.

When run-as is used, the vault must be provided for VisiConnect to use to extract the security information for the resource role. A resource role name and a set of credentials are written to this vault. When VisiConnect loads a Resource Adapter with a defined Security Map using run-as, it will read in the credentials for the defined role name(s) from the vault.

Authorization domain

The <authorization-domain> element in the ra-borland.xml descriptor file specifies the authorization domain associated with a specified user role. If <security-map> is set, you should set <authorization-domain> with its associated domain. If <authorization-domain> is not set, VisiConnect assumes the use of the **default** authorization domain. See “Getting started with security” in the *Security Guide* for more information on using authorization domains.

Default roles

In addition, the <security-map> element enables the definition of a default user role that can be associated with the appropriate resource role. This default role would be preferred to if the user role identified at run-time is not found in the mapping. The default user role is defined in the <security-map> element with a <user-role> element given a value of “*”. For example:

```
<user-role>*</user-role>
```

A corresponding <role-name> entry must be included in the <security-map> element. The following example illustrates the association between an AppServer user role and a resource role.

```
<security-map>
    <user-role>*</user-role>
    <run-as>
        <role-name>SHME_OPR</role-name>
    </run-as>
</security-map>
```

The default user role is also used at deployment time if the connection pool parameters indicate that the AppServer should initialize connections. The absence of a default user role entry or the absence of a <security-map> element may prevent the server from creating connections using container-managed security.

Generating a resource vault

To use run-as security mapping as described above, a resource role(s) must be defined in a vault which is provided to the AppServer. This is known as the resource vault.

VisiConnect provides a tool, ResourceVaultGen, to create a resource vault and to instantiate role objects in this vault. A role name and its associated security credentials are written to the resource vault by ResourceVaultGen. At this time only credentials of type Password Credential can be written to the resource vault. The usage of ResourceVaultGen is as follows:

```
java -Dborland.enterprise.licenseDir=<install_dir/var/domains/base/
configurations/<configuration_name>/mos/<partition_name>/adm> -
Dserver.instance.root=<install_dir/var/domains/base/configurations/
<configuration_name>/mos/<partition_name>/adm/properties/
management_vbroker.properties>
com.borland.enterprise.visiconnect.tools.ResourceVaultGen -rolename <role_name>
-username <user_name> -password <password> -vaultfile <full path to vault file>
-vpwd <vault_password>
```

where:

-rolename	Resource role name to store in the resource vault.
-username	Resource username to associate with the resource role.
-password	Resource password to associate with the resource role.

`-vaultfile` (optional)	Path to the vault file you write the resource role(s)to. If not specified, ResourceVaultGen will attempt to write to the default resource vault file <*install_dir*/var/domains/base/configurations/<configuration_name>/mos/<partition_name>/adm/properties/management_vbroker.properties>. If the vault file is does not already exist, a new vault file will be written to the specified location.
`-vpwd` (optional)	Password to assign to the vault for access authorization. If not specified, the vault will be created without a password.

When using ResourceVaultGen, ensure that the following jars are in your CLASSPATH:

- lm.jar
- visiconnect.jar
- vbsec.jar
- jsse.jar
- jaas.jar
- jce1_2_1.jar
- sunjce_provider.jar
- local_policy.jar
- US_export_policy.jar

Note

If you fail to include these jars in your CLASSPATH when you attempt to generate a vault, you may end up with a vault file which is invalid. If you attempt to reuse the invalid vault file, you will encounter an EOFException. To resolve, delete the invalid vault file and regenerate with ResourceVaultGen, ensuring that you have the proper jars in your CLASSPATH.

VisiConnect will use the vault if Security Map information is specified in at deployment time for a Resource Adapter. If the resource vault is password protected, VisiConnect will need to have the following property passed to it:

```
-Dvisiconnect.resource.security.vaultpwd=<vault_password>
```

If the resource vault is in a user specified location (-vaultfile ...), VisiConnect will need to have the following property passed to it:

```
-Dvisiconnect.resource.security.login=<path of specified vault file>
```

The following examples illustrate the use of ResourceVaultGen:

Example 1:

```
java -Dborland.enterprise.licenseDir=/opt/BES/var<install_dir/var/domains/base/
configurations/<configuration_name>/mos/<partition_name>/adm/properties/
management_vbroker.properties>
 -Dserver.instance.root=/opt/BES/var/servers/servername -
Dpartition.name=standard
 com.borland.enterprise.visiconnect.tools.ResourceVaultGen -rolename
administrator
 -username red -password balloon -vaultfile
 /opt/BES/var/servers/servername/adm/properties/partitions/standard/
resourcevault -vpwd
 lock
```

This usage generates a resource vault named `resourcevault` to /opt/BES/var/servers/servername/adm/properties/partitions/standard, with a role `administrator` associated with a Password Credential with username `red` and password `balloon`. The vault file itself is password protected, using the password `lock`. For VisiConnect to use this vault, the following properties must be set for it:

```
-Dvisiconnect.resource.security.vaultpwd=lock
-Dvisiconnect.resource.security.login=resourcevault
```

Example 2:

```
java -Dborland.enterprise.licenseDir=/opt/BES/var/domains/base/configurations/
<configuration_name>/mos/<partition_name>/adm/properties/
management_vbroker.properties>
-Dserver.instance.root=/opt/BES/var/domains/base/configurations/
<configuration_name>/mos/<partition_name>/adm/properties/
management_vbroker.properties>
-Dpartition.name=petstore
com.borland.enterprise.visiconnect.tools.ResourceVaultGen
-rolename manager accounts -username mickey daffy
-password mouse duck -vpwd goofy
```

This usage generates a default resource vault (named `resource_vault`) to /opt/BES/var/servers/servername/adm/properties/partitions/petstore, with a role `manager` associated with a Password Credential with username `mickey` and password `mouse`, and another role `accounts` associated with a Password Credential with username `daffy` and password `duck`. The vault file itself is password protected, using the password `goofy`. For VisiConnect to use this vault, the following properties must be set for it:

```
-Dvisiconnect.resource.security.vaultpwd=goofy
```

Example 3:

```
java -Dborland.enterprise.licenseDir=/opt/BES/var/servers/servername/adm -
Dserver.instance.root=/opt/BES/var/servers/servername
-Dpartition.name=standard
com.borland.enterprise.visiconnect.tools.ResourceVaultGen
-rolename OClone ENolco -username darkstar geraldo -password meteor rivera
```

This usage generates a default resource vault (named `resource_vault`) to /opt/BES/var/domains/base/configurations/<configuration_name>/mos/<partition_name>/adm/properties/management_vbroker.properties>, with a role `developer` associated with a Password Credential with username `darkstar` and password `meteor`, and a role `host` associated with a Password Credential with username `geraldo` and password `rivera`. The vault file itself is not password protected. VisiConnect requires no additional parameters to use this vault.

Note

`ResourceVaultGen` cannot be used to write vault information to an existing file containing invalid characters. For example, a file generated by 'touch', or a StarOffice or Word document. `ResourceVaultGen` can only write vault information to a new file that it itself generates, or a valid existing vault file.

Resource Adapter overview

According to the Connectors 1.5 specification, you must be able to deploy a Resource Archive (RAR) as part of an Enterprise Archive (EAR). With AppServer and VisiConnect you can also deploy a standalone RAR. Once the RAR is deployed, you must do the following:

- Write code to obtain a connection.
- Create an Interaction object.
- Create an Interaction Spec.
- Create record and/or result set instances.
- Run the execute command so the record objects become populated.

In addition to some introductory conceptual information, this chapter provides steps to help you understand the code you must write.

The J2EE Connector Architecture enables Enterprise Information System (EIS) vendors and third-party application developers to develop Resource Adapters that can be deployed to any J2EE 1.4 compliant application server. The Resource Adapter is the main component of the J2EE Connector Architecture (Connectors), providing platform-specific integration between J2EE application components and the EIS. When a Resource Adapter is deployed to the AppServer, it enables the development of robust J2EE applications which can access a wide variety of heterogeneous EISs. Resource Adapters encapsulate the Java components and, if necessary, the native components required to interact with the EIS.

Development overview

See “Developing the Resource Adapter” for more information.

Developing a Resource Adapter from scratch requires implementing the necessary interfaces and deployment descriptors, packaging these into a Resource Adapter Archive (RAR), and finally deploying the RAR to the AppServer. The following summarizes the main steps for developing a Resource Adapter:

1 Write Java code for the various interfaces and classes required by the Resource Adapter within the scope of the Connectors 1.5 specification.

2 Specify these classes in the ra.xml standard deployment descriptor file.

3 Compile the Java code for the interfaces and implementation into class files.

4 Package the Java classes into a Java Archive (JAR) file.

5 Create the Resource Adapter-specific deployment descriptors:

- `ra.xml`: describes the Resource Adapter-related attributes and deployment properties using the Sun standard DTD.
- `ra-borland.xml`: add additional AppServer-specific deployment information. This file contains the parameters for connection factories, connection pools, and security mappings.

1 Create the Resource Adapter Archive (RAR) file (that is, package the Resource Adapter)

2 Deploy the Resource Adapter Archive to the AppServer, or include it in an Enterprise Application Archive (EAR) file to be deployed as part of a J2EE application.

Editing existing Resource Adapters

If you have existing Resource Adapters you would like to deploy to the AppServer, it may only be necessary to edit the Borland-specific deployment descriptor described above and repackage the adapter. Doing so involves the following steps, with illustrative example:

1 Create an empty staging directory for the RAR:

```
mkdir c:/temp/staging
```

2 Copy the Resource Adapter to be deployed into the staging directory:

```
cp shmeAdapter.rar c:/temp/staging
```

3 Extract the contents of the Resource Adapter Archive:

```
jar xvf shmeAdapter.rar
```

The staging directory should now contain the following:

- a JAR containing Java classes that implement the Resource Adapter
- a `META-INF` directory containing the files `Manifest.mf` and `ra.xml`

1 Create the `ra-borland.xml` file using the Borland Deployment Descriptor Editor (DDEditor) and save it into the staging area's `META-INF` directory. See “Using the Deployment Descriptor Editor” in the *Management Console User’s Guide* for information on using the DDEditor.

2 Create the new Resource Adapter Archive

```
jar cvf shmeAdapter.rar -C c:/temp/staging
```

3 You may now deploy the Resource Adapter to the AppServer.

Resource Adapter Packaging

The Resource Adapter is a J2EE component contained in a RAR. Resource Adapters use a common directory format. The following is an example of a Resource Adapter's directory structure:

Resource Adapter Directory Structure:

```
.META-INF/ra.xml
.META-INF/ra-borland.xml
./images/shmeAdapter.jpg
./readme.html
./shmeAdapter.jar
./shmeUtilities.jar
./shmeEisSdkWin32.dll
./shmeEisSdkUnix.so
```

As shown in the structure above, the Resource Adapter can include documentation and related files not directly used by the Resource Adapter—for example, the image and readme files. Packaging the Resource Adapter means packaging these files as well.

Packaging a Resource Adapter includes the following steps:

1. Create a temporary staging directory.
2. Compile the Resource Adapter Java classes into the staging directory. (Or, as above, simply copy pre-compiled classes into the staging directory.)
3. Create a JAR file to store the Resource Adapter Java classes. Add this JAR to the top level of the staging directory.
4. Create a `META-INF` subdirectory in the staging area.
5. Create a ra.xml deployment descriptor in this subdirectory and add entries for the Resource Adapter. Refer to Sun Microsystems' documentation for information on the ra.xml document type definition, at `http://java.sun.com/dtd/connector_1_0.dtd`.
6. Create a ra-borland.xml deployment descriptor in this same META-INF subdirectory and add entries for the Resource Adapter. Refer to the DTD at the end of this document for details on the necessary entries.
7. Create the Resource Adapter Archive:

   ```
   jar cvf resource-adapter-archive.rar -C staging-directory
   ```

   This command creates a RAR file that you can deploy to the server. The `-C staging-directory` option instructs the JAR command to change to the `staging-directory` so that the directory paths recorded in the RAR file are relative to the directory where the Resource Adapters were staged.
   One or more Resource Adapters can be staged in a directory and packaged in a JAR file.

Deployment Descriptors for the Resource Adapter

The AppServer uses two XML files to specify deployment information. The first of these is `ra.xml`, based on Sun Microsystems' DTD for resource adapters. The second is Borland's proprietary `ra-borland.xml`, which includes additional deployment information necessary for AppServer.

Configuring ra.xml

If you do not already have an `ra.xml` file associated with your Resource Adapter, it is necessary to manually create a new one or edit an existing one. You can use a text editor or the Borland DDEditor to edit these properties. For the most up-to-date

information on creating an ra.xml file, refer to the Connectors specification at http://java.sun.com/j2ee/connector.

Configuring the transaction level type

It is of critical importance that you specify the transaction level type supported by your Resource Adapter in the ra.xml deployment descriptor. The following table shows the transaction levels supported and how they are rendered in XML.

Transaction Support Type	XML representation
None	<transaction-support>**NoTransaction**</transaction-support>
Local	<transaction-support>**LocalTransaction**</transaction-support>
XA	<transaction-support>**XA**</transaction-support>

Configuring ra-borland.xml

The ra-borland.xml file contains information required for deploying a Resource Adapter to the AppServer. Certain attributes need to be specified in this file in order to deploy the RAR file. This functionality is consistent with the equivalent .xml extensions for EJBs, EARs, WARs, and client components for the AppServer.

Until Borland-specific deployment properties are provided in the ra-borland.xml file, the RAR cannot be deployed to the server. The following attributes are required in ra-borland.xml for a deployable RAR:

- Resource Adapter instance name. This name must be unique among RARs deployed to the partition. It is used by the VisiConnect service to uniquely identify the deployed Resource Adapter. When a Resource Adapter supports inbound communication, this is the name used in the ejb-borland.xml descriptor for the endpoint MDB to identify the Resource Adapter from which the MDB expects to receive incoming messages.
- For each connection-definition in the ra.xml file:
  - Connection factory interface class name. This must be unique among all connection-definitions in the Resource Adapter. This class name is used to associate the specifications for a particular connection factory in the ra-borland.xml with the specifications for the corresponding factory in the ra.xml file.
  - Connection factory name. This must be unique among all Resource Adapters deployed to the partition.
  - Connection factory JNDI name. This must be unique among all Resource Adapters deployed to the partition.

The following optional attributes may also be specified in the ra-borland.xml file:

- Reference to a separately deployed connection factory that contains Resource Adapter components that should be shared with the current Resource Adapter.
- Directory where all shared libraries should be copied.
- Mapping of security principals for Resource Adapter/EIS sign-on processing. This mapping identifies resource principals to be used when requesting EIS connections for applications that use container-managed security and for EIS connections requested during initial deployment.
- For each connection-definition in the ra.xml file:
  - Connection factory description
  - Logging-required flag. This indicates whether logging should be done for the ManagedConnectionFactory and ManagedConnection classes.
  - Log file location
  - Connection pool properties:

- busy-timeout: the number of seconds to wait before a busy connection is released. The default is 600 seconds
- idle-timeout: a pooled connection remaining in an idle state for a period of time longer than this timeout value should be closed. All idle connections are checked for expiration every 60 seconds. The value of the idle-timeout is given in seconds. The default is 600 seconds.
- wait-timeout: the number of seconds to wait for a free connection. The default is 30 seconds.

Changes to the Deployment Descriptors for Connectors 1.5

BAS supports both Connectors 1.0 and Connectors 1.5. The following lists the significant changes in the deployment descriptor for Connectors 1.5:

- A new resource adapter implementation class needs to be specified in the ra.xml.

```
<resourceadapter>
   <resourceadapter-class>
      .ResourceAdapter.Implementation.Class
   </resourceadapter-class>
⋮
```

- Multiple connection definitions can be specified within the same RAR within the <outbound-resourceadapter>.
- Multiple message adapters can be specified within the <inbound-resourceadapter>
- Configuration properties can be specified at the Resource Adapter level using the config-property element. Note that the config-property type can only be objects and not primitives, for example, you need to use java.lang.Integer rather than int. Also, make sure that the set methods in the Resource Adapter implementation use the same types. For example, in this case: setCount(java.lang.Integer value):

```
<config-property>
   <description>Open User Name</description>
   <config-property-name>Count</config-property-name>
   <config-property-type>java.lang.Integer</config-property-type>
   <config-property-value>100</config-property-value>
</config-property>
```

  Config properties can also be set specific to each connection definition and also for each of the connection definitions.

Note

The Message Driven Bean that is configured as an Endpoint to receive messages through an inbound resource adapter has to implement the messagelistener-type interface specified in the ra.xml.

```
<inbound-resourceadapter>
   <messageadapter>
      <messagelistener>
         <messagelistener-type> </messagelistener-type>
```

The EJB must include the activation configuration required to activate the endpoint.

Resource Adapter Classloader Considerations

Borland AppServer provides a per-module classloader policy: each deployed module, whether .ear, .war, .rar, or ejb.jar, has its own classloader whose classpath includes all the classes packaged in that module. This allows each module to completely control the classes it will have access to; several modules can have their own copies of a library in different versions, or can use the same package names without clashing with

packages used by other modules. This is a powerful feature, but can lead to complications when multiple modules need to interoperate with each other and share some classes.

While each module has its own classloader, the partition in which VisiConnect runs has several additional classloaders of its own. Classes in these per-partition classloaders are available to all modules running in the partition. Generally, the per-module classloaders take precedence over partition-level classloaders, so if a class is found in both, the per-module copy is the one that will be used.

There are two types of situations where VisiConnect users may need to exercise some care in packaging their applications to account for classloader issues:

- Connection factories and connections for outbound communication
- Message listeners for inbound communication

Connection Factories and Connections

The connection factory and connection interface and implementation classes for outbound communication may be provided along with the resource adapter. In some cases, the interface classes may in fact be standards-based and supplied as part of the JDK or an extension, while the implementation classes are proprietary and specific to the resource adapter. For example, a JMS Resource Adapter may provide its own proprietary implementation of the standard `javax.jms.QueueConnectionFactory` and `javax.jms.QueueConnection` classes.

Classes in the `javax.jms` package are provided by Borland AppServer and will be present in the partition's classloader, and all modules will share the same class definitions for those classes. But the proprietary classes that implement these `javax.jms` interfaces will be supplied by the Resource Adapter, and each module may have its own copy of these proprietary class definitions.

In order to get a connection to a Resource Adapter, the client program does a JNDI lookup for one of the Resource Adapter's connection factories. The object returned by this JNDI lookup must be created using the class definitions available to the `.ear`, `.war`, or `.jar` module where the client resides; otherwise, when the client code attempts to make use of the connection factory object, it will receive a `ClassCastException`. The client next uses the connection factory to create Connection instances. These objects, too, must be created using the client module's classloader, so that the client code can manipulate the objects.

At the same time, the VisiConnect service, running inside the AppServer, needs to make use of some of the classes provided by the Resource Adapter. For example, 1.5 level Resource Adapters contain an implementation of `javax.resource.spi.ResourceAdapter`, which will be used by VisiConnect to start and stop of the Resource Adapter. Any use by VisiConnect of the Resource Adapter classes will be done using the classloader of the Resource Adapter. If the Resource Adapter was deployed as a standalone `.rar` (rather than being embedded in an `.ear` with its clients), it will have its own classloader, and thus its own copies of the class definitions for the Resource Adapter classes. In this situation there may be a potential for `ClassCastExceptions`.

One such problem area is the method `ManagedConnectionFactory.setResourceAdapter(javax.resource.spi.ResourceAdapter)`. The `ManagedConnectionFactory` instance is created using the client classloader, while the ResourceAdapter instance is created using the `.rar` classloader. If the implementation of this method casts the ResourceAdapter instance to a proprietary implementation class, a `ClassCastException` will be thrown.

Message Listeners

An inbound resource adapter must specify a class to be its message listener. This class will be implemented by any MDB which is to serve as an endpoint for inbound communications from this Resource Adapter. When the Resource Adapter has a message which is to be passed on to the MDB, it will invoke a method in the message listener class. For example, many JMS Resource Adapters will use `javax.jms.MessageListener` as their message listener class, and the `onMessage(javax.jms.Message)` method in this class to actually receive the incoming

messages. Since all these classes are provided by the `javax.jms` package, which is in the partition's classloader and therefore shared by both the Resource Adapter and the MDB client, there is no possibility here of `ClassCastExceptions`.

However, a Resource Adapter is free to provide its own proprietary message listener class, and this class may have any number of methods to actually deliver messages, all of which may use proprietary objects as arguments. This may be a source of `ClassCastExceptions`.

VisiConnect will ensure that even if the message listener class is proprietary, calls to the MDB will be properly handled such that the message delivery method in the MDB is invoked using the definition of the message listener from the MDB's own classloader. However, if that method takes an argument which is a proprietary object, VisiConnect cannot map from the Resource Adapter's class definition of that object to the MDB's class definition. This will lead to `ClassCastExceptions`.

For example, the Mail Resource Adapter provided as a sample with the product provides a message listener class called `com.borland.enterprise.ra.mail.api.MailListener`. This class contains a message delivery method called `onMessage(javax.mail.Message)`. Notice that the message listener class is proprietary, but its `onMessage()` method takes a non-proprietary object as an argument. This situation will NOT cause `ClassCastExceptions`. The message listener class itself may be proprietary. Only if one or more arguments on a message delivery method are proprietary objects will there be a classloader problem.

Correcting ClassCastExceptions

In any of the problem situations described above, you have two basic possibilities for a resolution:

- Create an `.ear` that contains the Resource Adapter `.rar` and its clients. Since the entire `.ear` will share a classloader, there can be no `ClassCastExceptions` moving objects between the Resource Adapter and the client.
- Remove the Resource Adapter classes from both the `.rar` and the client modules, and deploy these classes to the partition as a library `.jar`. A library `.jar` will be placed in the partition's classloader, shared by all deployed modules. Therefore, all modules will use the same class definitions for the Resource Adapter classes, and no `ClassCastExceptions` will be thrown.

Developing the Resource Adapter

This section describes how to develop a Connectors 1.5-compliant Resource Adapter. Resource Adapters must implement the following system contract requirements, discussed in detail below:

- Connection management
- Security management
- Transaction management
- Packaging and deployment

Connection management

The connection management contract for the resource adapter specifies a number of classes and interfaces necessary for providing the system contract. The resource adapter must implement the following interfaces:

- `javax.resource.spi.ManagedConnection`
- `javax.resource.spi.ManagedConnectionFactory`
- `javax.resource.spi.ManagedConnectionMetaData`

The ManagedConnection implementation provided by the Resource Adapter must, in turn, supply implementations of the following interfaces and classes to provide support to the application server. It is the application server which will ultimately be managing the connection and associated transactions.

Note

If your environment is non-managed (that is, not managed by the application server), you are not required to use these interfaces or classes.

- `javax.resource.spi.ConnectionEvent`
- `javax.resource.spi.ConnectionEventListener`

In addition, support for error logging and tracing must be provided by implementing the following methods in the Resource Adapter:

- `ManagedConnectionFactory.setLogWriter()`
- `ManagedConnectionFactory.getLogWriter()`
- `ManagedConnection.setLogWriter()`
- `ManagedConnection.getLogWriter()`

The resource adapter must also provide a *default implementation* of the `javax.resource.spi.ConnectionManager` interface for cases in which the Resource Adapter is used in a non-managed two-tier application scenario. A default implementation of `ConnectionManager` enables the Resource Adapter to provide services specific to itself. These services can include connection pooling, error logging and tracing, and security management. The default `ConnectionManager` delegates to the `ManagedConnectionFactory` the creation of physical connections to the underlying EIS.

In an application-server-managed environment, the Resource Adapter should not use the default `ConnectionManager` implementation class. Managed environments do not allow resource adapters to support their own connection pooling. In this case, the application server is responsible for connection pooling. A Resource Adapter can, however, have multiple `ConnectionManager` instances per physical connection transparent to the application server and its components.

Transaction management

Resource Adapters are easily classified based on the level of transaction support they provide. These levels are:

- **NoTransaction**: the Resource Adapter supports neither local not JTA transactions, and implements no transaction interfaces.
- **LocalTransaction**: the Resource Adapter supports resource manager local transactions by implementing the `LocalTransaction` interface. The local transaction management contract is specified in Section 6.7 of the Connectors 1.5 specification from Sun Microsystems.
- **XATransaction**: the Resource Adapter supports both resource manager local and JTA/XA transactions by implementing the LocalTransaction and XAResource interfaces, respectively. The XA Resource-based contract is specified in Section 6.7 of the Connectors 1.5 specification from Sun Microsystems.

The transaction support levels above reflect the major steps of transaction support that a Resource Adapter must implement to allow server-managed transaction coordination. Depending on its transaction capabilities and the requirements of its underlying EIS, a Resource Adapter can choose to support any one of the above levels.

Security management

The security management contract requirements for a Resource Adapter are as follows:

- The Resource Adapter is required to support the security contract by implementing the `ManagedConnectionFactory.createManagedConnection()` method.
- The Resource Adapter is not required to support re-authentication as part of its `ManagedConnection.getConnection()` method implementation.
- The Resource Adapter is required to specify its support for the security contract as part of its deployment descriptor. The relevant deployment descriptor elements are:
  - `<authentication-mechanism></authentication-mechanism>`
  - `<authentication-mechanism-type></authentication-mechanism-type>`
  - `<reauthentication-support></reauthentication-support>`
  - `<credential-interface></credential-interface>`

Refer to section 10.3.1 of the Connectors 1.5 specification for more details on these descriptor elements.

Packaging and deployment

The file format for a packaged Resource Adapter module defines the contract between a Resource Adapter provider and a Resource Adapter deployer. A packaged Resource Adapter includes the following elements:

- Java classes and interfaces that are required for the implementation of both the Connectors system-level contracts and the functionality of the Resource Adapter
- Utility Java classes for the Resource Adapter
- Platform-dependent native libraries required by the Resource Adapter

- Help files and documentation
- Descriptive meta information that ties the above elements together

For more information on packaging requirements, refer to Section 10.3 and 10.5 of the Connectors 1.5 specification, which discuss deployment requirements and supporting JNDI configuration and lookup, respectively.

Deploying the Resource Adapter

Deployment of Resource Adapters is similar to deployment of EJBs, Enterprise Applications and Web Applications. As with these modules, a Resource Adapter can be deployed as an archive file or as an expanded directory. A Resource Adapter can be deployed either dynamically using the AppServer Console or the `iastool` utilities, or as a part of an EAR. See the Borland AppServer *User's Guide* for deployment details.

When a Resource Adapter is deployed, a name must be specified for the module. This name provides a logical reference to the Resource Adapter deployment that, among other things, can be used to update or remove the Resource Adapter. The AppServer implicitly assigns a deployment name that matches the filename of the RAR file or deployment directory containing the Resource Adapter. This logical name can be used to manage the Resource Adapter after the server has started. The Resource Adapter deployment name remains active in the AppServer until the module is undeployed.

Application development overview

Developing application components

Common Client Interface (CCI)

The client APIs used by application components for EIS access can be categorized as follows:

- The standard common client interface (CCI) defined in Section 9 of the Connectors 1.5 specification.
- A general client interface specific to the type of Resource Adapter and its underlying EIS. For example, JDBC is one such interface for RDBMSs.
- A proprietary client interface specific to the particular Resource Adapter and its underlying EIS. For example, the CICS Java Gateway is one such interface for the IBM CICS transaction processor, and the JFC for the SAP R/3 enterprise resource planner is another.

The Connectors 1.5 specification defines the CCI for EIS access. The CCI is a standard client API for application components that enables these and EAI frameworks to drive interactions across heterogeneous EISs. The CCI is primarily targeted for Enterprise Application Integration (EAI), third-party enterprise tool vendors, and migration of legacy modules to the J2EE Platform.

In the CCI, a connection factory is a public interface that enables connection to an EIS instance. The ConnectionFactory interface is implemented by the Resource Adapter to provide this service. An application looks up a ConnectionFactory instance in the JNDI namespace, and uses it to request to obtain EIS connections.

The application then uses the returned Connection interface to access the EIS. To provide a consistent application programming model across both CCI and EIS-specific APIs, the ConnectionFactory and Connection interfaces comply to the Interface Template design pattern. This defines the skeleton of the connection creation and connection closing, deferring the appropriate steps to subclasses. This allows for these interfaces to be easily extended and adapted to redefine certain steps of connection creation and closing without changing these operations' structure. For more information on the application of the Interface Template design pattern to these interfaces, refer to Section 5.5.1 in the Connectors 1.5 specification (`http://java.sun.com/j2ee/connector`).

Managed application scenario

The following steps are performed when a *managed application* requests to obtain a connection to an EIS instance from a connection factory, as specified in the `res-type` variable:

1. The application assembler or component provider specifies the connection factory requirements for an application component by using a deployment descriptor:

```
res-ref-name: shme/shmeAdapter
res-type:javax.resource.cci.ConnectionFactory
res-auth: Application|Container
```

2. The Resource Adapter deployer sets the configuration information for the Resource Adapter.
3. VisiConnect uses a configured Resource Adapter to create physical connections to the underlying EIS.
4. The application component performs a JNDI lookup of a connection factory instance in the component's environment:

```
// obtain the initial JNDI Naming context
 javax.naming.Context ctx = new javax.naming.InitialContext();
 // perform the JNDI lookup to obtain the connection factory
 javax.resource.cci.ConnectionFactory cxFactory =
        (javax.resource.cci.ConnectionFactory)ctx.lookup(
            "java:comp/env/shme/shmeAdapterConnectionFactory");
```

5. The JNDI name passed in the context lookup is that same as that specified in the `res-ref-element` of the component's deployment descriptor. The JNDI lookup returns a connection factory instance of type `java.resource.cci.ConnectionFactory` as specified in the `res-type` element.
6. The application component invokes the `getConnection()` method on the connection factory to request to obtain an EIS connection. The returned connection instance represents an application level handle to an underlying physical connection. An application component requests multiple connections by invoking the `getConnection()` method on the connection factory multiple times.

```
javax.resource.cci.Connection cx = cxFactory.getConnection();
```

7. The application component uses the returned connection to access the underlying EIS. This is specific to the Resource Adapter.
8. After the component finishes with the connection, it closes it using the `close()` method on the connection interface.

```
cx.close();
```

9. If the application component fails to close an allocated connection after its use, that connection is considered an unused connection. The AppServer manages to cleanup of unused connections. When the container terminates a component instance, the container cleans up all the connections used by that component instance.

Non-managed application scenario

In the non-managed application scenario, a similar programming model must be followed in the application component. The non-managed application must lookup a connection factory instance, request to obtain an EIS connection, use the connection for EIS interactions, and close the connection when completed.

The following steps are performed when a *non-managed application* component requests to obtain a connection to an EIS instance from a connection factory:

1. The application component calls the `getConnection()` method on the `javax.resource.cci.ConnectionFactory` instance to get a connection to the underlying EIS instance.

2 The connection factory instance delegates the connection request to the default connection manager instance. The Resource Adapter provides the default connection manager implementation.

3 The connection manager instance creates a new physical connection to the underlying EIS instance by calling the `ManagedConnectionFactory.createManagedConnection()` method.

4 Invoking `ManagedConnectionFactory.createManagedConnection()` creates a new physical connection to the underlying EIS, represented by the `ManagedConnection` instance it returns. The `ManagedConnectionFactory` uses the security information from the JAAS `Subject` object, and `ConnectionRequestInfo`, and its configured set of properties (port number, server name, etc.) to create the new `ManagedConnection` instance.

5 The connection manager instance calls the `ManagedConnection.getConnection()` method to get an application-level connection handle. This method call does not necessarily create a new physical connection to the EIS instance; it produces a temporary handle that is used by an application to access the underlying physical connection, represented by the `ManagedConnection` instance.

6 The connection manager instance returns the connection handle to the connection factory instance; the connection factory in turn returns the connection to the requesting application component.

Code excerpts—programming to the CCI

The following code excerpts illustrate the application programming model based on the CCI requesting to obtain a connection, obtaining the connection factory, creating the interaction and interaction spec, obtaining a record factory and records, executing the interaction with the records, and performing the same using result sets and custom records.

```
// Get a connection to an EIS instance after lookup of a connection factory
// instance from the JNDI namespace. In this case, the component allows the
// container to manage the EIS sign-on
javax.naming.Context ctx = new javax.naming.InitialContext();
javax.resource.cci.ConnectionFactory cxFactory =
(javax.resource.cci.ConnectionFactory)ctx.lookup(
  "java:comp/env/shme/shmeAdapter" );
javax.resource.cci.Connection cx = cxFactory.getConnection();

// Create an Interaction instance
javax.resource.cci.Interaction ix = ct.createInteraction();

// Create a new instance of the respective InteractionSpec
com.shme.shmeAdapter.InteractionSpecImpl ixSpec = new
com.shme.shmeAdapter.InteractionSpecImpl();
ixSpec.setFunctionName( "S_EXEC" );
ixSpec.setInteractionVerb( javax.resource.cci.InteractionSpec.SYNC_SEND_RECEIVE
);
// ...
// Get a RecordFactory instance
javax.resource.cci.RecordFactory recFactory = // ... get a RecordFactory

// Create a generic MappedRecord using the RecordFactory instance. This record
// instance acts as an input to the execution of an interaction. The name of
the
// Record acts as a pointer to the metadata for a specific record type
javax.resource.cci.MappedRecord input = recFactory.createMappedRecord(
"ShmeExecRecord" );

// Populate the generic MappedRecord instance with input values. The component
// code adds values based on the metadata it has accessed from the metadata
// repository
```

```
input.put( "<key: element0>", new String( "S_APP01"     );
input.put( "<key: element1>", // ... );
// ...

// Create a generic IndexedRecord to hold output values that are set by the
// execution of the interaction
javax.resource.cci.IndexedRecord output =
    recFactory.createIndexedRecord( "ShmeExecRecord" );

// Execute the Interaction
boolean response = ix.execute( ixSpec, input, output );

// Extract data from the output IndexedRecord. Note that type mapping is done
// in the generic IndexedRecord by mean of the type mapping information in the
// metadata repository. Since the component uses generic methods on the
// IndexedRecord, the component code performs the required type casting
java.util.Iterator iter = output.iterator();

while ( iter != null && iter.hasNext() )
{
  // Get a record element and extract value ...
}

// Set up the requirements for the ResultSet returned by the execution of
// an Interaction. This step is optional. Default values are used if
// requirements are not explicitly set.
com.shme.shmeAdapter.InteractionSpecImpl rsIxSpec =
    new com.shme.shmeAdapter.InteractionSpecImpl();
rsIxSpec.setFetchSize( 20 );
rsIxSpec,setResultSetType( javax.resource.cci.ResultSet.TYPE_SCROLL_INSENSITIVE
);

// Execute an Interaction that returns a ResultSet
javax.resource.cci.ResultSet rSet =
    (javax.resource.cci.ResultSet)ix.execute( rsIxSpec, input );

// Iterate over the ResultSet. The example here positions the cursor on the
// first row and then iterates forward through the contents of the ResultSet.
// Appropriate get methods are then used to retrieve column values.
rSet.beforeFirst();

while ( rSet != null && rSet.next() )
{
// get the column values for the current row using the appropriate
// get methods
}

// This illustrates reverse iteration through the ResultSet
rSet.afterLast();

while ( rSet.previous() )
{
// get the column values for the current row using the appropriate
// get methods
}

// Extend the Record interface to represent an EIS-specific custom Record.
// The interface CustomerRecord supports a simple accessor/mutator design
// pattern for its field values. A development tool would generate the
// implementation class of the CustomerRecord
public interface CustomerRecord extends javax.resource.cci.Record,
    javax.resource.cci.Streamable
{
  public void setName( String name );
```

```
  public void setId( String custId );
  public void setAddress( String address );

  public String getName();
  public String getId();
  public String getAddress();
}

// Create an empty CustomerRecord instance to hold output from
// the execution of an Interaction
CustomerRecord customer = // ... create an instance

// Create a PurchaseOrderRecord instance as an input to the Interaction
// and set properties on this instance. The PurchaseOrderRecord is another
// example of a custom Record
PurchaseOrderRecord purchaseOrder = // ... create an instance
purchaseOrder.setProductName( "..." );
purchaseOrder.setQuantity( "..." );
// ...

// Execute an Interaction that populates the output CustomerRecord instance
boolean crResponse = ix.execute( rsIxSpec, purchaseOrder, customer  );

// Check the CustomerRecord
System.out.println( "Customer Name = [" + customer.getName() + "],
                     Customer ID = [" + customer.getId() + "],
                     Customer Address = [" + customer.getAddress() + "]" );
```

Deployment Descriptors for Application Components

The application component deployment descriptors need to specify connection factory information for the Resource Adapter which the component will use. Appropriate entries are required in:

1 In the component's Sun standard deployment descriptor. For example, in ejb-jar.xml, the following is required:

   – `res-ref-name: shme/shmeAdapter`

   – `res-type: javax.resource.cci.ConnectionFactory`

   – `res-auth: Application|Container`

1 In addition, any version specific entries can be included. For example, EJB 2.0's res-sharing-scope:

   – `res-sharing-scope: Shareable|Unshareable`

1 In the component's Borland-specific deployment descriptor. For example, in ejb-borland.xml, the following is required:

   – `res-ref-name: shme/shmeAdapter`

   – `res-type: javax.resource.cci.ConnectionFactory`

1 In addition, any version specific entries can be included. For example, EJB 1.1's cmp-resource:

   – `cmp-resource: True|False`

The following details example deployment descriptors for two EJBs—the first written to the EJB 2.0 spec, the second written to the EJB 1.1 spec. Both the standard and Borland-specific deployment descriptors are shown. In these examples, a hypothetical Resource Adapter is referenced.

EJB 2.x example

ejb-jar.xml deployment descriptor

This example uses container-managed persistence

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise JavaBeans
2.0//EN" "http://java.sun.com/dtd/ejb-jar_2_0.dtd">
<ejb-jar>
    <display-name>SHME Integration Jar</display-name>
    <enterprise-beans>
        <session>
            <description>Interface EJB for shmeAdapter Class /shme/test/
              shmeAdapter/schema/Customer</description>
            <display-name>customer_bean</display-name>
            <ejb-name>shme/customer_bean</ejb-name>
            <home>com.shme.test.shmeAdapter.schema.CustomerHome</home>
            <remote>com.shme.test.shmeAdapter.schema.CustomerRemote</remote>
            <ejb-class>com.shme.test.shmeAdapter.schema.CustomerBean
              </ejb-class>
            <session-type>Stateful</session-type>
            <transaction-type>Container</transaction-type>
            <env-entry>
                <description>SHME Repository URL for Connector configuration
                   </description>
                <env-entry-name>repositoryUrl</env-entry-name>
                <env-entry-type>java.lang.String</env-entry-type>
                <env-entry-value>s_repository://S_APP01</env-entry-value>
            </env-entry>
            <env-entry>
                <description>Location of Resource Adapter Configuration within
                   the SHME Repository</description>
                <env-entry-name>configurationUrl</env-entry-name>
                <env-entry-type>java.lang.String</env-entry-type>
                <env-entry-value>/shme/client</env-entry-value>
            </env-entry>
            <resource-ref>
                <description>Reference to SHME Resource Adapter</description>
                <res-ref-name>shme/shmeAdapter</res-ref-name>
                <res-type>com.shme.shmeAdapter.ConnectionFactory</res-type>
                <res-auth>Container</res-auth>
                <res-sharing-scope>Shareable</res-sharing-scope>
            </resource-ref>
    </session>
    </enterprise-beans>
    <assembly-descriptor>
        <container-transaction>
            <method>
                <ejb-name>customer_bean</ejb-name>
                <method-intf>Remote</method-intf>
                <method-name>s_exec_customer_query</method-name>
                <method-params/>
            </method>
            <trans-attribute>Required</trans-attribute>
        </container-transaction>
    </assembly-descriptor>
</ejb-jar>
```

This corresponds to the ejb-jar.xml above.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE ejb-jar PUBLIC "-//Borland Software Corporation//DTD Enterprise
   JavaBeans 2.0//EN" "http://www.borland.com/devsupport/appserver/dtds/
   ejb-jar_2_0-borland.dtd">
<ejb-jar>
    <enterprise-beans>
        <session>
            <ejb-name>shme/customer_bean</ejb-name>
            <bean-home-name>shme/customer_bean</bean-home-name>
            <resource-ref>
                <res-ref-name>shme/shmeAdapter</res-ref-name>
                <jndi-name>eis/shmeAdapter</jndi-name>
            </resource-ref>
        </session>
    </enterprise-beans>
</ejb-jar>
```

EJB 1.1 example

ejb-jar.xml deployment descriptor

This example uses bean-managed persistence.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE ejb-jar PUBLIC '-//Sun Microsystems, Inc.//DTD Enterprise JavaBeans
   1.1//EN' 'http://java.sun.com/j2ee/dtds/ejb-jar_1_1.dtd'>
<ejb-jar>
 <description />
     <display-name>ShmeAdapter Interface Jar</display-name>
     <small-icon />
  <large-icon />
     <enterprise-beans>
      <session>
  <description>Interface EJB for SHME Class /shme/test/shmeAdapter/schema/
      Customer</description>
   <display-name>customer_bean</display-name>
  <ejb-name>shme/customer_bean</ejb-name>
 <home>com.shme.test.shmeAdapter.schema.CustomerHome</home>
  <remote>com.shme.test.shmeAdapter.schema.CustomerRemote</remote>
   <ejb-class>com.shme.test.shmeAdapter.schema.CustomerBean</ejb-class>
   <session-type>Stateless</session-type>
   <transaction-type>Bean</transaction-type>
   <env-entry>
         <description>SHME Repository URL for Connector configuration
           </description>
      <env-entry-name>repositoryUrl</env-entry-name>
   <env-entry-type>java.lang.String</env-entry-type>
  <env-entry-value>s_repository://S_APP01</env-entry-value>
   </env-entry>
 <env-entry>
  <description>Location of Resource Adapter configuration within the SHME
Repository</description>
     <env-entry-name>configurationUrl</env-entry-name>
  <env-entry-type>java.lang.String</env-entry-type>
  <env-entry-value>/shme/client</env-entry-value>
  </env-entry>
 <resource-ref>
  <description>Reference to SHME Resource Adapter</description>
 <res-ref-name>shme/shmeAdapter</res-ref-name>
 <res-type>com.shme.shmeAdapter.ConnectionFactory</res-type>
```

```
 <res-auth>Container</res-auth>
 </resource-ref>
   </session>
</enterprise-beans>
<ejb-client-jar />
</ejb-jar>
```

ejb-inprise.xml deployment descriptor

This corresponds to the ejb-jar.xml above.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE inprise-specific PUBLIC '-//Inprise Corporation//DTD Enterprise
   JavaBeans 1.1//EN' 'http://www.borland.com/devsupport/appserver/dtds/
   ejb-inprise.dtd'>
<inprise-specific>
    <enterprise-beans>
        <session>
            <ejb-name>shme/customer_bean</ejb-name>
            <bean-home-name>shme/customer_bean</bean-home-name>
            <timeout>0</timeout>
            <resource-ref>
                <res-ref-name>shme/shmeAdapter</res-ref-name>
                <jndi-name>eis/shmeAdapter</jndi-name>
                <cmp-resource>False</cmp-resource>
            </resource-ref>
        </session>
    </enterprise-beans>
</inprise-specific>
```

Other Considerations

Working with Poorly Implemented Resource Adapters

Some commercially available Resource Adapters may be poorly implemented. As there does not yet exist any mechanism to test a Resource Adapter for compliance to the Connectors specs (as the J2EE Compatibility Test Suite (CTS) tests a Connectors implementation for spec compliance), it is currently not a simple task to recognize, but among the symptoms, you will find:

1. The Resource Adapter will exhibit strange errors during deployment
2. The Resource Adapter will exhibit strange errors during method invocation on the connection factory.

As VisiConnect strictly implements J2EE 1.4 and Connectors 1.5 requirements, it is often the only Connector Container which will detect poorly implemented Resource Adapters and not ignore the problem.

Examples of Poorly Implemented Resource Adapters

Generally, poorly implemented Resource Adapters are not compliant with the Connectors 1.5 specification. Examples of such Resource Adapters include:

- The Resource Adapter with a connection factory implementing only `java.io.Serializable`, and not both `java.io.Serializable` and `javax.resource.Referenceable` as per the Connectors specification (Section 10.5 “JNDI Configuration and Lookup”). The local JNDI context handlers of application servers such as AppServer can only register objects if they implement both interfaces. If a Resource Adapter implements a connection factory as `Serializable`, and doesn't implement `Referenceable`, you will see exceptions thrown when the application server attempts to deploy the connection factory to JNDI.
- The Resource Adapter with a connection factory which poorly implements `javax.resource.Referenceable` (which inherits `getReference()` from `javax.naming.Referenceable`). The J2SE 1.3.x and 1.4.x specs specify that for `javax.naming.Referenceable`, `getReference()` either:
  a. Returns a valid, non-null reference of the `Referenceable` object, or
  b. Throws an exception (`javax.naming.NamingException`).

  If the Resource Adapter implements `Referenceable` such that `getReference()` can (and will) return `null`, you will see exceptions thrown when a client attempts to invoke a connection factory method such as `getConnection()`.
- The Resource Adapter with a connection factory correctly implementing `Referenceable`, but which does not provide an implementation of `javax.naming.spi.ObjectFactory` (which is required by the Connectors specification (Section 10.5 “JNDI Configuration and Lookup”)). Although such a Resource Adapter can be deployed to an application server without incident, it cannot be deployed to JNDI outside the aegis of an application server, as a non-managed Connector. Also, including a `javax.naming.spi.ObjectFactory` implementation source Adapter with backup mechanism for JNDI `Reference`-based connection factory lookup.
- The Resource Adapter which specifies an connection factory or connection interface while not implementing that interface in its connection factory or connection class, respectively. Section 10.6 “Resource Adapter XML DTD” in the Connectors spec

discusses the related requirements. To illustrate, let's say that in the `ra.xml` of a particular Resource Adapter, you have the following elements:

```
//...
<connection-interface>java.sql.Connection</connection-interface>
<connection-impl-class>com.shme.shmeAdapter.ShmeConnection</connection-impl-
class>
//...
```

But your implementation of `ShmeConnection` is as follows:

```
package shme;
public class ShmeConnection
{
private ShmeManagedConnection mc;
  public ShmeConnection( ShmeManagedConnection mc )
  {
   System.out.println( "In ShmeConnection" );
   this.mc = mc;
  }
}
```

Any attempt to invoke `getConnection()` on this Resource Adapter's connection factory will result in a `java.lang.ClassCastException`, as you're indicating to the appserver in `ra.xml` that connection objects returned by the Resource Adapter are to be cast to `java.sql.Connection`.

Working with a Poor Resource Adapter Implementation

To work around a poor Resource Adapter implementation, perform the following:

Extend the connection factory and/or connection class of the Connector, and have the extension correctly implement the poorly implemented code. For example, when dealing with a connection factory which implements `Serializable`, and doesn't implement `Referenceable` the idea is to extend the original connection factory to implement `Referenceable`, which means implementing `getReference()` and `setReference()`.

To illustrate, if the connection factory is `com.shme.BadConnectionFactory`, extend the connection factory as `com.shme.GoodConnectionFactory`, and implement `Referenceable` as follows:

```
package com.shme.shmeAdapter;

public class GoodConnectionFactory
{
 private javax.naming.Reference ref;
// ...
 public javax.naming.Reference getReference()
{
// implement such that getReference() never returns null
// ...
  return ref;
    }
public javax.naming.Reference setReference( javax.naming.Reference ref )
 //  this.ref = ref;
  }
 //
```

Also, when dealing with a poorly behaving `getReference()`, there are various ways to accomplish this, but principally, the idea is to implement `getReference()` such that it never returns `null`. The best approach is to implement:

– A fallback mechanism in `getReference()` which sets the reference to be returned correctly if the connection factory's reference attribute is `null`—returning a registerable javax.naming.Reference object, and

– A helper class implementing `javax.naming.spi.ObjectFactory` to provide the fallback object to create the connection factory object from the valid Reference instance.

To illustrate, if the connection factory is `com.shme.BadConnectionFactory`, extend the connection factory as `com.shme.GoodConnectionFactory`, and override `getReference()` as follows:

```
package com.shme.shmeAdapter;

public class GoodConnectionFactory
{
  // ...

  public javax.naming.Reference getReference()
  {
    if ( ref == null )
    {
      ref = new javax.naming.Reference( this.getClass().getName(),
        "com.shme.shmeAdapter.GoodCFObjectFactory"
          /* object factory for GoodConnectionFactory references */,
            null );
      String value;
      value = managedCxFactory.getClass().getName();

      if ( value != null )
      {
        ref.add( new javax.naming.StringRefAddr(
          "managedconnectionfactory-class", value ) );
      }

      value = cxManager.getClass().getName();

      if ( value != null )
      {
        ref.add( new javax.naming.StringRefAddr(
          "connectionmanager-class", value ) );
      }
    }

    return ref;
  }

  // ...
}
```

Then implement the associated object factory class, in this case:

com.shme.shmeAdapter.GoodCFObjectFactory

```
package com.shme.shmeAdapter;

import javax.naming.spi.*;
import javax.resource.spi.*;

public class GoodCFObjectFactory implements ObjectFactory {
  public GoodCFObjectFactory() {};

  public Object getObjectInstance( Object obj,
                                   javax.naming.Name name,
                                   javax.naming.Context context,
                                   java.util.Hashtable env )
         throws Exception
  {
    if ( !( obj instanceof javinstance ofReference ) )
    {
      return null;
    }

    javax.naming.Reference ref = (javax.naming.Reference)obj;

    if ( ref.getClassName().equals(
"com.shme.shmeAdapter.GoodConnectionFactory" ) )
    {
      ManagedConnectionFactory refMcf = null;
      ConnectionManager refCm = null;

if ( ref.get( "managedconnectionfactory-class" ) != null )
      {
        String managedCxFactoryStr =
         (String)ref.get( "managedconnectionfactory-class" ).getContent();
        Class mcfClass = Class.forName( managedCxFactoryStr );
        refMcf = (ManagedConnectionFactory)mcfClass.newInstance();
      }

if ( ref.get( "connectionmanager-class" ) != null )
      {
        String cxManagerStr = (String)ref.get( "connectionmanager-class"
).getContent();
        Class cxmClass = Class.forName( cxManagerStr  );
        java.lang.ClassLoader cloader = cxmClass.getClassLoader();
        refCm = (ConnectionManager)cxmClass.newInstance();
      }

      GoodConnectionFactory cf = null;

      if ( refCm != null )
      {
        cf = new GoodConnectionFactory( refMcf, refCm  );
      }
      else
      {
        cf = new GoodConnectionFactory( refMcf );
      }

    return cf;
    }

    return null;
  }
}
```

Update the classes in the ra.xml standard deployment descriptor file. For example, before extending the implementation, the `ra.xml` may look something like this:

```
<managedconnectionfactory-class>com.shme.shmeAdapter.
LocalTxManagedConnectionFactory</managedconnectionfactory-class>
<connectionfactory-interface>javax.sql.DataSource</connectionfactory-interface>
<connectionfactory-impl-class>com.shme.shmeAdapter.BadConnnectionFactory</
connectionfactory-impl-class>
<connection-interface>java.sql.Connection</connection-interface>
<connection-impl-class>com.shme.Connection</connection-impl-class>
```

After extending the interfaces, the ra.xml may look something like this:

```
<managedconnectionfactory-class>com.shme.shmeAdapter.
LocalTxManagedConnectionFactory </managedconnectionfactory-class>
<connectionfactory-interface>javax.sql.DataSource</connectionfactory-interface>
<connectionfactory-impl-class>com.shme.shmeAdapter.GoodConnectionFactory</
connectionfactory-impl-class>
<connection-interface>java.sql.Connection</connection-interface>
<connection-impl-class>com.shme.shmeAdapter.Connection</connection-impl-class>
```

As this illustrates, this conversion only impacts the connection factory. No other Resource Adapter classes are affected by this conversion.

Compile the Java code for the extended implementation (and any helper classes) into class files.

Package these into the Resource Adapter's Java Archive (.jar) file.

Update the Resource Adapter Archive (.rar) file with this extended .jar.

Deploy the Resource Adapter Archive, or include it in an Enterprise Application Archive (.ear) file to be deployed as part of a J2EE application, to VisiConnect running standalone or as a Partition service in the AppServer.

You've now converted a badly behaving Resource Adapter into a well behaving one.

Sometimes the design of a Resource Adapter makes it impossible to extend the existing API implementation. In such cases you need to re-implement the offending class or classes, and set the elements in `ra.xml` to reference the re-implementation(s). Or better yet, choose another Resource Adapter, which is compliant with the Connectors specification to work with.

35

Borland AppServer Ant tasks and running AppServer examples

Many of the Borland AppServer (AppServer) examples now employ the Ant build script system. In addition to Ant's core functionality, the Borland AppServer version of Ant includes customized tasks for several of the AppServer command line tools, including commands of the following:

- `appclient`
- `iastool`
- `idl2java`
- `java2iiop`

These customized Ant tasks have the following advantages over using `exec` or `apply` directives:

- Customized Ant tasks run under the VM used to launch the Ant script, hence they run faster and use less memory compared to spawning new JVM's with the `exec`/`apply` commands.
- Customized tasks have a much simpler command syntax than the exec/apply version.
- Ant features such as filesets and patternsets are available in a more natural way.

General syntax and usage

The following table shows the currently defined Ant tasks and their relationship to the equivalent command line tools.

Ant task name	Equivalent command line	Function
appclient	appclient	Runs a client application.
idl2java	idl2java	Converts IDL to Java classes.
java2iiop	java2iiop	Executes the java2iiop command.
iastool	iastool	Runs iastool.

Generally the AppServer Ant task uses the same pattern as the command-line tool equivalent.

Name-value pair transformation

All name-value pair command line arguments can be transformed into Ant task attributes. The name-value pair command-line arguments should be translated into equivalent XML attributes. For example, the command line:

```
iastool -verify -src cart_beans_client.jar -role DEVELOPER
```

translates into the Ant task:

```
<iastool option="verify" src="cart_beans_client.jar" role="DEVELOPER" />
```

Name-only argument transformation

All name-only command-line arguments can be transformed into boolean type Ant task attributes. For boolean-style attributes, such as -nowarn, use the default usage of the attribute, according to the usage documented for each of the command line tools. See “iastool command-line utility” for more information about iastool command line attributes.

For example, the following command sets the warn attribute to false:

```
iastool -verify -src cart_beans_client.jar -role DEVELOPER -nowarn -nostrict
```

The equivalent Ant task is:

```
<iasverify src="cart_beans_client.jar" role="DEPLOYER" nowarn="true"
strict="false" />
```

Note

It is **not** valid to use “warn” as an attribute in the Ant task. For instance, the following line causes a syntax error:

```
****** INCORRECT SYNTAX!!! ******
<iasverify src="cart_beans_client.jar" role="DEPLOYER" warn="false"
strict="false" />
```

Multiple File Arguments

Many commands either act on multiple files or have options which can point to multiple files. There are several ways to achieve this functionality in the equivalent Ant task. For example, the `iastool -merge` command:

```
iastool -merge -target build\client.jar -type lib client\build\local_client.jar
build\local_stubs.jar
```

has the Ant equivalent:

```
<iastool option="merge" target="${build.dir}/client.jar" type="lib"
 jars="client/build/local_client.jar ; build/local_stubs.jar" />
```

Note

The files in the `jars` attribute must be separated by semi-colons (;) or colons (:)—spaces and commas are **not** valid separators.

Ant provides a convenient `<fileset>` task to include multiple files:

```
<iastool option="merge" target="build/client.jar" type="lib" >
    <fileset dir="client/build" includes="local_client.jar" />
    <fileset dir="build" includes="local_stubs.jar" />
</iastool>
```

The patternset feature of Ant can also be useful. The following alteration now includes all of the jar files contained in the build directory and all of its sub-directories:

```
<iastool option="merge" target="${build.dir}/client.jar" type="lib" >
    <fileset dir="${build.dir}" includes="**/*.jar" />
</iastool>
```

Class path attributes can include multiple paths separated by semicolons:

```
<iastool option="verify" src="cart_beans_client.jar" role="DEPLOYER"
classpath="alib.jar;blib.jar" />
```

or use the <classpath> element:

```
<iastool option="verify" src="cart_beans_client.jar" role="DEPLOYER" >
   <classpath>
       <pathelement location="alib.jar" />
       <pathelement location="blib.jar" />
   </classpath>
</iastool>
```

Syntax and usage for iastool

The iastool Ant tasks can use two different styles:

1 <iastool option="*myoption*" />

2 <ias*myoption* />

For example, the command line:

```
iastool -verify -src cart_beans_client.jar
```

translates into the Ant task:

```
<iastool option="verify" src="cart_beans_client.jar" />
```

or you can use the older Ant style for backward compatible Ant tasks:

```
<iasverify src="cart_beans_client.jar" />
```

The following table shows the Ant task styles for each of the iastool options.

iastool Ant task style 1	iastool Ant task style 2	Equivalent command line	Function	Fileset attribute
<iastool option= "compilejsp" />	<iascompilejsp />	iastool -compilejsp	Precompiles JSP's.	
<iastool option= "compress" />	<iascompress />	iastool -compress	Compresses a JAR file.	
<iastool option= "deploy" />	<iasdeploy />	iastool -deploy	Deploys a J2EE module.	jars
<iastool option= "dumpstack" />	<iasdumpstack />	iastool -dumpstack	Dumps the stack trace of a Partition process to the stdout.log, located in: <install_dir>/var/domains/ <domain-name>/configurations/ <configuration-name>/ mos/<partition-name>/adm/logs/ <partition_name>.stdout.log	
<iastool option= "genclient" />	<iasgenclient />	iastool -genclient	Generates a client library.	jars
<iastool option= "gendeployable" />	<iasgendeployable />	iastool -gendeployable	Generates a manually deployable module.	
<iastool option= "genstubs" />	<iasgenstubs />	iastool -genstubs	Generate a stub library.	

iastool Ant task style 1	iastool Ant task style 2	Equivalent command line	Function	Fileset attribute
<iastool option= "info" />	<iasinfo />	iastool -info	Displays system configuration information.	
<iastool option= "kill" />	<iaskill />	iastool -kill	Kills a Managed Object.	
<iastool option= "listhubs" />	<iaslisthubs />	iastool -listhubs	Lists available Hubs on a management port.	
<iastool option= "listpartitions" />	<iaslistpartitions />	iastool -listpartitions	Lists the Partitions running on a Hub.	
<iastool option= "listservices" />	<iaslistservices />	iastool -listservices	Lists the services running on a Hub.	
<iastool option= "manage" />	<iasmanage />	iastool -manage	Actively manage a Managed Object.	
<iastool option= "merge" />	<iasmerge />	iastool -merge	Merges a set of JAR files into a single JAR file.	jars
<iastool option= "migrate" />	<iasmigrate />	iastool -migrate	Migrates a module from J2EE 1.2 to J2EE 1.3.	
<iastool option = "newconfig" />	<iasnewconfig />	iastool –newconfig	Creates a new configuration.	
<iastool option= "patch" />	<iaspatch />	iastool -patch	Applies a patch to a JAR file.	
<iastool option= "ping" />	<iasping />	iastool -ping	Pings a Managed Object or Hub for its current state.	
<iastool option= "pservice" />	<iaspservice />	iastool -pservice	Enables, disables, or gets the state of a Partition service.	
<iastool option= "removestubs" />	<iasremovestubs />	iastool -removestubs	Removes stubs from a JAR.	
<iastool option= "restart" />	<iasrestart />	iastool -restart	Restarts a Managed Object.	
<iastool option= "setmain" />	<iassetmain />	iastool -setmain	Sets the main class of a client JAR or a JAR in an EAR.	
<iastool option= "stop" />	<iasstop />	iastool -stop	Stops a Managed Object.	
<iastool option= "uncompress" />	<iasuncompress />	iastool -uncompress	Uncompresses a JAR file.	
<iastool option= "undeploy" />	<iasundeploy />	iastool -undeploy	Undeploys a Managed Object.	
<iastool option= "unmanage" />	<iasunmanage />	iastool -unmanage	Removes a Managed Object from active management.	
<iastool option= "verify" />	<iasverify />	iastool -verify	Verifies a J2EE module.	

Note

The Fileset attributes column indicate attributes which can accept multiple file names. Such attributes can employ the Ant `<fileset>` element to designate these files. Techniques for including multiple files is explained in the “Multiple File Arguments”.

Omitting attributes

Omitting an attribute from the Ant task call has the same effect as omitting the option from the command line tool. Since some attributes are `true` by default, omitting an attribute does **not** necessarily set the attribute to `false`.

For more information on the default values of these options, see “iastool command-line utility”.

Examples of iastool Ant tasks

The following are a few examples to illustrate the usage details for iastool Ant tasks. For more details about the function of each iastool option and attribute, see “iastool command-line utility”.

deploy

```
<target name="deploy" description="Deploys the example to the server">
<iastool option="deploy" hub="${hub.name}" cfg="${cfg.name}"
 partition="${partition.name}" mgmtport="${default.mgmtport}"
 jars="${build.dir}/hello.ear;${bes.lib.dir}/../var/repository/archives/wars/
   bank_form.war"
 realm="${realm.name}" user="${server.user.name}" pwd="${server.user.pwd}"/>
</target>
```

merge

```
<iastool option="merge" target="${build.dir}/helloclient.jar" type="lib">
    <fileset dir="${build.dir}" includes="hello_stubs.jar" />
</iastool>
```

ping

```
<target name="ping">
<iastool option="ping" hub="${hub.name}" cfg="${cfg.name}"
 partition="${partition.name}" mgmtport="${default.mgmtport}"
 realm="${realm.name}" user="${server.user.name}" pwd="${server.user.pwd}" />
</target>
```

restart

```
<target name="iastoolrestart">
<iastool option="-restart" hub="${hub.name}" cfg="${cfg.name}"
 partition="${partition.name}" mgmtport="${default.mgmtport}"
 realm="${realm.name}" user="${server.user.name}" pwd="${server.user.pwd}" />
</target>
```

Syntax and usage for java2iiop

The `java2iiop` Ant task is very different from its command line tool. It is an exception to the Borland Ant task usage pattern. Ant task `java2iiop` takes classes in a directory instead of an individual file. The classpath attribute points to the directory containing the classes that need to be compiled by `java2iiop`. That classpath is Path-Like structure in Ant, and the usage of it is very flexible, but to use classpath with the `java2iiop` task you can only use one of the following styles:

1 Used as an attribute, its value only accepts colon- or semicolon-separated lists of locations:

```
<java2iiop classpath="${path1}:${path2}"/>
```

2 Used as nested classpath element.
This takes the general form of:

```
<java2iiop>
   <classpath>
      <pathelement path="${path1}"/>
      <pathelement location="lib/helper.jar"/>
   </classpath>
</java2iiop>
```

The location attribute specifies a single file or directory relative to the project's base directory (or an absolute filename), while the path attribute accepts colon- or semicolon-separated lists of locations. The path attribute is intended to be used with predefined paths. In any other case, multiple elements with location attributes should be preferred.

For details on the equivalent command line arguments for java2iiop, see “Programmer tools for Java” in the *VisiBroker for Java Developer’s Guide*.

Example of java2iiop Ant task

```
<target name="create_ejb_stubs" depends="home">
    <java2iiop root_dir="${stubsPath}" list_files="true"
classpath="${outputPath}" />
</target>
```

Syntax and usage for idl2java

The `idl2java` Ant task is similar to its equivalent command tool. It tasks nested Path-Like structure filesets which are equivalent to command line file inputs.

```
<idl2java>
    <fileset dir="server" includes="*.idl" />
</idl2java>
```

For details on the equivalent command line arguments for idl2java, see “Programmer tools for Java” in the *VisiBroker for Java Developer’s Guide*.

Attribute	Type	Required
classpath	Path	Yes
back_Compat_Mapping	boolean	No
bind	boolean	No
boa	boolean	No
comments	boolean	No
compile	boolean	No
compiler	String	No
destDir	Path	No

Attribute	Type	Required
dynamic_Marshal	boolean	No
examples	boolean	No
export_All	boolean	No
exported	String	No
gen_Included_Files	boolean	No
idl2package	String	No
idl_Strict	boolean	No
import	String	No
imported	String	No
include	File	No
invoke_Handler	boolean	No
line_Directives	boolean	No
list_Files	boolean	No
list_Includes	boolean	No
map_Keyword	String	No
narrow_Compliance	boolean	No
obj_Wrapper	boolean	No
object_Method	boolean	No
package	String	No
retain_Comments	boolean	No
root_Dir	File	No
sealed	String	No
servant	boolean	No
srcDir	Path	No
srcFile	String	No
stream_Marshal	boolean	No
strict	boolean	No
tie	boolean	No
undefine	String	No
VBJclassPath	Path	No
VBJdebug	String	No
VBJjavaVM	File	No
VBJprop	String	No
VBJquoteSpaces	String	No
VBJtag	String	No
version	String	No
warn_Missing_Define	String	No
warn_Unrecognized_Pragmas	boolean	No

Example of idl2java Ant task

```
<target name="idl2java" depends="init">
    <idl2java package="com.borland.examples.webservices.visibroker"
root_dir="${server-skel-src}">
        <fileset dir="server" includes="*.idl" />
    </idl2java>
    <javac srcdir="${server-skel-src}" destdir="${server-classes}"
classpathref="classpath"/>
</target>
```

Syntax and usage for appclient

```
<!-- Execute the example. -->
<target name="execute" description="Executes the Hello World example">
    <appclient jar="${build.dir}/hello.ear" uri="helloclient.jar" args="World"/
>
</target>
```

Building and running the Borland AppServer examples

Note

Many of the AppServer examples have their own `readme.html` files located in:

```
<install_dir>/examples
```

To build an AppServer example:

1. Open a command line window.
2. Set the current directory to an example directory. The “Hello World” example located at `<install_dir>/examples/j2ee/hello` is a good place to start.
3. On the command line, enter “ant”.

   The example should build automatically.

Note

The server does not have to be running to build the example. However, deployment and undeployment require that the server be operational. Executing an example requires that the Partition be running.

Deploying the example

1. Make sure that a server is running.
2. On the command line, enter `ant deploy`.

   This will deploy the example to the Hub, Configuration, and Partition set in the <install_dir>\examples\deploy.properties file.

If you wish to deploy to a different combination of Hub/Configuration/Partition, you can either edit the `deploy.properties` file to change the settings, or use `-D` options on the command line to override the `deploy.properties` settings.

For example, to use a Hub named “myhub", use the command:

```
ant -Dhub.name=myhub deploy
```

This will override the default Hub name in `deploy.properties` with the value `myhub`.

Running the example

1. Make sure that the Partition is running.
2. On the command line, enter `ant execute`.

   The precise response depends on the particular example.

Undeploying the example

1. Make sure that a server is running.

2 On the command line, enter `undeploy`.

Troubleshooting

1 Make sure that the `<appserver_install_dir>/bin` directory is on your path and precedes the path to any alternative Ant installations.

2 Before calling the `ant execute` command, make sure that the server and the Partition are running.

3 The `<appserver_install_dir>\examples\deploy.properties` contains default settings for the Hub, Configuration, Partition, and Management Port. These default properties include:

- `hub.name=`*`your_machine_name`*
- `cfg.name=j2ee`
- `partition.name=standard`
- `realm.name=ServerRealm`
- `server.user.name=admin`
- `server.user.pwd=admin`

where *your_machine_name* is the machine name designated at installation. You can reset these values as needed or specify them on the Ant command line using the `-D` option.

36

iastool command-line utility

This section describes the `iastool` command-line utility that you can use to manage your managed objects.

Using the iastool command-line tools

The `iastool` utility is a set of command-line tools for manipulating managed objects. The following table shows the command-line tools provided with the `iastool` utility:

Use...	To...
`-clonepartition`	Clone an existing partition and creates a new one.For more information, see "clonepartition".
`-compilejsp`	Precompile JSPs in a standalone WAR or all WARs in an EAR. For more information, see "compilejsp".
`-compress`	Compress a JAR file. For more information, see "compress".
`-deploy`	Deploy a J2EE module to the specified Partition. For more information, see "deploy".
`-dumpstack`	Dump the stack trace of a Partition process to the Partition stdout.log file. For more information, see "dumpstack".
`-genclient`	Generate a library containing client stubs, EJB interfaces, and dependent classes. For more information, see "genclient".
`-gendeployable`	Generate a manually deployable module. For more information, see "gendeployable".
`-gendeployableverify`	Generate a manually deployable module and verifies it. For more information, see "gendeployableverify".
`-genstubs`	Generate a library containing client or server stubs only. For more information, see "genstubs".
`-info`	Display system configuration information. For more information, see "info".
`-jndinamespace`	Display the list of objects bound to a context. For more information, see "jndinamespace".
`-kill`	Kill a managed object. For more information, see "kill".
`-listpartitions`	List the partitions on a hub. For more information, see "listpartitions".
`-listhubs`	List the available hubs on a management port. For more information, see "listhubs".
`-listservices`	List the services on a hub. For more information, see "listservices".

Use...	To...
-manage	Actively manage a managed object. For more information, see "manage".
-merge	Merge a set of JAR files into a single JAR file. For more information, see "merge".
-migrate	Migrate a module from J2EE 1.2, 1.3, or 1.4 to an alternative target J2EE version. For more information, see "migrate".
-newconfig	Create a new configuration from a configuration template. For more information see "newconfig".
-patch	Apply one or more patches to a JAR file. For more information, see "patch".
-ping	Ping a managed object or hub for its current state. For more information, see "ping".
-pservice	Enables, disables, or gets the state of a partition service. For more information, see "pservice".
-removestubs	Remove all stub files from a JAR file. For more information, see "removestubs".
-restart	Restart a hub or managed object. For more information, see "restart".
-setmain	Set the main class of a standalone Client JAR or a Client JAR in an EAR. For more information, see "setmain".
-start	Start a managed object. For more information, see "start".
-stop	Stop a hub or managed object. For more information, see "stop".
-uncompress	Uncompress a JAR file. For more information, see "uncompress".
-undeploy	Remove a J2EE module from a Partition. For more information, see "undeploy".
-unmanage	Remove a managed object from active management. For more information, see "unmanage".
-usage	Display the usage of command-line options. For more information, see "usage".
-verify	Verify a J2EE module. For more information, see "verify".

clonepartition

Clones an existing partition and creates a new one. This utility is a replica of the clone functionality in the Borland Management Console. The new partition name must be unique.

Syntax

```
-clonepartition [-hub <hub name>] [-cfg <configuration name>] [-mgmtport
<99999>]\n  [-bare] [-realm <realm>] [-user <username>] [-pwd <password>] [-
file <login_file>]\n  -osagent <osagent port number> [-partition<name partition
to be cloned>]   [-newpartition<name of new partition>]\n [-newdesc
<description of new partition>] -newdisplayname<display name of new partition>
-newgroup<group name for creating the new partition like {hubname}/
{configurationname}> -moagent<agent name> [-targetagent <target agent name>] -
jmxport<jmx port number default:8082> -tomcatport<Tomcat port number
default:8080>
```

Default Output

Creates a partition with default values.

Options

The following table describes the options available when using the `compilejsp` tool.

Option	Description
`-hub <hub>`	Specifies the hub name for which you want to list the running partitions.
`-cfg <configname>`	Specifies the name of the configuration on which to list partitions.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. The default value is 42424.
`-bare`	Suppresses the output information, other than the names of the running partitions.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, username, and password used to authenticate a user. You must specify the complete or relative path to this file.
`-osagent`	Specifies the osagent port.
`-partition`	Specifies the name of the partition that has to be cloned.
`-newpartition`	Specifies the name of the new partition.
`-newdesc`	Description of the new partition.
`-newdisplayname`	Display name for the name partition.
`-newgroup`	Specifies the group name where the partition needs to be created (hub/configuration).
`-moagent`	Specifies the agent with which the managed object is associated.
`-targetagent`	Specifies the target agent where the partition needs to be created.
`-jmxport`	Specifies the port number for java management console. The default value is 8082.
`-tomcatport`	Specifies the port number for the tomcat service. The default value is 8080.

Example

This example illustrates the usage of the options for clone partition. The example clones “MyPartition” and creates “VJ_partition” with specified port numbers under the given configuration and agent.

```
-clonepartition -hub gvijay -cfg j2eeSample -mgmtport 20001 -realm ServerRealm
-user admin -pwd admin -osagent 19999 -partition MyPartition -newpartition
VJ_partition  -newgroup gvijay/j2eeSample -moagent gvijay -jmxport 8084 -
tomcatport 8086
```

compilejsp

Use this tool to precompile JSP pages in a standalone WAR or in all WARs in an EAR. The JSP pages are compiled into Java servlet classes and saved in a WAR file. This operation enables the JSP pages to be served faster the first time they are accessed.

Note

When compiling JSPs using the iastool, you may encounter an out-of-memory error. Increase the size of the virtual memory on your system to solve this issue.

Syntax

```
-compilejsp -src <war_or_ear> -target <target_file> [-overwrite]
[-package <package_root>] [-excludefile <exclude_file>] [-loglevel <0-4>]
[-classpath <classpath>]
```

Default Output

By default, `compilejsp` reports if the operation was successful or not.

Options

The following table describes the options available when using the `compilejsp` tool.

Option	Description
`-src <war_or_ear>`	Specifies the WAR or EAR file you want to compile. The full or relative path to the file must be specified. There is no default.
`-target <target_file>`	Specifies the name of the target WAR or EAR archive file to be generated. If file name you specify already exists, use the `-overwrite` option to overwrite the previously existing target. The full or relative path to the file must be specified. There is no default.
`-overwrite`	Indicates that the <target_file> should be overwritten if it previously existed. If <target_file> exists but `-overwrite` is not used, you will get an error message saying that the target JAR must be different from the source JAR.
`-package <package_root>`	Specifies the base package name for the precompiled JSP servlet classes. The default is `com.bes.compiledjsp`.
`-excludefile <exclude_file>`	Specifies a text file containing a list of JSP files to exclude from the compile operation. See "Using the excludefile option" for more details.
`-loglevel <0-4>`	Specifies the amount of output diagnostic messages to be generated. A value greater than 2 will also leave the temporary servlet Java files for further inspection. The default is 2.
`-classpath <classpath>`	Specifies any additional libraries that may be required for compiling the JSP pages. There is no default.

Example

To precompile the JSP pages contained in a WAR file called `proj1.war` located in the current directory into a WAR file called `proj1compiled.war` in the same location:

```
iastool -compilejsp -src proj1.war -target proj1compiled.war
```

To precompile the JSP pages contained in an EAR file called `proj1.ear` located in the directory `c:\myprojects\` into an EAR file called `proj1compiled.ear` in the same location and generate the maximum amount of diagnostic messages:

```
iastool -compilejsp -src c:\myprojects\proj1.ear -target
c:\myprojects\proj1compiled.ear -loglevel 4
```

Using the excludefile option

The `compilejsp excludefile` option allows you to specify a text file containing a list of JSPs to exclude from the compile operation. The following list details the usage rules.

- Comment lines (with a leading '#') and blank lines are ignored.
- Leading and trailing blank spaces and white spaces on each line are trimmed.
- Each line in the exclude file represents one exclude pattern entry, which can be a string for exact matching or a Java Pattern regular expression.
- Each JSP exclude entry is used to perform an exact match against a JSP URL first. If there is no match, the JSP exclude entry will be used as a regular expression to match against the JSP URL again.

- A JSP URL is compared against each of the JSP exclude entries using the above algorithm. As soon as there is a match, the JSP is excluded from compilation. If the JSP URL does not match any of the JSP exclude entries, the JSP will be compiled.
- If a pattern entry is not a valid Java regular expression, a warning is shown. It will still be used to compare against JSP URLs for exact match.
- If the `iastool -compilejsp -loglevel` option is set to 3 or higher, the exclude pattern entries, the number of excluded JSP pages, and the excluded JSP URLs are be displayed.
- If all of the JSP files in the archive are excluded the `-compilejsp` command will fail.

The following are some example exclude patterns:

```
# This pattern excludes a specific JSP: /jsp/test/test.jsp
/jsp/test/test[.]jsp

# This pattern excludes the JSP /jsp/test/test.jsp or /jsp/test/test2jsp, etc.
# because the regular expression "." represents any character
/jsp/test/test.jsp

# This pattern excludes all the files under the /include URL path
/include/.*

# This pattern excludes all the include.jsp files
.*/include[.]jsp

# This pattern excludes all the JSP files that start with "tmp_" in the /jsp
URL path
# and all JSP files in any URL path that starts with "tmp_" under /jsp
/jsp/tmp_.*[.]jsp

# This pattern excludes all the JSP files that start with "tmp_" in the /jsp
URL path
/jsp/tmp_[^/]*[.]jsp
```

compress

Use this tool to compress a JAR file.

Syntax

```
-compress -src <srcjar> -target <targetjar>
```

Default Output

By default, `compress` reports if the operation was successful or not.

Options

The following table describes the options available when using the `compress` tool.

Option	Description
`-src <srcjar>`	Specifies the JAR file that you want to compress. The full or relative path to the file must be specified. There is no default.
`-target <targetjar>`	Specifies the name of the compressed JAR file to be generated. The full or relative path to the file must be specified. There is no default.

Example

To compress a JAR file, called `proj1.jar` and located in the current directory, into a file called `proj1compress.jar` in the same location:

```
iastool -compress -src proj1.jar -target proj1compress.jar
```

To compress a JAR file called `proj1.jar` located in the directory `c:\myprojects\` into a file called `proj1compress.jar` in the same location:

```
iastool -compress -src c:\myprojects\proj1.jar
-target c:\myprojects\proj1compress.jar
```

deploy

Use this tool to deploy a J2EE module to a specified Partition on the specified hub and configuration.

Syntax

```
-deploy -jars <jar1,jar2,...> <-hub <hub>|-host <host>:listener_port>>
-cfg <configname> -partition <partitionname> [-force_restart] [-cp <classpath>]
[-args <args>] [-javac_args <args>] [-noverify] [-nostubs] [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

By default, `deploy` reports if the operation was successful.

Options

The following table describes the options available when using the `deploy` tool.

Option	Description
`-jars <jar1,jar2...>`	Specifies the names of one or more JAR files to be deployed. To specify more than one JAR file, enter a comma (,) between each file name (no spaces). The full or relative path to the files must be specified. There is no default.
`-hub <hub>`	Specifies the name of the hub in which to deploy the JAR files.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the partition you are interested in is running. This option enables the `iastool` utility to locate a partition on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration containing the partition in which you want to load the JAR file.
`partition <partitionname>`	Specifies the name of the Partition in which you want to load the JAR file.
`-force_restart`	Restarts the specified Partition after deploying the module. If this option is not specified, you will need to restart the Partition manually to initialize the module.
`-cp <classpath>`	Specifies the classpath containing the class dependencies of the JAR file(s) to be deployed.
`-args <args>`	Specifies any arguments that are needed by the JAR file. For details, see “Programmer tools for Java” in the *VisiBroker for Java Developer’s Guide*.
`-javac_args <args>`	Specifies any Java compiler arguments that are needed by the JAR file.
`-noverify`	Turns off verification of the active connections to a Partition on a specified management port.
`-nostubs`	Prevents creation of client or server-side stub files for the deployed module.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. The default is 42424.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.

Option	Description
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

dumpstack

Use this tool to obtain diagnostic information about the threads running in a Partition. This tool causes the Partition to generate a stack trace of all threads, and the output is stored in the Partition's stdout.log, located in:

```
<install_dir>/var/domains/<domain-name>/configurations/<configuration-name>/
    mos/<partition-name>/adm/logs/<partition_name>.stdout.log
```

. The stack trace may be useful for diagnosing problems with the Partition. The log file is located in the directory:

```
<install_dir>\var\domains\<domain_name>\configurations\<config_name>\
    <partition_name>\adm\logs\partition_log.xml
```

Syntax

```
-dumpstack <-hub <hub>|-host <host>:<listener_port>> -cfg <configname>
-partition <partitionname> [-mgmtport <nnnnn>] [-realm <realm>]
[-user <username>] [-pwd <password>] [-file <login_file>]
```

Options

The following table describes the options available when using the `dumpstack` tool.

Option	Description
-hub <hub>|-host <hostname>:<listener_port>	Specifies the name of the hub or the host name and the listener port of the machine on which the partition process you are interested in is running. You must specify either a hub name or a host name and listener port. Specifying a listener port enables the `iastool` utility to locate a hub on a different subnet than the machine on which `iastool` is running.
-cfg <configname>	Specifies the name of the configuration containing the specified partition.
-partition <partitionname>	Specifies the name of the Partition that you want to diagnose. The name of a valid Partition must be specified.
-mgmtport <nnnnn>	Specifies the management port number used by the specified hub. The default is 42424
-realm <realm>	Specifies the realm used to authenticate a user when the user and password options are specified.
-user <username>	Specifies the user to authenticate against the specified realm.
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Examples

The following example shows how to perform a thread dump of the `standard` Partition in the `j2ee` configuration on the `BES1` hub:

```
iastool -dumpstack -hub BES1 -cfg j2ee -partition standard
```

The following example shows how to perform a thread dump of the `standard` Partition on a computer host on a specific listener port. Note that the `-host` option can be used regardless of whether `iastool` is executed on the same or a different host machine on which the partition is running.

```
iastool -dumpstack -host mymachine:1234 -cfg j2ee -partition standard
```

genclient

Use this tool to generate a library containing client stubs files, EJB interfaces, and dependent class files for one or more EJB JAR files, and to package them into one or more client JAR files. The client JAR is not an EJB, but is an EJB client.

If `genclient` fails for one of the EJB JARs in the argument list, an error is displayed and the `genclient` tool will continue to attempt to generate a client JAR on the remainder of the specified list.

The `genclient` tool will exit `0` (zero), for 100% success, or `1`, for any failure.

Syntax

```
-genclient -jars <jar1,jar2,...> -target <client_jar> [-cp <classpath>]
[-args <java2iiop_args>] [-javac_args <args>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `genclient` tool.

Option	Description
`-jars <jar1,jar2,...>`	Specifies one or more JAR files for which you want to generate one or more client JAR files. To specify more than one JAR file, enter a comma (,) between each file name (no spaces). The full or relative path to the JAR files must be specified. There is no default.
`-target <client_jar>`	Specifies the client-JAR files to be generated on the localhost. The full or relative path to the JAR files must be specified. There is no default.
`-cp <classpath>`	Specifies the classpath containing the class dependencies of the JAR file for which you want to generate a client JAR file. The default is none.
`-args <java2iiop_args>`	Specifies any arguments that are needed by the file. For details, see “Programmer tools for Java” in the *VisiBroker for Java Developer’s Guide*.
`-javac_args <args>`	Specifies any Java compiler arguments that are needed by the JAR file.

Example

The following example shows how to generate a manually deployable module client JAR file from each of the EJB JAR files: `proj1.jar`, `proj2.jar`, and `proj3.jar` into the EJB JAR `myproj.jar`.

```
iastool -genclient -jars proj1.jar,proj2.jar,proj3.jar -target myproj.jar
```

gendeployable

Use this tool to create a manually deployable server-side module. Server-side deployable JAR files are archives (EAR, WAR, or JAR beans only) that have been compiled to resolve all external code references by using stubs and are, therefore, ready for deployment.

For example, first use `gendeployable` to create the server-side deployable JAR file on a local machine, then use the `deploy` tool to copy and load it on the hub. The hub is

advised of the presence of the new JAR file and loads it automatically. Using the command-line tools lets you script a creation and deployment to several servers quite simply. You can also manually copy the server-side deployable JAR file to the correct location on each hub, but this requires restarting each hub to cause it to be recognized and loaded.

Syntax

```
-gendeployable -src <input_jar> -target <output_jar> [-cp <classpath>]
[-args <java2iiop_args>] [-javac_args <args>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `gendeployable` tool.

Option	Description
`-src <input_jar>`	Specifies the JAR file (or the directory of an expanded JAR) that you want to use to generate a new deployable JAR file. The full or relative path to the JAR file must be specified. There is no default.
`-target <output_jar>`	Specifies the deployable JAR files to be generated on the localhost. The full or relative path to the JAR files must be specified. There is no default.
`-cp <classpath>`	Specifies the classpath containing the class dependencies of the JAR file for which you want to generate a client JAR file. The default is none.
`-args <java2iiop_args>`	Specifies any arguments that are needed by the file. For details, see “Programmer tools for Java” in the *VisiBroker for Java Developer’s Guide*.
`-javac_args <args>`	Specifies any Java compiler arguments that are needed by the JAR file.

Example

The following example shows how to generate a server-side deployable module JAR file for `proj1.jar` into the file `server-side.jar`.

```
iastool -gendeployable -src proj1.jar -target serverside.jar
```

gendeployableverify

This command line property combines the utilities of `gendeployable` and `verify`. This combined option expedites the total stub generation and verification process. It creates and verifies manually deployable server-side modules.

Syntax

```
-gendeployableverify -src <input_jar>  -target <output_jar> [-args
<java2iiop_args>] [-javac_args <args>] [-role <DEVELOPER|ASSEMBLER|DEPLOYER>]
[client] [-nowarn] [-strict] [-cp <classpath>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `gendeployableverify` tool.

Option	Description
`-src <input_jar>`	Specifies the JAR file (or the directory of an expanded JAR) that you want to use to generate a new deployable JAR file. The full or relative path to the JAR file must be specified. There is no default.
`-target <output_jar>`	Specifies the deployable JAR files to be generated on the localhost. The full or relative path to the JAR files must be specified. There is no default.
`-cp <classpath>`	Specifies the classpath containing the class dependencies of the JAR file for which you want to generate a client JAR file. The default is none.
`-args <java2iiop_args>`	Specifies any arguments that are needed by the file. For details, see "Programmer tools for Java" in the *VisiBroker for Java Developer's Guide*.
`-javac_args <args>`	Specifies any Java compiler arguments that are needed by the JAR file.
`-role <DEVELOPER|ASSEMBLER|DEPLOYER>`	Specifies the level of error checking.
`-nowarn`	Specifies that the tool should only report errors that preclude deployment and not report warnings.
`-strict`	Specifies that the tool should report the most minute inconsistencies, many of which do not affect the overall integrity of the application.
`-classpath <classpath>`	Specifies the search path for application classes and resources. To enter more than one directory, ZIP, or JAR file entry, separate each entry with a semicolon (;).

Example

The following example shows how to generate a server-side deployable and verified module JAR file for `proj1.jar` into the file `server-side.jar`.

```
iastool -gendeployableverify  -src proj1.jar -target serverside.jar
```

genstubs

Use this tool to create a stubs library file containing client or server stubs.

Syntax

```
-genstubs -src <input_jar> -target <output_jar> [-client] [-cp <classpath>]
[-args <java2iiop_args>] [-javac_args <args>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `genstubs` tool.

Option	Description
`-src <input_jar>`	Specifies the JAR file (or the directory of an expanded JAR) for which you want to generate a stubs library. The full or relative path to the JAR file must be specified. There is no default.
`-target <output_jar>`	Specifies the name of the JAR file that will be generated on the localhost. The full or relative path to the JAR file(s) must be specified. There is no default.
`-client`	Specifies that you want to generate client-side stubs. If this option is not specified, the `genstubs` tool will generate server-side stubs.

Option	Description
-cp <classpath>	Specifies the classpath containing the class dependencies of the JAR file for which you want to generate a client JAR file(s). The default is none.
-args <java2iiop_args>	Specifies any arguments that are needed by the file. For details, see “Programmer tools for Java” in the *VisiBroker for Java Developer’s Guide*.
-javac_args <args>	Specifies any Java compiler arguments that are needed by the JAR file.

Examples

The following example shows how to generate a server-side stubs of the EJB JAR `proj1.jar` into the EJB JAR `server-side.jar`.

```
iastool -genstubs -src proj1.jar -target serverside.jar
```

The following example shows how to generate a client-side stub file for the EJB JAR `myproj.jar` into the EJB JAR `client-side.jar`.

```
iastool -genstubs -src c:\dev\proj1.jar -target
-client c:\builds\client-side.jar
```

info

Use this tool to display the Java system properties for the JVM the `iastool` is running in.

Syntax

```
-info
```

Default Output

The default output is the current Java system properties for the JVM the `iastool` is running in. For example, the first few lines of output look like the following partial listing:

```
application.home              : C:\Program Files\AppServer
awt.toolkit                   : sun.awt.windows.WToolkit
file.encoding                 : Cp1252
file.encoding.pkg             : sun.io
file.separator                : \
java.awt.fonts                :
java.awt.graphicsenv          : sun.awt.Win32GraphicsEnvironment
java.awt.printerjob           : sun.awt.windows.WPrinterJob
java.class.path               : C:\Program Files\AppServer\jdk\lib\tools.jar
.
.
.
```

Example

The following example shows how to display configuration information.

```
iastool -info|more
```

jndinamespace

Use this tool to view the objects bound in a context. This utility is a replica of the JNDI browser in the command line. It is an interactive utility that takes the user’s input to display the objects bound.

Syntax

```
-jndinamespace -osagent <osagent port number>
```

Default Output

Displays the objects bound for the context specified.

Options

`Osagent`

This option specifies the osagent port number. This is mandatory.

Example

The following example shows how to display the objects bound in a context:

```
iastool -jndinamespace -osagent 19999
```

The system displays the following output:

```
Retrieving the jndi name space details
The list of nodes under  19999
[0]order
[1]serial_provider_defaultname
[2]datasources
Enter the context [0,1,...] or X to exit
```

Enter `2`

The system displays the following output:

```
The list of nodes under  19999/datasources
[0]Jdatastore
[1]Oracle
Enter the context [0,1,...] or X to exit
```

Enter `X` to exit.

kill

Use this tool to kill a managed object on a specified hub and configuration.

Syntax

```
-kill <-hub <hub>|-host <host>:listener_port>> -cfg <configname>
-mo <managedobjectname> -moagent <managedobjectagent> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

By default, the `kill` tool lists the managed object that has been killed.

Options

The following table describes the options available when using the `kill` tool.

Option	Description
`-hub <hub>`	Specifies the name of the hub on which you want to kill a managed object.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the managed object you are interested in is running. The option is enables the `iastool` utility to locate a hub on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration containing the specified managed object.
`-mo <managedobjectname>`	Specifies the name of the managed object.
`-moagent <managedobjectagent>`	Specifies the managed object agent name. Use this option if the specified hub has more than one agent.

Option	Description
-mgmtport <nnnnn>	Specifies the management port number used by the specified hub. The default is 42424.
-realm <realm>	Specifies the realm used to authenticate a user when the user and password options are specified.
-user <username>	Specifies the user to authenticate against the specified realm.
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Examples

The following example kills the managed object j2ee-server using the default management port:

```
iastool -kill -hub AppServer1 -cfg j2ee -mo j2ee-server
```

listpartitions

Use this tool to list the partitions running on a specified hub, and optionally on a specified configuration or management port.

Syntax

```
-listpartitions <-hub <hub>|-host <host>:<listener_port>>
[-cfg <configname>] [-mgmtport <nnnnn>] [-bare] [-realm <realm>]
[-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

By default, the listpartitions tool displays the partitions running on a specified hub, or displays the partitions running on a specified hub on a specified configuration or management port.

Options

The following table describes the options available when using the listpartitions tool.

Option	Description
-hub <hub>	Specifies the hub name for which you want to list the running partitions.
-host <host>:<listener_port>	Specifies the host name and the listener port of the machine on which the partitions you are interested in are running. The option is enables the iastool utility to locate a hub on a different subnet than the machine on which iastool is running.
-cfg <configname>	Specifies the name of the configuration on which to list partitions.
-mgmtport <nnnnn>	Specifies the management port number used by the specified hub. The default is 42424.
-bare	Suppresses the output information, other than the names of the running partitions.
-realm <realm>	Specifies the realm used to authenticate a user when the user and password options are specified.
-user <username>	Specifies the user to authenticate against the specified realm.

Option	Description
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Examples

The following example lists the partitions running on the hub AppServer1 using the default management port:

```
iastool -listpartitions -hub AppServer1
```

The following example lists the partitions running on the hub AppServer1 using the management port 24410:

```
iastool -listpartitions -hub AppServer1 -mgmtport 24100
```

listhubs

Use this tool to list hubs running on a particular management port located on the same local area network.

Syntax

```
-listhubs [-mgmtport <nnnnn>] [-bare] [-realm <realm>] [-user <username>]
[-pwd <password>] [-file <login_file>]
```

Default Output

By default, the listhubs tool displays the hubs running in the default management port or on a specified management port.

Note

If a particular hub that is queried is down, it is not listed.

Options

The following table describes the options available when using the listhubs tool.

Option	Description
-mgmtport <nnnnn>	Specifies a management port number of the running hub you want to list. The default is 42424.
-bare	Suppresses output information, other than the names of the running hubs.
-realm <realm>	Specifies the realm used to authenticate a user when the user and password options are specified.
-user <username>	Specifies the user to authenticate against the specified realm.
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Examples

The following example lists the hubs running in the default management port:

```
iastool -listhubs
```

The following example lists the hubs running in the management port 24410:

```
iastool -listhubs -mgmtport 24100
```

listservices

Use this tool to list one or more services running on a hub.

Syntax

```
-listservices  <-hub <hub>|-host <host>:<listener_port>> [-cfg <configname>]
[-mgmtport <nnnnn>] [-bare] [-realm <realm>] [-user <username>]
[-pwd <password>] [-file <login_file>]
```

Default Output

By default, `listservices` displays a list of all partition services registered for the specified hub on a particular management port.

Options

The following table describes the options available when using the `listservices` tool.

Option	Description
`-hub <hub>`	Specifies the name of the hub for which you want to list the running services.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the services you are interested in are running. The option is enables the `iastool` utility to locate a hub on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration on which to list services.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. The default is 42424.
`-bare`	Suppresses output of information other than the names of the running services.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Example

The following example lists all services running on the `salsa` hub:

```
iastool -listservices -hub salsa
```

manage

Use this tool to actively manage a managed object in a configuration.

Syntax

```
-manage  (-hub <hub>|-host <host>:<listener_port>) [-cfg <configname>] -mo
<managedobjectname> [-moagent <managedobjectagent>] [-mgmtport <99999>] [-realm
<realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `manage` tool.

Option	Description
`-hub <hub>`	Specifies the name of the hub on which the managed object is running.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the managed object is running. This option enables the iastool utility to locate a hub on a different subnet than the machine on which iastool is running.
`-cfg <configname>`	Specifies the name of the configuration with which the managed object is associated.
`-mo <managedobjectname>`	Specifies name of the managed object.
`-moagent <managedobjectagent>`	Specifies the agent with which the managed object is associated.
`-mgmtport <99999>`	Specifies the port with which the managed object's agent is associated.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See "Executing iastool command-line tools from a script file" for more information.

Example

The following example puts the managed object j2ee-server into the actively managed mode using the default management port:

```
iastool -manage -hub AppServer1 -cfg j2ee -mo j2ee-server
```

merge

Use this tool to produce a single, new Java Archive file (EJB-JAR) containing the contents of a specified list of EJB-JARs. Multiple EJB 1.1 and EJB 2.0 deployment descriptors (if any) will be consolidated into a single deployment descriptor. If merging fails for one of the EJB-JARs in the argument list an error is displayed and the merge command will exit indicating failure.

Syntax

```
-merge -jars <jar1,jar2,...> -target <new_jar> -type <valid_type>
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `merge` tool.

Option	Description
`-jars <jar1,jar2,..>`	Specifies the JAR files to merge, comma separated and no spaces. The full or relative path to the JAR files must be specified. There is no default.

Option	Description
`-target <new_jar>`	Specifies the name of the new JAR file to be created containing the merged contents of the specified list of JAR files. The full or relative path to the new JAR file must be specified. There is no default.
`-type valid_type`	Specifies the type of the new archive file using one of the following supported formats: ■ `ejb2.0` – Version 2.0 Enterprise Java Bean ■ `ejb1.1` – Version 1.1 Enterprise Java Bean ■ `ear1.3` – Version 1.3 Enterprise Application Resource ■ `ear1.2` – Version 1.2 Enterprise Application Resource ■ `lib` – Library file ■ `war2.3` – Version 2.3 Web Application Archive ■ `war2.2` – Version 2.2 Web Application Archive ■ `rar1.0` – Version 1.0 Resource Adapter Archive ■ `client1.2` – Version 1.2 Client JAR ■ `client1.3` – Version 1.3 Client JAR ■ `jndi1.2` – Version 1.2 Java Naming and Directory Interface

Example

The following example merges the EJB-JAR files `proj1.jar`, `proj2.jar`, and `proj3.jar` into a new version 2.0 EJB-JAR file named `combined.jar`:

```
iastool -merge -jars proj1.jar,proj2.jar,proj2.jar
-target combined.jar -type ejb2.0
```

migrate

Use this tool to convert a JAR or XML file from one version of J2EE to another, for example, J2EE version 1.2 to J2EE version 1.3 or J2EE 1.4.

Note

The `migrate` command only converts the deployment descriptor for an EJB; as such, code changes may also be required to implement the conversion properly in your deployment.

If the conversion fails, an error is displayed.

Syntax

```
-migrate [-to[1.2|1.3|1.4]] -src <src-archive> -target <target-archive>
```

Default Output

The default returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `migrate` tool.

Option	Description
`-to[1.2|1.3|1.4]`	Specifies the target version that the source J2EE module should be migrated to. For instance, a J2EE 1.3 source module can be migrated to J2EE 1.2 or J2EE 1.4, and a J2EE 1.2 module can be migrated to a target J2EE 1.3 or J2EE 1.4 module. If you do not specify this option then it defaults to J2EE 1.4.
`-src <src-archive>`	Specifies the J2EE module to convert. The full or relative path to the archive file (or the directory of an expanded JAR) must be specified. There is no default.
`-target <target-archive>`	Specifies the name of the J2EE module to be created. The full or relative path to the archive file must be specified. There is no default.

Example

The following example migrates the file `myj1_2.jar` from J2EE version 1.2 to J2EE version 1.3 into new file called `myj1_3.jar`:

```
iastool -migrate -src myj1_2.jar -target myj1_4.jar //here -to 1.4 is implied
iastool -migrate -to 1.3 -src myj1_2.jar -target myj1_3.jar //here a 1.2 module
is converted to 1.3
```

newconfig

Use this tool to create a new configuration from a configuration template. The command takes the name of a new configuration, the filename and path of a template file relative to the installation's configuration template directory, and an optional property file that overrides the properties in the template used to create the new configuration.

Syntax

```
-newconfig (-hub <hub>|-host <host>:<listener_port>) -cfg <configname>
   -template <template_path> [-property <property_path>] [-mgmtport <99999>]
   [-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

Option	Description
`-hub <hub>`	Specifies the name of the hub in which to create the new configuration.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the hub is running to create the new configuration. This option enables the iastool utility to locate a hub on a different subnet than the machine on which iastool is running.
`-cfg <configname>`	Specifies the name of the new configuration.
`-template <template_path>`	Specifies the path where the configuration template is located. `template_path` can be either a full path of a configuration template XML file or a path relative to the configuration template directory (i.e. <install_dir>/var/templates/configurations/).
`-property <property_path>`	Specifies the path to an optional property file that overrides the properties used in the template to create the new configuration.
`-mgmtport <99999>`	Specifies the management port number used by the specified hub. The default is 42424.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Example

```
iastool -newconfig -hub myhub -cfg SimpleProcessConfig -template native.xml
   -property c:\simple.properties
```

patch

Use this tool to apply one or more patches to a JAR file and produce a new JAR file with the applied patches.

Syntax

```
-patch -src <original_jar> -patches <patch1_jar,...> -target <new_jar>
```

Default Output

The default output displays the patches that were applied.

Options

The following table describes the options available when using the `patch` tool.

Option	Description
`-src <original_jar>`	Specifies the JAR file to which you want to apply one or more patches. The full or relative path to the JAR file must be specified. There is no default.
`-patches <patch1_jar,...>`	Specifies one or more JAR files that contain the patches you want to apply. To specify more than one file, enter a comma (,) between each file name (no spaces). The full or relative path to the files must be specified. There is no default.
`-target <new_jar>`	Specifies the name of the new JAR file to be created. The full or relative path to the JAR file must be specified. There is no default.

Example

The following example applies the patches contained in the files `mypatch1.jar` and `mypatch2.jar` to the file `myold.jar` which are all located in the current directory and creates a new file called `mynew.jar` in the same location:

```
iastool -patch -src myold.jar -patches mypatch1.jar,mypatch2.jar
-target mynew.jar
```

ping

Use this tool to verify the current state of a hub or a managed object. The ping command will return nothing for a hub that is not running.

Syntax

```
-ping <-hub <hub>|-host <host>:<listener_port>> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

or

```
-ping <-hub <hub>|-host <host>:<listener_port>> -cfg <configname>
-mo <managedobjectname> -moagent <managedobjectagent> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output shows the name and status of the hub (and optionally the managed object) if the process is pinged and running. For example:

```
Pinging Hub xyz_corp1: Running
```

The `ping` tool returns one of the following states:

- Running
- Starting
- Stopping
- Not Running

- Restarting
- Cannot Load
- Cannot Start
- Terminated
- Unknown

Options

The following table describes the options available when using the `ping` tool.

Option	Description
`-hub <hub>`	Specifies the hub to ping or whose services to ping. There is no default.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the hub or managed object you are interested in is running. The option is enables the `iastool` utility to locate a managed object on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration on which to ping for managed objects.
`-mo <managedobjectname>`	Specifies the name of the managed object.
`-moagent <managedobjectagent>`	Specifies the managed object agent name. Use this option if the specified hub has more than one agent.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. The default is 42424.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Examples

The following example pings the hub `AppServer1` in the default management port:

```
iastool -ping -hub AppServer1
```

The following example pings the partition naming service running on the hub AppServer1 in the management port 24410:

```
iastool -ping -hub AppServer1 -cfg j2ee -mo standard_visinaming -mgmtport 24410
```

pservice

Use this tool to enable, disable, or get the state of a partition service.

Syntax

```
-pservice <hub <hub>|-host <host>:<listener_port>> -cfg <configname>
-partition <partitionname> -moagent <managedobjectagent>
-service <servicename> <-enable|-disable|-status> [-force_restart]
[-mgmtport <nnnnn>] [-realm <realm>] [-user <username>] [-pwd <password>]
[-file <login_file>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `pservice` tool.

Option	Description
`-hub <hub>`	Specifies the hub where the partition service you are interested in is located. There is no default.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the partition service you are interested in is running. The option is enables the `iastool` utility to locate a partition service on a different subnet than the machine on which `iastool` is running.
`-partition <partitionname>`	Specifies the name of the partition.
`-moagent <managedobjectagent>`	Specifies the managed object agent name. Use this option if the specified hub has more than one agent.
`-service <servicename>`	Specifies the name of the service.
`-enable|-disable|-status`	Specifies the operation you would like to perform on the partition service.
`-force_restart`	Restarts the specified Partition after completing the enable, disable, or status operation. If this option is not specified, you will need to restart the Partition manually to initialize the module.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. The default is 42424.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Example

The following example shows how to enable the partition naming service on the standard partition.

```
iastool -pservice -hub AppServer1 -cfg j2ee -partition standard
-service standard_visinaming -enable -force_restart -mgmtport 24431
```

removestubs

Use this tool to remove all stub files from a JAR file.

Syntax

```
-removestubs -jars <jar1,jar2,...> [-targetdir <dir>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `removestubs` tool.

Option	Description
`-jars <jar1,jar2...>`	Specifies the JAR file(s) from which you want to remove one or more stub files. To specify more than one JAR file, enter a comma(,) between each JAR file (no spaces). The full or relative path to the JAR file(s) must be specified. There is no default.
`-targetdir <dir>`	Specifies the directory in which the stub files that were removed will be stored. A full or relative path must be specified, if this option is specified. There is no default. If no target directory is specified, the stub files will be removed, but not saved.

Example

The following example shows how to remove stub files located from the EJB JAR files `proj1.jar`, `proj2.jar`, and `proj3.jar` located in the current directory and copy them to `c:\examples\proto`:

```
iastool -removestubs -jars proj1.jar,proj2.jar,proj3.jar
-targetdir c:\examples\proto
```

restart

Use this tool to restart a hub or managed object. The hub must already be running in order for the `restart` tool to work with a hub.

Syntax

```
-restart <-hub <hub>|-host <host>:<listener_port>> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

or

```
-restart <-hub <hub>|-host <host>:<listener_port>> [-cfg <configname>]
-mo <managedobjectname> -moagent <managedobjectagent> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output displays the hub or managed object that has been restarted.

If the `restart` tool fails (for example, when a managed object cannot be shutdown or restarted), an error is displayed with a status code which is returned to standard error output (stderr).

Options

The following table describes the options available when using the `restart` tool.

Option	Description
`-hub <hub>`	Specifies the name of the hub that you want to restart. Also used to locate a managed object on a particular hub.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the managed object you are interested in is running. The option is enables the `iastool` utility to locate a managed object on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration on which to locate managed objects.
`-mo <managedobjectname>`	Specifies the name of the managed object.
`-moagent <managedobjectagent>`	Specifies the managed object agent name. Use this option if the specified hub has more than one agent.

Option	Description
`-mgmtport nnnnn`	Specifies the management port number used by the specified hub. The default is 42424.
`-realm realm`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user username`	Specifies the user to authenticate against the specified realm.
`-pwd password`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See "Executing iastool command-line tools from a script file" for more information.

Examples

The following example restarts the hub `AppServer1` on the default management port:

```
iastool -restart -hub AppServer1
```

The following example restarts the partition naming service running on the hub `AppServer1` on the management port 24410:

```
iastool -restart -hub AppServer1 -cfg j2ee -mo standard_visinaming  -mgmtport
24410
```

setmain

Use this tool to set the main class of a standalone Client JAR or a Client JAR in an EAR file. Once the main class is set, the `java -jar jarfile` command will automatically invoke the main class that has been set for the JAR file.

Syntax

```
-setmain -jar <jar_or_ear> [-uri <client_jar_in_ear>] -class <main_classname>
```

Default Output

The default output displays the main class that has been set for the specified JAR file.

Options

The following table describes the options available when using the `setmain` tool.

Option	Description
`-jar <jar_or_ear>`	Specifies the name of the JAR or EAR file on which you want to set the main class.
`-uri <client_jar_in_ear>`	If you are setting the main class for an EAR file, you must use the `-uri` option to identify the URI (Uniform Resource Identifier) path of the client JAR in the EAR.
`-class <main_classname>`	Specifies the class name that will be set as the main class in the specified Client JAR. The class must exist in the client JAR file and contain a `main()` method.

Examples

The following example sets a main class for a standalone Client JAR:

```
iastool -setmain -jar myclient.jar -class com.bes.myjclass
```

The following example sets a main class for a Client JAR contained in an EAR file:

```
iastool -setmain -jar myapp.ear -uri base/myapps/myclient.jar
-class com.bes.myjclass
```

start

Use this tool to start a managed object on a specified hub and configuration.

Syntax

```
-start <-hub <hub>|-host <host>:<listener_port>> -cfg <configname>
-mo <managedobjectname> -moagent <managedobjectagent> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output displays the managed object that has been started.

Options

The following table describes the options available when using the `start` tool.

Option	Description
`-hub <hub>`	Specifies the name of the hub on which the managed object you want to start is located.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the managed object you are interested in is running. The option is enables the `iastool` utility to locate a hub on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration containing the managed object you are interested in.
`-mo <managedobjectname>`	Specifies the name of the managed object you are interested in.
`-moagent <managedobjectagent>`	Specifies the managed object agent name. Use this option if the specified hub has more than one agent.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. If not specified, the default is 42424.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Example

The following example starts the partition naming service running on the hub `AppServer1` in the `j2ee` configuration on management port 24410:

```
iastool -start -hub AppServer1 -cfg j2ee -mo standard_visinaming  -mgmtport
24410
```

stop

Use this tool to shut down a hub or managed object.

Syntax

```
-stop <-hub <hub>|-host <host>:<listener_port>> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

or

```
-stop <-hub <hub>|-host <host>:<listener_port>> [-mgmtport <nnnnn>]
-cfg <configname> -mo <managedobjectname> -moagent <managedobjectagent>
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output displays the process or processes that have been shut down.

If the `stop` tool fails (for example when a managed object cannot be shutdown), an error is displayed with a status code, which is returned to standard error output (stderr).

Options

The following table describes the options available when using the `stop` tool.

Option	Description
`-hub <hub>`	Specifies the name of the hub that you want to shut down, or the hub on which resides the managed object you want to shut down.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the hub or managed object you are interested in is running. The option is enables the `iastool` utility to locate a hub on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the name of the configuration containing the managed object you are interested in.
`-mo <managedobjectname>`	Specifies the name of the managed object you are interested in.
`-moagent <managedobjectagent>`	Specifies the managed object agent name. Use this option if the specified hub has more than one agent.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. The default is 42424.
`-realm <realm>`	Specifies the realm used to authenticate a user when the user and password options are specified.
`-user <username>`	Specifies the user to authenticate against the specified realm.
`-pwd <password>`	Specifies the user's password to authenticate against the specified realm.
`-file <login_file>`	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Example

The following example stops the partition naming service running on the hub `AppServer1` in the `j2ee` configuration on the management port 24410:

```
iastool -stop -hub AppServer1 -cfg j2ee -mo standard_visinaming -mgmtport 24410
```

uncompress

Use this tool to uncompress a JAR file.

Syntax

```
-uncompress -src <srcjar> -target <targetjar>
```

Default Output

By default, `uncompress` reports if the operation was successful or not.

Options

The following table describes the options available when using the `uncompress` tool.

Option	Description
`-src <srcjar>`	Specifies the JAR file that you want to uncompress. The full or relative path to the file must be specified. There is no default.
`-target <targetjar>`	Specifies the name of the uncompressed JAR file to be generated. The full or relative path to the file must be specified. There is no default.

Examples

The following example converts the compressed JAR file called `small.jar` located in the current directory into an uncompressed file called `big.jar` in the same location:

```
iastool -uncompress -src small.jar -target big.jar
```

The following example uncompresses a JAR file named `small.jar` located in the directory `c:\myprojects\` into a file named `big.jar` in the same location:

```
iastool -uncompress -src c:\myprojects\small.jar -target c:\myprojects\big.jar
```

undeploy

Use this tool to undeploy a J2EE module from a specified Partition on a specified hub and configuration.

Syntax

```
-undeploy -jar <jar> <-hub <hub>|-host <host>:<listener_port>>
-cfg <config_name> -partition <partitionname> [-mgmtport <nnnnn>]
[-realm <realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

By default, the `undeploy` tool reports if the operation was successful or not.

Options

The following table describes the options available when using the `undeploy` tool.

Option	Description
`-jar <jar>`	Specifies the name of the JAR file to be undeployed. The full or relative path to the file must be specified. There is no default.
`-hub <hub>`	Specifies the name of the hub from which to undeploy the JAR file.
`-host <host>:<listener_port>`	Specifies the host name and the listener port of the machine on which the deployed module you are interested in is located. The option is enables the `iastool` utility to locate a module on a different subnet than the machine on which `iastool` is running.
`-cfg <configname>`	Specifies the configuration name under which the partition is configured.
`-partition <partitionname>`	Specifies the name of the Partition that contains the JAR file.
`-mgmtport <nnnnn>`	Specifies the management port number used by the specified hub. If not specified, the default is 42424.

Option	Description
-realm <realm>	Specifies the realm used to authenticate a user when the user and password options are specified.
-user <username>	Specifies the user to authenticate against the specified realm.
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

unmanage

Use this tool to remove a managed object from the actively managed mode.

Syntax

```
-unmanage  (-hub <hub>|-host <host>:<listener_port>) [-cfg <configname>] -mo
<managedobjectname> [-moagent <managedobjectagent>] [-mgmtport <99999>] [-realm
<realm>] [-user <username>] [-pwd <password>] [-file <login_file>]
```

Default Output

The default output returns nothing to standard output (stdout).

Options

The following table describes the options available when using the `unmanage` tool.

Option	Description
-hub <hub>	Specifies the name of the hub on which the managed object is running.
-host <host>:<listener_port>	Specifies the host name and the listener port of the machine on which the managed object is running. This option enables the iastool utility to locate a hub on a different subnet than the machine on which iastool is running.
-cfg <configname>	Specifies the name of the configuration with which the managed object is associated.
-mo <managedobjectname>	Specifies name of the managed object.
-moagent <managedobjectagent>	Specifies the agent with which the managed object is associated.
-mgmtport <99999>	Specifies the port with which the managed object's agent is associated.
-realm <realm>	Specifies the realm used to authenticate a user when the user and password options are specified.
-user <username>	Specifies the user to authenticate against the specified realm.
-pwd <password>	Specifies the user's password to authenticate against the specified realm.
-file <login_file>	Specifies a login script file containing the realm, user name , and password used to authenticate a user. The full or relative path to this file must be specified. See “Executing iastool command-line tools from a script file” for more information.

Example

The following example removes the managed object j2ee-server from the actively managed mode using the default management port:

```
iastool -unmanage -hub AppServer1 -cfg j2ee -mo j2ee-server
```

usage

When invoked without arguments `usage` displays a list of recognized command-line options and a brief description of each. Invoking `usage` with one or more arguments provides a more detailed description of the specified commands and their arguments.

Syntax

```
-usage
-usage <tool>
-usage <tool1 tool2 tool3>
```

Note

Arguments to the `usage` command do not require a leading hyphen.

Default Output

By default, the `usage` tool displays a list with a brief description of each command-line tool.

Examples

The following example displays a list and a brief description of each the command-line tools:

```
iastool -usage
```

The following example displays detailed information on the `compress` tool:

```
iastool -usage compress
```

The following example displays detailed information on the `-start`, `-stop`, and `-restart` tools:

```
iastool -usage start stop restart
```

verify

Use this tool to check an archive file for correctness and consistency, and to check if all of the elements required for deploying your application are in place.

The verification process supports the following roles that correspond to a phase in the application's life cycle and the appropriate level of verification (similar to the J2EE role definitions):

- **Developer:** This is the lowest verification level. All xml is checked for syntax as well as standard and proprietary keywords relevant to the current archive type. Consistency across the archive is checked, but no external resources are verified at this level.
- **Assembler:** Once the archives are individually verified and are correct, other resources built into an application will start to be verified. For example, this level will verify the existence and correctness of URIs (Uniform Resource Identifiers), but not EJB or JNDI links.
- **Deployer:** (the default) All checks are turned on. EJB and JNDI links are checked at this level as well as the operational environment in which the application is to be deployed.

Supported archive types are EAR, EJB, WAR, JNDI, and Client JARs. The typical archive verification process includes the following checks:

- A pass over the XML, checking for correct XML syntax.
- Verification of the semantics of the standard and proprietary XML descriptors, and the compliance with their required descriptors for each supported archive type.

Verification always occurs in a hierarchical fashion, starting with the top module, then recursively working through its submodules, and finally checking for inter-archive links.

Syntax

```
-verify -src <srcjar> [-role <DEVELOPER|ASSEMBLER|DEPLOYER>] [-nowarn]
[-strict] [-classpath <classpath>]
```

Default Output

By default, `verify` reports nothing (for example, if no errors are found in the specified module).

Options

The following table describes the options available when using the `verify` tool.

Option	Description
`-src <srcjar>`	Specifies the JAR file (or the directory of an expanded JAR) that you want to verify. The full or relative path to the file must be specified. There is no default.
`-role <DEVELOPER|ASSEMBLER|DEPLOYER>`	Specifies the level of error checking to perform: ■ `DEVELOPER` ■ `ASSEMBLER` ■ `DEPLOYER` (default) For details, see the role descriptions above.
`-nowarn`	Specifies that the tool should only report errors that preclude deployment and not report warnings.
`-strict`	Specifies that the tool should report the most minute inconsistencies, many of which do not affect the overall integrity of the application.
`-classpath <classpath>`	Specifies the search path for application classes and resources. To enter more than one directory, ZIP, or JAR file entry, separate each entry with a semicolon (;).

Example

The following example performs a *developer* level verification of the JAR file `soap-client.jar` located in the `c:\examples\soap` directory:

```
-verify -src c:\examples\soap\soap-client.jar -role DEVELOPER
```

Executing iastool command-line tools from a script file

Several `iastool` utility tools require that you supply login information (realm, username, and password). You may, however, want to run `iastool` commands from a script file, but doing so would expose the realm, username, password information to anyone who has access to the script file. There are two methods you can use to protect this information:

- Enter the realm, username, and password information into a file and pipe that file to the command.
- Enter the realm, username, password information into a file and pass the file to the command with the -file option.

Piping a file to the iastool utility

The following example shows how to ping a hub named `east1` by piping the file `mylogin.txt` (located in the default Borland Deployment Platform installation directory) to the `iastool` utility:

```
iastool -ping -hub east1 < c:\AppServer\mylogin.txt
```

where the file `mylogin.txt` contains three lines that correspond to what you would enter for the realm, username, and password:

```
2
username
password
```

Note

The contents of the file are exactly what you would enter on the command-line. The first entry in the file is the `realm` option—not the realm name, but the number you would choose from the list presented to you if you run the `ping` tool without the `realm` option. The second line is the `username` and the third line is the `password`. This file can then be secured in such a way that it is readable by the `iastool` utility, but not by unauthorized users.

Passing a file to the iastool utility

The following command shows how to ping a hub named `east1` by passing a file to the `iastool` utility using the `-file` option:

```
iastool -ping -hub east1 -file c:\AppServer\mylogin.txt
```

where `mylogin.txt` has the following format:

```
Default Login
Smart Agent port number
username
password
false
ServerRealm
```

The `-file` option requires that you supply a fully qualified file name (the file name plus a relative or absolute path). When passing a file to the `iastool` utility, only the third, fourth, and sixth lines are used, which are the `username`, `password`, and `realm name`, respectively. The other lines must be present, but the information they contain is ignored by the `iastool` utility. For example:

```
Default Login
12448
myusername
mypassword
false
ServerRealm
```

37

Partition XML reference

This section describes the XML definition of a Partition's `partition.xml` configuration file that contains the core meta-data for a Partition's configuration.

<partition> element

The `partition` element is the root node of the schema which contains the attributes and sub-elements that define the settings that control the configuration of a Borland AppServer (AppServer) Partition.

Attribute	Description
`version`	Product version of the Partition.
`name`	The name of the Partition.
`description`	A description of the Partition.

Syntax

```
<partition version="version number" name="partition name"
description="description">
   ⋮
</partition>
```

The `partition` element contains the following sub-elements:

- `<jmx>`
- `<statistics.agent>`
- `<security>`
- `<container>`
- `<user.orb>`
- `<management.orb>`
- `<shutdown>`
- `<services>`
- `<archives>`

<jmx> element

The `jmx` element contains the sub-elements for configuring the JMX agent.

The jmx element contains the following sub-elements:

- <mbean.server>
- <mlet.service>
- <http.adaptor>
- <rmi-iiop.adaptor>

<mbean.server> element

The mbean.server element is used to enable or disable the JMX agent's MBean Server. The MBean Server is an interface and a factory object defined by the agent specification level of the JMX.

Attribute	Description
enable	Enables or disables the MBean Server. Valid values are true (default) or false.

<mlet.service> element

The mlet.service element configures the JMX agent's MLet service. The MLet service is useful for loading MBean classes and resources inside an MBean Server's JVM from a remote host and registering the MBeans in a single action.

Attribute	Description
enable	Enables or disables the MLet service. Valid values are true or false (default).

<http.adaptor> element

The http.adaptor element configures the JMX agent's HTTP adaptor. The HTTP adaptor is the adaptor for the HTTP protocol through which the Partition can be managed through any HTML 3.2 compliant browser or application.

Attribute	Description
enable	Enables or disables the HTTP adaptor. Valid values are true or false (default).
port	The port number at which the HTTP adaptor is listening. 8082 is the default port number.
host	Defines the host name which the server will be listening to. If left with the default setting (localhost) you cannot access the server from another computer. This is good for security reasons, forcing you to explicitly open the server to the outside. You can also use 0.0.0.0 which will open the server to all local interfaces. Default is localhost.
authentication.method	Sets the authentication method. Valid values are none, basic, and digest. Default is none.
socket.factory.name	Replaces the default socket factory with another using an ObjectName, the pointed MBean has to have a public ServerSocket createServerSocket(int port, int backlog, String host) throws IOException method.
processor.name	This sets the MBean's ObjectName to be used as XML processor. The MBean has to implement the mx4j.tools.adaptor.http.ProcessorMBean interface.

The http.adaptor element contains the following sub-element:

- <xslt.processor>

<xslt.processor> element

The xslt.processor element configures the XSLT processor for the HTTP adaptor. The XSLT processor transforms the raw XML into presentable HTML for the Web browser.

If this property is not enabled and you try to use the MX4J Web console you will get raw XML in your Web browser.

Attribute	Description
enable	Enables or disables the XSLT processor. Valid values are true (default) or false.
File	Determines where to look for XSL files. If the target file is a directory, it assumes that XSL files are located in the directory. Otherwise if it points to a JAR or ZIP file, it assumes that the files are located inside. Pointing to a file system is especially useful for testing.
PathInJar	Sets the directory in the JAR file where XSL files reside.
LocaleString	Sets the locale by using a String. Default is en.
UseCache	Indicates whether to cache the transformation objects. This speeds up the transformation process. It is usually set to true, but you can set it to false for easier testing. Default is true.

<rmi-iiop.adaptor> element

The rmi-iiop.adaptor element configures the JMX agent's RMI-IIOP adaptor. The RMI-IIOP adaptor is based on the client framework, which helps managers or managing applications to communicate with the MBean Server through RMI.

Attribute	Description
enable	Enables or disables the RMI-IIOP adaptor. Valid values are true (default) or false.
port	The port number to be allocated for the RMI-IIOP adaptor port. If you do not specify a port number or specify 0 as the port number, a random port number is assigned.

<statistics.agent> element

The statistics.agent element configures the Partition's statistics agent. The Partition statistics agent consists of two components:

- A statistics collector that periodically collects statistics data on the Partition and saves that data onto disk. These periodic data samples build up on disk enabling the product tools to provide current, and historical current statistical data on a Partition.
- A statistics reaper that periodically *reaps* (cleans up) the historical data from disk.

The Partition statistics agent is intended for collecting short term statistical data. However, it is only physically limited by the amount of disk space it is allowed to consume.

Attribute	Description
enable	Enables or disables the statistics agent. A disabled statistics agent will not collect or reap statistics data. Valid values are true (default) or false.
level	Sets the level of detail of statistics collected from a Partition. Valid values are: none, minimum (default), and maximum.
snapshot.period_secs	Specifies how often (in seconds) Partition statistics are collected and written to the disk. The default is 10 seconds.
reap.enable	Enables or disables the reaping (clean up) of Partition statistics data on disk. Valid values are true (default) or false.
reap.older_than_secs	If reap_enable is true, sets the threshold for the age (in seconds) of statistics data kept on disk before being deleted. The default is 600 seconds (10 minutes).
reap.period_secs	If reap_enable is true, sets the time period (in seconds) between sweeps to clean up statistics data older than reap.older_than_secs from disk. The default is 60 seconds (1 minute).

<security> element

The security element lets you configure the security settings for a given Partition. This empty element contains the attributes described in the following table.

Attribute	Description
enable	Enables or disables security for a Partition. Valid values are true (default) or false.
manager	Specifies the name of the security manager used by a Partition. Valid values are any available security provider names, for example com.borland.security.provider.CertificateWallet.
policy	Specifies the name of the security policy file that defines the security rules of a Partition. Valid values are any fully qualified security policy file names, for example <install_dir>/va/\security/profile/\management/java_security.policy

<container> element

The container element specifies how the Partition works with classloading.

Attribute	Description
system.classload.prefixes	This is a comma separated list of resource prefixes that the custom classloader will delegate to the system classloader prior to attempting to load itself.
verify.on.load	When true runs verify on JARs as they are loaded. Default is true.
classloader.policy	Determines the type of classloader to be used by the Partition. Valid values are per_module or container. The per_module classloader policy will create a separate application classloader for each deployed module. This policy is required if you want to be able to hot deploy. The container policy will load all deployed modules in the shared classloader. You cannot hot deploy if this policy is selected.
classloader.classpath	Contains a semicolon (;) separated list of JAR files to be loaded by each instance of the application classloader. This has the same logical effect as bundling these jars in every module.

<user.orb> element

The user.orb element controls the VisiBroker configuration used for the Partition's user domain ORB.

Attribute	Description
orb.propstorage	Path to the Partition's user ORB properties file. Relative paths are relative to the Partition's properties directory (the directory partition.xml is in).
use.default.smartagent.port	This property defines whether the Partition will use the SCU Smart Agent configuration to determine the Smart Agent port value.
use.default.smartagent.addr	This property defines whether the Partition will use the SCU Smart Agent configuration to determine the Smart Agent host address value.

<management.orb> element

The `management.orb` element controls the VisiBroker configuration for the Partition's management domain ORB.

Attribute	Description
`orb.propstorage`	Path to the Partition's management domain ORB properties file.
`required_roles.propstorage`	Path to the Partition's management domain ORB required roles configuration file.
`runas.propstorage`	Path to the Partition's management domain ORB `runas` configuration file.

All the paths are relative to the Partition's properties directory (the directory `partition.xml` is in).

<shutdown> element

The `shutdown` element determines the actions taken when a Partition stops. This empty element has no attributes.

Attribute	Description
`dump_threads`	Flag that causes the Partition to dump diagnostic information on threads still running late in Partition shutdown.
`dump_threads.count`	If defined, the value indicates the number of times to dump the thread states during shutdown. It is useful if you are trying to see if some threads are simply taking a long time to quit, but do quit eventually.
`delay.1`	Reserved for support use.
`garbage_collection.1`	Reserved for support use.
`delay.2`	Reserved for support use.
`runfinalizersonexit`	Reserved for support use.
`delay.3`	Reserved for support use.
`garbage_collection.2`	Reserved for support use.
`delay.4`	Reserved for support use.
`runfinalization`	Reserved for support use.

<services> element

The services element lets you configure the Partition's services. Each Partition service has a service sub-element with its specific configuration, the services element itself has the following attributes:

Attribute	Description
autostart	List of Partition's services to be started at Partition startup. The value is a space separated list of Partition service names.
startorder	The startup order to be imposed on the Partition services configured to be started by the autostart attribute. Partition services that are not specified are started after those specified. A valid value is a space separated list of Partition service names in their start order (left to right).
shutdownorder	The shutdown order to be imposed on the Partition services that are running at Partition shutdown. Partition services that are not specified are stopped before those specified. A valid value is a space separated list of Partition service names in their shutdown order (left to right).
administer	List of Partition services that are visible to the user. They appear in the tools when Partition services are listed.

The <services> element contains the following sub-element:

– service

<service> element

The <service> element provides the configuration for a Partition service. It contains attributes that govern the Partition's management of the service and a properties sub-element that contains the service's configuration metadata.

Attribute	Description
name	The Partition service's name.
version	The version of the Partition service.
description	The description for the Partition service.
vendor	The description of the vendor for the Partition service.
class	The class that implements the Partition's service plugin architecture and provides the management and control interface for the service.
in.management.domain	Flag that indicates if the service runs in the Partition's management domain or in the Partition's user domain.
startup.synchronization	The type of synchronization to be performed when the service is started. Valid values are: ■ service_ready—wait for the service to be ready for up to startup.service_ready.max_wait milliseconds. ■ delay—always wait for startup.delay milliseconds, do not monitor the service for it to become ready. Default is no synchronization.
startup.service_ready.max_wait	Limits the maximum time, in milliseconds, that the Partition waits for the service to start when the startup.synchronization value is service_ready. A value of 0 (zero) means no time limit is imposed. The default value is 0 (zero).
startup.delay	Defines the time, in milliseconds, that the Partition waits in order to give the service a chance to start when the startup.synchronization value is delay. A value of 0 (zero) means wait forever. Default is 0 (zero) .

Attribute	Description
shutdown.synchronization	The type of synchronization to be performed when the service is shutdown. Valid values are: ■ service_shutdown—wait for the service to stop for up to shutdown.service_shutdown.max_wait milliseconds. ■ delay—always wait for shutdown.delay milliseconds, do not monitor the service for it to stop. Default is no synchronization.
shutdown.service_shutdown.max_wait	Limits the maximum time, in milliseconds, that the Partition waits for the service to stop when the shutdown.synchronization value is service_shutdown. A value of 0 (zero) means no time limit is imposed. Default value is 0 (zero).
shutdown.delay	Defines the time, in milliseconds, that the Partition waits in order to give the service a chance to stop when the shutdown.synchronization value is delay. A value of 0 (zero) means wait forever. Default is 0 (zero).
shutdown.phase	This property governs which Partition shutdown phase the service is shutdown in. A Partition shuts down in 2 phases. In the first phase all services and components providing user facility are shutdown, and in the second phase the Partition's own infrastructure is shutdown. Valid values are 1 (default) and 2. It is not typical for any Partition service to be shutdown in phase 2.

<properties> element

The properties element lets you supply the specific service's configuration metadata.

<archives> element

The archives element contains configuration metadata for the archives that the Partition can host. A specific archive can have an archive sub-element with attributes specific to that archive. An archive does not have to have an archive sub-element.

Attribute	Description
ear.repository.path	Path to the Partition's EARs directory. All EARs found in that directory are loaded by the Partition on startup, unless specifically disabled with an archive element.
war.repository.path	Path to the Partition's WARs directory. All WARs found in that directory are loaded by the Partition on startup, unless specifically disabled with an archive element.
ejbjar.repository.path	Path to the Partition's EJB jars directory. All EJB jars found in that directory are loaded by the Partition on startup, unless specifically disabled with an archive element.
rar.repository.path	Path to the Partition's RARs directory. All RARs found in that directory are loaded by the Partition on startup, unless specifically disabled with an archive element.
dar.repository.path	Path to the Partition's DARs directory. All DARs found in that directory are loaded by the Partition on startup, unless specifically disabled with an archive element.
lib.repository.path	Path to the Partition's lib directory. All JAR files found in that directory are placed on the Partition's system classpath.
classes.repository.path	Path to the Partition's classes directory. All classes found in that directory are placed on the Partition's system classpath.

All the paths are relative to the Partition's root directory.

<archive> element

The `archive` element contains configuration metadata specific to an archive. Archives that are found in the Partition's archive repository directories do not need an archive element unless there is non-default configuration that need to be applied to them.

Attribute	Description
`name`	Name of the archive to which this element pertains. Is the filename of the archive.
`disable`	Flag for disabling the hosting of that archive in the Partition at startup. Valid values are `true` or `false` (default).
`path`	Path to an archive that exists outside of the Partition repositories. Use to get the Partition to host an archive from a specified path.

All the paths are relative to the Partition's root directory.

38

EJB, JSS, and JTS Properties

EJB Container-level Properties

Set EJB container properties in `partition.xml` (each Partition has its own properties file). This file is located in the following directory:

```
<install_dir>/var/domains/base/configurations/configuration_name/mos/
partition_name/adm/properties
```

Property	Description	Default
`ejb.copy_arguments=true|false`	This flag causes arguments to be copied in intra-bean in-process calls. By default, intra-bean calls use pass-by-reference semantics. Enable this flag to cause intra-bean calls to use pass-by-value semantics. **Note:** A number of EJBs will run significantly slower using pass-by-value semantics.	`false`
`ejb.use_java_serialization=true|false`	If set it overrides use of IIOP serialization with Java serialization for things like session passivation, and so forth.	`false`

Property	Description	Default
ejb.useDynamicStubs=true|false	This property is only relevant for CMP 2.0 entity beans that provide local interfaces. If set, the Container, which otherwise uses CORBA to dispatch calls, uses a dynamic proxy-based scheme to dispatch calls (creating custom lightweight, non-CORBA references). These local dynamic stubs provide many optimizations, especially due to the callers and callees being in the same VM, making a direct dispatch to the beans without going through the CORBA layer. Also, since the dynamic stubs are aware of the EJB container data structures, they access the target beans more quickly. Note that currently the stub generator, java2iiop (called using the iastool or directly) still generates the stubs for all the interfaces in the archive. When ejb.useDynamicStubs is active, the subset of stubs that correspond to the selected CMP 2.0 beans are ignored. This feature, when used, makes the whole dispatch mechanism dynamic, providing dynamic stubs for the client side as well as dynamic skeletons on the server side. Any statically generated stub and skeleton classes in the archive are ignored. You set the property in the bean. However, if there isn't an issue with using the property in all the entity beans, the easiest way is to set it at the EAR level in the deployment descriptor. **Important:** You must use this property in conjunction with ejb.usePKHashCodeAndEquals.	true
ejb.usePKHashCodeAndEquals=true|false	Data structures that support Active Cache (TxReady cache) and Associated Cache (Ready beans cache) use java.util.Hashtable, and java.util.HashMap. The values (entity bean instances) pooled in these data structures are keyed on the primary key values of the cached entity beans. As we know, the implementation of Hashtable relies on computing hashCode() and calling equals() methods of the keys to place and locate the values. These data structures are in the critical code path and are accessed frequently by the container while dispatching calls to methods in entity beans The default in Borland AppServer (AppServer) is a reflection-based computation. When this property is set, the container uses a user supplied implementation of the equals() and hashCode() methods.	true
ejb.no_sleep=true|false	Typically set from a main program that embeds a Container. Setting this property prevents the EJB container from blocking the current thread, thereby returning the control back to user code.	false
ejb.trace_container=true|false	Turns on useful debugging information that tells the user what the Container is doing. Installs debugging message interceptors.	false
ejb.xml_validation=true|false	If set, the XML descriptors are validated against its DTD at deployment time.	true
ejb.xml_verification=true|false	If set, J2EE archive is verified at deployment time.	false
ejb.classload_policy=per_module|container|none	Defines class loading behavior of standalone EJB container. Not applicable to the Partition. If set to per_module, the container uses a new instance of custom class loader with each J2EE archive deployed. If set to none, the container uses the system class loader. Hot-deployment and deployment of EARs does not work in this mode. If set to container, container uses single custom class loader. This enables deployment of EARs, but disables hot-deployment feature.	per_module
ejb.module_preload=true|false	Loads the entire J2EE archive into memory at deployment time, so the archive can be overwritten or rebuilt. This option is required by JBuilder running a standalone ejb container.	false
ejb.system_classpath_first=true|false	If set to true, the custom classloader will look at the system classpath first.	false

Property	Description	Default
ejb.sfsb.keep_alive_timeout=<num>	Defines the default value of the <timeout> element used in the ejb-borland.xml descriptor. This property affects an EJB whose <timeout> element is skipped or set to 0. The purpose of this property is to define a time interval in seconds how long to keep an inactive stateful session bean alive in the persistent storage (JSS) after it was passivated. After the time interval ends, JSS deletes the session's state from the persistent storage, so it becomes impossible to activate it later.	86400 (=24 hours)
ejb.cacheTimeout=<integer>	This property hints the container to invalidate the data fields of entity beans after a specified time-out period. Use the property by specifying the interval for which the container will not load a bean's state from the database, but uses the cached state instead. At the end of the expire period specified, the container marks the bean as dirty (but keeps its association with the primary key), forcing the instance to load its state from the database (not the cache) before it can be used in any new transactions. The property is expected to be used by entity beans that are not frequently modified. The property is a positive integer representing cache intervals in seconds. This is only valid for commit mode A. It is ignored if specified for any other commit mode.	0 (no timeout).
ejb.sfsb.aggressive_passivation=true|false	If set to true, stateful session bean is passivated no matter when it was used last time. This enables fail-over support, so if an EJB container fails, the session can be restored from the last saved state by one of EJB containers in the cluster. If set to false, only the beans which were not used since the last passivation attempt, are passivated to JSS. This makes the fail-over support less deterministic, but speeds things up. Use this setting, to trade performance for high-availability.	true
ejb.sfsb.factory_name=<string>	If set, makes the stateful session beans use a different JSS from the one that is running within the same EJB container or Partition. Specify the factory name of JSS to use. This is the name under which JSS is registered with Smart Agent (osagent).	none
ejb.logging.verbose=true|false	If set to true, the EJB container logs messages about unexpected situations which potentially could require user's attention. The messages are marked with >>>> EJB LOG <<<< header. Set it to false, to suppress these messages.	true
ejb.logging.doFullExceptionLogging=true|false	If set, the container logs all unexpected exceptions thrown in an EJB implementation.	false
ejb.jss.pstore_location=<path>	Use this to override the default name and location of the file used as a JSS backend storage. Applicable only for standalone ejb containers. This option is deprecated, use jss.pstore and jss.workingDir instead.	none
ejb.jdb.pstore_location=<path>	Use this to override the default name and location of the file used by the database service. Applicable only for standalone ejb containers.	none
ejb.interop.marshal_handle_as_ior=true|false	If set to true, each instance of javax.ejb.Handle is marshaled as a CORBA IOR. Otherwise, it is marshaled as a CORBA abstract interface. See CORBA IIOP spec for details.	false
ejb.finder.no_custom_marshal=true|false	When a multi-object finder returns a collection of objects, by default the EJB container does the following: ■ creates and returns a custom Vector implementation to the caller. ■ creates IORs (from the primary keys) lazily as the caller of the finder browses/iterates over the Vector returned. ■ compute IORs for the whole Vector, when result is to leave the JVM where it was created. If this property is set to true, the EJB container does not do any of the above.	false

Property	Description	Default
ejb.collect.stats_gather_ frequency=<num>	The time interval in seconds between printouts of container statistics. If set to zero, this disables stats gathering and no stats are displayed (since they are not collected). This means that a zero setting overrides whatever may be the value of ejb.collect.display_statistics, ejb.collect.statistics or ejb.collect.display_detail_statistics properties.	5
ejb.collect.display_ statistics=true|false	This flag turns on timer diagnostics, which allow the user to see how the Container is using the CPU.	false
ejb.collect.statistics= true|false	Same as the ejb.collect.display_statistics property, except this property does not write the timer value to the log.	false
ejb.collect.display_detail_ statistics=true|false	This flag turns on the timer diagnostics, as ejb.collect.display_statistics option does. In addition, it prints out method level timing information. This allows the developer to see how different methods of the bean are using CPU. Please note, that the console output of this flag will require you to widen your terminal to avoid wrapping of long lines.	false
ejb.mdb.threadMaxIdle=<num>	There is a VM wide thread pool maintained by the EJB container for message-driven bean execution. This pool has the same configurability as the ORB dispatcher pool for handling RMI invocations. This particular property controls the maximum duration in seconds a thread can idle before being reaped out.	300
ejb.mdb.threadMax=<num>	Maximum number of threads allowed in the MDB thread pool.	0 (no limit)
ejb.mdb.threadMin=<num>	Minimum number of threads allowed in the MDB thread pool.	0
ejb.allowNullsInFinders= true|false	This property is applicable only to CMP 2.x. If you set this property to true, it will allow the presence of NULLs as the return value of a finder or as part of a finder Collection. By default, this property is set to False.	False

EJB Customization Properties: Deployment Descriptor level

These properties customize the behavior of a particular EJB. Some of them are applicable only to a particular type of EJB (such as session or entity), others are applicable to any kind of bean. There are several places where these properties can be set. Below are these places in the order of precedence:

1 Property element defined on the EJB level in the ejb-borland.xml deployment descriptor of a JAR file. This setting affects this particular EJB only. For example, the following XML sets the ejb.maxBeansInPool property to 99 for the EJB named data:

```
<ejb-jar>
   ⋮
   <enterprise-beans>
      <entity>
         <ejb-name>data</ejb-name>
         <bean-home-name>data</bean-home-name>
         <property>
            <prop-name>ejb.maxBeansInPool</prop-name>
            <prop-type>Integer</prop-type>
            <prop-value>99</prop-value>
         </property>
      </entity>
   </enterprise-beans>
   ⋮
</ejb-jar>
```

2 Property element defined on the <ejb-jar> level in the ejb-borland.xml deployment descriptor of a JAR file. This setting affects all EJBs defined in this JAR. For example, the following XML sets the ejb.maxBeansInPool property to 99 for all EJBs in the particular JAR file:

```
<ejb-jar>
   ⋮
   <property>
      <prop-name>ejb.maxBeansInPool</prop-name>
      <prop-type>Integer</prop-type>
      <prop-value>99</prop-value>
   </property>
   ⋮
</ejb-jar>
```

3 Property element defined at the `<application>` level in the `application-borland.xml` deployment descriptor of an EAR file. This setting affects all EJBs defined in the all JARs located in this EAR file. For example, the following XML sets the `ejb.maxBeansInPool` property to 99 on the EAR level:

```
<application>
   ⋮
   <property>
      <prop-name>ejb.maxBeansInPool</prop-name>
      <prop-type>Integer</prop-type>
      <prop-value>99</prop-value>
   </property>
   ⋮
</application>
```

4 EJB property defined as an EJB container level property. This affects all EJBs deployed in this EJB container. For example, the following command sets the `ejb.maxBeansInPool` property to 99 for all beans deployed in the EJB container started standalone:

```
vbj -Dejb.maxBeansInPool=99 com.inprsie.ejb.Container ejbcontainer hello.ear
-jns -jss -jts
```

Complete Index of EJB Properties

Properties common for any kind of EJB

Property	Type	Description	Default
`ejb.default_transaction_attribute`	`Enumeration (NotSupported, Supports, Required, RequiresNew, Mandatory, Never)`	This property specifies a transaction attribute value for the methods which have no trans-attribute defined in the standard deployment descriptor. Note, that if this property is not specified, the EJB container does not assume any default transaction attribute. Thus, specifying this property, may simplify porting J2EE applications created with other application servers which assume some default transaction attribute.	None

Entity Bean Properties (applicable to all types of entities—BMP, CMP 1.1 and CMP 2)

Property	Type	Description	Default
ejb.maxBeansInPool	Integer	This option specifies the maximum number of beans in the ready pool. If the ready pool exceeds this limit, entities will be removed from the container by calling unsetEntityContext.	1000
ejb.maxBeansInCache	Integer	This option specifies the maximum number of beans in the cache that holds on to beans associated with primary keys, but not transactions. This is relevant for Option “A” and “B” (see ejb.transactionCommitMode below). If the cache exceeds this limit, entities will be moved to the ready pool by calling ejbPassivate.	1000
ejb.maxBeansInTransactions	Integer	A transaction can access any/large number of entities. This property sets an upper limit on the number of physical bean instances that EJB container will create. Irrespective of the number of database entities/rows accessed, the container will manage to complete the transaction with a smaller number of entity objects (dispatchers). The default for this is calculated as ejb.maxBeansInCache/2. If the ejb.maxBeansInCache property is not set, this translates to 500.	Calculated

Property	Type	Description	Default
ejb.transactionCommitMode	Enumeration (A|Exclusive, B|Shared, C|None)	This flag indicates the disposition of an entity bean with respect to a transaction. The values are: A or Exclusive: This entity has exclusive access to the particular table in the DB. Thus, the state of the bean at the end of the last committed transaction can be assumed to be the state of the bean at the beginning of the next transaction. For example, to cache the beans across transactions. B or Shared: This entity shares access to the particular table in the DB. However, for performance reasons, a particular bean remains associated with a particular primary key between transactions, to avoid extraneous calls to ejbActivate and ejbPassivate between transactions. This means the bean stays in the active pool. This setting is the default. C or None: This entity shares access to the particular table in the DB. A particular bean does not remain associated with a particular primary key between transactions, but goes back to ready pool after every transaction. This is generally not a useful setting.	Shared
ejb.transactionManagerInstanceName	String	Use this property to specify by name a particular transaction manager for driving the transaction started for method calls. This option is useful in cases where you need 2PC completion of a particular transaction but want to avoid the RPC overhead of using a 2PC transaction manager for all other transactions in the system. This is also supported for MDBs.	None

Property	Type	Description	Default
ejb.findByPrimaryKeyBehavior	Enumeration (Verify, Load, None)	This flag indicates the desired behavior of the findByPrimaryKey method. The values are: Verify: This is the standard behavior, for findByPrimaryKey to simply verify that the specified primary key exists in the database. Load: This behavior causes the bean's state to be loaded into the container when findByPrimaryKey is invoked, if the finder call is running in an active transaction. The assumption is that found objects will typically be used, and it is optimal to go ahead and load the object's state at find time. This setting is the default. None: This behavior indicates that findByPrimaryKey should be a no-op. Basically, this causes the verification of the bean to be deferred until the object is actually used. Since it is always the case that an object could be removed between calling find and actually using the object, for most programs this optimization will not cause a change in client logic.	ejb.checkExistenceBeforeCreate

Property	Type	Description	Default
Boolean	Most tables to which entity beans are mapped have a Primary Key Constraint. If the CMP engine attempts to create a bean that already exists, this constraint is violated and a DuplicateKeyException is thrown. Some tables, however, do not define Primary Key Constraints. In these cases, the checkExistanceBeforeCreate property can be used to avoid duplicate entities. When set to True, the CMP engine checks the database to see if the entity exists before attempting the insert operation. If the entity exists then the DuplicateKeyException is thrown.	False	

Message Driven Bean Properties

Property	Type	Description	Default
ejb.mdb.use_jms_threads	Boolean	Option to switch to using the JMS providers dispatch thread rather than the Container managed thread to execute the onMessage() method. For OpenJMS this value will be true, since the message will be delivered in the JMS providers dispatch thread.	false **Note:** This value will be true by default for OpenJMS.
ejb.mdb.local_transaction_optimization	Boolean	This property is currently used only with OpenJMS. It is used to attain atomicity without using the XAConnectionFactory. The same database is used for message persistence and application data.	true
ejb.mdb.maxMessagesPerServerSession	Integer	For JMS providers that support the option to batch load a ServerSession with multiple messages, use this property to tune performance.	5
ejb.mdb.max-size	Integer	This is the maximum number of connections in the pool.	None
ejb.mdb.init-size	Integer	When the pool is initially created, this is the number of connections AppServer populates the pool with.	None

Property	Type	Description	Default
ejb.mdb.wait_timeout	Integer	The number of seconds to wait for a free connection when maxPoolSize connections are already opened. When using the maxPoolSize property and the pool is at its max and can't serve any more connections, the threads looking for JDBC connections end up waiting for the connection(s) to become available for a long time if the wait time is unbounded (set to 0 seconds). You can set the waitTimeout period to suit your needs.	30
ejb.mdb.rebindAttemptCount	Integer	This is the number of times the EJB Container attempts to re-establish a failed JMS connection or a connection that was never established for the MDB. To make the Container attempt to rebind infinitely you need to explicitly specify ejb.mdb.rebindAttemptCount=0.	5
ejb.mdb.rebindAttemptInterval	Integer	The time in seconds between successive retry attempts (see above property) for a failed JMS connection or a connection that was never established.	60
ejb.mdb.maxRedeliverAttemptCount	Integer	This is the number of times a message will be re-delivered by the JMS service provider should the MDB fail to consume a message for any reason. The message will only be re-delivered five times. After five attempts, the message will be delivered to a dead queue (if one is configured).	5
ejb.mdb.unDeliverableQueueConnectionFactory	String	Should an MDB fail to consume a message for any reason, the message will be re-delivered by the JMS service. The message will only be re-delivered five times. After five attempts, the message will be delivered to a dead queue (if one is configured). This property looks up the JNDI name for the connection factory to create a connection to the JMS service. This property is used in conjunction with ejb.mdb.unDeliverableQueue.	None
ejb.mdb.unDeliverableQueue	String	Should an MDB fail to consume a message for any reason, the message will be re-delivered by the JMS service. The message will only be re-delivered five times. After five attempts, the message will be delivered to a dead queue (if one is configured). This property looks up the JNDI name of the queue. This property is used in conjunction with ejb.mdb.unDeliverableQueueConnectionFactory.	None
ejb.transactionManagerInstanceName	String	This property is currently supported only for MDBs that have the “Required” transaction attribute. Use this property to specify by name a particular transaction manager for driving the transaction started for the onMessage() call. This option is useful in cases where you need 2PC completion of this particular transaction but desire to avoid the RPC overhead of using the 2PC transaction manager for all other transactions in the system, for example, in entity beans. Please refer to the MDB chapter for more details.	None

Stateful Session Bean Properties

Property	Type	Description	Default
ejb.sfsb.passivation_timeout	Integer	Defines the time interval (in seconds) when to passivate inactive stateful session beans into the persistent storage (JSS).	5
ejb.sfsb.instance_max	Integer	Defines the maximum number of instances of a particular stateful session bean allowed to exist in the EJB container memory at the same time. If this number is reached and a new instance of a stateful session needs to be allocated, the EJB container throws an exception indicating lack of resources. 0 is a special value. It means no maximum set. Note, that this property is applicable only if the ejb.sfsb.passivation_timeout property is set to non-zero value.	0
ejb.sfsb.instance_max_timeout	Integer	If the max number of stateful sessions defined by the ejb.sfsb.instance_max property is reached, the EJB container blocks a request for an allocation of a new bean for the time defined by this property waiting if the number goes lower before throwing an exception indicating lack of resources. This property is defined in ms (1/1000th of second). 0 is a special value. It means not to wait and throw an exception indicating lack of resources immediately.	0
ejb.jsec.doInstanceBasedAC	Boolean	If set to true, the EJB container checks if the principal invoking an EJB's method is the same principal that created this bean. If this check fails, the method throws a java.rmi.AccessException (or javax.ejb.AccessLocalException) exception. This is applicable to stateful session beans only.	True

EJB Security Properties

Property	Type	Description	Default
ejb.security.transportType	Enumeration (CLEAR_ONLY, SECURE_ONLY, ALL)	This property configures the Quality of Protection of a particular EJB. If set to CLEAR_ONLY, only non-secure connections are accepted from the client to this EJB. This is the default setting, if the EJB does not have any method permissions. If set to SECURE_ONLY, only secure connections are accepted form the client to this EJB. This is the default setting, if the EJB has at least one method permission set. If set to ALL, both secure and non-secure connections are accepted from the client. Setting this property controls a transport value of the ServerQoPConfig policy.	None
ejb.security.trustInClient	Boolean	This property configures the Quality of Protection of a particular EJB. If set to true, the EJB container requires the client to provide an authenticated identity. By default, the property is set to false, if there is at least one method with no method permissions set. Otherwise, it is set to true. Setting this property controls a transport value of the ServerQoPConfig policy.	

Java Session Service (JSS) Properties

JSS can run as part of standalone EJB container (-jss option) or as part of a Partition.

As a “Partition service", JSS Configuration information is located in each Partition's data directory in the partition.xml file. By default, this file is located in the following directory:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>
/mos/<partition_name>/adm/properties/
```

For example, for a Partition named “standard", by default the JSS configuration information is located in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>
/mos/standard/adm/properties/partition.xml
```

For more information, go to the *partition.xml reference*, “<service> element”.

Otherwise, for the location of a Partition data directory, go to the configuration.xml file located in:

```
<install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/
```

and search for the Partition Managed Object directory attribute:

```
<partition-process directory=
```

The JSS supports two kinds of backend storage: JDataStore or a JDBC datasource. For more information, go to the Java Session Service (JSS) configuration section.

Property	Console Property Name	Description	Default
jss.workingDir=<path>	Working directory	The directory where the backend database (JDataStore) file is located. **Note:** this property is applicable only if the jss.pstore property is configured to use a JDataStore file as backend storage.	If not specified and the JSS runs in the Partition, then the Partition's working directory <install_dir>/var/domains/<domain_name>/configurations/<configuration_name>/mos/<partition_name> is used. If not specified and the JSS runs as part of a standalone EJB container, then the current directory where the container started is used.
jss.factoryName=<char_string>	Factory name	Name given to the JSS factory created by this service. The service gets registered with this name in the Smart Agent (osagent).	If not specified and the JSS runs in the Partition, the default value is: <server_name>:file:<install_dir>/ var/domains/<domain_name>/configurations/<configuration_name>/ mos/<partition_name>/ . If not specified and the JSS runs in a standalone EJB container, the default value is EJB/JSS[<container_name>].
jss.softCommit=true|false	Soft commit	If true, the JSS uses the JDataStore backend database with the Soft Commit mode enabled. Setting this property improves the performance of the Session Service, but can cause recently committed transactions to be rolled back after a system crash. **Note:** this property is applicable only if the jss.pstore property is configured to use a JDataStore file as backend storage. For more details, see the JDataStore documentation at: http://info.borland.com/techpubs/jdatastore/.	true
jss.maxIdle=<numeric value>	Max idle	The time interval in seconds between runs of JSS garbage collection job. The JSS garbage collection job is responsible for removing the state of expired sessions from the backend database. If set to 0, the garbage collection job never starts.	1800 (=30min)
jss.debug=true|false		Print debug information. If set to true, the JSS prints out debug traces.	false

Property	Console Property Name	Description	Default
jss.pstore=<char_string>	Persistent store	Specifies the JDatastore file to use for backend storage. If the file does not exist, JSS creates the file with the .jds extention, for example jss_factory.jds. For any compatible database supporting JDBC, specifies the JNDI name with the serial: prefix, for example serial://datasources/OracleDB to use for backend storage. In this case, JSS uses a datasource that is deployed in the Naming Service under the JNDI name specified.	If the JSS runs in the Partition, the JDataStore file named is used, such as jss_factory.jds. If the JSS runs in a standalone ejb container, the <container_name>_jss.jds is used.
jss.backingStoreType=<Dx|JDBC>		Specifies the type of persistent backend storage to use. Acceptable values are: Dx (indicating a local JDatastore database, or JDBC (indicating a JDBC datasource that is resolved from the JNDI name provided using the jss.pstore property value).	For a JSS that runs in a Partition, the default is Dx (local JDataStore database). If the JSS runs in a standalone ejb container, the default is Dx.
jss.userName=<char_string>	User name	User name JSS uses to open a connection with the JDataStore backend database. **Note:** this property is applicable only if the jss.pstore property is configured to use a JDataStore file as backend storage.	<default-user-name>
jss.passWord=<char_string>		When JSS persistent storage is defined through jss.backingStoretype=DX, use this property to specify the password value required for jss.userName access to the local JDataStore database. **Note:** This property is applicable only if jss.backingStoretype=Dx (configured to use JDataStore for persistent backend storage).	masterkey

Partition Transaction Service (Transaction Manager)

Listed below are properties that influence the behavior of the Partition Transaction Service (Transaction Manager). The properties can be specified when hosted by either a standalone EJB container or a Partition.

When configuring the Partition Transaction Service for a Partition, set the properties in the `partition.xml` file which is located in the `<install_dir>/var/domains/base/configurations/<configuration_name>/mos/<partition_name>/adm/properties`.

If running an EJB container standalone, they must be specified using system property names described below in section titled JTS System Properties. For example, when JTS is hosted by a standalone EJB Container property `jts.allow_unrecoverable_completion` must be specified using its system property equivalent:

```
prompt% vbj -DEJBAllowUnrecoverableCompletion com.inprise.ejb.Container
ejbcontainer beans.jar -jns -jts
```

Property	Description	Default
`jts.allow_unrecoverable_completion=true|false`	If set to `true`, this instructs the Container built-in JTS implementation to do a non-recoverable (that is, non two-phase) completion when there are multiple Resource registrations. Use at your own risk. It is provided only as a developer friendly feature. For OpenJMS, this property is set to true by default.	`False`
`jts.no_global_tids=true|false`	By default, JTS generates X/Open XA compatible transaction identifiers. By setting this property to `true`, the transaction key generation behavior changes to generate non-XA compliant tids. By generating XA compliant properties out of the box, the EJB container can work with JDBC2/XA drivers seamlessly.	`False`
`jts.no_local_tids=true|false`	There is an optimization where the EJB container detects that a transaction was started in the transaction service that lives in the same VM, and make the transaction comparison faster. Setting this property to `true` turns that off. This local transaction identifier (local tid) is a subset of the global transaction id hence makes the transaction comparisons faster.	`False`
`jts.timeout_enable=true|false`	By default, JTS transaction timeout facility is disabled. When enabled, each new transaction created by JTS will registered with a timeout in the JTS Timeout Manager. If the timeout expires before completion of the transaction, JTS will automatically rollback the transaction.	`False`
`jts.timeout_interval=<num>`	The JTS Timeout Manager examines registered transactions for timeout expiration at intervals in seconds controlled by the value of this property. Setting it to a value of `0` causes the interval to occur every 9999 seconds.	5

Property	Description	Default
jts.default_timeout=<num>	The timeout period for a Bean Managed transaction can be configured using JTA UserTransaction setTransactionTimeout() method. If not used or if transaction is a Container Managed transaction, the default transaction timeout value is applied. This can be configured upon JTS startup using the jts.default_timeout property value. The granularity of this property is 1 second.	600
jts.default_max_timeout=<num>	To prevent specification of an excessive timeout value for the jts.default_timeout property, the jts.default_max_timeout property controls the maximum time a transaction can be active before its expired. The granularity of this property is 1 second.	3600
jts.trace=true|false	Set this property to generate JTS debug messages.	False
jts.transaction_debug_timeout=<num>	If set, this property displays a list of active transactions maintained by JTS. Its value dictates the interval in seconds at which transactions are displayed.	None

39

Using LifeRay Portal with AppServer 6.7

This document describes the steps to prepare the liferay ear for deployment, create a liferay configuration using the Borland AppServer (AppServer) console, deploy the LifeRay portal, and deploy any custom portlets.

Liferay is an open-source portal that is designed to deploy portlets. It provides personalization, user/group management, web mail, message boards, content management all rolled into one package. It comes bundled with many portlet applications which are compliant with Java Portlet Specification, JSR-168.

In order to use the LifeRay Portal with AppServer do the following:

1. Download the LifeRay 4.0 EAR file from `http://www.liferay.com`.
2. Open a Borland Management Console and log in.
3. Create a configuration for the LifeRay Portal. When you create the configuration using the Management Console the configuration comes preconfigured with a LifeRay partition, JDataStore, and JMS.
   a. Right-click on Configurations node in the left pane and select Add Configuration... from the menu.
   b. Click on Portals in the Template Gallery.
   c. Select LifeRay Portal Configuration version 4.0 from the right pane in the Template Gallery and click on the Select button. The Create New Configuration dialog box will open.
   d. Enter a name for the new liferay configuration in the Name field.
   e. Change the Smart Agent Port in the Configuration Properties box by double-clicking on the value in the Value column.
   f. Click OK.
4. Run the configuration by right-clicking on the configuration name and selecting Start from the resulting menu.

5 Create a LifeRay server on which to host the LifeRay EAR file. Deploy the EAR file to the server.

   To create a LifeRay server:

   a Right-click on the Hosted Modules node under the LifeRay partition in the left pane of the Borland Management Console.

   b Select Host LifeRay module... from the menu. The Host LifeRay Portal dialog will open.

   c Enter the path to the LifeRay EAR file in the Liferay Ear Path box.

   d Enter the path to the directory where you want to host the LifeRay module in the Host Target Directory field. This directory must be on the same machine as the agent. Make sure that this directory already exists.

   e Make sure that the Generate Stub checkbox is checked.

   f Enter a name in the Module Name text box if you want to change the default module name. By default the module gets the same name as the directory in which you host the LifeRay module.

   g Click OK. You will see a status box. The AppServer will first generate the stub then extract the contents of the EAR file into the directory that you entered in step 3.

6 Test whether the LifeRay module has been deployed correctly by going to `http://localhost:8080` in a browser. This should open the LifeRay portal in the browser. The default login name for the LifeRay Portal is test@liferay.com and the default password is test.

Note

To change any of the LifeRay partition properties, right-click on the LifeRay partition name in the left pane of the Management Console and select Properties from the menu. Click on the desired tab to bring it forward and change the properties associated with that tab.

Using Other Databases

By default, LifeRay uses the JDataStore database to store data. You can use a database other than JDataStore. Refer to the `http://www.liferay.com/web/guest/documentation/development/databases` site for information on which databases are supported. If you would like to use a database other than JDataStore, you must do the following:

1 Create a new `liferay.dar` file with the JNDI information for the database you want to use in the `jndi-definitions.xml` file.

   For information on how to edit the `jndi-definitions.xml` file for the database you want to use, see `http://www.liferay.com/web/guest/documentation/development/databases` site.

   For information on how to create a DAR file, see “Creating a JNDI definitions archive (DAR)” in the *Management Console User's Guide*.

2 Replace the default `liferay.dar` that is included with the LifeRay portal configuration with the new one you created.

Deploying Portlet or J2EE modules to LifeRay module

You can add an EJB JAR, WAR, RAR or a library JAR to a hosted liferay module. To deploy any of these to a LifeRay module:

1. Open the Borland Management Console.
2. Expand the LifeRay partition node in the left pane.
3. Right-click on the LifeRay hosted module under the Hosted Modules node, and select Deploy Portlet from the menu. The LifeRay Portal Deployment Wizard will open.

   Proceed to step 4 only if you are using Liferay 3.6.
4. Click on the Add button to point the wizard to the portlet (WAR file) that you want to deploy.
5. Click on the Finish button.

To check whether the portlet is deployed successfully:

1. Open a web browser.
2. Go to the LifeRay portal by typing `http://localhost:8080` in a web browser
3. Log in to the portal using the default login which is test@liferay.com and password which is test.
4. if you are using Liferay 4.0, click Add Content.

   If you are using Liferay 3.6, Scroll down until you see the Add Portlet to Wide Column field. You should see the newly added portlet in the drop-down menu for field.
5. Select the portlet and click on the Add button to add the portlet to your portal.

40

Integrating Borland AppServer 6.7 with JBuilder 2006

This chapter explains how to install the Borland AppServer 6.7 plug-in for JBuilder 2006, configure the plug-in, and use Borland AppServer 6.7 with JBuilder 2006.

Configuring JBuilder 2006 for Borland AppServer 6.7

After you have installed the plug-in and restarted JBuilder, you need to configure JBuilder to use the plug-in.

To configure JBuilder settings for Borland AppServer 6.7,

1. Choose Enterprise|Configure Servers to display the Configure Servers dialog box. The right side of the dialog box displays the default settings for the server. The General page displays common fields, while the Custom page displays server-specific fields. In some cases, modifying a Custom setting will update a setting on the General page.
2. Select Borland Enterprise Server AppServer Edition 6.x from the User Home folder in the left pane.

   Note

   6.x refers to AppServer versions Borland Enterprise Server AppServer Edition 6.0RP1, Borland Enterprise Server AppServer Edition 6.5 (Patch 11), and Borland AppServer 6.7.
3. Select the Enable Server option at the top of the dialog box.
   Checking this option enables the fields for Borland AppServer 6.7. You won't be able to edit any fields until it is checked. The Enable Server check box also determines whether this server will appear in the list of servers when you select a server for your project using Project|Project Properties|Servers.
4. View and change (if required) fields on the General tab:
   - Home Directory: The directory where Borland AppServer 6.7 is installed. The default is `Borland/AppServer`. If the default directory is not correct, use the ellipsis (…) button to browse to the correct directory.

- Native Executable Launcher: The native executable used to run this server. The default is partition.exe in the <APPSERVER_HOME>/bin folder. If JBuilder is able to locate the native executable, this field is automatically filled in for you.
- VM Parameters: The parameters you want to pass to the virtual machine.
- Server Parameters: The parameters you want to pass to the server.
- Working Directory: The location of a working directory.

1 Click the Custom tab to view and change (if required) fields unique to the server. Change or fill in these fields:

- JDK Installation Directory: The directory where JDK v 1.5.0 is located. For Borland AppServer 6.7, this is automatically set to the <APPSERVER_HOME>/jdk/jdk1.5.0 folder. Your project will use this JDK to run the AppServer partition.
- Server Name: The Borland AppServer 6.7 hub name.
- Configuration Name: The name of the configuration that manages the partition. By default, this is set to jbuilder.
- Partition Name: The name of the partition in which the module is run. By default, this is set to jbpartition.
- Add A Management Agent Item To The Enterprise Menu: Adds a Management Agent item to the JBuilder Enterprise menu, so you can start the Management Agent quickly from the JBuilder IDE.
- Server Realm: The server realm name. For more information, see the *Borland Management Console User's Guide*. The default setting is ServerRealm.
- User Name: The name you use to identify yourself to the server. The default value is admin.
- User Password: The password you use to identify yourself to the server. The default value is admin.
- Advanced Settings: Click this button to display the Advanced Settings dialog box. Use this dialog box to change the port number used by the Management Agent and to select the Use Security option. The management port is used in JBuilder to detect the server during startup and deployment. Change the management port only if you are deploying to a remote server with a management port that is different from the default. When you change the port number, make sure that you entered the same port number as the server: the server won't start up without the correct port number. The port and security settings you choose must match the settings of your server. The values are read from the Borland AppServer property files, but the values will be changed automatically if you change the Home Directory setting. Note that changing the port number while you have the Management Agent started in JBuilder automatically shuts down the Management Agent.

1 Click OK to close the dialog box and save settings.
Any projects that were targeted for Borland Enterprise Server AppServer Edition 6.x are automatically updated to the new Borland AppServer home.

Displaying the Borland Management Console in JBuilder

After you have installed and configured JBuilder 2006 for Borland AppServer 6.7, you can display the Borland AppServer Management Console in the JBuilder message pane. To do so, you need to edit the jbuilder.config configuration file.

1 Save your project and exit JBuilder.

2 Open jbuilder.config in a text editor. This file is located in the <JBUILDER_HOME>/bin folder.

3 Add the following VM parameter to the configuration file:
`vmparam -Djava.endorsed.dirs=<APPSERVER_HOME>/lib/endorsed`

4 Restart JBuilder.

5 Choose View|Panes|BAS Console 6.7.
The Borland Management Console is displayed in the message pane.

VisiBroker development with JBuilder

You use the CORBA node of the Enterprise Setup dialog box (Enterprise|Enterprise Setup) to set up Borland AppServer for use with VisiBroker 7.0 in JBuilder 2006.

To make the ORB available to JBuilder,

1 Choose Enterprise|Enterprise Setup to display the Enterprise Setup dialog box. Select the CORBA page. The parameters in this dialog box allow a JBuilder user to develop CORBA applications.

2 Select the VisiBroker (Borland Enterprise Server AppServer Edition 6.x) option from the Configuration drop-down list. This option is automatically updated to point to the `<APPSERVER_HOME>/bin` folder.

3 To start the Smart Agent from the Tools menu, check the Add The VisiBroker Smart Agent Item To The Tools Menu option.

4 Enter the number of the SmartAgent port in the SmartAgent Port field.

5 Click OK to save the configuration.

Important

In the runtime configuration for your CORBA application, you need to add the following parameters to the VM Parameters field in the Edit Runtime Configuration dialog box (Run|Configurations|Edit):

```
-Dvbroker.agent.port=<your_osagent_port>
-Dborland.enterprise.licenseDir=<APPSERVER_HOME>/var
-Dborland.enterprise.licenseDefaultDir=<APPSERVER_HOME>/license
```

You've now completed setting up your system to use VisiBroker 7.0, installed with Borland AppServer 6.7. Before running your application, choose Tools|VisiBroker Smart Agent to start the Smart Agent.

Using the JBuilder Deployment Descriptor Editor to develop J2EE 1.4 applications

Borland AppServer 6.7 supports J2EE 1.4 applications. The JBuilder Deployment Descriptor Editor provides new editors for J2EE 1.4 applications targeted for Borland AppServer 6.7.

J2EE 1.4 requires that each deployment descriptor be validated against an XML schema, instead of a DTD. The Borland AppServer 6.7 plug-in for JBuilder 2006 supports these updates. When you right-click a J2EE deployment descriptor in the project pane and choose Validate, the descriptor is validated against an XML schema file. The J2EE modules for Borland AppServer deployment are validated against the following schema files:

- Application module: `application_1_4-borland.xsd`
- Application Client module: `application-client_1_4-borland.xsd`
- Connector module: `connector_1_5.xsd`
- EJB module: `ejb-jar_2_1-borland.xsd`
- Web module: `web-app_2_4-borland.xsd`

To view the standard Java commented J2EE 1.4 XML schemas, go to Java 2 Platform, Enterprise Edition (J2EE) : XML Schemas for J2EE Deployment Descriptors at:

`http://java.sun.com/xml/ns/j2ee/`

To support Borland AppServer 6.7, the JBuilder Deployment Descriptor Editor now contains updated editors or new editors for the following entities:

- Message Destinations page: Web module, EJB module, Application Client module
- Message Destination Reference page: Web module, Application Client module, session bean, entity bean, message-driven bean
- Message-Driven Bean page: Message-driven bean
- Resource Environment References page: Message-driven bean
- Admin Object and Admin Object Properties page: Message-driven bean
- Resource Adapter page: Connector module
- BES Connection Definition page: Connector module

Note

JBuilder provides a Deployment Descriptor editor for editing standard and server-specific J2EE 1.3 modules and deployment descriptors. For information on the editor, see “Editing EJB deployment descriptors” in *Developing Applications with Enterprise JavaBeans* in the JBuilder online Help.

Message Destinations page

The Messages Destinations page is a new DD Editor page for a Web module, EJB module, and Application Client module.

Use the Message Destinations page to set the new J2EE 1.4 deployment descriptor element `<message-destination-name>` for a Web module, EJB module, or Application Client module. This element specifies the name of a message destination reference. The Borland AppServer 6.7 value is a JNDI name that represents the message destination reference name used in the web module, EJB module, or application client module.

To display the Message Destinations page and set the standard and Borland AppServer-specific deployment descriptor elements,

1. Select the Web module, EJB module, or Application Client module in the project pane.
   Select the DD Editor tab at the bottom of the content pane.
2. Open the structure pane.
3. Expand the module node and select the Message Destinations node.
4. To add a message destination, right-click the node and choose Add.
5. Click the Standard tab in the DD Editor.
   a. Enter the name of the message destination in the Name field. The name must be unique among the names of message destinations within the deployment file.
   b. Enter the language with which to associate the Display Name, Description, and icons in the Language field. You can have one Display Name, Description, and large and small icon for each language. Use the Add and Remove buttons to add and remove languages.
   c. Enter the name to be displayed the Display Name field.
   d. Enter the description in the Description field.
   e. Enter the location of a large icon (32 x 32 pixels) in the Large Icon field. The icon must be contained in the module tree.

f Enter the location of a small icon (16 x 16 pixels) in the Small Icon field. The icon must be contained in the module tree.

6 Click the BES tab in the DD Editor. In the JNDI Name field, enter the JNDI name for the message destination. See “Using JMS” for more information.

Message Destination Reference page

The Message Destinations Reference page is a new DD Editor page for a Web module and Application Client module, as well as for entity beans, session beans, and message-driven beans.

Use the Message Destination Reference page to set the new J2EE 1.4 deployment descriptor element `<message-destination-ref>` for a Web module and Application Client module, as well as for entity beans, session beans, and message-driven beans. The message destination reference contains a declaration of the reference associated with a resource.

To display the Message Destination Reference page and set the Borland AppServer 6.7-specific deployment descriptor element,

1. Select the Web module, EJB module, or Application Client module in the project pane.

   Select the DD Editor tab at the bottom of the content pane.
2. Open the structure pane.
3. Expand the module or bean node and select the Message Destination References node.
4. To add a message destination reference, right-click the node and choose Add.
5. Click the Standard tab in the DD Editor. Set the following attributes:
   a. Enter the name of the message destination reference in the Name field. This is the name used in deployment component code.
   b. Specify the Java type of the message destination in the Type field. The type specifies the Java interface to be implemented by the destination.
   c. Choose how the message destination will be used in the Usage field. Choose Consumes if messages are consumed from the message destination; choose Produces if messages are produced for the destination, or choose ConsumeProduces if messages are both consumed and produced. The Assembler makes use of this information in linking producers of a destination with its consumers.
   d. Specify the message destination link in the Link field. This element links a message destination reference or message-driven bean to a message destination. The Assembler sets the value to reflect the flow of messages between producers and consumers in the application. The value must be the name of a message destination in the same deployment file or in another deployment file in the same J2EE application unit. Alternatively, the value may be composed of a path name specifying a deployment file containing the referenced message destination with the name of the destination appended and separated from the path name by `#`. The path name is relative to the deployment file containing a deployment component that is referencing the message destination. This allows multiple message destinations with the same name to be uniquely identified.
   e. Enter the language to which the description applies in the Description Language field. Use the Add and Remove buttons to add and remove languages. You can have one description per language.
   f. Enter the description of the message destination reference in the Description field.
6. Click the BES tab in the DD Editor.
7. Enter the JNDI name for the message destination. See “Using JMS” for more information.

Message-Driven Bean page

The Borland-specific Messages-Driven Bean page was updated for message-driven beans.

Use the Message-Driven Bean page of the DD Editor to set the deployment descriptor for a message-driven bean. Both the standard and Borland AppServer 6.7-specific pages have been updated for the J2EE 1.4 implementation.

To display the Message-Driven Bean page and set standard and BAS-specific deployment descriptor elements,

1. Select the EJB module in the project pane.
   The DD Editor is displayed in the content pane.
2. Open the structure pane.
3. Expand the EJB module, then the Message Driven Beans node. Select a message-driven bean.
   The Message-Driven Bean page is displayed in the DD Editor.
4. Click the Standard tab in the DD Editor. Set the following attributes:
   a. Enter the name of the message-driven bean in the Name field.
   b. Enter the fully-qualified name of the Java class that implements the bean's business methods in the EJB Class field. This information must be specified.
   c. Choose the bean's messaging type in the Messaging Type field.
   d. Select how the bean's transactions are managed in the Transaction Type drop-down list. Transactions can be managed by the bean itself or by the container.
   e. Select the message destination type in the Message Destination Type drop-down list. This is the actual topic or queue the message-driven bean is listening to.
   f. Enter the bean's message destination in the Message Destination Link field.
   g. Enter the language for the activation configuration description in the Description Language field. Use the Add and Remove buttons to add and remove languages. You can have one description per language.
   h. Enter a description of the bean in the Description field.
   i. Enter the language for the display information in the Language field. Use the Add and Remove buttons to add and remove languages. You can have one display description per language.
   j. Enter the name you want used to identify the bean for display purposes in the Display Name field.
   k. Enter a description for display purposes in the Description field.
   l. Enter the name of a large icon (32 x 32 pixels) you want associated with the bean in the Large Icon field.
   m. Enter the name of a small icon (16 x 16 pixels) you want associated with the bean in the Small Icon field.
5. Click the BES tab in the DD Editor. Use this page to set the Borland AppServer `<message-source>` element.
   a. Choose the message source type from the Message Source Type drop-down list. Choose jms-provider-ref to activate the message source through a JMS provider, using EJB 2.0 implementation. Choose adapter-ref to activate the message source through a JCA 1.5 resource adapter.
   b. Enter the message-driven bean destination in the Destination Name field. This is the actual topic or queue the message-driven bean is listening to. This field is available if jms_provider_ref is selected as the Message Source Type.

c Enter the resource connection factory used to connect to the JMS broker in the Connection Factory Name field. This field is available if jms_provider_ref is selected as the Message Source Type.

d Enter the initial number of connections in the Initial Pool Size field. This field is available if jms_provider_ref is selected as the Message Source Type.

e Enter the maximum number of connections in the Maximum Pool Size field. This field is available if jms_provider_ref is selected as the Message Source Type.

f Enter the length of time (in seconds) to wait for a connection in the Wait Timeout field. This field is available if jms_provider_ref is selected as the Message Source Type.

g Enter the name of the resource adapter instance that connects to a J2EE resource in the Instance Name field. This field is available if resource_adapter_ref is selected as the Message Source Type.

Resource Environment References page

The Borland-specific Resource Environment References page was updated for message-driven beans.

Use the Resource Environment References page to set the `<resource-environment-ref>` element for a message-driven bean. The resource environment reference can be set to either a JNDI name or an administered object. A resource environment reference maps a logical name used by the client application to the physical name of an object.

To display the Resource Environment Reference page and set the Borland AppServer 6.7-specific deployment descriptor element,

1 Select the EJB module in the project pane.
The DD Editor is displayed in the content pane.

2 Open the structure pane.

3 Expand the EJB module and select the Message Driven Beans node. Select a message-driven bean.

4 Right-click the Resource Environment References page and choose Add.

5 Click the BES tab in the DD Editor.

a Choose the type of reference from the Resource Environment References Type drop-down list. Choose JNDI name to select a JNDI reference. Choose Admin Object to select an administered object. If you select Admin Object, you need to set properties for the object. See “Admin Object and Admin Object Properties page” for more information.

b Enter the name of the JNDI bean that maps the logical name to the object name in the JNDI Name field. This field is only available if JNDI Name is selected as the Resource Environment Type. See “Using JMS” for more information.

Admin Object and Admin Object Properties page

The Admin Object and Admin Object Properties pages are new for resource environment references on message-driven beans.

Use the Admin Object page to add an administered object for a resource environment reference. Use the Admin Object Properties page to set the object properties. This page is only available if you select Admin Object as the Resource Environment References Type. Administered objects are specific to a messaging style or message provider.

To display the Admin Object page and set the Borland AppServer 6.7-specific deployment descriptor element,

1. Select the EJB module in the project pane.
   The DD Editor is displayed in the content pane.
2. Open the structure pane.
3. Expand the EJB module and select the Message Driven Beans node. Select a message-driven bean.
4. Right-click the Resource Environment References page and choose Add.
5. Click the BES tab in the DD Editor.
6. Choose Admin Object from the Resource Environment Reference Type drop-down list.
7. Expand the Resource Environment References node in the structure pane until you see the entry you just added.
8. Expand the node and select the Admin Object Properties node. Right-click the node and choose Add.
   The BES Admin Object Properties page is displayed in the DD Editor.
9. Enter properties:
   a. Enter the property name in the Name field.
   b. Choose the type of property from the Type drop-down list. Select one of java.lang.String, java.lang.Boolean, or Integer. You can also select <Unspecified>.
   c. Enter a value for the property in the Value field. The value must match the type of property.

Resource Adapter page

The Borland-specific Resource Adapter page is new for a Connector module.

Use the Resource Adapter page to set the Borland-specific JCA 1.5 deployment descriptor `<resourceadapter>` element for a Connector module. This element describes a resource adapter for a connector.

To display the Resource Adapter page and set the Borland AppServer 6.7-specific deployment descriptor element,

1. Select the Connector module in the project pane.
   The DD Editor is displayed in the content pane.
2. Open the structure pane.
3. Expand the Connector module and choose the Resource Adapter node.
4. Click the BES tab in the DD Editor.
   a. Enter the name of the connection factory in the Instance Name field.
   b. Enter the resource adapter link reference in the Resource Adapter Link Reference field. This allows you to associate multiple deployed resource adapters with a single deployed resource adapter. The link provides for linking

and reusing resources already configured in a base resource adapter to another resource adapter, modifying only a subset of attributes. Using this field avoids duplication of resources where possible. Any values defined in the base resource adapter deployment are inherited by the linked resource adapter unless otherwise specified.

c Enter the directory where all shared libraries should be copied in the Resource Adapter Library Directory field.

d Enter the authorization domain for the connection in the Authorization Domain field.

BES Connection Definition page

The BES Connection Definition is new for a Connector module.

Use the BES Connection Definition page to set the Borland-specific JCA 1.5 deployment descriptor `<outbound-resourceadapter>` element for a Borland Connector Module. The information includes the fully qualified names of classes and interfaces that are required as part of the connector architecture, the number of managed connections, and the connection timing intervals.

To display the BES Connection Definition page and set the Borland AppServer 6.7-specific deployment descriptor element,

1 Select the Connector module in the project pane.
The DD Editor is displayed in the content pane.

2 Open the structure pane.

3 Expand the Connector module and the Resource Adapter node.

4 Right-click the BES Connection Definitions node and choose Add.

5 Set attributes for the connection definition:

a Enter the name of the factory interface in the Factory Interface field.

b Enter the name of the factory class used to connect to the JMS broker in the Factory Name field.

c Enter the connection description in the Description field.

d Enter the name of the JNDI class to the connection factory in the JNDI Name field. See “Using JMS” for more information.

e Check the Enable Logging option to require logging for the ManagedConnectionFactory or ManagedConnection classes.

f Enter the name and location of the file where the logging results are to be written in the Log File Name field.

g Enter the initial number of managed connections the server attempts to allocate at deployment time in the Initial Capacity field.

h Enter the maximum number of managed connections the server allows to be allocated at any one time in the Maximum Capacity field.

i Enter the length of time in seconds to wait if the connection is busy in the Busy Timeout field.

j Enter the length of time in seconds to wait before timing out a connection in the Idle Timeout field.

k Enter the length of time in seconds to wait for a connection in the Wait Timeout field.

l Enter the number of managed connections the server attempts to allocate when fulfilling a request for a new connection in the Capacity Delta field.

m Select the Enable Cleanup option to force the server to attempt to claim unused managed connections to save system resources.

n Enter the time in seconds the server waits between attempts to claim used managed connections in the Cleanup Interval field.

6 To add connection definition properties, expand the node for the definition you just added, right-click the Properties node and choose Add. Enter properties:

a Enter the property name in the Name field.

b Choose the type of property from the Type drop-down list. Select one of java.lang.String, java.lang.Boolean, or Integer. You can also select <Unspecified>.

c Enter a value for the property in the Value field. The value must match the type of property.

Creating a run configuration for Borland AppServer 6.7 targeted projects

JBuilder uses a default Borland AppServer configuration called jbuilder and a default partition called jbpartition as a deployment target for Borland AppServer 6.7. If the configuration or partition doesn't exist, one will be created automatically when you configure the plug-in. The server name is the same as the hub name.

You can start multiple partitions using multiple JBuilder run configurations. To create multiple run configurations, follow these steps:

1 Choose Run|Configurations, then click New.

2 Change the Run Type to Server. The displayed server is the server selected for the project in Project Properties|Server.

3 In the Category list, select Server|Command Line.

4 Change the Partition and Configuration fields to the ones you want to use.

5 Click OK to close the New Runtime Configuration dialog box and save the configuration.

6 Repeat these steps to create additional run configurations to run other partitions.

7 To run multiple partitions, use the Management Agent (Enterprise|Borland Enterprise Server Management Agent).

To avoid naming service conflicts, make sure that you have the naming service enabled for only one of the partitions. To disable the naming service for a partition, edit the run configuration for the partition (Edit Runtime Configuration dialog box) and uncheck the Naming/Directory service in the Category list.

Before starting up the partition, make sure you have configured unique port numbers for the Tomcat and JDataStore services.

Important

You cannot change the Tomcat port setting from the JBuilder run configuration. You need to start the server, open the Borland AppServer Management Console, and set the port on the server side. To do this,

1 Start the Borland Management Agent from the command line: `<APPSERVER_HOME>/bin/scu.exe`

2 Open the Borland AppServer Management Console: `<APPSERVER_HOME>/bin/console.exe`

3 Log into the console.

4 Choose the Management Hubs node.

5 Expand the Management Hubs node until you see the running partition. Expand the partition node.

6 Right-click the Web Container node and choose Properties.
The Configure Web Container dialog box is displayed.

7 Expand the Service: HTTP node. Choose Connector.

8 Scroll through the Connector page until you see the Port Number field.

9 Change the port number to the one you want to use.

10 Choose File|Save.

Changing the management port

The Management Agent manages all partitions. The default management port, set on the Advanced Settings dialog box (Tools|Configure Servers|Advanced Settings) is 42424. You can change the management port used by JBuilder, or the one used by the server.

To change the port used by JBuilder,

1 Choose Enterprise|Configure Servers and chose Borland Enterprise Server AppServer Edition 6.x from the User Home Folder on the left.

2 Click the Custom tab, then the Advanced Settings button.

3 Change the port in the Management Port field. (The default is 42424.)

4 Click OK two times.

To change the management port used by the server,

1 Open the Borland Management Console embedded in JBuilder 2006 (View|Panes|BAS 6.7 Console).

Note

The embedded console has to be enabled first. See “Displaying the Borland Management Console in JBuilder”.

2 Double-click Installations.

3 Expand the server location node and select the server node.

Note

The server is also the machine ID.

4 Expand the server node until you see the server displayed as a sub-node of the Agents node.

5 Right-click the server node and choose Properties.

6 Change the port number in the Management Port field.

7 Click OK to save the settings.

If you change the management port, the management agent will shutdown if it was started in JBuilder.

Launching the partition in JBuilder 2006

When you launch the partition in JBuilder, the Management Agent is started by default and the server is launched. After the partition has been started, all deployable archives are automatically deployed. Startup output is displayed in the message pane.

If you want to start multiple partitions, or quickly redeploy partitions, you can start the Management Agent (Enterprise|Borland Enterprise Server Management Agent). The Management Agent refers to scu.

To launch the partition and configuration for the server, right-click the module in the project pane you want to run. Select Run Using <Configuration_Name>. Typically, this will be the name of the server runtime configuration.

Running the partition in JBuilder will:

- Create the partition and configuration, if not already present. The partition and configuration name are derived from the run configuration used to start up the server. If you are starting the configuration or server using the default configuration, the partition name as configured in the application server properties will be used.
- Deploy any resources defined in `jndi-definitions.xml` (if present) to the root of the partition directory, `<APPSERVER_HOME>\var\domains\base\configurations\<CONFIGURATION_NAME>\mos\<PARTITION_NAME>\dars\jbuilder.dar`. The `jndi-definitions.xml` file is packaged into this `.dar` file and deployed.

  Note

  The `jndi-definitions.xml` file is created if your project contains an EJB 2.0 module with data sources/messaging resources defined. This action can be turned off by unchecking the Deploy jndi-definitions.xml option on the Deployment|EJBs Service Properties page on the Server node of the Project Properties dialog box (Project Properties|Server|Services|Deployment|EJBs).
- Remove any archives deployed to the partition, if the Remove Archives Already Deployed To Server option is selected. You can set this option in the Server|Archives category on the Edit Runtime Configuration dialog box for the server run configuration (Run|Configurations|<Server_Config_Name>|Edit|Run page|Category list).
- Deploy the selected archives. By default, all deployable archives in the project are selected. You can choose the archives to deploy in the Server|Archives category on the Edit Runtime Configuration dialog box for the server run configuration (Run|Configurations|<Server_Config_Name>|Edit|Run page|Category list).
- Start the partition. When partition startup is complete, you should see the partitions listed in the message pane for Borland AppServer 6.7. Archives deployed at startup should be loaded and accessible.

Note

By default, all services associated with a partition are started.

Deploying

To deploy EJBs, WARs, and EAR modules to Borland AppServer 6.7, follow these steps.

1. Choose Enterprise|Configure Servers.
2. Select Borland Enterprise Server AppServer Edition 6.x from the left side of the dialog box.
3. Click the Custom tab and set the Server, Configuration, and Partition names to match that of the server. (The server can be running remotely or on the local machine.)
4. Click the Advanced Settings button and make sure the Management Port setting matches the port set for the server.
5. Click OK two times.

Now you are ready to deploy. There are two ways to deploy EJBs, WARs, and EAR modules using JBuilder. You can deploy using the Deployment wizard or the context menu.

To deploy using the Deployment wizard,

1. Build your project (Project|Make).
2. Start the Management Agent (Enterprise|Borland Enterprise Server Management Agent).
3. Open the Server Deployment wizard (Enterprise|Server Deployment).
4. On Page 1 of the wizard, select the modules to deploy. Set the Restart Partitions On Deploy (Cold Deploy) option if the partition has already been started. Click Next.
5. On Page 2, select the partition you want to deploy modules to from the list.

   Important

   If the selected partition has already been started, restart it to access the deployed modules.
6. Click Finish to deploy.

To deploy from the context menu,

1. Build your project (Project|Make).
2. Start the Management Agent (Enterprise|Borland Enterprise Server Management Agent).
3. In the project pane, right-click any deployable node.
4. Choose Deploy Options|Deploy.

   Note

   To deploy more than one module, you can select multiple deployable nodes in the project pane, right-click them, and use the Deploy Options context menu.

Remote debugging

Before you can debug your application remotely, you need to configure the partition. Choose one of these topics for more information:

– Preparing to remote debug partitions that are not managed in JBuilder

– Preparing to remote debug partitions with JBuilder

Once the configuration, partition, and server are started, follow the instructions in "Remote debugging from JBuilder". For a remote debugging tutorial, see "Tutorial: Remote debugging with the Borland Enterprise Server AppServer Edition 6.0" in *Developing Enterprise JavaBeans* in the JBuilder online help. The steps are the same for both the 6.x and 6.7 versions of the Borland application server.

Preparing to remote debug partitions that are not managed in JBuilder

To prepare to remote debug partitions that are not managed in JBuilder,

1 Start the Borland Management Agent from the command line: `<APPSERVER_HOME>/bin/scu.exe`

2 Open the Borland AppServer Management Console: `<APPSERVER_HOME>/bin/console.exe`
Log into the console.

3 Choose the Management Hubs node. Expand the node until you see the partition you want to debug.

4 Right-click the partition name and choose Properties.
The Partition Properties dialog box is displayed.

5 Select the Partition Process Settings tab.

6 Check the Enable JPDA Remote Debugging option.

7 Set the JPDA Debugging Transport Address to 3999.

8 Uncheck the Suspend Partition Until Debugger Attaches option.

9 Click OK.

Preparing to remote debug partitions with JBuilder

To prepare to remote debug partitions with JBuilder,

1 Shut down the server.

2 Open the file: `<APPSERVER_HOME>/var/domains/base/configurations/<CONFIGURATION_NAME>/configuration.xml`

3 Look for the JPDA element and edit attribute values as follows:

```
 enable-jpda-debug="true"
jpda-transport-address="3999"
jpda-suspend="false"
```

Remote debugging from JBuilder

Once the server, partition, and Management Agent have been started, follow these steps from the JBuilder IDE:

1. In the project from which you want to launch the remote debug session, choose Run|Configurations.
2. Select the Server run configuration and choose Edit.
3. Select the Debug|Connection node.
4. Select the Remote Attach option.
5. Set the Transport Type to dt_socket and the localhost value to 3999.
6. Click OK two times to close the Run Configuration dialog boxes.
7. Set a breakpoint in the process you want to debug.
8. Click the down arrow next to the Debug Project button on the toolbar and select the Server configuration you just created or edited. The debugger launches, attaches to the partition running remotely, and stops at the breakpoint.

41

Managing the Borland AppServer

This section contains an overview of the Borland AppServer (BAS) management architecture. The key concepts and terminology you need to understand in order to successfully manage BAS are covered as well.

Understanding the SCU process

At the heart of AppServer is a process called *SCU*. The SCU contains both the Hub logic and the Agent logic, making it one process with two capabilities. As a result, the SCU can support the following modes:

- Hub/Local Agent, where the both the Hub and the Agent are located on the same host
- Hub/Agent, where the Hub and the Agent are located on different hosts

By default, when you install AppServer you install the Hub. However, because the SCU contains both the Hub logic and the Agent logic, you have the flexibility to change the mode of any Hub/Agent post installation using the Management Console.

The following graphic illustrates the relationship between the various entities:

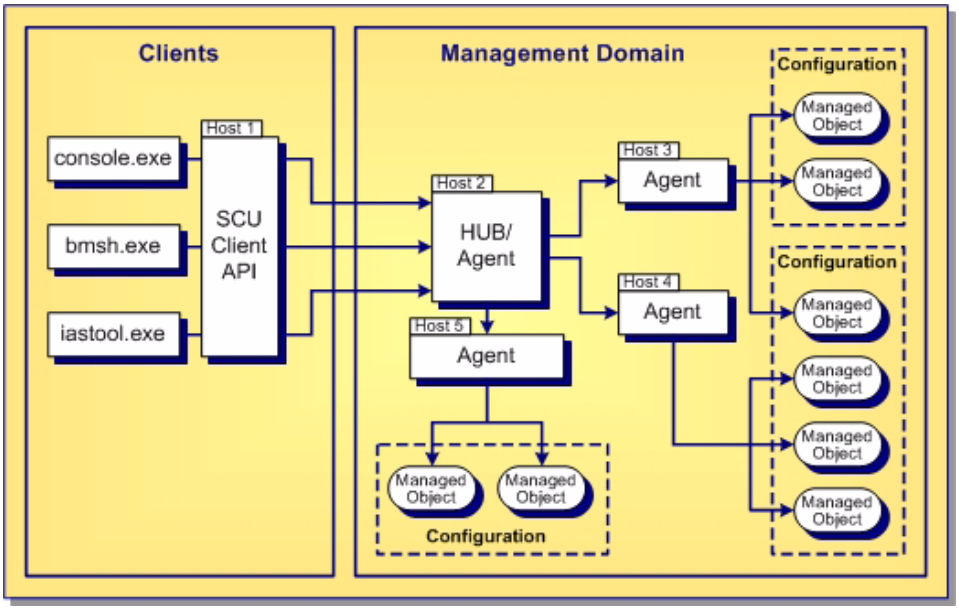

For information about starting the SCU process (Hubs and Agents), go to the Starting the Hub and Agents section.

Hub

A *Hub* is the administrative, control center of your AppServer management system. A Hub can also act as an agent, because the Hub and Agent are part of the same SCU process.

By default, the Hub is assigned the name of the host on which you are installing it - unless you enter a different name during install.

For information about starting a Hub, go to the Starting the Hub and Agents section.

A Hub can manage local resources as well as distributed resources (using Agents). By default, the BAS installation installs a Hub. Once you have installed BAS, you can later change a Hub to an Agent using the Management Console.

Agent

As the administrator of your management system, the Hub delegates actions to the *Agents* which in turn implement those actions.

Because the Hub and Agent are part of the same SCU process, a Hub has at least one Agent - the *Local Agent*. The Local Agent is co-located on the same host as a Hub.

You can install Agents on hosts other than the host on which the Hub is installed. By default, when you install a Agent, it is assigned the name of the host on which you are installing it.

By default, when you install AppServer you install the Hub. However, because the SCU contains both the Hub logic and the Agent logic, you have the flexibility to change the mode of any Hub/Agent post installation using the Management Console. To change a Hub to an Agent:

1. In the Management Console, on the tree structure expand the Agents node.
2. Right-click on the agent whose properties you want to change.
3. Choose Properties from the resulting menu.
4. In the Properties dialog, make sure you are viewing the General tab.

5 From the Agent type drop-down menu, select Agent.

6 Click OK.

Important

The Management Console prompts you to restart the Agent - you must restart the Agent in order for this change to be applied.

Starting the Hub and Agents

The Hub and Agents are started by starting the SCU executable (`scu.exe`).

Starting a Hub and its Local Agent

To start a Hub and its Local Agent:

1 On the machine where the Hub/Local Agent is installed, open a Command Prompt window.

2 Migrate to: `<install_dir>/bin`

3 Type the following command and press Enter:

```
scu
```

or

On a Windows machine where the Hub/Local Agent is installed, on the Start menu, choose:

```
Start|Programs|Borland AppServer|Borland Management Agent
```

Starting an Agent

To start a Agent:

1 On the machine where the Agent is installed, open a Command Prompt window.

2 Migrate to: `<install_dir>/bin`

3 Type the following command and press Enter:

```
scu
```

or

On a Windows machine where the Agent is installed, on the Start menu, choose:

```
Start|Programs|Borland AppServer|Borland Management Agent
```

Note

If you have BAS installed on an NFS filesystem and it is being run across NFS, you could potentially get an error message saying "There may be an agent running in this footprint" when you attempt to start the SCU process. This problem only manifests itself on certain combinations of client servers on different vendor's UNIX implementations. To avoid this problem we recommend that you avoid installing BAS on an NFS.

Management Port

A Hub and its Agents must be set to the same *Management Port* to find each other. By default, when you install a Hub or Agent, the Management Port is set to `42424`.

However, you can later change the Management Port for a Hub or Agent using the Management Console.

To change the Management port number:

1. Open the Borland Management Console.
2. On the tree structure, right-click the name of the Agent whose port you want to change and select Properties.
3. In the Properties dialog, make sure you are viewing the General tab.
4. In the Management Port field, type port number.
5. Click OK.

Important

The Management Console prompts you to restart the Agent - you must restart the Agent in order for this change to be applied.

Management Domain

A Management Domain is defined by the set of Hubs and Agents that share a single Management Port. You can have multiple Hubs and Agents use the same Management Port and in turn be in the same Management Domain.

Configurations

A Configuration is the central container of the information a Hub references in order to manage the entities of a Management Domain. For each Hub, you can create multiple Configurations. A single configuration can span across multiple agents. It is within a Configuration that you define the collections of resources that you want to monitor and control as a single entity.

You create Configurations for your Hubs, not for your Agents. Only a Hub can understand and interpret the information within a Configuration. Based on the Configurations you create, the Hub delegates actions to the Agents. Each Agent implements the actions on its host.

For more information about Configurations, go to the Configurations section.

Managed Resources

The Managed Resources are the actual entities of your system that you want to manage. For example:

- a text editor such as Notepad
- a web server such as Apache
- a database such as JDataStore
- a web page
- a pinger

Managed Objects

In BAS, you use a *Managed Object* to represent each Managed Resource you want to configure, control and/or simply monitor. For each Managed Resource that you want a Hub and its Agents to manage, a Managed Object must exist within a Configuration.

You use Managed Objects to model your Managed Resources in one of the following two ways:

- to control the Managed Resource's lifecycle, or
- to solely determine status.

BAS management of Managed Objects is implemented by Agents per instructions from a Hub. Each Managed Object must be assigned to a single Agent within the Management Domain. However, each Agent can manage multiple Managed Objects.

When you create a Configuration using a Configuration template, for each of the Managed Objects the template includes, you specify an Agent. For each Managed Object assigned to an Agent, the Hub sends the Managed Object Configuration information to the Agent when:

1. The Managed Object is added, and
2. each time you update the Managed Object information in the Configuration.

For more information about Managed Objects, go to the Managed Objects section.

For more information about Configuration Templates, go to the Configuration Templates section.

Borland Management Console

The Borland Management Console is a graphical user interface. It provides both logical and physical graphical depictions of your Configurations, as well as access to the `configuration.xml` file for each Configuration. You should use the Management Console for all stages of BAS management - from creating and modifying your Configurations to monitoring and diagnosing configuration states.

42

Configurations

Using Configurations, Borland AppServer (BAS) enables you to group and manage distributed resources based on functional dependencies. Furthermore, through Configurations, BAS enables you to control, monitor and manage these defined resources as a single entity. This section describes the important role Configurations play in managing your system resources and the stages of the Configuration lifecycle.

Creating Configurations

You should use the Management Console to create new Configurations. Remember, you always create a Configuration within the context of a Hub, and each of the Managed Objects you define within a Configuration you assign to an Agent.

When you create a Configuration, you use a Configuration Template which includes all or a subset of the following:

- Configuration-wide properties, elements, and attributes,
- the Managed Objects appropriate for that particular Configuration, and
- the Managed Object specific properties, elements, and attributes.

BAS includes the capability to add, remove and otherwise modify Managed Objects and properties after you create a Configuration. However, after you first create your Configurations, we do not recommend modifying the Managed Object-to-Agent assignments.

For more information, go to the Managed Objects section.

Configuration Templates

Configuration Templates are models and Configurations are the real instances you create using these models. BAS includes a variety of Configuration Templates that help you quickly and accurately create common Configurations. While some of the Configuration Templates provide good starting points upon which you can build your Configurations, others are samples of common Configurations which you can use as is or modify to better match your distributed environment. You access the Configuration Templates and create your Configuration using the Management Console.

To create a Configuration:

1. In the Management Console, on the tree structure locate the Hub for which you want to create a Configuration and expand the tree.

2 Right-click Configurations and choose Add Configuration.

The Template Gallery dialog for Configuration Templates appears. Use this dialog to locate the Template Category and choose a Configuration Template that best meets your needs.

Note: The Categories of Configuration Templates that appear by default are based on the Borland product selection you have installed. To view all categories and access all available Configuration Templates: click the Favorite Categories drop-down and choose All categories.

configuration.xml files

The Configuration information is stored in a `configuration.xml` file. Hubs are coded to look for the `configuration.xml` file name. For this reason, it is important that you **do not** rename any of your `configuration.xml` files.

Implementing Configurations

In addition to creating and modifying your Configurations, you use the Management Console to start, monitor, and stop them as well. When you run a Configuration, Agents implement all or a sub-set of the following *actions* on their Managed Objects. The basic actions that Agents are capable of implementing are:

- Start
- Stop
- Ping
- Kill

For each Managed Object within your Configurations, defaults are set for each basic action. However, using the Management Console, you can override these settings. For more information, go to the Managed Objects, Managed Object actions section.

Starting Configurations

Use the Management Console to start a Configuration. When you start a Configuration, the Hub communicates the start to each Agent. Each Agent in turn starts its Managed Resource(s) based on start order defined for each of its Managed Objects in the Configuration.

Running Configurations

As your Managed Resources run, each Agent regularly checks status by executing the Managed Object Ping action based on the properties defined for it in the Configuration. If a ping returns a state change, the Agent notifies the Hub.

If a ping determines that the Managed Resource has failed, by default the Agent independently invokes the BAS auto-recovery feature by immediately implementing the Managed Object start strategy defined in the Configuration. Then, when the Managed Resource restarts, the Agent communicates the state change to the Hub.

Stopping Configurations

Configurations run until you stop them. You use the Management Console to stop a Configuration. When you stop a Configuration, the Hub communicates the stop to the Agents. Each Agent in turn stops its Managed Resources(s) based on the stop order defined for its Managed Object(s) in the Configuration.

In the event an Agent is unable to stop a Managed Resource, you can use the Management Console to manually perform a kill action on a Configuration's Managed Objects.

Scheduling Configuration tasks

You can add one or more scheduled tasks to a Configuration to start processes that are independent of the MOs in the Configuration. A scheduled task can be started and aborted manually, or it can be started automatically based on the time rule(s) specified for the task. For more information about configuring a scheduled task for a Configuration, see Scheduling tasks for a Configuration.

Stopping the BAS infrastructure

It is important to note that if you stop a Hub or Slave Agent SCU process, or either of these unintentionally goes down, that the state of your Managed Resource(s) is unaffected. Stopping the SCU process on a Hub or a Slave Agent host, does not stop the resources being managed.

Important

While a Hub and/or Slave Agent SCU process is not running, the Managed Resources are not actively being managed.

For example:

1. You run Configuration A.
2. Per Configuration A, the Hub instructs Slave Agent 1 to start Managed Resource 1.
3. Slave Agent 1 starts and then pings Managed Resource 1 based on its Managed Object start and ping strategies.
4. While Managed Resource 1 is still running, the Slave Agent 1 SCU process goes down.
5. Managed Resource 1 continues to run - its state is unaffected by its Agent going down.
6. However, as long as Slave Agent 1 is down, it is unmanaged and auto-recovery no longer exists. For example, if Managed Resource 1 stops running while Slave Agent 1 is down:
   - The state change is not communicated to the Hub, and
   - the Managed Object start strategy is not implemented.

Stopping a Local Hub/Agent

By default, if you stop a Hub/Agent SCU process, the state of your Managed Resources is affected; before the SCU process shuts down, all Managed Resources are stopped. However, if the SCU process unintentionally goes down, the state of your Managed Resources is not affected.

Important

While the Hub/Agent is not running, the Managed Resources are not actively being managed.

The `agent.shutdown.policy` property controls this behavior for the Hub/Agent. This property is located in the Agent properties file, which by default is located in:
`<install_dir>\var\domains\<domain_name>\adm\properties`

The following lists and describes the acceptable values for this property:

Property	Values	Description
`agent.shutdown.policy`	■ `stop` - (applicable to Hub only) for a hub, stops all local configurations which have all the MOs on the same agent. ■ `stop-and-escalate` - (applicable to Hub only) for a hub, stops all local configurations. If they don't stop, escalates to kill. ■ `leave` (Default) - leave all configurations in their current state when exiting ■ `local-stop` - attempts to stop all managed objects on an agent when the agent shuts down regardless of configuration/group rules and without regard to group ordering ■ `local-stop-and-escalate` - for any scu, stop all MOs on this agent, regardless of group orders, etc.. If they don't stop, escalate to kill.	When a Local Hub/Agent SCU process is intentionally stopped, controls how the running Configuration's Managed Resources are affected.

Restarting a Hub/Agent

After you stop a Hub/Agent SCU process, when you restart it, by default any Configurations that were running when the SCU process stopped are automatically started. This behavior applies to a Hub/Agent for which the `agent.shutdown.policy` property is set to `stop-and-escalate` or to `stop`.

After you stop a Hub/Agent for which the `agent.shutdown.policy` property is set to `leave` (the Default value), when you restart it the Hub/Agent resumes management of the Configurations that were running when the shut down occurred. The Hub/Agent determines the state of each Managed Resource and based on the Configuration information, takes action where needed.

43

Managed Objects

You use a *Managed Object* to represent a Managed Resource, or a group of Managed Resources, that you want to configure, control, and/or simply monitor.

Creating and Adding Managed Objects

You use the Management Console to create and add new Managed Objects to your Configurations. When you create your Configurations, the Configuration Template you choose includes the Managed Objects pertinent to that Configuration.

If you want to add Managed Objects to an existing Configuration, the Management Console also provides a selection of Managed Object Templates for you to use. Each Managed Object Template contains predefined information specific to the type of Managed Object that Template creates and adds to your Configurations.

BAS includes the capability to add, remove and otherwise modify Managed Objects that you have added to a Configuration. However, after you first create your Configurations, we do not recommend modifying the Managed Object-to-Agent assignments.

For more information, go to the Managed Objects section.

Managed Object Templates

Managed Object Templates are models and the Managed Objects you create using these Templates are the real instances. BAS includes a built-in set of Managed Object Templates that you must use to quickly and accurately create the Managed Objects for your Configurations.

Unlike some of the Configuration Templates, none of the Managed Object Templates are samples; they are all models of each of the available types of Managed Objects available to you to use to represent your Managed Resources. You access the Managed Object Templates and create your Managed Objects using the Management Console.

To create and add a Managed Object to a Configuration:

1. In the Management Console, on the tree structure locate the Configuration for which you want to create and add a Managed Object and expand the tree.
2. Right-click the Configuration name and choose Add Managed Object.

The Template Gallery dialog for Managed Object Templates appears. Use this dialog to locate the Managed Object Category and choose a Managed Object Template that best meets your needs.

Note: To access all available Managed Object Templates: click the Favorite Categories drop-down and choose All categories.

Note

When creating an Apache Web Server managed object, do not include spaces in the Managed Object name. If spaces are included in the name, for example, "Apache Managed Object", the Apache Web Server will fail to start.

Note

When a UNIX text file is downloaded as part of a Managed Object Template, it must follow the UNIX platform rules for Ctrl-M. If the file is edited using a Windows text editor, you must be sure that it does not contain any extra Ctrl-M. Watch for this when including .sh files as part of the Custom Executable Managed Object Template.

Managed Object types

BAS uses the required Managed Object `type` attribute to determine how to handle controlling a Managed Object and what Managed Object specific information to look for within the Configuration. BAS hard-codes the type attribute value for each Managed Object in your Configuration - either when you add a Configuration that includes the Managed Object(s) or when you add a Managed Object to an existing Configuration.

Managed Object Type	Managed Object Sub-Type	Applies to...
Ordered Group Redundancy Group State Proxy		Pure Management Groups (does not represent a Managed Resource)
Process		Simple Process Managed Resources
Java Process	VBJ Process	Java Process Managed Resources
Custom JavaScript Custom Executable		Extensibility Managed Resources
osagent VisiNaming process Apache Web Server process OTS Tibco Partition Partition Service		BAS-specific Managed Resources

Note

To allow the Management Agent to manage an IIS server, using the Windows management console, you must set the IISAdmin and WWW services to: Startup Type = Manual. Also be sure all the Recovery actions are set to: Take No Action

Pure management Managed Object types

The Managed Object types you use for pure management of your Managed Resources within a Configuration are:

- Ordered Group
- Redundancy Group
- State Proxy

These types of Managed Objects are simply containers for other Managed Objects. You use these types of Managed Objects in your Configurations to represent groupings

of functionally dependent Managed Resources. Additionally, you can nest these Managed Object groups.

Ordered Group

You use an Ordered Group Managed Object to group a collection of Managed Objects within a Configuration for the following purposes:

- To define the order that BAS will use to start and stop the Managed Objects within the context of the group,
- To aggregate the status of the Managed Resources represented by the Managed Objects within the group.

For example, you configure BAS to manage a Web server that depends on access to a database in order to server its Web requests. In this case, you would want to be sure that the database was available before the Web server started.

1. Create and add an Ordered Group Managed Object to your Configuration.
2. Create and add the following to the Ordered Group:
   - A web server Managed Object.
   - A database Managed Object.
1. Set the start orders so that BAS starts the database first and the Web server next.

   By defining these start orders, BAS ensures that the Web server is not started until the database is up and running.

When you run this Configuration:

1. When the Configuration starts the Ordered Group Managed Object, the database is started.
2. Only when the database it is up and running, does BAS attempt to start the Web server.

And when you stop this Configuration:

- When the Configuration stops the Ordered Group Managed Object, the web server is stopped first and then the database.

While you can configure BAS to start and stop your Managed Resources sequentially, you can also assign the same start order and/or stop order to multiple Managed Objects and have BAS start and/or stop them concurrently.

Note

Each of your Configurations is itself an Ordered Group. A Configuration contains a collection of Managed Objects, for which you define the start and stop order within the context of the Configuration, and the status of the Managed Objects is aggregated and ultimately represented by the status of the Configuration.

Note

Refer to the `readme` files found in the `start_dependency` and `fail_dependency` folders under the `<install_dir>/var/templates/configurations/examples` directory for information about using the Op-Center Example Configurations to demonstrate how start and failure dependencies work for Ordered Group MOs.

For more information about creating a Configuration and adding Ordered Group Managed Objects, go to the Creating Configurations and Creating and Adding Managed Objects sections.

Redundancy Group

You use a Redundancy Group Managed Object to group a collection of homogeneous Managed Resources. The Redundancy Managed Object type enables you to model failover across the members of the group.

By default, when you run the Configuration and the Redundancy Group is started, all the members are started. However, using the Management Console, you can customize your Redundancy Group in the following ways:

- To specify the number of members that you require BAS to start and run.
- To specify the desired minimum number of members you want BAS to start and run.
- To specify the desired maximum number of members you want BAS to start and run.

Each of these values can be a subset of the total number of members in the group or all the members (the default).

When you stop a Configuration, in order for a Redundancy Group Managed Object to stop, all running members are stopped at the same time.

For example, if you want BAS to manage 3 Web servers which are configured to perform the same functions and be used for failover, you can:

1. Create and add a Redundancy Group Managed Object to your Configuration.
2. Create and add the following members to the group:
   - Web server A Managed Object
   - Web server B Managed Object
   - Web server C Managed Object
1. Define the following for the Redundancy Group:
   - 1 required.
   - a desired minimum of 2.
   - a desired maximum of 2.

When you run the Configuration:

1. When the Configuration starts the Redundancy Group Managed Object, BAS attempts to start and run Web server A and Web server B at the same time.

Note

BAS bases the member start priority on the order in which the members are listed (top to bottom) in the Configuration.

2. In this example, as soon as 1 of these Web servers is running, the required value is met for the group and the running status is aggregated to the Redundancy Group Managed Object.
3. However, BAS will continue to start members until the desired minimum of 2 running is achieved. For example, if Web server B starts, but Web server A does not, BAS will try to start Web server C.
4. As long as the Configuration is running, and members are available, BAS will continue to maintain the desired minimum of 2 members running.

And when you stop this Configuration:

- Before the Redundancy Group Managed Object can stop, all running Web server members must be stopped. BAS will attempt to stop all that are running at the same time.

Note

Refer to the `readme` file found in the `<install_dir>/var/templates/configurations/examples/fault_tolerance/datadir` directory for information about using the Fault Tolerance example in the Op-Center Example Configurations to demonstrate how a Redundancy Group can be used for failover and fault tolerance.

For more information about creating a Configuration and adding Redundancy Group Managed Objects, go to the Creating Configurations and Creating and Adding Managed Objects sections.

State Proxy

You use a State Proxy Managed Object to represent the state of another MO that is defined elsewhere in a Configuration. A State Proxy:

- Can be used anywhere within the same Configuration as the original MO for which it is a proxy.
- Will affect the state and behavior of the group to which it belongs, just like any other MO.

A State Proxy reflects the current "run" state of the MO it represents (for example, Running or Stopped), but not what it is doing (for example, starting or stopping). A State Proxy does not control the lifecycle of the MO it represents.

The State Proxy is useful in cases where different Ordered Groups share dependency on the same MO. When the proxied MO goes down, its State Proxy likeness also goes down. Each of the Ordered Group MOs can be configured to appropriately respond to this state change.

For example, if you create separate Ordered Group MOs but want both to depend on the same instance of a Naming Service MO, you can:

1. Create the first Ordered Group, GroupA.
2. Add a Naming Service MO to GroupA.
3. Create the second Ordered Group, GroupB.
4. Add a State Proxy MO to Group B.
5. Set the logical name in the State Proxy MO to that of the Naming Service MO created in GroupA, using the syntax `${agent.name}/mo-name`, for example, `${agent1.name}/nsdb_jds`

   When the State Proxy MO in Group B starts, it pings the Naming Service MO in Group A. If the Naming Service MO is stopped, the State Proxy MO will also stop and the states of both groups will reflect this state change. If Group B is configured to shut down when a child fails, it will initiate the shutdown sequence. The behavior of Group B has no affect on the Naming Service MO itself or on Group A.

Note

Refer to the `readme` found in the `<install_dir>/var/templates/configurations/examples/state_proxy/datadir` directory for information about using the example State Proxy MO included with the Op-Center Example Configurations to demonstrate State Proxy behavior.

Process Managed Object type

Process Managed Objects are directly associated with Managed Resources. You use a Process Managed Object to represent a Managed Resource that is a single process (including a VisiBroker for Java process)- that is started by issuing a command to the Operating System and is pinged and stopped using the Operating System.

For more information about creating a Configuration and adding Process Group Managed Objects, go to the Creating Configurations and Creating and Adding Managed Objects sections.

UNIX: Managing Ownership of a Process Managed Object

Processes which access privileged ports on UNIX hosts must have appropriate permissions; the process must be started with the permissions of user `root`. Typically, only some processes in a particular configuration need to be started with `root` permissions. The `setuser` script helps configure BAS to allow Process Managed Objects to start as `root` or with `root` permissions. See “Using the setuser tool to manage ownership” in the installation guide for information about using the `setuser` tool and multi-user mode.

Java Process Managed Object type

Java Process Managed Objects are directly associated with Managed Resources. You use a Java Process Managed Object to represent a Managed Resource that is a single java process - that is started by issuing a command to the Operating System and is pinged and stopped using the Operating System.

For more information about creating a Configuration and adding Java Process or VBJ Process Managed Objects, go to the Creating Configurations and Creating and Adding Managed Objects sections.

Extensibility Managed Object types

These Managed Object types are directly associated with Managed Resources. They give you the flexibility to extend BAS by providing you the capability to custom code your Managed Object action strategies. As such, these types enable you to manage a wide variety of resources.

Custom JavaScript Managed Object type

You use the Custom JavaScript Managed Object type to represent a Managed Resource for which you want to implement the start, ping, stop and kill action strategies in JavaScript. BAS uses a JavaScript interpreter to execute theses strategies.

Important

(UNIX only) JavaScript execution is typically not allowed when a Managed Object is started by an Agent configured for Multi-user Mode (MUM). If any of the Managed Objects that will run in Multi-user Mode (MUM) contain JavaScripts, you need to modify the `agent.config` to allow those JavaScripts to run.

Custom Executable Managed Object type

You use the Custom Executable Managed Object type to represent a Managed Resource for which you want to execute a process for each of the action strategies.

AppServer-specific Managed Object types

These Managed Object types are directly associated with AppServer Managed Resources in Configurations.

For more information about adding a AppServer Configuration and AppServer Managed Objects, go to the Creating Configurations and Creating and Adding Managed Objects sections.

AppServer Managed Object Type	Use to represent a single instance of...
osagent	the VisiBroker Smart Agent - a simple directory service used to facilitate communication and RPCs between CORBA objects.
Naming Service process	the VisiBroker Naming Service - a distributed directory service.
Apache web server process	an Apache web server.
OTS	VisiTransact - a 2-phase commit transaction service.
Tibco	the Tibco Java Messaging Service (JMS).

AppServer Managed Object Type	Use to represent a single instance of...
Partition	the AppServer Partition - the basic unit of object hosting which provides services for J2EE and CORBA components.
Partition Service	any of the following AppServer Partition services: ■ EJB container ■ JDataStore database ■ Tomcat web container ■ Java Session Service ■ VisiBroker Interceptor ■ VisiConnect ■ OpenJMS

Managed Object actions

For each Managed Object that controls the lifecycle of a Managed Resource, an Agent can perform all of the following basic actions:

- Start
- Stop
- Ping
- Kill

For each Managed Object that is set up to just perform status determination for a Managed Resource (pingers), the Agent will execute only the Ping action.

Note

For UNIX only: The Management Agent will be unable to stop a JDataStore server as a Managed Object process unless it is started using the -ui none option. For an example, see the Pet Store cluster.

Strategies

How an Agent executes a Managed Object action is determined by a "strategy". The selection of strategies available for each action are determined by the Managed Object `type`. For each Managed Object in your Configurations, a default strategy is set for each action. However, you can override any of the defaults by specifying another available strategy using the Management Console or by directly editing the `configuration.xml`. Within the `control-overrides` element, you specify all action parameter and action strategy default overrides.

Scheduling tasks for a Configuration

You can add one or more scheduled tasks to a Configuration to start processes that are independent of the MOs in the Configuration. For example, you can add a scheduled task that starts a backup process during non-peak processing times at the start and end of each business week.

Op-Center starts scheduled tasks at the configured time, even if the Configuration is not currently running. A scheduled task does not impact the state of the Configuration, nor does the state of the Configuration impact the scheduled task's operation. The logical depiction of a Configuration with one or more scheduled tasks displays the task as independent of the rest of the Configuration's structure.

Each scheduled task:

- Has a name.
- Has an Agent on which it should run.
- Is configured to run a process.

– Can have one or more rules associated with the task's schedule.

Note

A scheduled task can only be configured to run the process it is associated with: it cannot be configured to stop the process.

To configure a Scheduled Task:

1. Right-click on the Configuration, then select Add Scheduled Task.
2. Enter the Object information, then click Next.
3. Enter the Process Settings for the process you want to run, and configure any Advanced Process Settings before clicking Next.
4. To add one or more rules to the Task Schedule, click New.
5. To configure Advanced Options, click Next.
6. When you have completed making entries on the Task Schedule or Advanced Options tabs, click Finish.

Creating an availability schedule for a Managed Object

Note

Availability schedules cannot be configured for AppServer Partition Services. You can, however, configure scheduled tasks for a Configuration that contains AppServer Partition Services (for a list of these Services, see the Partition Service section.) For information about scheduling tasks for a Configuration, see the Scheduling tasks for a Configuration section.

Each Managed Object can have an optional set of time-dependent rules that dictate when an MO's state should change to Running or Stopped. These rules are expressed with `cron`-style time specifications to indicate the start time when the rule applies and, in the case of a rule for an MO, a duration to indicate how long after the start time the rule continues to apply. If multiple rules are configured for the same time slot, the last rule configured for that time slot is used.

You can add rules to control an MO's running state through the Availability Schedule tab, which is available in the MO's Properties editor. The Configuration must be running in order for the rules in an MO's Availability Schedule to be implemented. Unlike rules entered for a scheduled task, a rule configured for a specific MO can specify the MO's state to be either Stopped or Running, and can have a configurable duration that controls how long the configured state is in effect. For example, you can configure a rule that stops a specific MO during a regular backup or maintenance cycle, or during holiday shutdowns, without shutting down the entire Configuration. Once the configured duration is reached, the MO reverts to its default state.

If an MO's Configuration dictates that the default state for the MO is Running but a rule configured for the MO dictates that the MO's state should change to Stopped, the rule's state change overrules the MO's default state and the MO is stopped as configured in the rule. However, if an MO's Configuration dictates that the default state for the MO is Stopped and a rule configured for the MO dictates that the MO's state should change to Running, the Configuration's default state for the MO overrules the Running state configured for the rule and the MO remains in the Stopped state.

Each rule in an availability schedule either has a description that is generated from the configured rules (for example, "Stop for 1 minute, at midnight, every day"), or a custom name that you specify (for example, "Holiday Shutdown"). The rules are defined with standard `cron`-style syntax, which is configured by editing the MO's Properties using the entry mechanisms found in the availability schedule's Time Rule page, or by specifying the actual `cron` syntax where indicated in the Time Rule page. The rules are applied in the order in which they were entered, however, you can rearrange the order anytime after configuring the rules.

Note

Configuring a Running rule on a root managed object will not start its Configuration. The Configuration must be started by an administrator.

44

Getting started with the Borland Management Shell

The Borland Management Shell (BMSH) is a scripting environment you can use to write scripts to configure and run AppServer both locally and remotely. BMSH lets you write scripts that automate operations that would otherwise involve manual, error-prone tasks that require the use of a combination of the Borland Management Console, a text editor, and other Borland AppServer tools.

BMSH scripts are text files that you can write using any text editor and require no compiling or preprocessing. BMSH is built on the open-source Mozilla Rhino project. The keywords, statements, expressions, and operators used by BMSH are identical to the JavaScript used on web pages. BMSH extends the built-in standard Mozilla Rhino commands for the purpose of scripting a AppServer installation and it includes objects that provide a scriptable API to Borland AppServer. For more information on JavaScript, go to Mozilla Rhino.

Using BMSH

Some example uses of BMSH include:

- Performing general file system operations.
- Starting and stopping a AppServer.
- Managing and configuring Partitions located in an AppServer instance.
- Configuring and running AppServer services.

Running BMSH

There are two ways to execute a BMSH script:

- Interactively by typing commands at the `bmsh>` prompt.
- Executing a script file from the command prompt (Windows) or the shell prompt (UNIX). If your path has not been updated, you will need to navigate to the `bin` directory of your installation.

Note

BMSH scripts run on all platforms.

Using BMSH interactively (interactive mode)

To start BMSH in interactive mode, at the command prompt, type:

```
bmsh
```

The `bmsh>` prompt appears. (BMSH assumes that the path to your AppServer installation is included in the PATH environment variable.)

To exit the interactive shell, at `bmsh>`, type:

```
quit()
```

Important

The closing parentheses are required.

Interactive mode examples:

To display the Management Port for a local installation, type:

```
bmsh>bes.getManagementPort();
```

The Management Port is returned:

For example:

```
42424
```

To set the Management Port for the local installation, type:

```
bmsh>bes.setManagementPort(portnumber);
```

To display the Management Console Management Port used for discovery, type:

```
bmsh>bes.getConsolePort();
```

To set the Management Console Management Port, type:

```
bmsh>bes.setConsolePort("portnumber");
```

In the shell you can execute any BMSH API by typing it at the prompt. Any variable you use will remain defined until you exit the shell session. This method is a convenient way to interactively test script snippets or perform simple one or two line operations. The interactive behavior of the BMSH is unchanged from the Mozilla-Rhino release. For more complex operations that you want to perform repeatedly, you will want to write script files as described in the following sections.

Using BMSH script files (batch mode)

Use the BMSH batch mode to execute script files. At the command prompt (Windows) or the shell prompt (UNIX), specify a file that contains JavaScript commands and its arguments. The example shown below is located in your Borland AppServer installation at `<install_dir>/examples/bmsh/`.

Navigate to the BMSH example directory (located in `<install_dir>/examples/bmsh`), then type:

```
bmsh setManagementPort.js ?
```

This script above displays the usage for the `setManagementPort.js` script.

Executing the following script sets all the management ports for the local installation of AppServer. For example, at the command prompt or shell prompt:

```
bmsh setManagementPort.js -p 33333
```

BMSH files and folders

This section describes the directories that contain the script files associated with the BMSH component in your AppServer installation. Users can place their own BMSH scripts in any directory location.

BMSH scripts supplied with the product are installed in the following locations:

```
<install_dir>/bin/bscript/autoload/
<install_dir>/bin/bscript/scripts/
<install_dir>/examples/bmsh/autoload/
<install_dir>/examples/bmsh/scripts
```

BMSH files and folders in the bin directory

The `<install_dir>/bin` directory contains files and folders that are globally applicable to an AppServer installation. The BMSH scripts here should be treated the same as other binary files included with the product and should not be altered by the user. Files and folders in this directory may depend on assumptions regarding installation default names of directories and the location and content of files. Files and folders in this directory, however, should not contain information that is specific to a particular BMS configuration or any user defined variables.

BMSH files and folders in the /bin/bscript/autoload subdirectory

The `<install_dir>/bin/bms/autoload/` subdirectory contains the scripts that are loaded when BMSH starts. For more information on the BMSH autoload feature, see the Autoload feature.

Files in /bin/bscript/scripts

The `<install_dir>/bin/bscript/scripts/` subdirectory contains scripts written and supported by Borland. The scripts placed in this directory are used as command line tools for frequently executed and common BMSH function such as enabling and disabling BMS debugging. They are set with a read-only attribute to help safeguard against accidental modification or deletion. BMSH includes a script search path feature that looks in this directory location to find a script when it is entered on the command line.

Note

The scripts in this location serve as examples of using BMSH. You can copy and tailor them for custom purposes, but Borland strongly recommends that you do not save your own customized scripts in this location.

BMSH subdirectories in the examples directory

The `<install_dir>/examples/bmsh` subdirectory contains sample BMSH scripts and supporting documentation. You can use these scripts as templates when creating your own customized scripts.

Files in /examples/bmsh

The `<install_dir>/examples/bmsh` subdirectory contains scripts and folders that may be globally applicable to an AppServer installation, to a particular BMS configuration, or a to specific AppServer application. The primary intent of the files in this subdirectory is for illustrative purposes only and Borland does not guaranty their functionality and operability.

Files in /examples/bmsh/scripts

The <install_dir>/examples/bmsh/scripts subdirectory contains additional sample scripts that are expected to be used less frequently than those in the <install_dir>/examples/bmsh/scripts subdirectory. The scripts in this location typically have dependencies on a particular BMS configuration or application and may not be as fully developed or well documented. Borland does not guaranty their functionality and operability.

Files in /examples/bmsh/autoload

The <install_dir>/examples/bmsh/autoload subdirectory contains example scripts that demonstrate how to use the BMSH autoload feature.

BMSH features

BMSH includes several features that extend the Mozilla-Rhino environment.

Autoload feature

This BMSH extension mechanism can be used to define global BMSH variables, objects and functions. You can use this mechanism to create a BMSH environment that is highly customized to your installation and deployment preferences. When the BMSH starts it executes the scripts in an autoload folder. The default autoload folder is <install_dir>/bin/bscript/autoload. BMSH runs the scripts in the autoload folder before executing interactive commands and before executing command-line script files. Note that the BMSH startup time can be affected because it must load and run the scripts.

After installation, you can the built-in autoload file and two examples. The built-in file is called mgmt_utils.js and contains simple wrappers for management domain functions for setting the port. Note that some of the command-line scripts that are installed assume that this file as been autoloaded, using a leading underscore prefix for autoloaded variables and functions.

There are two examples in the BMSH examples folder. Simply copy them to the autoload folder to try them out. They are:

- javaarray.js — A utility routine that simplifies the creation of Java arrays in JavaScript.
- sendmail.js — An example of using LiveConnect. This example uses only javax.mail interfaces, but no JavaScript objects. This script requires use of the javaarray.js file.

Note

The AppServer object and Hub class are extended using JavaScripts extension mechanisms—bes_ext.js and hub_ext.js. These extensions are located in:

```
<install_dir>/bin/bscript/autoload
```

These extensions are not formally documented, however you can view these API extensions by opening these files.

Search path feature

BMSH looks in a specific path location to find a script when it is entered on the command line in addition to the current directory. By default, the BMSH search path is set to:

```
<install_dir>/bin/bscript/scripts/
```

BMSH classpath add JARs feature

You can add JARs to the BMSH classpath by editing the `bmsh.config` file in your installation's `bin` directory. Find the section headed

```
# Set up the list of jars in the product system directory
```

and add your paths to the JARs. Note that your paths must be absolute, but you can use the `$var(installRoot)` variable for portability if your paths are relative to the product installation.

BMSH Objects

The Rhino JavaScripting shell has objects defined by the ECMA JavaScript standard (such as `Date` and `RegExp`). BMSH adds objects specific to AppServer. BMSH objects fall into the following categories:

- Pre-instantiated objects — such as those that contain methods used for file system interactions or to extract and set values in property files.
- User-instantiated objects — such as those that contain methods used to interact with the Borland Management Hub, substitute property values in place holders, or interact with HTTP web pages.
- Static objects — such as those that contain methods to start, stop and ping processes using the functionality of the BMS agent, or methods that parse and save XML files.
- Factory objects — objects that are created by other objects (usually the pre-instantiated objects) such as those that contain methods to extract and set values in property files or to read and write text files.

BMSH pre-instantiated objects

Pre-instantiated objects do not require you to perform a "`new`". Each time BMSH is started, the pre-instantiated objects are created. You can begin using them without instantiation. The detailed API documentation can be found in the API documentation. BMSH has two pre-instantiated objects:

- `bes` — Provides an API to the local AppServer installation where BMSH is executed. This object was demonstrated in the interactive example above.
- `fso` — Provides general purpose functions for accessing the local file system. The `fso` has some AppServer specific methods for getting and setting properties of Borland AppServer local property files.

BMSH user-instantiated objects

BMSH extends the built-in Mozilla Rhino user-instantiated objects (`Date`, `RegExp`). BMSH user-instantiated objects require the script author to create instances using "`new`". For this release there is only one BMSH user-instantiated object: `XDOM`, a scriptable API to access xml files.

BMSH factory objects

BMSH includes objects that are created by other objects (factory objects). The most common is the `TextStream` object that reads and writes to a file. The `fso` object acts as a factory for these objects. You cannot use "`new`" to create a factory object.

Example

```
var inputStream = fso.openTextStream("a_filename",1) // 1 = for reading
```

The `inputStream` is a BMSH `TextStream` object.

Java LiveConnect

The Rhino JavaScript interpreter has a LiveConnect mechanism that will compile Java on the fly. This mechanism lets you include Java code in a script file. This feature is ideal for simple functionality that BMSH does not provide. For example, you could use the Java runtime object to execute another process or to sleep for a prescribed amount of time. It also could be useful to provide a simple user dialog, such as the `joptionpane.js` dialog box. Use of this mechanism to access functionality in AppServer JAR files, however, is not supported.

See the example: `<install_dir>/examples/bmsh/autoload/sendmail.js` which demonstrates how to include Java code in a script file.

Writing your own Rhino scriptable objects

BMSH objects are all derived from Rhino scriptable objects. This is the standard Rhino mechanism for adding objects to the shell. Because BMSH is based on Rhino, AppServer customers also have this mechanism available to them to add any objects to BMSH that they choose. This requires the use of a Java compiler and an advanced understanding of Java.

To extend BMSH with additional objects, you must edit the `bmsh.config` file:

1. Open the `bms.config` file located in the `<install_dir>/bin` directory.
2. Locate the section that starts with the comment: `# Set up the list of jars in the product system directory`
3. Add a line similar to the following settings that points to the fully qualified path location of the jar file containing the objects you want to add to the shell:

```
addpath $var(installRoot)/lib/dom4j.jar
addpath $var(installRoot)/lib/scu.jar
addpath $var(installRoot)/lib/scu_client.jar
...
```

   For example:

```
addpath $var(installRoot)/lib/my.jar
```

   or

```
addpath c:/myfolder/my.jar
```

Environment variables

Environment variables are available in BMSH using the Rhino environment JavaScript array. For example:

```
for(i in environment)
  print(i +"=" + environment[i])
```

Cut and paste this script fragment into the BMSH shell. Then run it at the `bmsh` prompt to display all the BMSH environment variables.

JavaScript arrays

JavaScript arrays are implemented as associative arrays (hash tables). The index is a string or a number. For example, to extract a single environment variable value, run the following script fragment at the `bmsh` prompt:

```
environment["user.dir"]
```

The example above displays the current working directory. It is equivalent to:

```
print( environment["user.dir"] )
```

Alternatively you can assign the value to a variable:

```
var cwd = environment["user.dir"];
print( cwd );
```

Java string arrays

Java string arrays are handled differently than JavaScript arrays . For example:

```
var files = fso.findFiles(environment["user.dir"],"*",false);
for(var i=0; i < files.length; i++ )
   print("files["+ i + "]=" + files[i])
```

The `findFiles()` routine returns a Java `String[]`. String arrays can be accessed only by number. They are not hashed.

JavaScript built-in functions

The JavaScript regular expression processor is available in BMSH. The `Date()` object is also available.

BMSH interactive behavior

The BMSH interactive behavior follows the same rules used by Mozilla-Rhino operations:

- When you type a line at the `bmsh` prompt, the script executes when you press the *Enter* key.
- If the last character of a line in a script is an open curly-bracket (`{`), it will not execute until you enter a closing curly-bracket (`}`). (This rule follows the C language convention for curly-brackets.)
- A semicolon (`;`) is optional when a script statement is followed by a new line.

BMSH Help

The information in the Borland Management Shell API Reference is also available from the BMSH shell using the global `help()` function. The following table describes the `help()` command options.

Option	Description
`help();`	Displays: Rhino help. This is unchanged from the Mozilla-Rhino release. It displays the intrinsic global functions of the Rhino interpreter.
`help(`*`jsObject`* `|` *`jsClass`* `);`	`jsClass` displays the functions and properties of a BMSH JavaScript class name. Only BMSH classes are valid. For example: `help(Dom4j)` `jsObject` displays: the functions and properties of a BMSH JavaScript object. Only BMSH objects are valid. For example: `bmsh>help(appserver);`

Option	Description
help(jsObject | *jsClass*, "String *function*" | "*property*");	jsObject | jsClass JavaScript is a BMSH JavaScript object or JavaScript class. function | property displays the function parameter names and types, and the return type of a property in a jsObject. This parameter must be quoted.
Help(*jsObject* | *jsClass*, String *function* | *property*, boolean *verbose*);	jsObject | jsClass - is a BMSH JavaScript object instance or JavaScript class. function | property - is the name of a function or property in a jsObject. This parameter must be quoted. For example: help(appserver,"installRoot"); verbose - displays detailed descriptions including the parameters and return values of the function or property. This information is the same as that contained in the API documentation. For example: bmsh>help(appserver,"getManagementPort",true);

45

Configurations and configuration.xml files

All your Configuration information is stored in a `configuration.xml` file. For each Configuration you create, a single `configuration.xml` is created and stored on the Hub host. Because the Hub is coded to look for the `configuration.xml` file, it is important that you do not rename this file.

configuration element

Within each `configuration.xml` file, all your Configuration information is contained within the `configuration` element.

```
<configuration ...attributes...>
        <configuration-id ...attributes.../>
        <main-root ...attributes.../>
        <managed-objects ...attributes...>
                ...
        </managed-objects>
        <properties>
                ...
        </properties>

</configuration>
```

configuration element attributes

The following are the `configuration` element attributes:

Attribute	Required	Default	Description
`revision`	Yes	n/a	The `configuration.xml` version.
`display-name`	Yes	n/a	A unique name you give the Configuration when you create it.

Attribute	Required	Default	Description
description	No	n/a	A brief description of your Configuration. This description appears in some views within the Management Console.
small-icon	No	n/a	For internal use.

```
<configuration revision="revisionnumber"
 description="string" display-name="string">
       <configuration-id ...attributes.../>
       <main-root ...attributes.../>
       <managed-objects ...attributes...>
              ...
       </managed-objects>
       <properties>
              ...
       </properties>

</configuration>
```

configuration element sub-elements

The following are the `configuration` element sub-elements:

Element	Description
configuration-id	Contains the Configuration name and version information. For more information, go to the configuration-id element section.
main-root	Contains a reference to the Managed Object through which Configuration start and stop actions are implemented. Additionally, the Configuration state is represented by the Managed Object identified as the `main-root`. For more information, go to the main-root sub-element section.
managed-objects	Contains the Managed Objects you define for the configuration and configuration information specific to each Managed Object. For more information, go to the Managed Object elements and attributes section.
properties	Contains the Configuration-wide properties you define. For more information, go to the Working with Configuration properties section.

```
<configuration ...attributes...>
       <configuration-id ...attributes.../>
       <main-root ...attributes.../>
       <managed-objects ...attributes...>
             ...
       </managed-objects>
       <properties>
             ...
       </properties>

</configuration>
```

configuration-id element

In each `configuration.xml`, the following configuration identification information is stored within the `configuration-id` element:

Attribute	Required	Default	Description
`name`	Yes	n/a	The name of your Configuration. This name is used to identify a Configuration throughout the Management Console displays. Following your initial creation of a Configuration, this name cannot be changed.
`version`	n/a	1	The version of your Configuration. Each time you save any modifications to a Configuration, BAS increases this value by 1.

```
<configuration ...attributes...>
      <configuration-id name="string" version="integer"/>
      <main-root ...attributes.../>
      <managed-objects ...attributes...>
            ...
      </managed-objects>
      <properties>
            ...
      </properties>

</configuration>
```

main-root sub-element

The `main-root` sub-element identifies the one Managed Object that is the main root of a Configuration.

The main root Managed Object:

- Starts when you start its Configuration.
- Stops when you stop its Configuration.
- Represents the state of its Configuration.

By default, when you create a Configuration using a template, a Configuration Ordered Group Managed Object is created. By default, all other Managed Objects within the Configuration are children of this Configuration Managed Object and it is the main root.

The `main-root` element has the following attribute.

Attribute	Required	Default	Description
`mo-ref`	Yes	`${HUB.name}/${CONFIG.name}` - The Configuration Ordered Group Managed Object. The logical name of the Configuration Managed Object which consists of the name of the Hub to which the Configuration belongs and the name of the Configuration following the replacement string syntax: `${HUB.name}/${CONFIG.name}`.	The logical name of the Managed Object that is the one main root within the Configuration using the syntax: `<Configuration_hub_name>/<managed_object_name>` **Note:** For the purpose of porting a Configuration to another Hub, you can use the replacement string `${HUB.name}` for the Hub.

```
<configuration ...attributes...>
       <configuration-id ...attributes.../>
       <main-root mo-ref="managed_object_logical_name"/>
       <managed-objects ...attributes...>
             ...
       </managed-objects>
       <properties>
             ...
       </properties>
</configuration>
```

46

Managed Object elements and attributes

This section lists and describes `configuration.xml` `managed-objects` element, its attributes and sub-elements. Additionally, all the elements, sub-elements, and attributes available for each Configuration Managed Object type are described.

Generic XML definition

Managed Objects are listed within the `managed-objects` element of the `configuration.xml` file. Each individual Managed Object is defined within the `managed-object` sub-element.

```
<configuration>
    ...
    <managed-objects>
        <managed-object />
        ...
    </managed-objects>
    ...
</configuration>
```

managed-object element attributes

For each Managed Object you define within your Configurations, all or a subset of the following attributes are available. These attributes appear within the `managed-object` element of the `configuration.xml`.

Attribute	Required	Default	Description
`agent`	Yes	The Local Agent which resides on the same host as the Hub for which the Configuration is created. By default, this Hub/Local Agent name is represented by the string variable `${HUB.name}`.	Name of the Agent to which the Managed Object is assigned. Possible values are: ■ `${HUB.name}` (default) ■ *`hub_name`* (by default, the name of the host on which the Hub is installed) ■ *`agent_name`* (by default, the name of the host on which the Slave Agent is installed)
`name`	Yes	Generic Managed Resource name or null.	The logical name you give a Managed Object. After you initially create and add a Managed Object to a Configuration, you cannot change the value of this attribute. However, you can change the name displayed throughout the Management Console at any time by using the `display-name` attribute.
`display-name`	No	The Managed Object `name` attribute value represented by the string variable `${MO.name}`.	A name that you want to appear in the Management Console displays in place of the Managed Object `name` attribute value.
`initial-manage`	No	`true`	For Managed Resources you want AppServer to actively manage, set to `true` (default). If set to `true`, the Managed Object is "managed"; when you run the Configuration, AppServer will perform actions upon the Managed Resource in order to maintain its state. If you want AppServer to only monitor the object for state changes, but implement no actions, set to `false`.
`initial-monitor`	No	`true`	For Managed Resources you want AppServer to monitor state during pings from its Agent, set to `true` (default). For Managed Resources that you do not want AppServer to monitor state during pings from its Agents, set to `false`.
`type`	Yes	Determined by the template you use to create a Configuration or create and add a Managed Object to a Configuration.	AppServer uses this attribute to determine how to handle controlling the Managed Object, and what `type`-specific information to look for within the Configuration. For more information, go to the Managed Object type attribute section.

Attribute	Required	Default	Description
initial-desired-state	No	Stopped	When the Managed Object is first added to a Configuration, by default the Managed Object state is Stopped. If you want the Managed Object running when it is first created and added, set to Running.
data-directory	No	${CONFIG.path}/mos/${MO.name} Where the replacement string variables ${CONFIG.path} represents the fully-qualified path to the configuration.xml directory, and ${MO.name} represents the Managed Object name value.	Internal use.
version	No	n/a	Used to specify the Managed Resource version.
vendor	No	n/a	Used to specify the Managed Resource vendor.
description	No	n/a	Used to describe the Managed Resource.
deploy-data	No	false	Internal use.

```
<configuration>
    ...
    <managed-objects>
          <managed-object agent="{$HUB.name}|hub_name|agent_name"
 name="managedobject_name"
 [display-name="${MO.name}|display_name"
 initial-desired-state="stopped|running"
 data-directory="${CONFIG.path}/mos/${MO.name}|fullyqualifiedpath"
 deploy-data="true|false" initial-manage="true|false"
 initial_monitor="true|false"]
 type="ordered-group|redundancy-group|state-proxy|process|java-process|
userdefined-jscript|userdefined-executable|visientity|visiserver|osagent|
apache-process|ots|tibco|partition|">
 [vendor="vendorname" version="version"
 description="description"]
              ...
          </managed-object>
          ...
    </managed-objects>
    ...
</configuration>
```

Managed Object type attribute

The following describes the acceptable values for the `type` attribute.

Value	Description
`ordered-group`	Use to define a set of related Managed Objects that can be managed as one. This allows you to build object hierarchies. Note that the `ordered-group` itself is just a placeholder - a logical construction. The actual Managed Objects that are controlled by AppServer are the *members* of the group. For general information about the Ordered Group Managed Object, go to the *Borland AppServer Developer's Guide*, Ordered group section. For more `ordered-group` Managed Object `type` information, go to the ordered-group Managed Object type section.
`redundancy-group`	Similar to an `ordered-group` in that it represents a collection of Managed Objects. The Managed Objects in a `redundancy-group` are functionally equivalent to one another. For general information about the Redundancy Group Managed Object, go to the *Borland AppServer Developer's Guide*, Redundancy group section. For more `redundancy-group` Managed Object `type` information, go to the redundancy-group Managed Object type section.
`state-proxy`	Use to represent the state of another MO that is defined elsewhere in a configuration, without impacting the lifecycle of the proxied MO. For general information about the State Proxy Managed Object, go to the *Borland AppServer Developer's Guide*, State Proxy section. For more `state-proxy` Managed Object `type` information, go to the state-proxy Managed Object type section.
`process`	Use to represent a Managed Resource that is a single process - that is started by issuing a command to the Operating System and is pinged and stopped using the Operating System. For general information about the Process Managed Object, go to the *Borland AppServer Developer's Guide*, Process Managed Object type section. For more `process` Managed Object `type` information, go to the process Managed Object type section.
`java-process`	Use to represent a Managed Resource that is a single Java process - that is started by issuing a command to the Operating System and is pinged and stopped using the Operating System. For general information about the Java Process Managed Object, go to the *Borland AppServer Developer's Guide*, Java Process Managed Object type section. For more `java-process` Managed Object `type` information, go to the java-process Managed Object type section.
`custom-javascript`	Use to represent a single instance of a Managed Resource for which you want to implement the start, ping, stop and kill actions in JavaScript. For general information about the Custom JavaScript Managed Object, go to the *Borland AppServer Developer's Guide*, Extensibility Managed Object types section. For more `custom-javascript` Managed Object `type` information, go to the custom-javascript Managed Object type section.
`custom-executable`	Use to represent a single instance of a Managed Resource for which you want to execute a process for each of the start, ping, stop and kill actions. For general information about the Custom Executable Managed Object, go to the *Borland AppServer Developer's Guide*, Extensibility Managed Object types section. For more `custom-executable` Managed Object `type` information, go to the custom-executable Managed Object type section.
`osagent`	Use to represent a single instance of the VisiBroker Smart Agent - a simple directory service used to facilitate communication and RPCs between CORBA objects. For general information about the osagent Managed Object, go to the *Borland AppServer Developer's Guide*, AppServer-specific Managed Object types section. For more `osagent` Managed Object `type` information, go to the osagent Managed Object type section.
`apache-process`	Use to represent a single instance of an Apache web server. For general information about the Apache web server Managed Object, go to the *Borland AppServer Developer's Guide*, AppServer-specific Managed Object types section. For more `apache-process` Managed Object `type` information, go to the apache-process Managed Object type section.

Value	Description
ots	Use to represent a single instance of VisiTransact - a 2-phase commit transaction service. For general information about the OTS Managed Object, go to the *Borland AppServer Developer's Guide*, AppServer-specific Managed Object types section. For more ots Managed Object type information, go to the ots Managed Object type section.
tibco	Use to represent a single instance of the Tibco Java Messaging Service (JMS). For general information about the Tibco JMS Managed Object, go to the *Borland AppServer Developer's Guide*, AppServer-specific Managed Object types section. For more tibco Managed Object type information, go to the tibco Managed Object type section.
partition	Use to represent a single instance of a AppServer Partition - the basic unit of object hosting in AppServer, which provides services for J2EE and CORBA components, among other features. For general information about the Partition Managed Object, go to the *Borland AppServer Developer's Guide*, AppServer-specific Managed Object types section. For more partition Managed Object type information, go to the partition Managed Object type section.

ordered-group Managed Object type

This section lists and describes the ordered-group Managed Object type sub-elements and sub-element attributes.

Note

For information about the ordered-group Managed Object attributes, go to the managed-object element attributes section.

ordered-group sub-elements

The following describes the ordered-group sub-elements:

ordered-group sub-element

Sub-Element	Description
ordered-group	Contains the ordered-group Managed Object type-specific Configuration information.

```
<configuration>
    ...
   <managed-objects>
          <managed-object ...type="ordered-group"...>
               <ordered-group>
                      ...
               </ordered-group>
          </managed-object>
           ...
    </managed-objects>
    ...
</configuration>
```

Additionally, the managed-object sub-element time-rules, which is not type-specific, can be used to define cron-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
time-rules	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="ordered-group"...>
               [<time-rules attributes >
                     ...
               </time-rules>]
               <ordered-group>
                     ...
               </ordered-group>
          </managed-object>
           ...
    </managed-objects>
...
```

Additionally, the `managed-object` sub-element `control-overrides`, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="ordered-group"...>
               [<control-overrides>
                     ...
               </control-overrides>]
               <ordered-group>
                     ...
               </ordered-group>
          </managed-object>
           ...
    </managed-objects>
...
```

ordered-group sub-element attributes

The following are the `ordered-group` sub-element attributes:

Attribute	Required	Default	Description
`fail-policy`	No	`none`	Defines the action to take when a member of the `ordered-group` fails. The available values are as follows: ■ `none` — The ordered group does not attempt to stop any group members. ■ `ordered-stop` — The ordered group stops all members whose `stop` sequence value is lower than the member that failed, in the order specified by the `stop` attribute specified in the `member` sub-element. After all members have been stopped, the ordered group attempts to restart its members in the ordered specified by the member's `start` attribute. ■ `stop-all` — The ordered group attempts to stop all members. Once all members have successfully stopped, the ordered group's state changes to `Failed`. Typically, the ordered group then attempts to restart its members in the order specified by the members' `start` attribute. ■ `shutdown` — The ordered group attempts to stop all members in the ordered specified by their `stop` attributes. Once all members have been stopped, the ordered group's state is set to `Cannot Start`.
`cannot-start-policy`	No	`none`	Defines the action to take when a member of the `ordered-group` fails to start. The available values are as follows: ■ `none` — The ordered group does not attempt to stop any group members. ■ `ordered-stop` — The ordered group stops all members whose `stop` sequence value is lower than the member that failed to start, in the order specified by the `stop` attribute specified in the `member` sub-element. After all members have been stopped, the ordered group attempts to restart its members in the ordered specified by the member's `start` attribute. ■ `stop-all` — The ordered group attempts to stop all members. Once all members have successfully stopped, the ordered group's state changes to `Failed`. Typically, the ordered group then attempts to restart its members in the order specified by the members' `start` attribute. ■ `shutdown` — The ordered group attempts to stop all members in the ordered specified by their `stop` attributes. Once all members have been stopped, the ordered group's state is set to `Cannot Start`.

```
...
<managed-objects>
       <managed-object ...type="ordered-group"...>
            <ordered-group [fail-policy="null|ordered-stop|stop-all|shutdown"
            cannot-start-policy="null|ordered-stop|stop-all|shutdown"]>
                ...
           </ordered-group>
       </managed-object>
        ...
```

```
</managed-objects>
...
```

member sub-element

Sub-Element	Required	Default	Description
`member`	Yes	n/a	Use to designate each Managed Object member of the group and the start and stop order for each within the context of the group.

```
...
 <managed-objects>
        <managed-object ...type="ordered-group"...>
           <ordered-group>
                 [<member attributes/>]
                 [<member attributes/>]
                 [<member attributes/>]
                 [...]
           </ordered-group>
        </managed-object>
        ...
 </managed-objects>
...
```

member sub-element attributes

The following are the `member` sub-element attributes.

Attribute	Required	Default	Description
`mo-ref`	Yes	n/a	Logical AppServer name of the Managed Object representing the object grouped here, following the syntax: *`agent-name/mo-name`*.
`start`	No	`1` - if no members already exist in group. *highest existing member start value +1* - when adding member(s) to a group with existing members.	The order in which the Managed Resource is to be started with regard to the other members of the `ordered-group`. Must be a positive integer not greater than the total number of members in the group. The `start` order ascends beginning with "1". To set a member to start first, set this attribute to "1". To set multiple members to start concurrently, set this attribute to the same number for those members.
`stop`	No	Reverse of the `start` order. `1` - if no start order.	The order in which this object is to be stopped with regard to other members of the `ordered-group`. Must be a positive integer not greater than the total number of members in the group. The `stop` order ascends beginning with "1". To set a member to stop first, set this attribute to "1". To set multiple members to stop concurrently, set this attribute to the same number for those members.

Attribute	Required	Default	Description
stop-policy	No	require-stop	The following are the acceptable values: ■ ignore: the state of the member is ignored when stopping the group. This is the default for pingers. ■ require-stop: (default except for pingers) the member is required to stop for the group to be considered stopped.
group-policy-level	No	all	Indicates the level to which this member participates in group policies (see fail-policy and cannot-start-policy). The following are the acceptable values: ■ all: this member triggers ordered stops and shutdowns as indicated by the group's policies. ■ start-only: This member only triggers a shutdown (if indicated by the group's policies) and only during a normal start (not on a failure after the group is already running).

```
...
<managed-objects>
     <managed-object ...type="ordered-group"...>
            <ordered-group>

                  [<member mo-ref="${agent.name}/managedobject_name"
 start="start_order" stop="stop_order"
 stop-policy="require-stop|ignore"
 group-policy-level="all|start-only"/>]
                  [...]
              </ordered-group>
      </managed-object>
        ...
</managed-objects>
...
```

redundancy-group Managed Object type

This section lists and describes the redundancy-group Managed Object type sub-elements and sub-element attributes.

Note

For information about the redundancy-group Managed Object attributes, go to the managed-object element attributes section.

redundancy-group sub-elements

The following are the redundancy-group type sub-elements:

Sub-Element	Description
redundancy-group	Contains the redundancy-group Managed Object type-specific Configuration information.

```
...
<managed-objects>
       <managed-object ...type="redundancy-group"...>
```

```
            <redundancy-group>
                 ...
            </redundancy-group>
        </managed-object>
        ...
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="redundancy-group"...>
               [<time-rules ... >
                     ...
               </time-rules>]
               <redundancy-group>
                     ...
               </redundancy-group>
          </managed-object>
           ...
    </managed-objects>
...
```

Additionally, the `managed-object` sub-element `control-overrides`, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="redundancy-group"...>
               [<control-overrides>
                     ...
               </control-overrides>]
               <redundancy-group>
                     ...
               </redundancy-group>
          </managed-object>
           ...
    </managed-objects>
...
```

redundancy-group sub-element attributes

The following are the `redundancy-group` sub-element attributes:

Attribute	Required	Default	Description
`require`	No	`1`	The actual number of members that have to be running for the group to be considered "up". To set to all members in the group, enter `-1`.
`desired-range-min`	No	`1`	The minimum number of members that you desire to be up and running when the group is run. To set to all members in the group, enter `-1`.
`desired-range-max`	No	`1`	The maximum number of members that you want up and running. If more are running, AppServer will try to shutdown members to meet this value. This value cannot exceed the total number of members in the group. To set to all members, enter `-1`.
`count-member-unknown-as-stopped`	No	`false`	When calculating how many members to start, AppServer subtracts the number of running members from the `desired-range-min` value. By default, AppServer includes the members with unknown state in the count of running members. If you want AppServer to skip the members with unknown state in the count of running members (treats them as stopped), set `count-member-unknown-as-stopped` to `true`.
`count-member-problem-as-stopped`	No	`false`	When calculating how many members to start, AppServer subtracts the number of running members from the `desired-range-min` value. By default, AppServer skips the members with state equal to "malfunctioning" in the count of running members (treats them as stopped). If you want AppServer to include the members with state equal to "malfunctioning" in the count of running members, set `count-member-problem-as-stopped` to `true`.
`excess-running-ok`	No	`true`	By default, if the number of members running exceeds the `desired-range-max` value, AppServer considers the group as running. If you want AppServer to consider the group as malfunctioning when the number of members running exceeds the `desired-range-max` value, set `excess-running-ok` to `false`.
`clear-member-start-failures`	No	`false`	When the required number of members reaches the running state, clears any failure states of the other members (for example, change `Cannot Start` to `Stopped`). This makes the other members available for failover, even though they were originally in a `Cannot Start` state. The default value is `false`.

```
...
<managed-objects>
       <managed-object ...type="redundancy-group"...>
```

```
            <redundancy-group [require="integer"
 desired-range-min="integer" desired-range-max="integer"
 count-member-unknown-as-stopped="true|false"
 count-member-problem-as-stopped="true|false"
 excess-running-ok="true|false"
 clear-member-start-failures="true|false"]>
                  ...
              </redundancy-group>
       </managed-object>
        ...
</managed-objects>
...
```

member sub-element

Sub-Element	Required	Default	Description
member	No	n/a	Use to designate each Managed Object member of the group.

```
...
<managed-objects>
      <managed-object ...type="redundancy-group"...>
            <redundancy-group>
                  [<member attributes/>]
                  [<member attributes/>]
                  [<member attributes/>]
                  [...]
             </redundancy-group>
      </managed-object>
      ...
</managed-objects>
...
```

member sub-element attributes

The following are the member sub-element attributes:

Attribute	Required	Default	Description
mo-ref	Yes	n/a	Logical AppServer name of the Managed Object referenced, following the syntax: ${*agent-name*}/*mo-name*.
stop-policy	No	require-stop	The following are the acceptable values: ■ ignore: the state of the member is ignored when stopping the group. This is the default for pingers. ■ require-stop: (default except for pingers) the member is required to stop for the group to be considered stopped.

```
...
<managed-objects>
      <managed-object ...type="redundancy-group"...>
            <redundancy-group>
                  [<member
 mo-ref="${agent1_name}/managedobject1-name"
 stop-policy="ignore|require-stop"/>]
                 [<member
 mo-ref="${agent2_name}/managedobject2-name"
 stop-policy="ignore|require-stop"/gt;]
                   [...]
            </redundancy-group>
```

```
      </managed-object>
       ...
</managed-objects>
...
```

state-proxy Managed Object type

This section lists and describes the `state-proxy` Managed Object `type` sub-elements and attributes and sub-elements.

Note

For information about the `state-proxy` Managed Object attributes, go to the managed-object element attributes section.

state-proxy sub-elements

The following are the `state-proxy` type sub-elements.

Sub-Element	Description
`state-proxy`	Contains the `state-proxy` Managed Object `type`-specific Configuration information.

```
...
<managed-objects>
       <managed-object ...type="state-proxy"...>
             <state-proxy>
                  ...
             </state-proxy>
       </managed-object>
       ...
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="state-proxy"...>
               [<time-rules ... >
                     ...
               </time-rules>]
               <state-proxy>
                     ...
               </state-proxy>
          </managed-object>
           ...
    </managed-objects>
...
```

Additionally, the `managed-object` sub-element `control-overrides`, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="state-proxy"...>
               [<control-overrides>
                     ...
               </control-overrides>]
               <state-proxy>
                     ...
               </state-proxy>
          </managed-object>
           ...
    </managed-objects>
...
```

state-proxy sub-element attributes

The following are the `state-proxy` sub-element attributes:

Attribute	Required	Default	Description
`mo-ref`	Yes	n/a	Logical AppServer name of the Managed Object referenced, following the syntax: `${agent-name}/mo-name`.

```
...
<managed-objects>
      <managed-object ...type="state-proxy"...>
            [<state-proxy> mo-ref="${agent_name}/managedobject1-name"/>]
      </managed-object>
       ...
</managed-objects>
...
```

process Managed Object type

This section lists and describes the `process` Managed Object `type` sub-elements and attributes.

Note

For information about the `process` Managed Object attributes, go to the managed-object element attributes section.

process sub-elements

The following are the `process` Managed Object type sub-elements.

Sub-Element	Description
`process`	Contains the `process` Managed Object type-specific configuration information.

```
...
<managed-objects>
```

```
      <managed-object ...type="process"...>
            <process ...attributes...>
                ...
            </process>
             ...
      </managed-object>
  </managed-objects>
  ...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
  ...
     <managed-objects>
            <managed-object ...type="process"...>
                 [<time-rules ... >
                       ...
                 </time-rules>]
                 <process>
                       ...
                 <process>
            </managed-object>
             ...
      </managed-objects>
  ...
```

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
  ...
     <managed-objects>
            <managed-object ...type="process"...>
                       ...
                 [<control-overrides>
                       ...
                 </control-overrides>]
            </managed-object>
             ...
      </managed-objects>
  ...
```

process sub-element attributes

The following are the `process` sub-element attributes:

Attribute	Required	Default	Description
`command`	Yes	n/a	Full path to the program file. Windows users may omit the extension, in which case "`.exe`" is assumed.
`directory`	No	`${CONFIG.path}/mos/${MO.name}/` - Where the replacement string variable `${CONFIG.path}` represents the path to the Managed Object `configuration.xml` file.	The full path and initial working directory of the process.
`data-id`	No	`1`	A unique identifier for the `process` element - used by AppServer to identify multiple `process` entries within the same Managed Object definition. Acceptable values are positive integers.

```
...
<managed-objects>
      <managed-object ...type="process"...>
           <process command="fullpath" directory="fullpathdirectory>
 data-id="integer">
            ...
           </process>
      </managed-object>
</managed-objects>
...
```

process sub-elements

The following are the java-process sub-element elements.

arguments sub-element

Sub-Element	Required	Default	Description
`arguments`	No	n/a	Use to supply command-line arguments you want referenced when the process is started. Each command-line argument you enter appears within a `argument` sub-element. The arguments are added to the command-line string in the order in which they are listed in the `configuration.xml`. **Note:** Each argument sub-element cannot contain more than a single command-line argument. For example, `-user Joe` is two arguments and must be entered as such within the `arguments` element: `<argument>-user</argument> <argument>Joe</argument>`

```
...
<managed-objects>
      <managed-object ...type="process"...>
           <process ... >
              [<arguments>
                  [<argument>argument1</argument>]
                  [<argument>argument2</argument>]
                  [...more arguments...]
              </arguments>]
           </process>
```

```
      </managed-object>
</managed-objects>
...
```

env-vars sub-element

Sub-Element	Required	Default	Description
env-vars	No	If no path, library-path, or env-vars attributes are set, uses the Agent SCU environment.	Use to set environment variables for the process. Enter a sequence of env-var elements whose name and value attributes define the variable to be set. Additionally, you can use the following env-vars element attributes to point to standard environment variable sources. This enables you to use the same environment variables for several Managed Resources. **Note:** If the same environment variable is set in use-vbroker-env, then the variable set using the env-vars element overrides it.

env-vars sub-element attributes

Attribute	Required	Default	Description
use-current-env	No	false	If you want the process to use the entire Agent SCU process environment, set to true.
use-default-env	No	true	If you do not want the process to use the common host system environment variables required to run a process, set to false. **Note:** If the same variable is set in use-user-login-env, then use-default-env overrides it.
use-vbroker-env	No	false	For a VisiBroker process, set to true. **Note:** If the same variable is set in use-default-env, then the use-vbroker-env variable overrides it.
use-user-login-env	No	false	UNIX only. If you want the process to use the environment the Managed Object process owner has when that user logs in to their default shell, set to true. **Note:** If the same variable is set in use-current-env, then use-user-login-env overrides it. The login shell must support running its command line argument as a shell script when the shell is invoked as a login shell. Also, when an absolute path to a program file is specified as the only content on a line, the shell's script language must invoke a child process, for example, normal KornShell, C Shell and so on semantics.

```
...
<managed-objects>
      <managed-object ...type="process"...>
           <process ... >
              [<env-vars use-current-env="true|false" use-default-env="true|
false" use-vbroker-env="true|false" use-login-env="true|false">]
                  [<env-var name="variable1"
 value="value1" />]
                  [<env-var name="variable2"
 value="value2" />]
                  [...more env-vars...]
              </env-vars>]
```

```
            </process>
        </managed-object>
    </managed-objects>
    ...
```

library-path sub-element

Sub-Element	Required	Default	Description
library-path	No	None	For UNIX only. Use to place items in the library path of the process. Each path you enter appears in the configuration.xml within a directory sub-element within the library-path element. Note: If the same environment variable is set using the env-vars element, then the library-path setting overrides it.

```
...
<managed-objects>
      <managed-object ...type="process"...>
            <process ... >
                  [<library-path>
                   <directory>path1</directory>
                  [<directory>path2</directory>]
                  [...more paths...]
                </library-path>]
            </process>
      </managed-object>
</managed-objects>
...
```

path sub-element

Sub-Element	Required	Default	Description
path	No	n/a	Non-platform specific. Use to place items in the process executable path. Each path you enter appears in the configuration.xml within a directory sub-element within the path element. Note: If the same environment variable is set using the env-vars element, then the path setting overrides it.

```
...
<managed-objects>
      <managed-object ...type="process"...>
            <process ... >
                  [<path>
                   <directory>path1</directory>
                  [<directory>path2</directory>]
                  [...more paths...]
                </path>]
            </process>
      </managed-object>
</managed-objects>
...
```

stdin|stdout|stderr sub-elements

Sub-Element	Required	Default	Description
`stdin`	No	Inherited from the assigned Agent.	Use to specify process standard-in settings.
`stdout`	No	Inherited from the assigned Agent.	Use to specify process standard-out settings.
`stderr`	No	Inherited from the assigned Agent.	Use to specify process standard-error settings.

stdin|stdout|stderr sub-element attributes

Attribute	Required	Default	Description
`path`	No	Inherited from the assigned Agent.	Use to set the process path location for standard-in, standard-out, or standard-error.
`append`	No	false	Applies to `stdout` and `stderr`. If you want the file to be opened for append, set to `true`. If set to `false`, the file is truncated.

```
...
<managed-objects>
      <managed-object ....type="process"...>
           <process ... >
               [<stdin path="path1/>]
               [<stdout path="path2"  append="true|false"/>]
               [<stderr path="path3"
 append="true|false"/>]
           </process>
      </managed-object>
</managed-objects>
...
```

platform-specific sub-element

Sub-Element	Required	Default	Description
`platform-specific` (Windows)	No	Attributes: `show-gui="false"` `start-minimized="true"` `window-title`: name of the executable file.	Used to specify Windows platform-specific settings for a process that has a GUI. The following are the Windows attributes for this element: `show-gui`: If you want the process window launched when the process is started, set to `true`. `window-title`: Use to specify a title for the process window. `start-minimized`: If you want the window maximized upon process start, set to `false`.
`platform-specific` (UNIX)	No	Attributes: `group-name="agentgroupname"` `user-name="agentusername"` `stdout-mode="644"` `stderr-mode="644"` `umask="agentumask"` `nice-value="agentnicevalue"` `root-directory="agentrootdirectory"`	Used to specify UNIX platform-specific settings. The following are the UNIX attributes for this element: `group-name`: group name of the process. For information on how to set `user-name`: user name of the process. `stdout-protection-mode`: octal file access mode for the stdoutput file. `stderr-protection-mode`: octal file access mode for the stderrput file. `umask`: octal umask for the process. `nice-value`: scheduling priority for the process - must be a positive integer. `root-directory`: the root directory of the process.

Important

The UNIX platform-specific settings can only be specified when the Agent SCU process is started with root privileges.

Windows

```
...
<managed-objects>
      <managed-object ...type="process"...>
            <process ... >
                  [<platform-specific show-gui="true|false"
 window-title="string" start-minimized="true|false"/>
            </process>
      </managed-object>
</managed-objects>
...
```

UNIX

```
...
<managed-objects>
      <managed-object ...type="process"...>
            <process ... >
                  [<platform-specific group-name="groupname"
 user-name="username" stdout-mode="octalfileaccessmode"
 stderr-mode="octalfileaccessmode" umask="octalumask"
 nice-value="integer"/>]
            </process>
```

```
      </managed-object>
</managed-objects>
...
```

java-process Managed Object type

This section lists and describes the `java-process` Managed Object `type` sub-elements and attributes.

Note

For information about the `java-process` Managed Object attributes, go to the managed-object element attributes section.

java-process sub-elements

The Configuration information specific to a `java-process` Managed Object is contained within the following `java-process` sub-elements:

- `java-process`
- `VBJ-process`

The following are all the available `java-process` type sub-elements.

Sub-Element	Description
`java-process`	Used for Java classes on which a `main()` method is invoked.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
            <java-process ...attributes...>
             ...
            </java-process>
      </managed-object>
</managed-objects>
...
```

Sub-Element	Description
`VBJ-process`	Use to define a VisiBroker for Java process (that is, a process started with `vbj` rather than `java`).

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
            <VBJ-process ...attributes...>
             ...
            </VBJ-process>
      </managed-object>
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="java-process"...>
               [<time-rules ... >
                     ...
               </time-rules>]
               <java-process>
                     ...
               </java-process>
          </managed-object>
           ...
    </managed-objects>
...
```

Additionally, the `managed-object` sub-element `control-overrides`, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="java-process"...>
                     ...
               [<control-overrides>
                     ...
               </control-overrides>]
          </managed-object>
           ...
    </managed-objects>
...
```

java-process sub-element attributes

The following are the `java-process` sub-element attributes:

Attribute	Required	Default	Description
`command`	Yes	n/a	Full path to the Java process program file. Windows users may omit the extension, in which case "`.exe`" is assumed.
`main-class`	Yes	n/a	The main class of the Java process Managed Resource in the format: *`MyPackage.MyClass`*.
`directory`	No	`${AGENT.root}/var/domains/`*`domain-name/`* - Where the replacement string variable `${AGENT.root}` represents the root of the Agent to which the Managed Object is assigned.	The full path and initial working directory of the process.

Attribute	Required	Default	Description
vm-type	No	n/a	Acceptable values are: ■ classic ■ hotspot ■ client ■ server
data-id	No	1	A unique identifier for the java-process element - used by AppServer to identify multiple java-process entries within the same Managed Object definition. Acceptable values are positive integers.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <process command="fullpath" main-class="MyPackage.MyClass"
directory="fullpathdirectory"
vm-type="server|client|classic"
data-id="integer">
            ...
           </java-process>
      </managed-object>
</managed-objects>
...
```

java-process sub-elements

The following are the java-process sub-element elements:

Sub-Element	Required	Default	Description
java-properties	No	n/a	Use to provide VM properties for the Managed Object. Enter a sequence of java-property elements containing a name and value attribute for each property to be set (*name* must not include -D). Results in java -D command-line syntax. **Note:** For a vbj-process type Managed Object - results in VBJ -VBJprop command-line syntax.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
              [<java-properties>
                    [<java-property name="name1" value="value1" />]
                    [<java-property name="name2" value="value2"  />]
                    [...more_properties...]
              </java-properties>]
           </java-process>
      </managed-object>
</managed-objects>
...
```

Sub-Element	Required	Default	Description
classpath	No	n/a	Use to specify classpath entries for the java -classpath. Each classpath entry appears within a classpath-entry sub-element within the classpath element. **Note:** For vbj-process type Managed Objects, we recommend using -vbj classpath instead. For more information, go to the vbj-process sub-elements table.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ...attributes...>
               [<classpath>
                   [<classpath-entry>classpath1</classpath-entry>[
                   [<classpath-entry>classpath2</classpath-entry>]
                   [...more classpaths...]
               </classpath>]
               [...]
           </java-process>
      </managed-object>
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`options`	No	n/a	Use to specify options to precede the main class on the command-line. Each option must appear within an `option` sub-element within the `options` element. **Note:** Each `option` sub-element cannot contain more than a single option. For example, `-Xms -Xmx` are two options and must be entered as such within the `options` element: `<option>-Xms</option> <option>-Xmx>/option>`

```
...
<managed-objects>
      <managed-object ...type="process"...>
           <java-process ...attributes...>
              [<options>
                  [<option>option1</option>]
                  [<option>option2</option>]
                  [...more options...]
              </options>]
           </java-process>
      </managed-object>
</managed-objects>
...
```

arguments sub-element

For information about this sub-element, go to the process arguments sub-element section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
              [<arguments>
                  [<argument>argument1</argument>]
                  [<argument>argument2</argument>]
                  [...more arguments...]
              </arguments>]
           </java-process>
      </managed-object>
</managed-objects>
...
```

env-vars sub-element

For information about this sub-element, go to the process env-vars sub-element section.

env-vars sub-element attributes

For information about this sub-element, go to the process env-vars sub-element attributes section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
              [<env-vars use-current-env="true|false" use-default-env="true|
false" use-vbroker-env="true|false" use-login-env="true|false">
                   <env-var name="variable1"
 value="value1" />
                   [<env-var name="variable2"
 value="value2" />]
                   [...more env-vars...]
              </env-vars>]
           </java-process>
      </managed-object>
</managed-objects>
...
```

library-path sub-element

For information about this sub-element, go to the process library-path sub-element section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
                  [<library-path>
                   <directory>path1</directory>
                  [<directory>path2</directory>]
                  [...more paths...]
              </library-path>]
           </java-process>
      </managed-object>
</managed-objects>
...
```

path sub-element

For information about this sub-element, go to the process path sub-element section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
                 [<path>
                  <directory>path1</directory>
                 [<directory>path2</directory>]
                 [...more paths...]
              </path>]
           </java-process>
      </managed-object>
</managed-objects>
...
```

stdin|stdout|stderr sub-elements

For information about this sub-element, go to the process stdin|stdout|stderr sub-elements section.

```
...
<managed-objects>
      <managed-object ....type="java-process"...>
           <java-process ... >
               [<stdin path="path1"/>]
               [<stdout path="path2"
 append="true|false"/>]
               [<stderr path="path3"
 append="true|false"/>]
          </java-process>
      </managed-object>
</managed-objects>
...
```

platform-specific sub-elements

For information about this sub-element, go to the process platform-specific sub-element section.

Windows

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
                [<platform-specific show-gui="true|false"
 window-title="string" start-minimized="true|false"/>
          </java-process>
      </managed-object>
</managed-objects>
...
```

UNIX

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <java-process ... >
                [<platform-specific group-name="groupname"
 user-name="username" stdout-mode="octalfileaccessmode"
 stderr-mode="octalfileaccessmode" umask="octalumask"
 nice-value="integer"/>]
          </java-process>
      </managed-object>
</managed-objects>
...
```

VBJ-process sub-element attributes

The `vbj-process` element inherits all of the attributes of the `java-process` element. For more information, go to the java-process sub-element attributes section.

There are no additional attributes specific to the `VBJ-process` element.

VBJ-process sub-elements

The following are the VBJ-process element-specific sub-elements:

Sub-Element	Required	Default	Description
vbj-java-options	No	n/a	Use to add options for the vbj -J command-line syntax - must not include -J. Each option must appear within the option sub-element within the vbj-java-options element. **Note:** Each option sub-element cannot contain more than a single option. For example, -Xms -Xmx are two options and must be entered as such within the options element: <option>-Xms</option> <option>-Xmx>/option>

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <VBJ-process ... >
              [<vbj-java-options>]
                  [<option>option1</option>]
                  [<option>option2</option>]
                  [...more options...]
              </vbj-java-options>]
           </VBJ-process>
      </managed-object>
</managed-objects>
...
```

Sub-Element	Required	Default	Description
vbj-classpath	No	n/a	Use to specify classpath entries for vbj -VBJClasspath classpath. Each classpath entry you enter must appear within the classpath-entry sub-element within the vbj-classpath element.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <VBJ-process ... >
              [<vbj-classpath>
                  [<classpath-entry>classpath1</classpath-entry>]
                  [<classpath-entry>classpath2</classpath-entry>]
                  [...more classpaths...]
              </vbj-classpath>]
           </VBJ-process>
      </managed-object>
</managed-objects>
...
```

Sub-Element	Required	Default	Description
vbj-add-jars	No	n/a	Use to provide a list of JARs to be used with the vbj -VBJaddJar command-line option. Each JAR you want to add to the command-line must appear within the add-jars-entry sub-element within the vbj-add-jars element.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <VBJ-process ... >
              [<vbj-add-jars>
                  [<add-jars-entry>jar1</add-jars-entry>]
                  [<add-jars-entry>jar2</add-jars-entry>]
```

```
                [...more JARs...]
              </vbj-add-jars>]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

For information about this sub-element, go to the java-process java-process sub-elements section.

Note

For `vbj-process` Managed Objects - results in `VBJ -VBJprop` command-line syntax.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <VBJ-process ... >
             [<java-properties>
                  [<java-property name="name1" value="value1" />]
                  [<java-property name="name2" value="value2"  />]
                  [...more_properties...]
             </java-properties>]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

For information about this sub-element, go to the java-process java-process sub-elements section.

Note

For `VBJ-process` Managed Objects, we recommend using `-vbj classpath` instead. For more information, go to the vbj-process sub-elements table.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <VBJ-process ...attributes...>
              [<classpath>
                  [<classpath-entry>classpath1</classpath-entry>[
                  [<classpath-entry>classpath2</classpath-entry>]
                  [...more classpaths...]
              </classpath>]
              [...]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

options sub-element

For information about this sub-element, go to the java-process java-process sub-elements section.

```
...
<managed-objects>
      <managed-object ...type="process"...>
          <VBJ-process ...attributes...>
             [<options>
                 [<option>option1</option>]
                 [<option>option2</option>]
                 [...more options...]
             </options>]
```

```
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

arguments sub-element

For information about this sub-element, go to the process arguments sub-element section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <java-process ... >
             [<arguments>
                 [<argument>argument1</argument>]
                 [<argument>argument2</argument>]
                 [...more arguments...]
             </arguments>]
          </java-process>
      </managed-object>
</managed-objects>
...
```

env-vars sub-element

For information about this sub-element, go to the process env-vars sub-element section.

env-vars sub-element attributes

For information about this sub-element, go to the process env-vars sub-element attributes section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <VBJ-process ... >
             [<env-vars use-current-env="true|false" use-default-env="true|
false" use-vbroker-env="true|false" use-login-env="true|false">
                  <env-var name="variable1"
 value="value1" />
                  [<env-var name="variable2"
 value="value2" />]
                  [...more env-vars...]
             </env-vars>]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

library-path sub-element

For information about this sub-element, go to the process library-path sub-element section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <VBJ-process ... >
                 [<library-path>
                  <directory>path1</directory>
                 [<directory>path2</directory>]
```

```
                [...more paths...]
              </library-path>]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

path sub-element

For information about this sub-element, go to the process path sub-element section.

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <VBJ-process ... >
                [<path>
                 <directory>path1</directory>
                [<directory>path2</directory>]
                [...more paths...]
              </path>]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

stdin|stdout|stderr sub-elements

For information about this sub-element, go to the process stdin|stdout|stderr sub-elements section.

```
...
<managed-objects>
      <managed-object ....type="java-process"...>
          <VBJ-process ... >
              [<stdin path="path1"/>]
              [<stdout path="path2"
 append="true|false"/>]
              [<stderr path="path3"
 append="true|false"/>]
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

platform-specific sub-elements

For information about this sub-element, go to the process platform-specific sub-element section.

Windows

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
          <VBJ-process ... >
               [<platform-specific show-gui="true|false"
 window-title="string" start-minimized="true|false"/>
          </VBJ-process>
      </managed-object>
</managed-objects>
...
```

UNIX

```
...
<managed-objects>
      <managed-object ...type="java-process"...>
           <VBJ-process ... >
                [<platform-specific group-name="groupname"
 user-name="username" stdout-mode="octalfileaccessmode"
 stderr-mode="octalfileaccessmode" umask="octalumask"
 nice-value="integer"/>]
           </VBJ-process>
      </managed-object>
</managed-objects>
...
```

custom-javascript Managed Object type

This section lists and describes the `custom-javascript` Managed Object `type` sub-elements and attributes.

Note

For information about the `custom-javascript` Managed Object attributes, go to the managed-object element attributes section.

Important

(UNIX only) JavaScript execution is typically not allowed when a Managed Object is started by an Agent configured for Multi-user Mode (MUM). If any of the Managed Objects that will run in Multi-user Mode (MUM) contain JavaScripts, you need to modify the `agent.config` to allow those JavaScripts to run.

custom-javascript sub-elements

The following are the `custom-javascript` type sub-elements:

Elements	Description
`jscript`	Contains the Configuration information specific to a `custom-javascript` type Managed Object.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             ...
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
         <managed-object ...type="custom-javascript"...>
              [<time-rules ... >
                    ...
              </time-rules>]
              <jscript>
                    ...
              </jscript>
         </managed-object>
          ...
    </managed-objects>
...
```

Elements	Description
control-overrides	A required element that contains the action strategy and action parameters information for the Managed Object. For more information, go to the Managed Object control-overrides element section.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
            <control-overrides>
                 ...
            </control-overrides>
            ...
       </managed-object>
       ...
</managed-objects>
...
```

control-overrides element sub-elements

The following are the control-overides element sub-elements:

Sub-Element	Required	Default	Description
start	Yes (not required for pinger Managed Objects)	n/a	Used to specify the start action strategy and to reference the jscript sub-element containing the start action specific information.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
            <control-overrides>
                  <start ...attributes.../>
                  ...
            </control-overrrides>
            <jscript ...attributes...>
                 ...
            </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

Sub-Element	Required	Default	Description
ping	Yes	n/a	Used to specify the ping action strategy and to reference the jscript sub-element containing the ping action specific information.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    ...
                    <ping ...attributes.../>
                    ...
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`stop`	Yes (not required for pinger Managed Objects)	n/a	Used to specify the stop action strategy and to reference the `jscript` sub-element containing the stop action specific information.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    ...
                    <stop ...attributes.../>
                    ...
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`kill`	No	n/a	Used to specify the kill action strategy and to reference the `jscript` sub-element containing the kill action specific information.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    ...
                    <kill ...attributes.../>
                    ...
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
```

```
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`parameters`	No	n/a	Used to specify action strategy additional parameters. For more information, go to the Managed Object `control-overides` element, parameters sub-element attributes section.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    <start ...attributes.../>
                    <ping ...attributes.../>
                    <stop ...attributes.../>
                    <kill ...attributes.../>
                    <parameters ...attributes.../>
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

start sub-element attributes

The following are the `start` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`jscript-control`	Runs a JavaScript to start the Managed Resource. To disable this action, set to `null`.
`data-id-ref`	No	`1`	A unique identifier for the `jscript` start action strategy - used by AppServer to reference the `jscript` sub-element which contains the reference to the JavaScript to run for the start action. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    <start strategy="jscript-control|null" [data-id-
ref="integer"]/>
                    <ping ...attributes.../>
                    <stop ...attributes.../>
                    <kill ...attributes.../>
                    <parameters ...attributes.../>
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

ping sub-element attributes

The following are the `ping` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`jscript-ping`	Runs a JavaScript to ping the Managed Resource. To disable this strategy, set to `null`.
`data-id-ref`	No	`2`	A unique identifier for the `jscript` ping action strategy - used by AppServer to reference the `jscript` sub-element which contains the reference to the JavaScript to run for the ping action. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    <start strategy="jscript-control|null"
[data-id-ref="integer"]/>
                    <ping strategy="jscript-ping|null" [data-id-ref="integer"]/
>
                    <stop ...attributes.../>
                    <kill ...attributes.../>
                    <parameters ...attributes.../>
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

stop sub-element attributes

The following are the `stop` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`jscript-control`	Runs a JavaScript to stop the Managed Resource. To disable this strategy, set to `null`.
`data-id-ref`	No	`3`	A unique identifier for the `jscript` stop action strategy - used by AppServer to reference the `jscript` sub-element which contains the reference to the JavaScript to run for stop. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                    <start strategy="jscript-control|null"
[data-id-ref="integer"]/>
                    <ping strategy="jscript-ping|null" [data-id-ref="integer"]/
>
                    <stop strategy="jscript-control|null"
[data-id-ref="integer"]/>
                    <kill ...attributes.../>
                    <parameters ...attributes.../>
             </control-overrrides>
             <jscript ...attributes...>
                  ...
```

```
            </jscript>
        </managed-object>
        ...
</managed-objects>
...
```

kill sub-element attributes

The following are the `kill` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`jscript-control`	Runs a JavaScript to kill the Managed Resource. To disable this strategy, set to `null`.
`data-id-ref`	No	`4`	A unique identifier for the `jscript` kill action strategy - used by AppServer to reference the `jscript` sub-element which contains the reference to the JavaScript to run for kill. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <control-overrides>
                   <start strategy="jscript-control|null"
[data-id-ref="integer"]/>
                    <ping strategy="jscript-ping|null" [data-id-ref="integer"]/
>
                    <stop strategy="jscript-control|null"
[data-id-ref="integer"]/>
                    <kill strategy="jscript-control|null"
[data-id-ref="integer"]/>
                    <parameters ...attributes.../>
             </control-overrrides>
             <jscript ...attributes...>
                  ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

jscript sub-element attributes

The following is the only `jscript` sub-element attribute:

Attribute	Required	Default	Description
`data-id`	Yes	`1`	A unique identifier for the `jscipt` element - used by AppServer to identify multiple `jscipt` entries within the same `custom-javascript` Managed Object definition. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <jscript data-id="integer1">
                  ...
             </jscript>
             [<jscript data-id="integer2">
                  ...
             </jscript>]
```

```
            [<jscript data-id="integer3">
                ...
            </jscript>]
        </managed-object>
        ...
</managed-objects>
...
```

jscript sub-elements

The following are the `jscript` sub-elements:

Sub-Element	Description
`run`	Contains the reference to the JavaScript to run. The JavaScripts are executed in the order they appear in the `configuration.xml`. Each `run` can contain only one script. You can define multiple run elements - each containing its own JavaScript and `arguments` defined for a Managed Object.

```
...
<managed-objects>
        <managed-object ...type="custom-javascript".../>
            <jscript data-id="integer1">
                <run ...attributes...>
                    ...
                </run>
            </jscript>
            [<jscript data-id="integer2">
                <run ...attributes...>
                    ...
                </run>
            </jscript>]
        </managed-object>
        ...
</managed-objects>
...
```

Sub-Element	Description
`classpath`	For information about this sub-element, go to the java-process sub-elements section.

```
...
<managed-objects>
        <managed-object ...type="custom-javascript".../>
            <jscript ...attributes... >
              [<classpath>
                  [<classpath-entry>classpath1</classpath-entry>[
                  [<classpath-entry>classpath2</classpath-entry>]
                  [...more classpaths...]
              </classpath>]
                ...
            </jscript>
            [<jscript ...attributes... >
              [<classpath>
                  [<classpath-entry>classpath1</classpath-entry>[
                  [<classpath-entry>classpath2</classpath-entry>]
                  [...more classpaths...]
              </classpath>]
                ...
            </jscript>]
        </managed-object>
```

```
        ...
</managed-objects>
...
```

run sub-element attributes

The following are the `run` sub-element available attributes:

Attribute	Required	Default	Description
`script`	Yes	n/a	The fully-qualified path to a JavaScript file. If you are embedding the JavaScript, **do not** use this attribute.
`name`	No	`scriptn`	A name for the JavaScript that is used for Management Console and Logging display purposes only - not a referenceable name.

```
...
<managed-objects>
        <managed-object ...type="custom-javascript".../>
              <jscript data-id="integer1">
                   <run script="fullpathfilename" [name="scriptn"|string1"]>
                        ...
                   </run>
                   ...
              </jscript>
              [<jscript data-id="integer2">
                   <run script="fullpathfilename" [name="scriptn+1"|string2"]>
                        ...
                   </run>
                   ...
              </jscript>]
        </managed-object>
        ...
</managed-objects>
...
```

run sub-element

The following is the only `run` sub-element:

Sub-Element	Required	Default	Description
`arguments`	No	n/a	For a description of this element, go to the process sub-elements section.

```
...
<managed-objects>
        <managed-object ...type="custom-javascript".../>
              <jscript data-id="integer1">
                   <run script="fullpath_filename" [name="scriptn"|string1"]>
                        [<arguments>
                           [<argument>argument1</argument>]
                           [<argument>argument2</argument>]
                        </arguments>]
                   </run>
              </jscript>
              [<jscript data-id="integer2">
                   <run script="fullpath_filename" [name="scriptn+1"|string2"]>
                        [<arguments>
                           [<argument>argument1</argument>]
                           [<argument>argument2</argument>]
                        </arguments>]
```

```
                    </run>
                </jscript>]
            </managed-object>
            ...
    </managed-objects>
    ...
```

![CDATA[section

To embed a JavaScript in the XML, use the `<![CDATA[` section within the `run` sub-element.

Important

In the embedded JavaScipt, AppServer **does not** support the following:

- C++ type comments.
- `try {`
  `} catch(Exception` *`name`*`) {`
- `return` statement - we recommend you use `quit(`*`returncode`*`)` instead.

![CDATA[attribute

Within the `![CDATA[` section you can specify command-line arguments using the `arguments` element. For a description of this element, go to the process sub-elements section.

```
...
<managed-objects>
       <managed-object ...type="custom-javascript".../>
             <jscript data-id="integer1">
                  <run script="fullpathfilename" [name="scriptn"|string1"]>
                       [<arguments>
                           [<argument>argument1</argument>]
                           [<argument>argument2</argument>]
                        </arguments>]
                  </run>
             </jscript>
             <jscript data-id="integer2">
                  <run name="string" [name="scriptn+1"|string1"]>
                        <![CDATA[
                               an-embedded-javascript
                         [<arguments>
                            [<argument>argument1</argument>]
                            [<argument>argument2</argument>]
                            [...more arguments...]
                          </arguments>]
                        ]]>
                  </run>
              ...
             </jscript>
       </managed-object>
       ...
</managed-objects>
...
```

custom-executable Managed Object type

This section lists and describes the `custom-executable` Managed Object `type` sub-elements and attributes.

Note

For information about the `custom-executable` Managed Object attributes, go to the managed-object element attributes section.

custom-executable sub-elements

The following is the `custom-executable` type sub-elements:

Elements	Description
`process`	Contains the information for the process(es) you want to run for the Managed Resource start, stop, ping, and/or kill action. Typically, a `custom-executable` Managed Object definition contains a `process` element for each action.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             ...
             <process ...attributes...>
                  ...
             </process>
             [<process ...attributes...>
                  ...
             </process>]
             [<process ...attributes...>
                  ...
             </process>]
             ...
       </managed-object>
       ...
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="custom-javascript"...>
               [<time-rules ... >
                     ...
               </time-rules>]
               <process>
                     ...
               </process>
          </managed-object>
           ...
    </managed-objects>
...
```

Elements	Description
`control-overrides`	A required element that contains the action strategy and parameters information for the Managed Object. For more information, go to the Managed Object control-overrides element section.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             <control-overrides>
                  ...
             </control-overrides>
             ...
       </managed-object>
       ...
</managed-objects>
...
```

control-overrides sub-elements

The following are the `control-overides` sub-elements:

Sub-Element	Required	Default	Description
`start`	Yes (not required for pinger Managed Object)	n/a	Used to specify the start action strategy and to reference the `process` element containing the start action specific information.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             <control-overrides>
                    [<start ...attributes.../>]
                    ...
             </control-overrrides>
             <process ...attributes...>
                  ...
             </process>
             [<process ...attributes...>
                  ...
             </process>]
             [<process ...attributes...>
                  ...
             </process>]
             ...
       </managed-object>
       ...
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`ping`	Yes	n/a	Used to specify the ping action strategy and to reference the `process` element containing the ping action specific information.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             <control-overrides>
                    ...
                    <ping ...attributes.../>
                    ...
             </control-overrrides>
             <process ...attributes...>
                  ...
             </process>
```

```
            [<process ...attributes...>
                  ...
            </process>]
            [<process ...attributes...>
                  ...
            </process>]
            ...
      </managed-object>
      ...
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`stop`	Yes (not required for pinger Managed Object)	n/a	Used to specify the stop action strategy and to reference the `process` element containing the stop action specific information.

```
...
<managed-objects>
      <managed-object ...type="custom-executable".../>
            <control-overrides>
                  ...
                  <stop ...attributes.../>
                  ...
            </control-overrrides>
            <process ...attributes...>
                  ...
            </process>
            [<process ...attributes...>
                  ...
            </process>]
            [<process ...attributes...>
                  ...
            </process>]
            ...
      </managed-object>
      ...
</managed-objects>
...
```

Sub-Element (cont.)	Required	Default	Description
`kill`	No	n/a	Used to specify the kill action strategy and to reference the `process` element containing the kill action specific information.

```
...
<managed-objects>
      <managed-object ...type="custom-executable".../>
            <control-overrides>
                  ...
                  [<kill ...attributes.../>]
                  ...
            </control-overrrides>
            <process ...attributes...>
                  ...
            </process>
            [<process ...attributes...>
                  ...
            </process>]
```

```
            [<process ...attributes...>
                ...
            </process>]
            ...
        </managed-object>
        ...
</managed-objects>
...
```

Sub-Element	Required	Default	Description
parameters	No	n/a	Used to specify action strategy additional parameters. For more information, go to the Managed Object control-overides element, parameters sub-element attributes section.

```
...
<managed-objects>
        <managed-object ...type="custom-executable".../>
            <control-overrides>
                [<start ...attributes.../>]
                <ping ...attributes.../>
                [<stop ...attributes.../>]
                [<kill ...attributes.../>]
                <parameters ...attributes.../>
            </control-overrrides>
            <process ...attributes...>
                ...
            </process>
            [<process ...attributes...>
                ...
            </process>]
            [<process ...attributes...>
                ...
            </process>]
            ...
        </managed-object>
        ...
</managed-objects>
...
```

start sub-element attributes

The following are the start sub-element attributes:

Attribute	Required	Default	Description
strategy	Yes	run-custom-executable	Runs a process to start the Managed Resource. To disable this action, set to null.
data-id-ref	Yes	1	A unique identifier for the custom-executable start action strategy - used by AppServer to reference the process sub-element which contains the configuration information specific to the process you want to execute for the start action. Acceptable values are positive integers.

```
...
<managed-objects>
        <managed-object ...type="custom-executable".../>
            <control-overrides>
                <start strategy="run-custom-executable|null" data-id-
ref="integer"/>
```

```
                    <ping ...attributes.../>
                    <stop ...attributes.../>
                    <kill ...attributes.../>
                    <parameters ...attributes.../>
              </control-overrrides>
              <process ...attributes...>
                    ...
              </process>
              <process ...attributes...>
                    ...
              </process>
              <process ...attributes...>
                    ...
              </process>
              ...
        </managed-object>
        ...
</managed-objects>
...
```

ping sub-element attributes

The following are the `ping` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`ping-custom-executable`	Runs a process to ping the Managed Resource. To disable this action, set to `null`.
`data-id-ref`	Yes	`2`	A unique identifier for the `custom-executable` ping action strategy - used by AppServer to reference the `process` sub-element which contains the configuration information specific to the process you want to execute for the ping action. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             <control-overrides>
                    [<start strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    <ping strategy="ping-custom-executable|null" data-id-
ref="integer"/>
                    [<stop ...attributes.../>]
                    [<kill ...attributes.../>]
                    [<parameters ...attributes.../>]
             </control-overrrides>
             <process ...attributes...>
                    ...
             </process>
             [<process ...attributes...>
                    ...
             </process>]
             [<process ...attributes...>
                    ...
             </process>]
             ...
       </managed-object>
       ...
</managed-objects>
...
```

stop sub-element attributes

The following are the `stop` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`run-custom-executable`	Runs a process to stop the Managed Resource. To disable this action, set to `null`.
`data-id-ref`	Yes	`3`	A unique identifier for the `custom-executable` stop action strategy - used by AppServer to reference the `process` sub-element which contains the configuration information specific to the process you want to execute for the stop action. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             <control-overrides>
                    [<start strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    <ping strategy="ping-custom-executable|null" data-id-
ref="integer"/>
                    [<stop strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    [<kill ...attributes.../>]
                    [<parameters ...attributes.../>]
             </control-overrrides>
             <process ...attributes...>
                  ...
             </process>
             [<process ...attributes...>
                  ...
             </process>]
             [<process ...attributes...>
                  ...
             </process>]
             ...
       </managed-object>
       ...
</managed-objects>
...
```

kill sub-element attributes

The following are the `kill` sub-element attributes:

Attribute	Required	Default	Description
`strategy`	Yes	`run-custom-executable`	Runs a process to kill the Managed Resource. To disable this action, set to `null`.
`data-id-ref`	Yes	`4`	A unique identifier for the `custom-executable` kill action strategy - used by AppServer to reference the `process` sub-element which contains the configuration information specific to the process you want to execute for the kill action. Acceptable values are positive integers.

```
...
<managed-objects>
       <managed-object ...type="custom-executable".../>
             <control-overrides>
```

```
                    [<start strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    <ping strategy="ping-custom-executable|null" data-id-
ref="integer"/>
                    [<stop strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    [<kill strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    [<parameters ...attributes.../>]
               </control-overrrides>
               <process ...attributes...>
                    ...
               </process>
               [<process ...attributes...>
                    ...
               </process>]
               [<process ...attributes...>
                    ...
               </process>]
               ...
          </managed-object>
          ...
</managed-objects>
...
```

process sub-element attributes

For a list and description of the `process` element attributes and sub-elements, go to the process sub-element attributes section.

```
...
<managed-objects>
          <managed-object ...type="custom-executable".../>
               <control-overrides>
                    [<start strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    <ping strategy="ping-custom-executable|null" data-id-
ref="integer"/>
                    [<stop strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    [<kill strategy="run-custom-executable|null"
data-id-ref="integer"/>]
                    [<parameters ...attributes.../>]
               </control-overrrides>
               <process data-id="integer" command="fullpath"
directory="fullpathdirectory">
                    ...
               </process>
               [<process ...attributes...>
                    ...
               </process>]
               [<process ...attributes...>
                    ...
               </process>]
               ...
          </managed-object>
          ...
</managed-objects>
...
```

osagent Managed Object type

This section lists and describes the `osagent` Managed Object `type` sub-elements and attributes.

Note

For information about the `osagent` Managed Object attributes, go to the managed-object element attributes section.

osagent sub-elements

Configuration information specific to the `osagent` type Managed Object is contained within the following sub-elements:

- `process`
- `osagent`

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="osagent"...>
               [<time-rules ... >
                     <time-rule ... />
                     ...
               </time-rules>]
          </managed-object>
           ...
    </managed-objects>
...
```

Additionally, the `managed-object` sub-element `control-overrides`, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="osagent"...>
                     ...
               [<control-overrides>
                     ...
               </control-overrides>]
          </managed-object>
           ...
    </managed-objects>
...
```

Sub-Element	Description
`process`	Contains the Smart Agent (osagent) Managed Object process configuration information.

```
...
<managed-objects>
      <managed-object ...type="osagent"...>
            <process ...attributes...>
                ...
            </process>
            ...
     </managed-object>
</managed-objects>
...
```

Sub-Element	Description
osagent	Contains the Smart Agent (osagent) location information.

```
...
<managed-objects>
      <managed-object ...type="osagent"...>
            <process ...attributes...>
                ...
            </process>
            <osagent ...attributes.../>
     </managed-object>
</managed-objects>
...
```

process attributes and sub-elements

For a list and description of the process element attributes and sub-elements, go to the following sections:

- process sub-element attributes
- process sub-elements

```
...
<managed-objects>
      <managed-object ...type="osagent"...>
            <process command="fullpath"  directory="fullpathdirectory">
                 [<arguments>
                     [<argument>argument1<argument>
                     [<argument includePlatforms="Windows|
UNIX"argument3<argument>
                     [<argument>argument2<argument>]
                     [...more arguments...]
                 </arguments>]
                 [<library-path>
                     [<directory>path1</directory>]
                     [<directory>path2</directory>]
                     [...more librarypaths...]
                 </library-path>]
                 [<env-vars use-default-env="true|false" use-current-
env="true|false" use-vbroker-env="true|false">
                     [<env-var name="name" value="value"/>]
                     [...more env-vars...]
                 <env-vars>]
            </process>
             ...
      </managed-object>
</managed-objects>
...
```

osagent sub-elements

The following are the `osagent` sub-elements:

Sub-Element	Required	Default	Description
`port`	Yes	n/a	The Smart Agent (osagent) port.
`host`	No	n/a	The physical location of this Smart Agent (osagent), either by IP address or hostname, which AppServer uses to execute the ping action. If not provided, AppServer looks for any Smart Agent (osagent) on the network with the proper port during the ping.
`logdir`	No	`${CONFIG.path}/mos/${MO.name}`	Directory to which you want the Smart Agent (osagent) log file written.

```
...
<managed-objects>
      <managed-object ...type="osagent"...>
            <process attributes>
             ...
            </process>
            <osagent port="port_number"
host="string" logdir="fullpathdirectory"/>
              ...
      </managed-object>
</managed-objects>
...
```

apache-process Managed Object type

This section lists and describes the `apache-process` Managed Object `type` sub-elements and attributes.

Note

For information about the `apache-process` Managed Object attributes, go to the managed-object element attributes section.

Note

When creating an Apache Web Server managed object, do not include spaces in the Managed Object name. If spaces are included in the name, for example, "Apache Managed Object", the Apache Web Server will fail to start.

apache-process sub-elements

The following are the `apache-process` type sub-elements:

Sub-Element	Description
`apache-data`	Contains the `apache-process` Managed Object type-specific configuration information.

```
...
<managed-objects>
      <managed-object ...type="apache-process"...>
            <apache-data ...attributes...>
                ...
            </apache-data>
```

```
                ...
        </managed-object>
    </managed-objects>
    ...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="apache-process"...>
               [<time-rules ... >
                     ...
               </time-rules>]
          </managed-object>
           ...
    </managed-objects>
...
```

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="apache-process"...>
               <apache-data>
                     ...
               </apache-data>
               [<control-overrides>
                     ...
               </control-overrides>]
          </managed-object>
           ...
    </managed-objects>
...
```

apache-data sub-element attributes

The following lists the `apache-data` sub-element attributes:

Attribute	Required	Default	Description
`command`	Yes	`${AGENT.root}/bin/<apache_managedobject_name>` - where the replacement string `${AGENT.root}` represents the root of the host where the assigned Agent resides.	The fully-qualified path to the command used to start the Apache web server Managed Resource.
`httpd-conf`	Yes	`${CONFIG.path}/mos/${MO.name}/conf/httpd.conf` - where the replacement strings: ■ `${CONFIG.path}` represent the fully-qualified path to the `configuration.xml`, ■ `${MO.name}` represent the Managed Object name (per the Managed Object `name` attribute value).	The fully-qualified path to the Apache `httpd.conf` file.
`nt-service`	`Yes` (for Windows platforms)	`false`	This attribute affects how Op-Center interacts with the Apache Managed Object that runs as an Microsoft NT service. If this Managed Object runs as an Microsoft NT service, set to `true`, otherwise, set to false.

```
...
<managed-objects>
      <managed-object ...type="apache-process"...>
            <apache-data command="fullpath" httpd-conf="fullpath" nt-
service="true|false">
                ...
            </apache-data>
             ...
      </managed-object>
</managed-objects>
...
```

apache-data sub-elements

The following are the `apache-data` sub-elements:

Sub-Element	Required	Default	Description
`arguments`	No	n/a	For a description of this element, go to the process sub-elements section.

```
...
<managed-objects>
      <managed-object ...type="apache-process"...>
            <apache-data attributes>
                [<arguments>
                     <argument>argument1</argument>
                    [<argument>argument2</argument>]
```

```
                    [...more arguments...]
                </arguments>]
                ...
            </apache-data>
             ...
        </managed-object>
    </managed-objects>
    ...
```

Sub-Element	Required	Default	Description
env-vars	No	n/a	For a description of this element, go to the process sub-elements section.

```
    ...
    <managed-objects>
        <managed-object ...type="apache-process"...>
            <apache-data attributes>
                [<env-vars>
                    <env-var name="variable1"
    value="value1"/>
                   [<env-var name="variable2"
    value="value2"/>]
                   [...more env-vars...]
                </env-vars>]
                ...
            </apache-data>
             ...
        </managed-object>
    </managed-objects>
    ...
```

Sub-Element	Required	Default	Description
library-path	No	n/a	For UNIX only. For a description of this element, go to the process sub-elements section.

```
    ...
    <managed-objects>
        <managed-object ...type="apache-process"...>
            <apache-data attributes>
                [<library-path>
                    <directory>path1</directory>
                   [<directory>path2</directory>]
                   [...more paths...]
                </library-path>]
                ...
            </apache-data>
             ...
        </managed-object>
    </managed-objects>
    ...
```

Sub-Element	Required	Default	Description
path	No	n/a	For a description of this element, go to the process sub-elements section.

```
    ...
    <managed-objects>
        <managed-object ...type="apache-process"...>
            <apache-data attributes>
```

```
                    [<path>
                        <directory>path1</directory>
                       [<directory>path2</directory>]
                       [...more paths...]
                    </path>]
                    ...
                </apache-data>
                 ...
        </managed-object>
    </managed-objects>
    ...
```

Sub-Element	Required	Default	Description
`stdin path`	No	n/a	For a description of this element, go to the process sub-elements section.
`stdout path`	No	n/a	For a description of this element, go to the process sub-elements section.
`stderr path`	No	n/a	For a description of this element, go to the process sub-elements section.

```
    ...
    <managed-objects>
          <managed-object ...type="apache-process"...>
                <apache-data attributes>
                    <stdin path="path1"/>
                    <stdout path="path2"
     append="true|false"/>
                    <stderr path="path3"
     append="true|false" />
                    ...
                </apache-data>
                 ...
          </managed-object>
    </managed-objects>
    ...
```

ots Managed Object type

This section lists and describes the `ots` Managed Object `type` sub-elements and attributes.

Note

For information about the `ots` Managed Object attributes, go to the managed-object element attributes section.

ots sub-elements

The following are the `ots` type sub-elements:

Sub-Element	Description
`process`	Configuration information specific to an `ots` Managed Object is contained within the `process` sub-element. For a list and description of the `process` element attributes and sub-elements, go to the process sub-element attributes section.

```
    ...
    <managed-objects>
          <managed-object ...type="ots"...>
```

```
                <process command="fullpath">
                   [<arguments>
                       <argument>argument1</argument>
                       [...more arguments...]
                   [<library-path>
                       <directory>path1</directory>
                       [...more paths...]
                   </library-path>
                   [<env-vars use-default-env="true|false"
use-current-env="true|false" use-vbroker-env="true|false"/>
                      [<env-var name="variable1" value="value1"/>]
                      [<env-var name="variable2" value="value2"="/>]
                      [...more env-vars...]
                   </env-vars>]
                   [<stdout path="path1"/>]
                   [<stderr path="path2"/>]
                </process>
                ...
        </managed-object>
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="ots"...>
               [<time-rules ... >
                     ...
               </time-rules>]
          </managed-object>
           ...
    </managed-objects>
...
```

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
<managed-objects>
      <managed-object ...type="ots"...>
               [<control-overrides>
                     ...
               </control-overrides>]
               <process command="fullpath">
                  [<arguments>
                      <argument>argument1</argument>
                      [...more arguments...]
                  [<library-path>
                      <directory>path1</directory>
                      [...more paths...]
                  </library-path>
```

```
                  [<env-vars use-default-env="true|false"
use-current-env="true|false" use-vbroker-env="true|false"/>
                     [<env-var name="variable1" value="value1"/>]
                     [<env-var name="variable2" value="value2"="/>]
                     [...more env-vars...]
                  </env-vars>]
                  [<stdout path="path1"/>]
                  [<stderr path="path2"/>]
               </process>
            ...
      </managed-object>
</managed-objects>
...
```

tibco Managed Object type

This section lists and describes the `tibco` Managed Object `type` sub-elements and attributes.

Note

For information about the `tibco` Managed Object attributes, go to the managed-object element attributes section.

tibco sub-elements

The Configuration information specific to a `tibco` Managed Object is contained within the following sub-elements:

- `tibco-data`
- `process`

Sub-Element	Description
`tibco-data`	Contains the `tibco` Managed Object type Java Messaging Service-specific configuration information.

```
...
<managed-objects>
      <managed-object ...type="tibco"...>
            <tibco-data attributes/>
            <process attributes>
                ...
            </process>
      </managed-object>
      ...
</managed-objects>
...
```

Sub-Element	Description
`process`	Contains the Tibco JMS Managed Object process configuration information.

```
...
<managed-objects>
      <managed-object ...type="tibco"...>
            <tibco-data attributes/>
            <process attributes>
                ...
            </process>
      </managed-object>
```

```
      ...
</managed-objects>
...
```

Additionally, the `managed-object` sub-element `time-rules`, which is not type-specific, can be used to define `cron`-style rules for starting or stopping the Managed Resource processes.

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules. For more information, go to the Managed Object time-rules element section.

```
...
   <managed-objects>
          <managed-object ...type="tibco"...>
               [<time-rules ... >
                     ...
               </time-rules>]
          </managed-object>
           ...
    </managed-objects>
...
```

Additionally, the `managed-object` sub-element `control-overrides`, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
`control-overrides`	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
          <managed-object ...type="tibco"...>
                     ...
               [<control-overrides>
                     ...
               </control-overrides>]
          </managed-object>
           ...
    </managed-objects>
...
```

tibco-data sub-element attributes

The following are the `tibco-data` sub-element attributes:

Attribute	Required	Default	Description
`jms-home`	Yes	`${AGENT.root}jms/tibco` - where the replacement string variable `${AGENT.root}` represents the root of the Agent host to which this `tibco` Managed Object is assigned.	The full path and directory of the Tibco JMS provider root.
`server-url`	Yes	`tcp://`*`localhost`*`:7222`	The address used to programmatically connect to the Tibco JMS provider, given in the form: `tcp://`*`hostname`*`:`*`portnumber`*. **Note:** The *portnumber* must be different from the `partner-url` attribute *`portnumber`*.
`server-name`	Yes	n/a	A unique name that identifies the Tibco JMS server. For a fault tolerant group, both Tibco JMS Managed Objects in the group must have the same `server-name`.
`partner-url`	Yes, if Managed Resource is in a fault tolerant group.	n/a	The address used by both Tibco JMS servers in the same fault tolerant group to communicate with each other, in the format `tcp://`*`localhost`*`:`*`portnumber`* **Note:** The *portnumber* must be different from the `server-url` attribute *`portnumber`*.
`shared-data-storage`	Yes	n/a	Used for persist messages. Fully qualified path and directory where the Tibco server has full read/write access. **Note:** For a fault tolerant group, both Tibco JMS Managed Objects in the group must use the same `shared-data-storage` location.

```
...
<managed-objects>
      <managed-object ...type="tibco"...>
            <tibco-data jms-home="fullpathdirectory"
server-url="tcp://hostname:portnumber1" server-name="string"
 partner-url="tcp://hostname:portnumber2"  shared-data-
storage="fullpathdirectory"/>
            <process attributes>
                ...
            </process>
      </managed-object>
      ...
</managed-objects>
...
```

process sub-element attributes

For a list and description of the process element attributes and sub-elements, go to the process sub-element attributes section.

```
...
<managed-objects>
      <managed-object ...type="tibco"...>
            <tibco-data attributes/>
            <process command="fullpath" directory="fullpathdirectory">
                 <stdout path="path1"/>
                 <stderr path="path2"/>
            </process>
      </managed-object>
      ...
</managed-objects>
...
```

partition Managed Object type

This section lists and describes the partition Managed Object type sub-elements and attributes.

Note

For information about the partition Managed Object attributes, go to the managed-object element attributes section.

partition-process sub-elements

The Configuration information specific to a partition type Managed Object is contained within the following partition sub-elements:

- partition-process
- partition-services
- jmx

Sub-Element	Description
partition-process	Contains the AppServer Partition Managed Object process configuration information.
partition-services	Contains the list of Managed Object references to the AppServer Partition services Managed Objects.
jmx	Contains the information needed to connect to the JMX server.

Additionally, the managed-object sub-element control-overrides, which is not type-specific, can be used to override the default action strategy and/or action parameter(s).

Sub-Element	Description
control-overrides	Contains the action strategy and/or action parameter(s) you specify to override the defaults. For more information, go to the Managed Object control-overrides element section.

```
...
   <managed-objects>
         <managed-object ...type="partition"...>
                  ...
              [<control-overrides>
                  ...
              </control-overrides>]
         </managed-object>
```

```
            ...
    </managed-objects>
...
```

partition-process sub-element attributes

The `partition-process` sub-element inherits all the `process` sub-element attributes. For more information about these attributes, go to the process sub-element attributes section.

partition-process sub-elements

The `partition-process` sub-element inherits all the `process` sub-elements. For more information about these sub-elements, go to the process sub-elements section.

The following are the additional `partition-process` element-specific sub-elements:

Sub-Element	Required	Default	Description
`jpda`	No	n/a	Determines how debugging functions with the Partition Managed Resource.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
            <partition-process command="fullpath">
                [<java-properties>
                   <java-property name="name1"
value="value1"/>
               </java-properties>]
               [<jpda ...attributes.../>]
               [<env-vars use-default-env="true|false"
use-current-env="true|false" use-vbroker-env="true|false"/>
                  [<env-var name="variable1" value="value1"/>]
                  [<env-var name="variable2" value="value2"="/>]
                  [...more env-vars...]
               </env-vars>]
               [<stdout path="path1"/>]
               [<stderr path="path2"/>]
            </partition-process>
            ...
      </managed-object>
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`optimizeit`	No	n/a	Contains all Configuration information for managing AppServer running with Borland Optimizeit ServerTrace or Profiler.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              [<optimizeit ...attributes...>]
                 ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

optimizeit sub-element attributes

The following are the `optimizeit` sub-element attributes:

Attribute	Required	Default	Description
`enable`	Yes	`false`	To enable a AppServer Partition Managed Object to run with Borland Optimizeit ServerTrace or Profiler, set to `true`.
`mode`	Yes - if `enable` is set to `true`.	`servertrace`	When the `enable` element is set to `true`, specifies which Optimizeit product to enable. If you want the AppServer Partition Managed Object to run with Optimizeit Profiler, set to `profiler`.
`sthome`	Yes - if enable is set to `true`.	n/a	Specifies the full path and directory to the Optimizeit ServerTrace or Profiler installation.
`jdkpath`	no	Template JDK path	Used to specify the full path to the JDK home.
`xmlpath`	no	The AppServer Template Optimizeit xml configuration file.	Used to specify the full path to the Optimizeit xml configuration file.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              [<optimizeit enable="false|true"
 mode="servertrace|profiler" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">]
                   ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

optimizeit element sub-elements

The following are the available `optimizeit` element sub-elements.

Sub-Element	Required	Default	Description
`strace`	No	n/a	Used to run AppServer Partition with Borland Optimizeit ServerTrace.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit ...attributes...>
                 [<strace ...attributes...>]
                 ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

Sub-Element	Required	Default	Description
`profiler`	No	n/a	Used to run a AppServer Partition with Borland Optimizeit Profiler.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit ...attributes...>
                 [<profiler ...attributes...>]
                 ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

strace sub-element attributes

The following are the `strace` sub-element attributes:

Attribute	Required	Default	Description
`sthome`	Yes - if the optimizeit attributes: `enable="true"` and `mode="servertrace"`.	n/a	Specifies the full path and directory to the Optimizeit ServerTrace installation.
`jdkpath`	no	Template JDK path	Used to specify the full path to the JDK home.
`xmlpath`	no	The AppServer Template Optimizeit xml configuration file.	Used to specify the full path to the Optimizeit xml configuration file.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="servertrace" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <strace sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 ...
                </strace>
           </partition-process>
      </managed-object>
</managed-objects>
...
```

profiler sub-element attributes

The following are the `profiler` sub-element attributes:

Attribute	Required	Default	Description
`sthome`	Yes - if the optimizeit attributes: `enable="true"` and `mode="profiler"`.	n/a	Specifies the full path and directory to the Optimizeit Profiler installation.
`jdkpath`	no	Template JDK path	Used to specify the full path to the JDK home.
`xmlpath`	no	The AppServer Template Optimizeit xml configuration file.	Used to specify the full path to the Optimizeit xml configuration file.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
```

```
            <optimizeit enable="true"
 mode="profiler" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                <profiler sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                ...
                   ...
                </profiler>
           </partition-process>
      </managed-object>
</managed-objects>
...
```

strace element sub-elements

The following are the available `strace` element sub-elements.

strace element classpath sub-element

For information about this sub-element, go to the java-process sub-elements section.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="servertrace" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <strace ...attributes...>
                    [<classpath>
                       [<classpath-entry>classpath1</classpath-entry>]
                       [<classpath-entry>classpath2</classpath-entry>]
                       [...more classpaths...]
                    </classpath>]
                    ...
                 </strace>
                 ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

strace element bootclasspath sub-element

Sub-Element	Required	Default	Description
`bootclasspath`	Yes	n/a	Used to specify the classpath to the ServerTrace oibcp.jar.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="strace" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <strace ...attributes...>
                    [<bootclasspath>
                      [<classpath-entry>classpath-to-oibcp.jar
</classpath-entry>]
                     </bootclasspath>]
```

```
                    ...
                  </strace>
                  ...
            </partition-process>
      </managed-object>
</managed-objects>
...
```

strace element options sub-element

For information about this sub-element, go to the java-process sub-elements section.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <java-process ...attributes...>
              <optimizeit enable="true"
 mode="servertrace" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <strace ...attributes...>
                    [<options>
                      [<option>option1</option>]
                      [<option>option2</option>]
                      [...more options...]
                    </options>]
                    ...
                 </strace>
            </partition-process>
      </managed-object>
</managed-objects>
...
```

strace element path sub-element

For information about this sub-element, go to the process path sub-element section.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="servertrace" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <strace ...attributes...>
                    [<path>
                      <directory>path1</directory>
                     [<directory>path2</directory>]
                     [...more paths...]
                    </path>]
                    ...
                 </strace>
            </partition-process>
      </managed-object>
</managed-objects>
...
```

profiler element sub-elements

The following are the available `profiler` element sub-elements.

profiler element classpath sub-element

For information about this sub-element, go to the java-process sub-elements section.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="profiler" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <strace ...attributes...>
                    [<classpath>
                       [<classpath-entry>classpath1</classpath-entry>]
                       [<classpath-entry>classpath2</classpath-entry>]
                       [...more classpaths...]
                    </classpath>]
                    ...
                 </profiler>
                 ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

profiler element bootclasspath sub-element

Sub-Element	Required	Default	Description
bootclasspath	Yes	n/a	Used to specify the classpath to the Optimizeit Profiler oibcp.jar.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="profiler" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                 <profiler ...attributes...>
                    [<bootclasspath>
                      [<classpath-entry>classpath-to-oibcp.jar
</classpath-entry>]
                     </bootclasspath>]
                    ...
                 </profiler>
                 ...
           </partition-process>
      </managed-object>
</managed-objects>
...
```

profiler element options sub-element

For information about this sub-element, go to the java-process sub-elements section.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="profiler" sthome="fullpath"
```

```
jdkpath="fullpath" xmlpath="fullpath">
                <profiler ...attributes...>
                   [<options>
                     [<option>option1</option>]
                     [<option>option2</option>]
                     [...more options...]
                   </options>]
                   ...
                </profiler>
           </partition-process>
      </managed-object>
</managed-objects>
...
```

profiler element path sub-element

For information about this sub-element, go to the process path sub-element section.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
           <partition-process ...attributes...>
              <optimizeit enable="true"
 mode="profiler" sthome="fullpath"
jdkpath="fullpath" xmlpath="fullpath">
                <profiler ...attributes...>
                   [<path>
                     <directory>path1</directory>
                    [<directory>path2</directory>]
                    [...more paths...]
                   </path>]
                   ...
                </profiler>
           </partition-process>
      </managed-object>
</managed-objects>
...
```

jpda sub-element attributes

The following are the `jpda` element attributes.

Attribute	Required	Default	Description
`enable-jpda-debug`	No	`false`	To enable JPDA debugging of the applications running in the Partition, set to `True`. When set to `false` (default) JPDA debugging of the applications running in the Partition is disabled.
`jdpa-transport-address`	No	Dynamically allocated by jdpa.	Use to configure the port that the JPDA debugger must use when attaching to this application - using the format `host:port`. If not specified, JPDA will dynamically allocate a port and print the port number. The user must read this printed number and use it when attaching the debugger.
`jdpa-suspend`	No	`true`	If you do not want the Partition's operations suspended until the debugger attaches, set to `false`.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
            <partition-process command="fullpath">
                [<java-properties>
                   <java-property name="name1"
value="value1"/>
               </java-properties>]
               [<jpda enable-jpda-debug="true|false"
jpda-transport-address="host:port"
jpda-suspend="true|false"/>]
               [<env-vars use-default-env="true|false"
use-current-env="true|false" use-vbroker-env="true|false"/>
                  [<env-var name="variable1" value="value1"/>]
                  [<env-var name="variable2" value="value2"="/>]
                  [...more env-vars...]
               </env-vars>]
               [<stdout path="path1"/>]
               [<stderr path="path2"/>]
            </partition-process>
            ...
      </managed-object>
</managed-objects>
...
```

partition-services element sub-elements

The following is the only `partition-services` element sub-element.

Sub-Element	Required	Default	Description
`member`	Yes	n/a	Used to list and reference each of the Partition Managed Object Partition Services.

```
...
<managed-objects>
      <managed-object ...type="partition"...>
            <partition-process attributes>
                  ...
            </partition-process>
            <partition-services>
                  [<member attribute/>]
                  [...more members...]
                  ...
            </partition-services>
             ...
      </managed-object>
</managed-objects>
...
```

member sub-element attribute

The `member` sub-element has only the following attribute:

Attribute	Required	Default	Description
`mo-ref`	Yes	n/a	For each Partition Service of the Partition Managed Resource, the Logical AppServer name of Partition service represented in the following format: `${agentname}/partition_managedobject_name_partitionservice_managedobjectname`. The replacement string `${mo.agent}` is used to represent the Agent to which the Partition Managed Object is assigned, and `${MO.name}` is used to represent the Partition Managed Object `name` attribute value. For example: `${MO.agent}/${MO.name}_jss`

```
...
<managed-objects>
      <managed-object ...type="partition"...>
            <partition-process attributes>
                  ...
            </partition-process>
            <partition-services>
                  <member
mo-ref="agentname/managedobjectname_partitionservicename/>
                  [...more members...]
            </partition-services>
             ...
      </managed-object>
</managed-objects>
...
```

Important

The `mo-ref` naming convention used within a `partition-services` element is strict and you should not modify it.

```
<managed-object [...] type="partition-process">
    <partition-process [...] />
    <partition-services>
        <member mo-ref="${MO.agent}/${MO.name}__interceptor"/>
        <member mo-ref="${MO.agent}/${MO.name}_jdatastore"/>
        <member mo-ref="${MO.agent}/${MO.name}_jss"/>
        <member mo-ref="${MO.agent}/${MO.name}_jts"/>
        <member mo-ref="${MO.agent}/${MO.name}_tomcat4"/>
        <member mo-ref="${MO.agent}/${MO.name}_visiconnect"/>
        <member mo-ref="${MO.agent}/${MO.name}_ejbcontainer"/>
        [<member mo-ref="${MO.agent}/${MO.name}_jms"/>]
    </partition-services>
</managed-object>
```

jmx sub-element attributes

The following are the `jmx` sub-element attributes.

Attribute	Required	Default	Description
`jmxserver`	Yes	`bes`	Used to identify the JMX server type. For AppServer configurations the valid value is `bes`.
`classpath`	Yes	n/a	Used to identify the classpath used to connect to the JMX server.

```
...
<managed-objects>
```

```
        <managed-object ...type="partition"...>
                <partition-process attributes>
                        ...
                </partition-process>
                <partition-services>
                        [<member attribute/>]
                        [...more members...]
                        ...
                </partition-services>
                <jmx attributes />
                 ...
        </managed-object>
  </managed-objects>
  ...
```

Partition Services Managed Object types

For a `partition` type Managed Object, each of the Partition Services is defined within the `configuration.xml` as a Managed Object. The following describes each of the `partition-services` Managed Object types in AppServer.

Partition Service Managed Object type	Description
`ps-interceptor`	Used to represent a AppServer *Partition Lifecycle Interceptor* Partition service. Partition Lifecycle Interceptors allow you to listen for Partition events and respond to them as needed.
`ps-jdatastore`	Used to represent an instance of the JDataStore (Borland's all-Java relational database) Partition service.
`ps-jss`	Used to represent an instance of the Java Session Service (JSS) Partition service.
`ps-jts`	Used to represent an instance of the Java Transaction Service (JTS) Partition service.
`ps-tomcat4`	Used to represent an instance of the Tomcat web container Partition service.
`ps-ejbcontainer`	Used to represent an instance of the EJB container Partition service.
`ps-visiconnect`	Used to represent an instance of the VisiConnect (Borland's Java Connector Architecture (JCA) implementation) Partition service.

47

Managed Object control-overrides element

`control-overrides` is a sub-element of the `managed-object` element. Within the `control-overrides` element, you can specify Managed Object default overrides for the action parameters and/or the action strategies.

control-overrides sub-elements

The following are the `control-overrides` sub-elements:

Element	Description
`parameters`	Contains action parameter settings overrides you have specified.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    ...
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

Elements	Description
`start`	Contains the start action strategy you have specified.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
```

```
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start ...attributes.../>]
                    ...
                </control-overrides>]
                ...
            </managed-object>
             ...
     </managed-objects>
...
```

Elements	Description
stop	Contains the stop action strategy you have specified.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start ...attributes.../>]
                    [<stop ...attributes.../>]
                    ...
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

Element	Description
kill	Contains the kill action strategy you have specified.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start ...attributes.../>]
                    [<stop ...attributes.../>]
                    [<kill ...attributes.../>]
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

Element	Description
ping	Contains the ping action strategy you have specified.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start ...attributes.../>]
```

```
                    [<stop ...attributes.../>]
                    [<kill ...attributes.../>]
                    [<ping ...attributes.../>]
                </control-overrides>]
                ...
            </managed-object>
            ...
        </managed-objects>
...
```

parameters sub-element attributes

The following are the `parameters` sub-element attributes:

Attribute	Required	Default	Description
`local-restart`	No	`false`	Boolean attribute. If set to `true`,when the Hub is unavailable, allows for agent-based restarts of the Managed Resource.
`escalate-stop`	No	`false`	Boolean attribute. If set to `true`, if a stop operation times-out, the stop operation is escalated to a kill.
`ping-policy`	No	`not-when-stopped`	The Managed Object is not pinged when its state is stopped. To ping, regardless of the state of the Managed Object, set to `always`.
`ping-interval`	No	`5` (seconds)	The time, in seconds, between checks during normal monitoring of a Managed Resource. 0 (zero) and positive integers are acceptable values.
`start-timeout`	No	`90` (seconds)	The time, in seconds, for a start attempt to return before the start operation times out. 0 (zero) and positive integers are acceptable values.
`start-ping-interval`	No	Managed Object's `ping-interval` value	The time, in seconds, between checks on the progress of a start operation. 0 (zero) and positive integers are acceptable values.
`start-first-ping-delay`	No	Managed Object's `first-ping-delay` value	The time, in seconds, before BAS starts to check on the progress of a start operation. 0 (zero) and positive integers are acceptable values.
`stop-timeout`	No	`90` (seconds)	The time, in seconds, for a stop attempt to return before the stop operation times out. 0 (zero) and positive integers are acceptable values.

Attribute	Required	Default	Description
stop-retry-interval	No	blank - never retry	The time, in seconds, between retry attempts for stop operations once the previous operation has timed-out. 0 (zero) and positive integers are acceptable values.
stop-ping-interval	No	Managed Object's ping-interval value	The time, in seconds, between checks on the progress of a stop operation. 0 (zero) and positive integers are acceptable values.
kill-timeout	No	90 (seconds)	The time, in seconds, for a kill attempt to return before the kill operation times out. 0 (zero) and positive integers are acceptable values.
kill-retry-interval	No	blank - never retry	The time, in seconds, between retry attempts for kill operations once the previous operation has timed-out. 0 (zero) and positive integers are acceptable values.
kill-ping-interval	No	Managed Object's ping-interval value	The time, in seconds, between checks on the progress of a kill operation. 0 (zero) and positive integers are acceptable values.
kill-first-ping-interval	No	Managed Object's first-ping-interval value	The time, in seconds, before BAS starts to check on the progress of a stop operation. 0 (zero) and positive integers are acceptable values.

```
...
   <managed-objects>
         <managed-object ...attributes...>
               ...
               [<control-overrides>
                   [<parameters local-restart="true|false"
escalate-stop="true|false" ping-policy="not-when-stopped|always"
ping-retry-interval="seconds" start-timeout="seconds"
start-ping-interval="seconds" start-first-ping-interval=="seconds"
stop-timeout="seconds"
stop-retry-interval="seconds" stop-ping-interval="seconds"
kill-timeout="seconds" kill-retry-interval="seconds"
kill-ping-interval="seconds"
kill-first-ping-interval="seconds"/>]
                   ...
               </control-overrides>]
               ...
         </managed-object>
          ...
    </managed-objects>
...
```

start sub-element attribute

The following is the only `start` sub-element attribute:

Attribute	Required	Default	Description
`strategy`	Yes	Default is Managed Object type-specific. For more information, go to the start sub-element strategy attribute section.	Specifies the strategy BAS uses for the Managed Object start action.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start strategy="None|managedobjecttypestrategy"/>]
                    [<stop ...attributes.../>]
                    [<kill ...attributes.../>]
                    [<ping ...attributes.../>]
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

stop sub-element attribute

The following is the only `stop` sub-element attribute:

Attribute	Required	Default	Description
`strategy`	Yes	Default is Managed Object type-specific. For more information, go to the stop sub-element strategy attribute section.	Specifies the strategy BAS uses for the Managed Object stop action.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start strategy="None|managedobjecttypestrategy"/>]
                    [<stop strategy="None|managedobjecttypestrategy"/>]
                    [<kill ...attributes.../>]
                    [<ping ...attributes.../>]
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

kill sub-element attribute

The following is the only `kill` sub-element attribute:

Attribute	Required	Default	Description
`strategy`	Yes	Default is Managed Object type-specific. For more information, go to the kill sub-element strategy attribute section.	Specifies the strategy BAS uses for the Managed Object kill action.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start strategy="None|managedobjecttypestrategy"/>]
               [<stop strategy="None|managedobjecttypestrategy"/>]
                    [<kill strategy="None|managedobjecttypestrategy"/>]
                    [<ping ...attributes.../>]
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

ping sub-element attribute

The following is the only `ping` sub-element attribute:

Attribute	Required	Default	Description
`strategy`	Yes	Default is Managed Object type-specific. For more information, go to the ping sub-element strategy attribute section.	Specifies the strategy BAS uses for the Managed Object ping action.

```
...
   <managed-objects>
          <managed-object ...attributes...>
                ...
                [<control-overrides>
                    [<parameters ...attributes.../>]
                    [<start strategy="None|managedobjecttypestrategy"/>]
               [<stop strategy="None|managedobjecttypestrategy"/>]
                    [<kill strategy="None|managedobjecttypestrategy"/>]
                    [<ping strategy="None|managedobjecttypestrategy"/>]
                </control-overrides>]
                ...
          </managed-object>
           ...
    </managed-objects>
...
```

start sub-element strategy attribute

The following are the `start` sub-element `strategy` attribute available values for each Managed Object type:

Managed Object `type`	**Default Start Strategy**	**Description**
`process`	`run-process`	■ `run-process` - executes a command to start the process. ■ `null` - BAS does not execute a start action for the Managed Object.
`custom-javascript`	`jscript-control`	■ `jscript-control` - launches a JavaScript that carries out the start action. ■ `null` - BAS does not execute a start action for the Managed Object.
`custom-executable`	`run-custom-executable`	■ `run-custom-executable` - runs a process that carries out the start action. ■ `null` - BAS does not execute a start action for the Managed Object.
`ordered-group`	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
`redundancy-group`	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
`state-proxy`	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
`osagent`	`run-process`	■ `run-process` - executes a command to start the process. ■ `null` - BAS does not execute a start action for the Managed Object.
`apache-process`	`run-apache-process`	■ `run-apache-process` - executes a command to start the Apache process. ■ `null` - BAS does not execute a start action for the Managed Object.
`ots`	`run-process`	■ `run-process` - executes a command to start the process. ■ `null` - BAS does not execute a start action for the Managed Object.
`tibco`	`run-process`	■ `run-process` - executes a command to start the process. ■ `null` - BAS does not execute a start action for the Managed Object.
`partition`	`start-partition`	■ `start-partition` - Starts the Partition Managed Resource. ■ `always-success` - internal use only. ■ `null` - BAS does not execute a start action for the Managed Object.

stop sub-element strategy attribute

The following are the `stop` sub-element `strategy` attribute available values for each Managed Object type:

Managed Object type	Default Stop Strategy	Description
`ordered-group`	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
`redundancy-group`	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
`state-proxy`	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
`process`	`stop-process-term-signal`	UNIX: For all strategies, sends the process a `sigterm` signal. Windows: ■ `stop-process-term-signal` - sends a `Ctrl+c` event to the process. ■ `stop-gui-process` - sends a `WM_CLOSE` message to all windows owned by the process. ■ `stop-bes-launcher` - internal use only. ■ `stop-process-windows-c-exit` - causes the process to call the C runtime `exit` function. ■ `stop-process-windows-exit-process` - causes the process to call the Windows `ExitProcess` function. ■ `kill-signal` - performs a "`kill -9`" or its platform-specific functional equivalent on the process. ■ `null` - BAS does not execute a stop action for the Managed Object.
`custom-javascript`	`jscript-control`	■ `jscript-control` - launches a JavaScript that carries out the stop action. ■ `null` - BAS does not execute a stop action for the Managed Object.
`custom-executable`	`run-custom-executable`	■ `run-custom-executable` - runs a process that carries out the stop action. ■ `null` - BAS does not execute a stop action for the Managed Object.
`osagent`	`stop-process-windows-c-exit`	■ `stop-process-windows-c-exit` - UNIX: sends the process a `sigterm` signal. Windows: causes the process to call the C runtime `exit` function. ■ `kill-signal` - performs a "`kill -9`" or its platform-specific functional equivalent on the process. ■ `null` - BAS does not execute a stop action for the Managed Object.
`apache-process`	`shutdown-apache-process`	■ `shutdown-apache-process` - executes the command to shutdown the Apache process. ■ `null` - BAS does not execute a stop action for the Managed Object.

Managed Object type	Default Stop Strategy	Description
ots	stop-process-term-signal	■ stop-process-term-signal - UNIX: sends the process a sigterm signal. Windows: sends a Ctrl+c event to the process. ■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a stop action for the Managed Object.
tibco	stop-process-term-signal	■ stop-process-term-signal - UNIX: sends the process a sigterm signal. Windows: sends a Ctrl+c event to the process. ■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a stop action for the Managed Object.
partition	stop-partition	■ stop-partition - stops the Partition Managed Resource. ■ null - BAS does not execute a stop action for the Managed Object.

ping sub-element strategy attribute

The following are the ping sub-element strategy attribute available values for each Managed Object type:

Managed Object type	Default Ping Strategy	Description
process	check-pid	■ check-pid - checks whether the process ID exists in the system process table and verifies its start time. ■ jscript-ping - launches a JavaScript that carries out the ping action. The JavaScript must return 0 for success or 1 for failure. ■ ping-custom-executable - runs a process that carries out the ping action. The process must return 0 for success or 1 for failure.
ordered-group	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
redundancy-group	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
state-proxy	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
custom-javascript	jscript-ping	Launches a JavaScript that carries out the ping action. The JavaScript must return 0 for success or 1 for failure.
custom-executable	ping-custom-executable	Runs a process that carries out the ping action. The process must return 0 for success or 1 for failure.
osagent	check-osagent	Based on the property entries in the configuration.xml for the osagent Managed Object, checks whether the process ID exists in the system process table and/or checks if the osagent Managed Resource is available.

Managed Object type	Default Ping Strategy	Description
apache-process	ping-apache-process	Checks for the process ID found in the Apache's httpd.pid file in the process table.
ots	check-pid	Checks whether the process ID exists in the system process table and verifies its start time.
tibco	ping-tibco-server	Checks whether the process ID exists in the system process table, then checks if the Tibco Managed Resource is available.
partition	ping-partition	Checks whether the process ID exists in the system process table, then requests the Partition state from the Managed Resource.

kill sub-element strategy attribute

The following are the kill sub-element strategy attribute available values for each Managed Object type:

Managed Object type	Default Kill Strategy	Description
ordered-group	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
redundancy-group	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
state-proxy	n/a	This type of Managed Object is an organizational construct only - action strategies do not apply.
process	kill-signal	■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a kill action for the Managed Object.
custom-javascript	null	■ jscript-control - launches a JavaScript that carries out the kill action. ■ null - BAS does not execute a kill action for the Managed Object.
custom-executable	null	■ run-custom-executable - runs a process that carries out the kill action. ■ null - BAS does not execute a kill action for the Managed Object.
osagent	kill-signal	■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a kill action for the Managed Object.
apache-process	kill-signal	■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a kill action for the Managed Object.
ots	kill-signal	■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a kill action for the Managed Object.

Managed Object type	Default Kill Strategy	Description
tibco	kill-signal	■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a kill action for the Managed Object.
partition	kill-signal	■ kill-signal - performs a "kill -9" or its platform-specific functional equivalent on the process. ■ null - BAS does not execute a kill action for the Managed Object.

48

Managed Object time-rules element

`time-rules` is a sub-element of the `managed-object` element. Within the `time-rules` element, you can add one or more configured `time-rule` sub-elements that contain `cron`-style rules defining when a Managed Object's processes should be started or stopped.

time-rules Sub-element

Sub-Element	Description
`time-rules`	Contains an attribute describing the base state of the Managed Object before running the configured timer rules.

```
...
   <managed-objects>
          <managed-object ...attributes...>
           ...
               [<time-rules attributes >
                     ...
               </time-rules>]
                     ...
          </managed-object>
           ...
    </managed-objects>
...
```

time-rules Sub-element attributes

The following are the `time-rules` sub-element's attributes.

Attribute	Required	Default	Description
`base-state`	Yes	`Running`	The base operational state for the Managed Object. If you want to control the MO to only run during a configured time slot, set the `base-state` to `Stopped`, then configure one or more `time-rule` sub-elements containing instructions for the time slot(s) during which the MO should run. Conversely, if you want to control the MO to stop during a configured time slot, set the `base-state` to `Running`, then configure one or more `time-rule` sub-elements containing instructions for the time slot(s) during which the MO should stop.

```
...
<managed-objects>
      <managed-object ...attributes...>
                ...
           [<time-rules base-state="Running|Stopped">
                          ...
           </time-rules>]
                ...
      </managed-object>
       ...
</managed-objects>
...
```

time-rule Sub-element

By default, a Configuration's `time-rules` sub-elements are evaluated in the order in which they are created, starting with the first entered `time-rule`. You can rearrange the evaluation order any time after a `time-rule` is created. The `time-rule` actions are based on the system clock of the Hub's host.

The `time-rule` attributes use a `cron`-style specification to indicate start and stop times, as well as duration to indicate how long after the specified time the `time-rule` continues to apply. For additional information about valid `cron` syntax, please refer to your UNIX system's `crontab` man page.

Note

Time rules specified for an MO that is monitored but not managed are ignored.

Sub-Element	Description
`time-rule`	Contains attributes that define the rules for a specific timed event for an MO.

```
...
   <managed-objects>
          <managed-object ...attributes...>
           ...
               <time-rules ... >
                    [<time-rule attributes>
                           ...
                    </time-rule>]
               </time-rules>
                    ...
          </managed-object>
           ...
    </managed-objects>
...
```

time-rule Sub-element attributes

The following are the `time-rule` sub-element attributes.

Attribute	Required	Default	Description
`rule-name`	No	Automatically generates a rule name for display in the Management Console Time Rules tab based on configured time rules (for example, "Stop for 1 minute, at midnight, every day"). The automatically-generated rule name is not represented in `configuration.xml`.	Assigns a customized name to the rule for display in the Management Console Time Rules tab.
`type`	Yes	`standard-time-rule`	The type of time rule. The only currently-supported type is `standard-time-rule`, which is the default.
`state`	Yes	■ For Scheduled Tasks time rules: `Running` ■ For all other time rules: `Stopped`	The MOs's run state in effect at the time(s) configured for the rule. For a Scheduled Task time rule, the only valid `state` is `Running`. For a time rule added to an MO within a Configuration's hierarchy, the valid values are `Stopped` or `Running`.
`hour`	Yes	`0`	The `cron`-style specification for the start time hour. Hours are expressed as integer values from 0 to 23. Entries that span a range of these values or present a comma-delimited list of specific hours are all valid. For example, `12-14` stands for 12 p.m., 1 p.m., and 2 p.m, while `12,14,15` stands for 12 p.m., 2 p.m., and 3 p.m. An asterisk (`*`) is a valid wildcard that means any hour. The default value for `hour` is `0` (12 a.m.).
`minute`	Yes	`0`	The `cron`-style specification for the start time minute. Minutes are expressed as integer values from 0 to 59. Entries that span a range of these values or present a comma-delimited list of specific minutes are all valid. For example, `0-29` stands for the first 30 minutes of the hour. An asterisk (`*`) is a valid wildcard `*` that means every minute of the hour. The default value for `minute` is `0`, the first minute of the hour.
`day`	Yes	`*`	If `day-of-week` is not specified, the `cron`-style specification for the day of the month to run the rule. Days are expressed as integer values from `1` to `31`. Entries that span a range of values, or a comma-delimited list of values are valid. For example, an entry of `1,15` means that the rule only runs on the 1st and 15th day of the month. The default is an asterisk (`*`), which means the rule runs every day of the month.

Attribute	Required	Default	Description
day-of-week	Yes	*	The cron-style specification for the day(s) of the week to run the rule. Integer values from 0 to 6 are valid. Entries that span a range of values, or a comma-delimited list of values are valid. For example, 1-5 means the rule runs only on weekdays (Monday through Friday), while 0,6 means the rule runs only on weekends (Saturday and Sunday). The default is an asterisk (*), which means that the rule runs every day. Specific day values are as follows: ■ 0=Sunday ■ 1=Monday ■ 2=Tuesday ■ 3=Wednesday ■ 4=Thursday ■ 5=Friday ■ 6=Saturday
month	Yes	*	The cron-style specification for the month(s) to run the rule. Entries that span a range of values, or a comma-delimited list of values are valid. For example, a month entry of 1,3,5 means that the rule only runs at the appointed time in the months of January, March, and May. The default is an asterisk (*) which means the rule runs every month. Specific month values are as follows: ■ 1=January ■ 2=February ■ 3=March ■ 4=April ■ 5=May ■ 6=June ■ 7=July ■ 8=August ■ 9=September ■ 10=October ■ 11=November ■ 12=December
duration	No	1	Specifies the number of duration-unit iterations for the rule. For example, the default value of 1 when combined with the default duration-unit of minute means that the rule applies for 1 minute at the scheduled time.
duration-unit	No	minute	The duration measurement unit for applying the rule. Valid values are minute, hour, day, or month.

```
...
<managed-objects>
          <managed-object ...attributes...>
                ...
                 <time-rules ... >
                     [<time-rule rule-name="<name>"
```

```
                    type="standard-time-rule"
                    state="Running|Stopped"
                    hour="*|0-23"
                    minute="*|0-59"
                    day="*|1-31"
                    month="*|1-12"
                    day-of-week="*|0-6"
                    duration-unit="minute|hour|day|month">
                    </time-rule>
                        ...

                </time-rules>
                ...
            </managed-object>
        ...
</managed-objects>
...
```

time-rule Examples

The following examples demonstrate the use of the `time-rule` sub-element to add timed rules to an MO:

- The following `time-rule` stops the MO for 3 minutes every day at 10:30 p.m.:

  ```
  <time-rule state="Stopped" minute="30" hour="22" duration="3" duration-
  unit="minute"/>
  ```

- The following `time-rule` stops the MO at all times. This is a useful first rule in a multi-rules set if you want to specify when the MO should run in subsequent `time-rule` rules.

  ```
  <time-rule state="Stopped" minute="0" hour="0" duration="24" duration-
  unit="hour"/>
  ```

- The following `time-rule` runs the MO from 9a.m. to 5p.m. on weekdays (Monday through Friday):

  ```
  <time-rule state="Running" minute="0" hour="9" day-of-week="1-5" duration="8"
  duration-unit="hour"/>
  ```

- The following `time-rule` means the MO does not run on the January 1st New Year's Day holiday:

  ```
  <time-rule state="Stopped" minute="0" hour="0" day="1" month="1"
  duration="24" duration-unit="hour"/>
  ```

- The following `time-rule` means that the MO does not run during the December 24th through December 31st winter holiday:

  ```
  <time-rule state="Stopped" minute="0" hour="0" day="24-31" month="12"
  duration="24" duration-unit="24"/>
  ```

49

Working with Configuration properties

Borland AppServer (BAS) ships with a selection of built-in, pre-defined properties whose values are set at the time of installation. Additionally, you can define sets of Configuration-wide properties within the `configuration.xml` file. Configuration properties are used as a string-substitution mechanism within the `configuration.xml` file.

properties element

You define Configuration properties within the `configuration` element, `properties` sub-element. Multiple `properties` sub-elements can appear in the same `configuration.xml` file.

```
<configuration>
    ...
    <properties>
        <property ...attributes.../>
        [<property ...attributes.../>]
        [...more properties...]
    </properties>
    <properties>
        <property ...attributes.../>
        [<property ...attributes.../>]
        [...more properties...]
    </properties>
    ...
<configuration>
```

Defining a property

Each Configuration-wide property must appear within a `property` sub-element within the `properties` element. This is done using the syntax:

```
<property name="name" value="value"/>
```

where *name* is the name of the property you are defining, and *value* is its value. The quotes ("") are required.

Referencing a defined property

Once a property is defined, it can be referenced elsewhere in the Configuration using the string-replacement syntax, for example `${property-name}`, where *property-name* is the name of a previously-defined property. The Hub looks up the property's name and substitutes its value in place of the replacement-string, such as `${property-name}`.

For example, if you define

```
<properties>
    <property name="install-root" value="c:/apache" />
    <property name="apache-command" value="apachect1" />
</properties>
```

you can then reference those properties within an appropriate Managed Object block, such as

```
<managed-object type="process" />
    <process command="${install-root}/bin/apache2/${apache-command}">
```

The Hub interprets the `command` attribute of `<process>` as

```
command="c:/apache/bin/apache2/apachect1"
```

by substituting the values of the `install-root` and `apache-command` properties in the appropriate places.

Defining properties in terms of other properties

Configuration-wide property values can also be defined in terms of other defined properties. For example:

```
<properties>
    <property name="appHome" value="c:/windows" />
    <property name="appDir" value="system32" />
    <property name="theApp" value="${appHome}/${appDir}/notepad" />
</properties>
```

Wherever you use the `theApp` property in the Configuration, it is processed by the Hub as: `c:/windows/system32/notepad`.

Note

Properties and macros cannot be nested. For example, `${appHome${appDir}}` is illegal.

Hub and Agent properties

Listed below are some of the Hub and Agent properties:

Property	Description
agent.heartbeat.startup.delay.max	Specifies the maximum time an agent will wait for initial states of its managed objects to be calculated before first heartbeating the hub. Defaults to 20 (seconds).
hub.heartbeat.misses.allowed	On a hub, specifies how many heartbeat intervals can go by before the hub considers an agent to be down. Defaults to 2, which means that 2 missed heartbeat is OK but the hub will consider the agent to be down if a third heartbeat is missed. By default, the heartbeat rate is 30 seconds.

Property	Description
agent.to-hub.ior.file	If set, a slave agent uses the IOR found in the file indicated by this property (the "hub internal IOR") for contacting the hub. The "hub internal IOR" is registered with the Smart agent with rep-id IDL:borland.com/Enterprise/Management/Internal/HubInternal:1.0 and object name <hub-name>_InternalHub
agent.output.iors	If this property is set to true, the scu will output its management IORs to files named <interface-name>.ior in the scu's working directory (by default <install-dir>/var/domains/base).

Built-in, pre-defined properties

BAS ships with built-in properties for which the values are set during installation. You can use these properties within your Configurations including within the Configuration-wide properties that you define.

For example, instead of using:

```
<property name="theApp" value="c:/windows/system32/notepad" />
```

you could use an AGENT property instead:

```
<property name="theApp" value="${AGENT.env.windows.SystemRoot}/system32/
notepad" />
```

In a distributed installation, where the Hub and slave Agents are not on the same host, the built-in properties can be used to communicate information known to the Agent but not its remote Hub, such as the host name the slave Agent resides on.

Note

The AGENT built-in properties are Java system properties, and are found in the following files: `<install_dir>/bin/scu.config`, `<install_dir>/var/domains/base/adm/properties/agent.config` and `agent.properties`. These Java system properties are all available to use as Configuration properties values.

Important

Built-in properties from releases prior to 6.5 are still valid, but are deprecated as of this release. The new and deprecated forms of the built-in properties are listed in the following tables.

Important

Although built-in Configuration properties can be used for substituting element values and attributes, the `AGENT`, `MO`, and `%platform%` properties cannot be used as values for the `name` or `agent` attributes.

Note

You cannot make any changes to the following built-in properties.

Property	Deprecated as of 6.5	Description
`HUB.name`	`hub.name`	The name of the "Master Hub" or "Local Hub", which by default is the name of the host on which you install the Hub.

Property	Deprecated as of 6.5	Description
`AGENT.install.root`	`agent.root`	The root directory of the BAS Hub Agent installation.
`AGENT.instance.root`	`agent.instance.root`	The directory of the Agent's domain. For example: `${AGENT.root}/var/domains/<domain-name>`

Property	Deprecated as of 6.5	Description
AGENT.host	agent.host	The host name of the physical resource on which the Agent is installed.
AGENT.address	n/a	The IP address of the Agent's host.
AGENT.name	agent.name	The name of the Agent. The default AGENT.name value is the host's name.
AGENT.default.smartagent.port	agent.default.smartagent.port	The port on the local subnet on which the Agent's osagent is active.
AGENT.default.smartagent.addr	agent.default.smartagent.addr	The IP address or host name of the host running the Agent's osagent.
AGENT.env.windows.SystemRoot	agent.env.windows.SystemRoot	The value of the Windows SYSTEMROOT environment variable. **Windows only.**
AGENT.env.windows.userprofile	agent.env.windows.userprofile	The value of the Windows USERPROFILE environment variable. **Windows only.**
AGENT.env.unix.DISPLAY	agent.env.unix.DISPLAY	The value of the DISPLAY environment variable. **UNIX only.**
AGENT.env.unix.TZ	agent.env.unix.TZ	The value of the TZ environment variable. **UNIX only.**
AGENT.env.unix.LANG	agent.env.unix.LANG	The value of the LANG environment variable. **UNIX only.**
AGENT.env.unix.LOCPATH	agent.env.unix.LOCPATH	The value of the LOCPATH environment variable. **UNIX only.**
AGENT.env.unix.NLSPATH	agent.env.unix.NLSPATH	The value of the NLSPATH environment variable. **UNIX only.**

Property	Deprecated as of 6.5	Description
CONFIG.name	config.name	The name of the Configuration.
CONFIG.path	config.path	The path to the configuration.xml file, for example, ${AGENT.root}/var/domains/base/configurations/${CONFIG.name}. Valid on the Hub and Agent.

Property	Deprecated as of 6.5	Description
MO.name	mo.name	The Managed Object name attribute value. This is valid only within the context of each Managed Object.
MO.agent	mo.agent	The Managed Object agent attribute value. This is valid only within the context of each Managed Object.

%platform% property

A special property called %platform% is used for platform checks. Unlike other properties which can be any property defined on the system, the %platform% property is natively defined. This property resolves to the following values:

- SunOS
- WinNT
- Win2000
- WinXP

- Win2003
- MacOS
- Linux
- AIX
- HPUX
- Windows
- UNIX

Using Property Expressions

There may be times when a property's value might depend on an external condition. For example, you might want to define the command attribute of a process, but the command used to start that process might vary depending on the operating system hosting it. BAS provides a mechanism to perform these types of checks called *Property Expressions*.

Property Expressions are used in the `value` attribute of the `property` element. Using a Property Expression, you can set the value of a property based on conditions you define. This provides a quick way to make many aspects of your Configurations more portable. Property expressions come in two varieties:

- `if` statements
- `switch` statements

if statements

You can use an `if` statement as a substitute for a property's `value` attribute. The basic syntax is:

```
if(condition) {
    "result"
    }
    [else {
    "result"
    }]
```

where *condition* is a comparison computed with a valid operator that evaluates either `true` or `false`, and *result* is the value of the property if the condition evaluates to `true`. If the condition evaluates to `false`, another result in this eventuality can be provided by using the `else` operator.

if statement

Conditions used in `if` statements are exclusively used to evaluate the value of properties using operators. The valid operators are:

- = (equals)
- ! (not equals)
- # (has or contains)

There is also a unary operator, `?`, which only tests for the existence of a property key and does not require a value for comparison.

You can also use multiple conditions in the same check using the `&` (AND) operator and the `|` (OR) operator.

For example, the condition

```
${AGENT.default.smartagent.port} = "14000"
```

would be `true` if the Agent's Smart Agent port was `14000` and `false` otherwise.

Note

All values must be in quotes ("").

An `if` statement using this condition could look like the following:

```
if(${AGENT.default.smartagent.port} = "14000") {
    "default_port"
    } else {
    "user_defined_port"
}
```

The `if` statement can then be substituted for the `value` attribute of a property using proper XML syntax:

```
<property name="isDefault" value=
    "if(${AGENT.default.smartagent.port} = &quot;14000&quot;) {
        &quot;default_port&quot;
        } else {
        &quot;user_defined_port&quot;
    }" />
```

Warning

Because `if` statements appear only in XML files, they must use proper XML grammar in order to be parsed properly. This means that quotation marks cannot be nested. Instead of using quotation marks around your condition values and results, you must use `"`.

Basic if statement examples

– A statement to check whether or not the property `agent.name` exists:

```
if(${AGENT.name} ?) {
    "named_agent"
    } else {
    "no_name"
    }
```

The same statement as a property value:

```
<property name="isNamed" value=
    "if(${AGENT.name} ?) {
        &quot;named_agent&quot;
        } else {
        &quot;no_name&quot;
        }" />
```

– A statement to check if the `agent.host` property contains the string "nett2":

```
if(${AGENT.name} # "nett2") {
    "bad_name"
    } else {
    "ok"
    }
```

The same statement as a property value:

```
<property name="nameCheck" value=
    "if(${AGENT.name} # &quot;nett2&quot;) {
        &quot;bad_name&quot;
        } else {
        &quot;ok&quot;
        }" />
```

– A statement that checks whether the platform is Windows NT and the `AGENT.host` property exists, illustrating the use of the `&` (AND) operator and the `%platform%` property.

```
if(${%platform%} = "WinNT" & ${AGENT.host} ?) {
    "ready"
    }
```

The same statement as a property value:

```
<property name="sysCheck" value=
    "if(${%platform%} = &quot;WinNT&quot; & ${AGENT.host} ?) {
        &quot;ready&quot;
        }" />
```

Nested if statements

You can nest `if` statements by replacing *result* in the default syntax with another `if` statement. The finally-resolved `if` statement's result becomes the property's value. The syntax is:

```
if(condition) {
    if(condition) {
        "result"
        }
        [else {
        "result"
        }]
    [else {
    "result"
    }]
```

You can continue nested substitutions of `if` statements for any other *result* field provided they can be resolved.

Example of a nested if statement

A statement that checks whether the platform is UNIX (returning "other_platform" if not), returning a command that depends on whether the user is named "root":

```
if ($%platform% = "Unix") {
    if(${user.name} = "root" {
        "rm -rf *"
        } else {
        "rm -rf ." }
    else {
    "other_platform"
    }
```

The same statement as a property value:

```
<property name="unixCommand" value=
    "if(${user.name} = &quot;root&quot; {
        &quot;rm -rf *&quot;
        } else {
        &quot;rm -rf .&quot; }
    else {
    &quot;other_platform&quot;
    }" />
```

switch statements

You can also use a `switch` statement as a substitute for a property's *value* attribute. This allows you to evaluate several cases and provide unique results for each, as well as a default value should none of your cases evaluate. The syntax is:

```
switch(${property-name}) {
    case "value" : "result";
    [case "value" : "result";
    ...]
    [default : "default-result"]
}
```

where *value* is a possible value for the property called *property-name*, and *result* is the value of the property being defined using the `switch` statement.

For example, the following `switch` statement checks for a variety of operating systems in order to set the value for a command:

```
switch (${%platform%}) {
    case "WinNT" : "dd";
    case "WinXP" : "dd -x";
    case "Win2000" : "";
    default : "dd1"
}
```

When used as a property value, it would appear as:

```
<property name="command" value=
    "switch (${%platform%}) {
        case &quot;WinNT&quot; : &quot;dd&quot;;
        case &quot;WinXP&quot; : &quot;dd -x&quot;;
        case &quot;Win2000&quot; : &quot;&quot;;
        default : &quot;dd1&quot;
    }" />
```

50

Borland Management Shell (BMSH) API Specification

This document is the API specification for the Borland Management Shell.

There are four kinds of BMSH objects:

- Pre-instantiated objects
  - fso
  - bes
- User-instantiated objects
  - Hub
  - PropertyContext
  - HttpPage
- Static objects
  - BmsExec
  - Dom4j
  - BmsState
- Factory objects - objects that are created by other objects (usually the pre-instantiated objects)
  - PropertyFileSO
  - TextStream

fso

The `fso` object has methods and properties for interacting with the file system. It has convenience methods for getting and setting values in property files (.properties). It is pre-instantiated as variable "fso" in the BMSH..

Property Summary

String classpath [get]

Description:
Returns the classpath

String jars [get]

Description:
Returns the jars in the classpath

String pwd [get]

Description:
Returns the current working directory.

Method Summary

void copyFile(String sourceFile, String destinationFile, boolean overwrite)

Description:
Copies a file from source to destination, overwrites destination if desired.

Parameters:
sourceFile - the file to copy
destinationFile - the destination file
overwrite - if `true`, overwrite destination file if it exists

Returns:
none

Boolean createArchive(String arFile, String archiveAs, String fileToArchive)

Description:
Create an archive with name `arFile`. Any existing archive file with that name will be and replaced.

Parameters:
arFile - name of the destination archive file created. This file will contain `fileToArchive` as an entry with name `archiveAs`. A manifest will be created. This routine does not update archive files. It is limited to archives with one file, eg. .dar files.
archiveAs - the path of the file as it will be archived.
fileToArchive - the path to the file to be archived.

Returns:
true if successfull else false

void createFolder(String folderName)

Description:
Creates a directory path. Will create full path. Uses java.io.File(folderName).mkdirs()

Parameters:
folderName - the path of the folder to be created.

Returns:
none.

TextStream createTextFile(String filename, boolean overwrite)

Description:
Creates a text file and returns a TextStream object so the file can be written to. CAUTION: the file is open upon return.

Parameters:
filename - the name of the text file
overwrite - if `true`, overwrite the existing text file

Returns:
TextStream interface

boolean deleteFile(String filename)

Description:
Delete a file.

Parameters:
filename - the file to delete

Returns:
`true` if successful

boolean deleteFolder(String folderName)

Description:
Deletes a directory recursively.

Parameters:
folderName - the path of the folder to be deleted (recursively).

Returns:
true if successful, false otherwise.

boolean exists(String filename)

Description:
Returns whether or not the file exists.

Parameters:
filename - file to check

Returns:
`true` if the file exists.

Boolean extractFile(String jarFile, String jarEntry, String destDir, String usePath)

Description:
Extract a file from a jar (or zip) file

Parameters:
jarFile - name of the jar file containing the file to extract
jarEntry - the name of the file in jarFile to extract. Important the file seperators and case of `jarEntry` must be exactly as the in the archive in order for the entry to be found.
destDir - the name of the directory the jarFile will be extracted to. The jarEntry path will be rooted here when extracted. If you want to extract to a different location set usePath=="useDestDir"
usePath - - selects where to extract the file to. If usePath == "useDestDir" the file's is extracted to `destDir`. if usePath == "useZipPath" or undefined the file's path is same as in the jarFile rooted at `destDir`.

Returns:
true if successfull else false

int fileLength(String filename)

Description:
Return the length of the file specified in the argument

Parameters:
filename - name of file to find the length of

Returns:
-1 if file cannot be found else integer >= 0 length of file

String[] files(String dir)

Description:
Lists the files in a directory. The list is alphabetized. To filter files or recursively find files use `fso.findFiles()`

Parameters:
dir - the directory whose files are to be listed.

Returns:
a String[] of the files.

String[] findFiles(String dir, String matchName, boolean recurse)

Description:
Searches for files and directories that match the mathchName parameter. The search begins in the dir parameter and will recurse if the recurse parameter is true. The search is not case sensitive. It is not optimized, searching a multi-gigabyte disk will take awhile. The list is alphabetized. Each file returned has a full path.

Parameters:
dir - the directory whose files are to be listed. Recursion traverses subfolders starting here.
matchName - the string to match file names against, can include wild cards.
recurse - if true will do a recursive search into subdirectories.

Returns:
a String[] of the files and directories

String[] folders(String dir)

Description:
Lists the subdirectories of a directory. The list is alphabetized.

Parameters:
dir - the root directory whose child folders are to be listed.

Returns:
a String[] of the child folders.

String getBaseName(String path) (obsolete)

Description:
Returns the base file name given a full path. Use LiveConnect:`java.io.File(path).getName();`

Parameters:
path - the full path to a file

Returns:
the base file name

String getProperty(String propFile, String propName)

Description:
Gets the specified property. This is a convenience routine. If there will many accesses to the same file the recommended approach is to open the property file with `fso.loadPropFile()` and use the methods of `PropertiesSO` to access and modify properities.

Parameters:
propFile - the property file
propName - the property to get

Returns:
the property value or null

boolean isDirectory(String filename)

Description:
Returns whether or not the file is a directory.

Parameters:
filename - file to check

Returns:
`true` if the file is a directory.

Properties loadProperties(String fileName) (experimental)

Description:
Loads a `java.util.Properties` object from a properties file

Parameters:
fileName - the properties file

Returns:
the `java.util.Properties` Object loaded from file

PropertyFileSO loadPropFile(String filename)

Description:
Loads a scriptable java properties file. The scriptable properties file retains comments upon modification. A properties file open using this method will is open on return. T

Parameters:
filename - the property file to load.

Returns:
the scriptable interface into the property file.

boolean moveFile(String currentName, String newName)

Description:
Moves a file. Use to rename a file.

Parameters:
currentName - the current file name
newName - the relocated file name

Returns:
`true` if successful

TextStream openTextFile(String filename, int ioMode, boolean okToCreate)

Description:
A factory method for open file as a TextStream. Example:

```
var stream = fso.openTextFile( "fileFoo", 1, true)
while( !stream.atEndOfStream) {
   line = stream.readLine();
   print(line);
}
stream.close();
```

Parameters:
filename - the name of the file to open
ioMode - 1 for reading only, 2 - for writing only
okToCreate - true if it is ok to create the file if does not exist

Returns:
a TextStream object.

boolean setProperty(String filename, String propname, String propvalue)

Description:
Sets the specified property. This is a convenience routine. If there will many accesses to the same file the recommended approach is to open the property file with `fso.loadPropFile()` and use the methods of `PropertiesSO` to access and modify properties. This routine will preserve the comments in the property file. The property will be added to the file if does not already exist. Example:

```
fso.setProperty("vbroker.properties","vbroker.security.disabled"
, "false");
```

Parameters:
filename - the property file
propname - the property to set
propvalue - the new value for the property. Add property if it does not exist.

Returns:
the true if successfull

bes

The `bes` object is the scriptable interface to the functions that act on the **local** BES installation. No login is required to use these functions. This object is pre-instantiated by the BMSH as `bes`.

Property Summary

String installRoot [get]

Description:
The BES installation root folder property. This will be the name of the parent of the BES /bin folder of bmsh.exe and scu.exe

String tempFolder [get]

Description:
The root path for temporary files

Method Summary

String getConsolePort()

Description:
Returns the management port of the console

Returns:
the port of the console

String getManagementPort(String domain)

Description:
Returns the management port of the domain

Returns:
the port

void setConsolePort(String port)

Description:
Sets the management discovery port of the console.

Returns:
the port of the console

void setManagementPort(String port, String domain)

Description:
Sets the management port of the given domain

Parameters:
port - new value for the port
domain - domain name

Hub

The `HubSO` object has core methods and properties for interacting with the Borland Management Hub. ${installRoot}/bin/bscript/autoload/hub_ext.js contains a set of convenience API for configurations and managed objects.

Method Summary

void addArchivedConfiguration(String archiveFile)

Description:
Add a configuration to the Hub.

void addConfiguration(String file)

Description:
Add a configuration to the Hub.

Parameters:
file - a full path (local to the script) to a configuration.xml file to be added to the configuration.

String addConfigurationFromTemplate(String templatePath, String configName, Object templateProperties)

Description:
Add a configuration to the Hub based on a template.

Parameters:
templatePath - path to the template file. If not absolute, assumes relative to the installRoot/var/templates/configurations folder
configName - name for the new configuration. If empty, the root file name of the template is used.
templateProperties - java.util.Properties object with the template property overides to be used when instantiating the template. (See the properties section in the template itself for the property names)

Returns:
string with either "Success" for successful, or an error message.

boolean connect(String osAgentPort, String hubName)

Description:
Connect to a hub. You must connect to a hub before calling any Hub API. This creates a VisiBroker ORB.

Parameters:
osAgentPort - the management osagent port
hubName - the name of the management hub

Returns:
true if connect is successful false otherwise

boolean corbalocConnect(String host, String port)

Description:
Connect to a hub. You must connect to a hub before calling any Hub API. This creates a VisiBroker ORB.

Parameters:
host - host hub is on
port - iiop port of the management hub

Returns:
true if connect is successful false otherwise

void disconnect()

Description:
Disconnects from the hub. You are not required to disconnect but explicitly disconnecting will free resources on the hub more quickly then allowing the connection to time-out. This will stopthe hub from sending notifications. TODO: login w/name&password.

String[] getAgents(String config)

Description:
Get the agents belonging to a configuration

Parameters:
config - configuration name

Returns:
JavaScript array of strings for each agent involved in the configuration

String[] getAgentsBelongingToHub()

Description:
Get the names of the agents belonging to a Hub

Returns:
JavaScript array of strings for each agent belonging to a Hub

Object getConfiguration(String cfgName)

Description:
Get a configuration object

Returns:
Java Configuration object

Scriptable getConfigurations()

Description:
Get the configurations being managed by a Hub. Example:

```
var h = new Hub();
h.connect(42424)
var cfgs = h.getConfigurations();
 for( c in cfgs )
  print(cfgs[c]);
```

Returns:
JavaScript array of strings for each configuration the Hub is managing

String getHost()

Description:
Get the host name of the hub

Returns:
host name of the hub

String getName()

Description:
Returns the name of the hub

Returns:
the string hub name.

Object getNamedAgent(String name)

Description:
Get an interface to an Agent

Returns:
Java ManagementAgent object

String getStateName(int state)

Description:
Translates a managed-object state integer value to a string. The state is returned as an integer in notifications and other methods of the API.

Parameters:
state - a state value returned in notification.

Returns:
the string for the state.

boolean iorConnect(String ior)

Description:
Connect to a hub. You must connect to a hub before calling any Hub API. This creates a VisiBroker ORB.

Parameters:
ior - ior string for hub

Returns:
true if connect is successful false otherwise

void removeConfiguration(String configName, boolean stopFirst)

Description:
Remove the configuration from the hub

Parameters:
configName - String representing the configuration name.
stopFirst - boolean flag indicating config should be stopped before removing (recommended)

void updateConfiguration(String file)

Description:
Update a configuration to the Hub.

Parameters:
file - a full path (local to the script) to a configuration.xml file to be updated.

Object waitForEvent()

Description:
Wait for an hub event. A blocking call that wait until a event is recieved from the hub. The caller responsibility to filter the event. Example of monitoring all events and displaying them.

```
var hubName = "myHub";
var hub = new Hub();
hub.connect(42424, hubName);

while(1) {
    print("Sleeping... Waiting for notifications...");
    var event = hub.waitForEvent();
    print(event.toString());
}
```

Returns:
An Event object.

Object waitForStartEvent(int timeoutInSec)

Description:
Wait for a start event. This is the same as waitForEvent() but returns only when the event is state == RUNNING.

Parameters:
timeoutInSec - the time to wait for a start.

Returns:
An Event object.

Object waitForStopEvent(int timeout)

Description:
Wait for a hub start event. This is the same as waitForEvent() but returns only when the event is state == STOPPED.

Returns:
An Event object.

PropertyContext

The `PropertyContext` class represents a scriptable interface that is used to substitute property values into a String containing placeholders for property names. The `PropertyContext` is initialized from a Dom4j containing elements representing the properties and their values. The static create method allows control over the xpath string used to locate the property elements, their names, and their values within the root Node.

Method Summary

void addProperties(Object propertiesContainingNode)

Description:
Given a wrapped XML Node containing child elements adds those properties to the existing set of properties.

Parameters:
propertiesContainingNode -

static PropertyContext	create(org.w3c.dom.Node root, String propElementXPath, String nameAttr, String valueAttr) **Description:** Creates a PropertyContext object initialized from a root org.w3c.dom.Node. If the Node passed in is a org.w3c.dom.Document, the actual node used is the result of getDocumentElement(). The optional parameters allow control over the xpath strings used to locate the property element, the property name attribute in that element, and the property value attribute in that element. **Parameters:** root - the root node to be used for locating properties propElementXPath - xpath used to locate elements holding property values nameAttr - used to locate the attribute representing the property name valueAttr - used to locate the attribute representing the property value **Returns:** PropertyContext instance initialized from root Node using optional xpath strings
Object	getProperties() **Description:** Returns the underlying properties container in a wrapped NativeJavaObject **Returns:** java.lang.Object containing the properties
String	getProperty(String property) **Description:** Retrieve an individual property value **Parameters:** property - the name of the property **Returns:** the property value or `null` if the property is undefined.
void	properties() (experimental) **Description:**
String	setProperty(String property, String value) **Description:** Set an individual property value **Parameters:** property - the name of the property value - the new value of the property **Returns:** the old value of the property or null if the property was undefined.
String	substituteProperties(String source) **Description:** Replace instances of known properties in the source String with the stored values. **Parameters:** source - the String containing properties of the form "${property.name}" **Returns:** the resultant String with property values substituted.

HttpPage

The `HttpPage` object has methods and properties for interacting with the web pages (http urls). Example:

```
var p = new HttpPage("http://www.google.com");
var e = p.getPage();
if( e != null ) {
   var s = fso.openTextFile("google.html",2,true);
   s.write(e);
   s.close();
} else {
   print("Page:" + p.getUrl() + " is NOT available.");
}
```

Notes: 1. The connection (HTTP request) will follow redirects (HTTP 300). 2. The most likely exception is MalformedURLException which can occur when creating the object 3. There is no control of the timeout used to wait for a HTTP response. 4. No javascript on a page is executed. 5. URL references on a page are not followed - ie. a .jpeg referenced on a page will not be retrieved.

Method Summary

int getHttpResponseCode()

Description:
Returns the Http Response code for the url.

Returns:
positive value, the HTTP Response Code from HTTP header. negative value if no page was recieved.

String getHttpResponseMessage()

Description:
Returns the Http Response Message for the url.

Returns:
the HTTP response message from HTTP header. null if no page was recieved.

String getPage()

Description:
Returns the contents of a web page. Hyperlinks on the page are not retrieved. For example if hyperlink refers to a .jpg the file will not be retrieved. Redirects using Javascript will not be followed. No Javascriopt is executed,

Returns:
the web page as String, null if not received.

int getPageLength()

Description:
Returns the content length of a web page

Returns:
positive value, the content length of the page. negative value if no page was recieved.

static Object httpRequest(String urlStr)

Description:
Make an http request to the specified url. Returns Java HttpURLConnection object. The HttpURLConnection api is defined at http://java.sun.com/j2se/1.4.1/docs/api/

Returns:
HttpURLConnection or null if connection fails.

static int ping(String urlStr, int desiredRespCode)

Description:
Ping a url and check the response code against the argument. This is a convience method for pinging a web page without instantiating HttpPage object. It is designed for pinging web page as a BMS managed object. If you need more control over what success means you must instantiate an HttpPage object and use its methods.

Parameters:
urlStr - the http url to ping
desiredRespCode - the HTTP response code desired

Returns:
0 if httpResponseCode == httpRespCode, page is running 1 if httpResponseCode != httpRespCode, page is NOT running

BmsExec

Methods for starting, stopping and pinging process using the functionality of the BMS agent.

Method Summary

static boolean killProcess(String processToken)

Description:
Terminates a process specified by the argument string in the form pid/starttime. This is the same form returned by `startProcess`. If the process is successfully killed, true is returned, false otherwise. This is the most brutal method to end a process. Use stopProcess() to more gracefully stop a process.

Returns:
true if process is killed, false otherwise.

static boolean pingProcess(String processToken)

Description:
Returns whether a process specified by the argument string in form pid/starttime is running. The argument is in the same form returned by `startProcess`.

Returns:
true if process is running, false otherwise.

static String startProcess(String cmdLine)

Description:
Starts a process returning pid and start time as a string in form: pid/starttime. If process couldn't start, null is returned. If start time couldn't be obtained, -1 is used for its value.

Parameters:
cmdLine - command line to execute.

Returns:
String containing pid and start time.

static String startProcessX(Object processNode, Object propertyContext)

Description:
Starts a process referenced by the Dom4j processNode argument.

Parameters:
processNode - a Dom4j <process> node from a configuration file.
propertyContext - a PropertyContext node containing the ${} property substitution values that are in <process> node.

Returns:
String pid and start time as a string in form: pid/starttime. null if unable to start process

static boolean stopProcess(String process, String strategy)

Description:
Attempts to stop a process specified by the argument string in the form pid/starttime. This is the same form returned by `startProcess`. If the process is successfully stopped, true is returned, false otherwise. If this does not stop the process use killProcess().

Parameters:
process - pid and start time of the process to stop.
strategy - one of:
On Windows:
`"ctrl-c"` sends ctrl-c to the process
`"wm-close"` sends WM_CLOSE to all windows owned by the process
`"c-exit"` calls the C runtime exit() function
`"exit-process"` calls the Windows ExitProcess() API, equivalent to kill
`"bes-launcher-exit"` calls an API in the BES launcher that causes it to call the defined Java cleanup method and exit
On Unix:
All of these strategies result in a SIGTERM signal being sent to the process.

Returns:
true if process is stopped, false otherwise.

Dom4j

The `Dom4j` object has static methods for conveniently parsing and saving XML files. The Dom4j API is described here: `http://www.dom4j.org/apidocs/index.html`. This API can be used with the object returned from parseAXmlFile object using the LiveConnect feature of JavaScript. Example:

```
var dom = Dom4j.parseXmlFile("logConfiguration.xml");

var rn = dom.selectSingleNode("/log4j:configuration/root");
var an = rn.selectSingleNode("appender-ref[@ref='SOCKET']");

 if( an == null ) {
       an = rn.addElement("appender-ref");
       an.addAttribute("ref","SOCKET");
}

Dom4j.save(dom,"logConfiguration.xml");
```

Method Summary

static String getAttribute(String xmlFile, String attrXPath)

Description:
Gets an attribute value from an xml file. A convenience routine for getting a single attribute value.

Parameters:
xmlFile - - name of file to use.
attrXPath - - xpath of attribute to set.

Returns:
the attribute value

static Object parseXmlFile(String xmlFile)

Description:
Parses an XML file and returns a Dom4j document. The document is a native Dom4j document. It is normalized();

Returns:
- an org.dom4j.Document

static boolean save(Object node, String xmlFile)

Description:
Saves a node of a dom to a file. It saves in a "pretty" xml format.

Parameters:
node - - a dom node. Can be a document node.
xmlFile - - name of the file to save the node to.

Returns:
boolean - true if saved, false otherwise

static boolean setAttribute(String xmlFile, String attrXPath, String attrValue)

Description:
Opens an xml file, sets the attribute value and saves the xml. A convenience routine for setting a single attribute value.

Parameters:
xmlFile - - name of file to use.
attrXPath - - xpath of attribute to set.
attrValue - - the value to set the attribute to.

Returns:
boolean - true if saved, false otherwise

BmsState

The `state` object is the class representing possible states of a managed object.

Property Summary

int STATE_DEPENDENCY_DOWN [get]

Description:
A group dependency may be down

int STATE_ERROR_STARTING [get]

Description:
Error while starting

int STATE_ERROR_STOPPED [get]

Description:
The entity stopped with an error

int STATE_ERROR_STOPPING [get]

Description:
Error stopping the service

int STATE_ERROR_WAIT_START [get]

Description:
Error occured while waiting to start

int STATE_ERROR_WAIT_STOP [get]

Description:
Error occured while waiting to stop

int STATE_MALFUNCTIONING [get]

Description:
Malfunctioning state

int STATE_RESTARTING [get]

Description:
About to restart (on restart request)

int STATE_RUNNING [get]

Description:
Up and Running

int STATE_STARTING [get]

Description:
About to start

int STATE_STOPPED [get]

Description:
Stop completed

int STATE_STOPPING [get]

Description:
About to Stop

int STATE_WAITING_TO_RESTART [get]

Description:
Waiting for a restart

int STATE_WAITING_TO_START [get]

Description:
Waiting for a start

int STATE_WAITING_TO_STOP [get]

Description:
Waiting to stop

PropertyFileSO

The `PropertyFileSO` object has methods and properties for accessing a .properties file. It differs from `java.util.Properties` object in that a `PropertyFileSO` will preserve the comments (#) in a .properties file. Example:

```
var pf = fso.loadPropFile("foo.properties");
var val = pf.getValue(propname);
pf.setValue(propname,"propValue");
pf.write()
pf.close()
```

If you are getting or setting only one property you may find it more convenient to use `fso.get/setProperty()`.

Method Summary

void close()

Description:
Closes a property file.

String getValue(String propname)

Description:
Get a property value

Parameters:
propname - the name of the property to get

void setValue(String propname, String propvalue)

Description:
Set a property value

Parameters:
propname - the name of the property to get
propvalue - the value to set the property to

void write()

Description:
Writes a property file. You must write a property file to preserve any changes you may have made.

TextStream

The `TextStream` object that methods and properties for reading and writing text files. It is a factory object of `fso`. It can be created by `fso.openTextFile() or fso.createTextFile()`.

Property Summary

boolean atEndOfStream [get]

Description:
Indicates when the end of a file stream has been reached. Example:

```
while( !stream.atEndOfStream )
  line = readline();
```

Method Summary

void close()

Description:
close the stream

String readAll()

Description:
Reads all of a file into a buffer.

String readLine()

Description:
Reads the next line of a file into a buffer.

boolean reset()

Description:
Resets the stream to it first character

Returns:
true if stream points to first character.

void write(String text)

Description:
Writes text to the end of the stream. The stream must have been opened forWriting. See `fso.openTextStream()`

Parameters:
text - the text to write to the stream

void writeLine(String text)

Description:
Writes the text to the stream and appends a cr/lf after the text. The stream must have been opened forWriting.

Parameters:
text - the text to write to the stream

Index

Symbols

! (not equals) operator 509
(has or contains) operator 509
%platform% property 508
& (AND) operator 509
= (equals) operator 509
? (unary) operator 509
[ ] square brackets 3
| (OR) operator 509
| vertical bar 3

A

action strategy
  kill 496
  ping 495
  start 492
  stop 494
actions
  Kill 396
  Ping 396
  Start 396
  Stop 396
adding time rules to an MO 499
ADLoginModule, using 247
Agent 389
  Management Domain 392
  Management Port 391
  starting 391
agent attribute 420
Agent properties
  referencing 507
agent.shutdown.policy property 397, 398
always-success
  partition start action strategy 493
Ant 305
  building AppServer examples 312
  customized tasks 305
  deploying AppServer examples 312
  running AppServer examples 312
  troubleshooting AppServer examples 313
  undeploying AppServer examples 312
Ant tasks
  iastool examples 309
  ommitting attributes 309
  syntax 305
  usage 305
Apache Ant 305
  building AppServer examples 312
  deploying AppServer examples 312
  running AppServer examples 312
  troubleshooting AppServer examples 313
  undeploying AppServer examples 312
  web services 73
Apache Axis
  Axis Toolkit libraries 72
  web service samples 73
  web services 68, 69
  web services Admin tool 74
Apache web server 7, 27
  clustering 55, 58
  configuration 27
  configuration syntax 27
  connecting to CORBA 61
  connecting to web container 32
  CORBA server 63
  directory structure 29
  .htaccess files 29
  HTTP sessions 59
  httpd.conf file 27, 37
  IIOP configuration 39
  IIOP connector 35
  IIOP connector configuration 37
  IIOP module 35
  Managed Object 422
  privileged port 28
Apache web server Managed Object 400, 404, 422, 467
Apache, httpd.conf configuration 28
apache-data element
  apache-process type 467
  arguments sub-element 469
  command attribute 469
  env-vars sub-element 470
  httpd-conf attribute 469
  library-path sub-element 470
  path sub-element 470
  stderr path sub-element 471
  stdin path sub-element 471
  stdout path sub-element 471
apache-process
  Managed Object type 400, 422, 467
  Managed Object type kill strategy 496
  Managed Object type ping strategy 495
  Managed Object type start strategy 493
  Managed Object type stop strategy 494
  process start action strategy 493
apache-process type
  control-overrides element 468
  kill action strategy 496
  ping action strategy 495
  start action strategy 493
  stop action strategy 494
  time-rules element 468
append attribute
  stderr element 437
  stdout element 437
applications
  managed 268
  non-managed 268
AppServer examples
  building 312
  deploying 312
  running 312
  troubleshooting 313
  undeploying 312
AppServer object
  in BMSH 411
AppServer Partition
  Managed Object 423
AppServer Partition Managed Object 423, 476
AppServer Partition Service Managed Object 484
AppServer web components 27
AppServer web server 27

directory structure 29
archive
deploying to a Partition 15
hosting by a Partition without deploying 17
verifying at deployment 16
archives, deploying in JBuilder 386
arguments element
jscript element run sub-element 456
arguments sub-element
apache-data element 469
java-process element 442
process element 434
VBJ-process element 447
arrays
Java string in BMSH 413
JavaScript in BMSH 412
attributes
time-rule sub-element 500
authentication, VisiConnect 266
autoload feature
in BMSH 410
Axis Toolkit libraries, web services 72

B

BAS
action strategies 487
Agent 389
Borland Management Console 393
Configurations 392, 395, 415
Configuration-wide properties 505
configuration.xml 415
configuration.xml file 396
configuring Managed Objects 419
creating Configurations 395
customizing Managed Object behavior 487
Hub 389
killing Configurations 396
Managed Object types defined 422
Managed Objects 392, 399, 419
Managed Resources 392
Management Domain 392
Management Port 391
property if statements 509
property switch statements 509
running Configurations 396
starting Configurations 396
stopping SCU 397
xml for Managed Objects 419
base-state
time-rules attributes 500
batch mode
in BMSH 408
BLOB 147
BMSH
autoload feature 410
batch mode 408
executing 407
execution method 413
exiting the interactive shell 408
file 409
interactive behavior 413
interactive mode 408
running 407
scripts 409
search path 410
starting the interactive shell 408
uses of 407
BMSH (Borland Management Shell)
defined 407
bootclasspath sub-element
profiler element 482
strace element 480
Borland AppServer
Agent 389
architecture 6
Configurations 415
configuration.xml 415
Connector service 8
EJB container 8
examples, running 305
Hub 389
J2EE APIs 10
JDataStore 8
JMS services 7
Managed Objects 399
Management Agent 375
Naming service 9
Partition Services 8
Partitions 8
partitions, starting Borland 384
services 6
session service 9
Smart Agent 7
transaction manager 9
Transaction Service 7
web container 9
web server 7
Borland AppServer 6.6
configuring in JBuilder 373
deploying to remotely 386
remote debugging 387
starting 384
Borland Developer Support, contacting 4
Borland Management Console, in JBuilder 374
Borland Management Shell (see BMSH) 407
Borland Technical Support, contacting 4
Borland virtual directory, IIS/IIOP redirector 46
Borland web container 29
adding environment variables 31
clustering 55, 58
configuration files 29
connecting to JSS 32
ENV variables 31
IIOP configuration 35
IIOP connector 35
JavaServer Pages 30
JSS and failover 58
server.xml 29, 35
servlets 30
Borland Web site 4
Borland-specific web DTD 31
brackets 3
built-in functions
for JavaScript 413

C

cannot-start-policy attribute
ordered-group element 425
cascade delete 128
database 128

Index

cascade delete database 128
! 457
arguments attribute
  ! 457
custom-javascript Managed Object
  ! 457
jscript element
  run element ! 457
run element
  457
custom-javascript Managed Object
  ! 457
run element
  457
CGI-bin Apache directory 29
check-osagent
  osagent ping action strategy 495
check-pid
  custom-javascript ping action strategy 495
  ots ping action strategy 496
  process ping action strategy 495
classloading
  support 271
  VisiConnect 271
classpath attribute
  jmx element 485
classpath sub-element
  profiler element 481
  strace element 480
clear-member-start-failures attribute
  redundancy-group element 429
client
  definition of 75
  get bean information 81
  initialization of 75
  invoke enterprise bean methods 78
  locate home interface 75
  manage transaction 80
  obtain remote interface 76
  use bean handle 79
client j2ee, running 92
client-side stub file, generating 324
CLOB 147
clustering
  Apache web server 55
  Borland web container 55
  Java Session Service 58
  JSS 58
  message-driven beans 175
  Session Service 58
  web components 55
clusters
  deploying JAR files 320
  IIOP connector 39
  IIOP redirector 48
  undeploying JAR files 340
CMP 2.x 117, 119
  and entity beans 117
  Borland implementation 120
  CMP mapping 123
  coarse-grained fields 123
  configuring database tables 122
  configuring datasources 122
  container-managed relationships 117
  many-to-many 127
  mapping fields to multiple tables 124
  one-to-many 126
  one-to-one 125
  optimistic concurrency 120
  persistence manager 118, 119
  schema 122
  specifying relationships 125
command attribute
  apache-data element 469
  java-process element 440
  process element 434
command line tools
  compilejsp 317
  compress 319
  deploy 320
  dumpstack 321
  genclient 322
  gendeployable 322, 323
  genstubs 324
  info 325
  kill 326
  listhubs 328
  listpartitions 327
  listservices 329
  manage 329
  merge 330
  migrate 331
  newconfig 332
  patch 333
  ping 333
  pservice 334
  removestubs 335
  restart 336
  start 337, 338
  stop 339
  uncompress 339
  undeploy 340
  unmanage 341
  usage 342
  verify 342
commands
  conventions 3
compilejsp, iastool command 317
component managed sign-on 266
compress, iastool command 319
conf Apache directory 29
conf IIS directory 32
Configuration
  adding Managed Objects 399
  Managed Object types 400
  properties 505
  scheduled tasks, about 396
configuration element 415
  attributes 415
  description attribute 415
  display-name attribute 415
  revision attribute 415
  small-icon attribute 415
Configuration properties 416

using if statements 509
using switch statements 509
valid if statement operators 509
configuration-id element 416
configuration-id element attributes 416
Configurations 392, 415
adding using a Template 395
configuration.xml files 396
creating 395
Hub 392
killing 396
Managed Objects 392
Managed Resources 392
Master Hub 392
models 395
overview 395
running 396
scheduled tasks 405
starting 396
stopping 396
Templates 395
configurations
Managed Objects 419
xml elements 419
configurations (Borland), starting in JBuilder 384
Configuration-wide properties
defining 505
referencing 506
using 506, 507
configuration.xml
configuration element 415
configuration-id element 416
main-root element 416, 417
Managed Object types 422
managed-object element 420
managed-objects element 416, 419
properties element 416
configuration.xml file 396, 415
configuring
JBuilder for Borland AppServer 6.6 373
JNDI objects for OpenJMS 226
connection
leak detection 271
management 264
recovery, JMS 176
Connection Pool, Partition 20
Connector service 8
connector, IIOP 35
connectors, connection management 264
Console
creating a Configuration 395
creating a Managed Object 399
Container-Managed Persistence 2.0
automatic table creation 148
CMP engine properties 133
column properties 134, 137, 138
data access support 146
entity bean properties 131
entity properties 133, 135
fetching special data types 146
Oracle Large Objects (LOBs) 147
table properties 134, 136
control-overides element
kill sub-element 451, 460
parameters sub-element 452, 461
ping sub-element 450, 459
start sub-element 450, 459
stop sub-element 451, 460
control-overrides element
apache-process type 468
custom-executable type 458
custom-javascript type 450
java-process type 440
kill sub-element 488
kill sub-element attribute 491
kill sub-element strategy attribute 491
ordered-group type 424
osagent type 465
ots type 472
parameters escalate-stop attribute 489
parameters kill-first-ping-interval attribute 490
parameters kill-ping-interval attribute 490
parameters kill-retry-interval attribute 490
parameters kill-timeout attribute 490
parameters local-restart attribute 489
parameters ping-interval attribute 489
parameters ping-policy attribute 489
parameters start-first-ping-delay attribute 489
parameters start-ping-interval attribute 489
parameters start-timeout attribute 489
parameters stop-ping-interval attribute 490
parameters stop-retry-interval attribute 489
parameters stop-timeout attribute 489
parameters sub-element 487
parameters sub-element attributes 489
ping sub-element 488
ping sub-element attribute 492
ping sub-element strategy attribute 492
process type 433
redundancy-group type 428
start sub-element 487
start sub-element attribute 490
start sub-element strategy attribute 490
state-proxy type 431
stop sub-element 488
stop sub-element attribute 491
stop sub-element strategy attribute 491
control-overrides kill sub-element
data-id-ref attribute 454, 463
strategy attribute 454, 463
control-overrides ping sub-element
data-id-ref attribute 453, 462
strategy attribute 453, 462
control-overrides start sub-element
data-id-ref attribute 452, 461
strategy attribute 452, 461
control-overrides stop sub-element
data-id-ref attribute 453, 463
strategy attribute 453, 463
control-overrrides element
kill sub-element attributes 454, 463
ping sub-element attributes 453, 462
start sub-element attributes 452, 461
stop sub-element attributes 453, 463
CORBA
connecting to web server 61
distribution mapping 82
IIOP connector 61
mapping to EJB 81
naming mapping 83
object instances and IIOP connector 64
security mapping 84

Index

transaction mapping 83
web server connection 61
CORBA methods, urls 61
CORBA servant, implemeting ReqProcessor IDL 62
CORBA server
implementing ReqProcessor IDL 62
ReqProcessor IDL 61, 62
web-enabling 61
corbaloc load balancing 56
count-member-problem-as-stopped attribute
redundancy-group element 429
count-member-unknown-as-stopped attribute
redundancy-group element 429
Custom Executable Managed Object 400, 404, 422
Custom JavaScript Managed Object 400, 404, 422
custom-executable
Managed Object type 400, 422, 457
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 493
Managed Object type stop strategy 494
custom-executable type
control-overrides element 458
kill action strategy 496
ping action strategy 495
process element 458, 464
process element attributes 464
process element sub-elements 464
start action strategy 493
stop action strategy 494
time-rules element 458
custom-javascript
Managed Object type 400, 422, 449
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 493
Managed Object type stop strategy 494
custom-javascript type
control-overrides element 450
jscript element 449, 454
kill action strategy 496
ping action strategy 495
start action strategy 493
stop action strategy 494
time-rules element 449

D

Data Archive (DAR) 182
creating and deploying 183
jndi-definitions module 182
migrating to 183
packaging 184
databases, connecting 181
data-directory attribute 421
DataExpress 51
data-id attribute
java-process element 440
jscript element 454
process element 434
data-id-ref attribute
control-overrides kill sub-element 454, 463
control-overrides ping sub-element 453, 462
control-overrides start sub-element 452, 461
control-overrides stop sub-element 453, 463
datasources. *See* Data Archive (DAR)
Date object
for JavaScript 413
in BMSH 411
debugging, Partitions and 19
debugging, remote 387
deploy, iastool command 320
deploy-data attribute 421
deploying
archives in JBuilder 386
deployment
how to deploy to a Partition 15
Partitions and 15
stub generator options 16
verifying 16
deployment descriptor, customization properties 356
deploy.wssd file 69
description attribute 421
configuration element 415
desired-range-max attribute
ordered-group element 425
redundancy-group element 429
desired-range-min attribute
ordered-group element 425
redundancy-group element 429
Developer Support, contacting 4
diagnostic tools, dumpstack (iastool) 321
directory attribute
java-process element 440
process element 434
Dispatcher Pool, Partition 20
display-name attribute 420
configuration element 415
distributed transaction, two-phase commit 155
DOCTYPE declaration 89
documentation 2
Borland AppServer Developer's Guide 2
Borland AppServer Installation Guide 2
Borland Security Guide 2
Management Console User's Guide 2
platform conventions used in 3
type conventions used in 3
VisiBroker for Java Developer's Guide 2
VisiBroker VisiTransact Guide 2
domain
Management 392
DTD, XML 89, 90
dump, generating 321
dumpstack, iastool command 321
dynamic queries, EJB-QL 144

E

EIS integration 261
EJB
mapping to CORBA 81
web services 68
EJB Container
ejb.classload_policy property 354
ejb.collect.display_detail_statistics property 356
ejb.collect.display_statistics property 356

ejb.collect.statistics property 356
ejb.collect.stats_gather_frequency property 356
ejb.copy_arguments property 353
ejb.finder.no_custom_marshal property 355
ejb.interop.marshal_handle_as_ior property 355
ejb.jdb.pstore_location property 355
ejb.jss.pstore_location property 355
ejb.logging.doFullExceptionLogging property 355
ejb.logging.verbose property 355
ejb.mdb.threadMax property 356
ejb.mdb.threadMaxIdle property 356
ejb.mdb.threadMin property 356
ejb.module_preload property 354
ejb.no_sleep property 354
ejb.sfsb.aggressive_passivation property 355
ejb.sfsb.factory_name property 355
ejb.sfsb.keep_alive_timeout property 355
ejb.system_classpath_first property 354
ejb.trace_container property 354
ejb.use_java_serialization property 353
ejb.useDynamicStubs property 354
ejb.usePKHashCodeAndEquals property 354
ejb.xml_validation property 354
ejb.xml_verification property 354
EJB container 8
properties 353
ejb.classload_policy for EJB Container 354
ejb.collect.display_detail_statistics for EJB Container 356
ejb.collect.display_statistics for EJB Container 356
ejb.collect.statistics for EJB Container 356
ejb.collect.stats_gather_frequency for EJB Container 356
ejb.copy_arguments for EJB Container 353
ejb.default_transaction_attribute for EJBs 357
EJBException 164
ejb.findByPrimaryKeyBehavior for Entity Beans 360
ejb.finder.no_custom_marshal for EJB Container 355
ejb.interop.marshal_handle_as_ior for EJB Container 355
ejb.jdb.pstore_location for EJB Container 355
ejb.jsec.doInstanceBasedAC for Stateful Session Beans 363
ejb.jss.pstore_location for EJB Container 355
ejb.logging.doFullExceptionLogging for EJB Container 355
ejb.logging.verbose for EJB Container 355
ejb.maxBeansInCache for Entity Beans 358
ejb.maxBeansInPool for Entity Beans 358
ejb.maxBeansInTransactions for Entity Beans 358
ejb.mdb.init-size for Message Driven Beans 361
ejb.mdb.local_transaction_optimization, for Message Driven Beans 361
ejb.mdb.maxMessagesPerServerSession for Message Driven Beans 361
ejb.mdb.max-size for Message Driven Beans 361
ejb.mdb.rebindAttemptCount for Message Driven Beans 362
ejb.mdb.rebindAttemptInterval for Message Driven Beans 362
ejb.mdb.threadMax for EJB Container 356
ejb.mdb.threadMaxIdle for EJB Container 356
ejb.mdb.threadMin for EJB Container 356
ejb.mdb.unDeliverableQueue for Message Driven Beans 362
ejb.mdb.unDeliverableQueueConnectionFactory for Message Driven Beans 362
ejb.mdb.use_jms_threads for Message Driven Beans 361
ejb.mdb.wait_timeout for Message Driven Beans 362
ejb.module_preload for EJB Container 354
ejb.no_sleep for EJB Container 354
EJB-QL 139
dynamic queries 144
GROUP BY extension 143
optimizing SQL 145
ORDER BY extension 141
return types for aggregate functions 140
selecting a cmp-field 139
selecting a collection of cmp-fields 139
selecting a ResultSet 140
specifying custom SQL 145
sub-queries 143
using aggregate functions 140
ejb-ref-name 89, 90
ejb-refs 90
ejb.security.transportType for EJB Security 364
ejb.security.trustInClient for EJB Security 364
ejb.sfsb.aggressive_passivation for EJB Container 355
ejb.sfsb.factory_name for EJB Container 355
ejb.sfsb.instance_max for Stateful Session Beans 363
ejb.sfsb.instance_max_timeout for Stateful Session Beans 363
ejb.sfsb.keep_alive_timeout for EJB Container 355
ejb.sfsb.passivation_timeout for Stateful Session Beans 363
ejb.system_classpath_first for EJB Container 354
ejb.trace_container for EJB Container 354
ejb.transactionCommitMode for Entity Beans 359
ejb.transactionManagerInstanceName for Message Driven Beans 359, 362
ejb.use_java_serialization for EJB Container 353
ejb.useDynamicStubs for EJB Container 354
ejb.usePKHashCodeAndEquals for EJB Container 354
ejb.xml_validation for EJB Container 354
ejb.xml_verification for EJB Container 354
enable attribute
optimizeit sub-element 478
enable_loadbalancing attribute 65
enable-jpda-debug attribute
jpda sub-element 483
enterprise bean
bean-managed transaction 160
container-managed transaction 160
get information about 81
home interface, locate 75
metadata 81
remote interface, reference to 76
remove instances of 79
transaction management 159
enterprise bean methods, to invoke 78
Enterprise JavaBeans
common properties 357
ejb.default_transaction_attribute property 357
Entity Bean properties 358
MDB properties 361
properties index 357
security properties 364
Stateful Session Bean properties 363
entity bean
create methods 78

Index

find methods 77
remote interface
create methods 78
find methods 77
reference to 77
remove methods 78
remove instances of 79
remove methods 78
Entity Beans
EJB 2.0 117
ejb.findByPrimaryKeyBehavior property 360
ejb.maxBeansInCache property 358
ejb.maxBeansInPool property 358
ejb.maxBeansInTransactions property 358
ejb.transactionCommitMode property 359
entity beans
interfaces 118
packaging requirements 118
primary keys 149
re-entrancy 119
ENV variables
Borland web container 31
Tomcat-based web container 31
web container 31
environment variables
in BMSH 412
environment variables web container 31
env-vars element
attributes 443, 447
use-current-env attribute 443, 447
use-default-env attribute 443, 447
use-vbroker-env attribute 443, 447
env-vars sub-element
apache-data element 470
java-process element 442
VBJ-process element 447
error recovery, JMS 176
escalate-stop attribute
parameters element 489
examples 305
building 312
deploying 312
running 312
troubleshooting 313
undeploying 312
web services 72, 73
excess-running-ok attribute
redundancy-group element 429
executing iastool from a script 344
existing applications 93
exiting the BMSH interactive shell 408
extensibility
custom-executable Managed Object type 422, 457
custom-javascript Managed Object type 422, 449

F

factory objects
in BMSH 411
failover
IIOP connector 55, 57
JSS 58
web component clustering 55
fail-policy attribute
ordered-group element 425
fault tolerance
IIOP connector 55, 57
MDB 176
web component clustering 55
-file option, executing iastool from a script 344
find methods 77
fso object
in BMSH 411
functions
built-in JavaScript 413

G

genclient, iastool command 322
gendeployable, iastool command 322, 323
genstubs, iastool command 324
group
ordered Managed Object 422
redundancy Managed Object 422
group-name attribute
platform-specific sub-element 438, 444, 448
groups
Managed Objects 401

H

handle 79
heap size, setting for a Partition 20
help
in BMSH 413
hidden attribute
stderr element 437
stdin element 437
stdout element 437
home interface, locate 75
host sub-element
osagent element 467
hosting, Partitions and 17
.htaccess files 29
htdocs Apache directory 29
HTTP sessions, Apache web server 59
httpd.conf 27
IIOP and CORBA 63
location 27
httpd-conf attribute
apache-data element 469
httpd.conf file
configuration syntax 28
IIOP connector configuration 37
Hub 389
Configurations 392
Management Domain 392
Management Port 391
restarting 398
starting 391
HUB.name property 507
hub.name property 507

I

iastool
compilejsp 317

compress 319
deploy 320
dumpstack 321
executing from a script 344
genclient 322
gendeployable 322, 323
genstubs 324
info 325
kill 326
listhubs 328
listpartitions 327
listservices 329
manage 329
merge 330
migrate 331
newconfig 332
patch 333
ping 333
pservice 334
removestubs 335
restart 336
start 337, 338
stop 339
uncompress 339
undeploy 340
unmanage 341
usage 342
verify 342
icons Apache directory 29
if statement
conditions used in a Configuration property 509
in a Configuration property 509
IIOP
adding new CORBA objects 64
CORBA 63, 64
plugin 35
IIOP connector 35
adding a CORBA instance 64
adding clusters 39
adding CORBA instances 64
adding web applications 41
Apache configuration 37
Apache configuration files 64
Apache web server 35
clustering 55
configuration files 39
CORBA 61
failover 55, 57
fault tolerance 55, 57
load balancing 55, 56
mapping CORBA URLs 64
mapping URIs 39
mapping URIs to CORBA servers 66
server.xml 35
smart session handling 55, 57
UriMapFile.properties 41, 66
web components 55
web container 35
web server 35
WebClusters.properties file 39, 64
IIOP redirector 32
adding clusters 48, 49
adding web applications 50
configuration 48
configuration files 48
IIS web server 46
mapping URIs 48
UriMapFile.properties 50
WebClusters.properties file 49
IIS
adding new clusters 49
adding new web applications 50
IIS redirector 32
directories 32
IIS web server
connecting to web container 46
IIOP redirector 32, 46
IIOP redirector configuration 46, 48
IIOP redirector directory structure 32
versions supported 32
IIS/IIOP redirector
ISAPI filter 46
virtual directory 46
Windows 2000 configuration 46
Windows 2003 configuration 46
Windows XP configuration 46
info, iastool command 325
initial-desired-state attribute 420
initial-manage attribute 420
initial-monitor attribute 420
interactive behavior
BMSH 413
interactive mode
in BMSH 408
internet, accessing CORBA 61
ISAPI filter, IIS/IIOP redirector 46

J

J2EE
APIs supported 10
connector architecture 261
VisiClient 87
VisiClient environment 87
JACC
authorization 243
configuring external providers 245
configuring provider 244
contracts 243
enabling/disabling provider 245
using 243
JAR files
deploying 320
server-side deployable 322, 323
undeploying 340
Java
APIs for XML Registries 251
Server Pages, precompiling 317
Transaction API 162
types mapped to SQL types 113, 148
Java arrays
in BMSH 413
Java LiveConnect
with BMSH 412
Java process
Managed Object 422
Java Process Managed Object 400, 403
Java Session Service 51
automatic storage 58
configuration 54
JDataStore 54
JDBC datasource 54

Index

jss.backingStoreType property 366
jss.debug property 365
jss.factoryName property 365
jss.maxIdle property 365
jss.passWord property 366
jss.pstore property 365
jss.softCommit property 365
jss.userName property 366
jss.workingDir property 365
programmatic storage 58
properties 54, 364
session management 51
storage implementation 58
web components 58
See also JSS
Java Transaction Service 154
jts.allow_unrecoverable_completion property 367
jts.default_max_timeout property 368
jts.default_timeout property 368
jts.no_global_tids property 367
jts.no_local_tids property 367
jts.timeout_enable property 367
jts.timeout_interval property 367
jts.trace property 368
jts.transaction_debug_timeout property 368
properties 367
Java2WSDL tool, web services 74
java-process
Managed Object type 400, 439
VBJ-process element 444
java-process element
(UNIX) platform-specific sub-element 444
(Windows) platform-specific sub-element 444
arguments sub-element 442
command attribute 440
data-id attribute 440
directory attribute 440
env-vars sub-element 442
java-process type 439
java-properties sub-element 441
library-path sub-element 443
main-class attribute 440
path sub-element 443
stderr sub-element 443
stdin sub-element 443
stdout sub-element 443
vm-type attribute 440
java-process type
control-overrides element 440
java-process element 439
time-rules element 439
VBJ-process element 439, 444
java-properties sub-element
java-process element 441
JavaScript
built-in functions,in BMSH 413
Date object 413
embedding in configuration.xml 457
embedding xml restrictions 457
regular expression processor 413
JavaScript arrays
in BMSH 412
JavaServer Pages (JSPs) 30
JAXR 251
JBuilder 373
JDataStore 8
DataExpress 51
JDBC 185
API modifications 162, 163
configuring datasources 186
configuring properties 189
connecting from deployed modules 197
Connection Pooling 88
datasource and JSS 54
datasources 185
debugging 193
deployment descriptor contruction 195
enabling and disabling datasources 184
JDBC 1.x drivers 194
JDK
customizing per Partition 20
Partition options 20
jdkpath attribute
optimizeit sub-element 478
profiler sub-element 479
strace sub-element 479
jdpa-suspend attribute
jpda sub-element 483
jdpa-transport-address attribute
jpda sub-element 483
JMS 199, 225
configuring 201
configuring connection factories 201
connecting from deployed modules 204
connection factories 199
connection recovery 176
deployment descriptor elements 212
error recovery 176
OpenJMS 225
provider, clustering 175
security 212
security enabling for Tibco 224
security Tibco 224
transactions 209
jms-home attribute
tibco-data element 475
JMX
agent, Partition options 19
jmx
partition type 476
jmx element
classpath attribute 485
jmxserver attribute 485
jmxserver attribute
jmx element 485
JNDI support 81
jndi-definitions module 182
JPDA debugging 19
jpda sub-element
enable-jpda-debug attribute 483
jdpa-suspend attribute 483
jdpa-transport-address attribute 483
partition-process element 477
jpda sub-element attributes

partition-process element 483
jscript element
custom-javascript type 449, 454
data-id attribute 454
run element sub-elements 456
run sub-element 455
run sub-element attributes 456
jscript-control
custom-javascript kill action strategy 496
custom-javascript start action strategy 493
custom-javascript stop action strategy 494
jscript-ping
custom-javascript ping action strategy 495
process ping action strategy 495
JSP, definition 30
JSS 51
automatic storage 58
configuration 54
connecting to web containers 32
failover 58
JDataStore 54
JDBC datasource 54
programmatic storage 58
properties 54
session management 51
storage implementation 58
web components 58
See also Java Session Service
jss.backingStoreType for Java Session Service 366
jss.debug for Java Session Service 365
jss.factoryName for Java Session Service 365
jss.maxIdle for Java Session Service 365
jss.passWord for Java Session Service 366
jss.pstore for Java Session Service 365
jss.softCommit for Java Session Service 365
jss.userName for Java Session Service 366
jss.workingDir for Java Session Service 365
JTA 162
JTS, two-phase commit 155
jts.allow_unrecoverable_completion for Java Transaction Service 367
jts.default_max_timeout for Java Transaction Service 368
jts.default_timeout for Java Transaction Service 368
jts.no_global_tids for Java Transaction Service 367
jts.no_local_tids for Java Transaction Service 367
jts.timeout_enable for Java Transaction Service 367
jts.timeout_interval for Java Transaction Service 367
jts.trace for Java Transaction Service 368
jts.transaction_debug_timeout for Java Transaction Service 368
JVM, advanced editing for Partitions 20

K

key cache size 152
Kill action 396
kill action strategy
apache-process Managed Object type 496
apache-process type 496
custom-executable Managed Object type 496
custom-executable type 496
custom-javascript Managed Object type 496
custom-javascript type 496
ordered-group Managed Object type 496
ordered-group type 496
osagent Managed Object type 496
osagent type 496
ots Managed Object type 496
ots type 496
partition Managed Object type 497
partition type 497
process Managed Object type 496
process type 496
redundancy-group Managed Object type 496
redundancy-group type 496
state-proxy Managed Object type 496
state-proxy type 496
tibco Managed Object type 497
tibco type 497
kill element
strategy attribute 491, 496
kill sub-element
control-overides element 451, 460
control-overrides element 488
kill sub-element attribute
control-overrides element 491
kill sub-element attributes
control-overrrides element 454, 463
kill sub-element strategy attribute
control-overrides element 491
kill, iastool command 326
kill-first-ping-interval attribute
parameters element 490
kill-ping-interval attribute
parameters element 490
kill-retry-interval attribute
parameters element 490
kill-signal
apache-process kill action strategy 496
osagent kill action strategy 496
osagent stop action strategy 494
ots kill action strategy 496
ots stop action strategy 495
partition kill action strategy 497
process kill action strategy 496
process stop action strategy 494
tibco kill action strategy 497
tibco stop action strategy 495
kill-timeout attribute
parameters element 490

L

library-path sub-element
apache-data element 470
java-process element 443
VBJ-process element 447
LifeRay
creating MySQL database 370
deploying custom portlets 371
using with AppServer 369
listhubs, iastool command 328
listpartitions, iastool command 327
listservices, iastool command 329
load balancing 56
corbaloc-based 56
IIOP connector 55, 56
osagent-based 56
web component clustering 55
Local Agent
restarting 398

Index

SCU process 397
stopping 397
Local Hub
SCU process 397
stopping 397
local-restart attribute
parameters element 489
log settings, Partition 21
logdir sub-element
osagent element 467
login information
protecting 344
running from a script file 344
logs Apache directory 29
logs IIS directory 32

M

main-class attribute
java-process element 440
main-root element 416, 417
mo-ref attribute 417
main-root element attributes 417
manage, iastool command 329
Managed Object
agent attribute 420
data-directory attribute 421
deploy-data attribute 421
description attribute 421
display-name attribute 420
initial-desired-state attribute 420
initial-manage attribute 420
initial-monitor attribute 420
name attribute 420
time rules 406
type attribute 420
vendor attribute 421
version attribute 421
Managed Object strategies 405
Managed Object types 400, 420
Managed Objects 392, 399, 419
adding to Configuration 399
adding using a Template 399
apache-process type 400, 467
apache-process type kill strategy 496
apache-process type ping strategy 495
apache-process type start strategy 493
apache-process type stop strategy 494
configuring 419
control-overrides element 487
creating 399
creating time rules for 499
custom-executable type 400, 457
custom-executable type kill strategy 496
custom-executable type ping strategy 495
custom-executable type start strategy 493
custom-executable type stop strategy 494
customizing behavior 487
custom-javascript type 400, 449
custom-javascript type kill strategy 496
custom-javascript type ping strategy 495
custom-javascript type start strategy 493
custom-javascript type stop strategy 494
element attributes 419
elements 419
generic xml definition 419
java-process type 439
Kill action 396
Managed Resources 392
namingservice type 400
ordered-group type 400, 423
ordered-group type kill strategy 496
ordered-group type ping strategy 495
ordered-group type start strategy 493
ordered-group type stop strategy 494
osagent type 400, 465
osagent type kill strategy 496
osagent type ping strategy 495
osagent type start strategy 493
osagent type stop strategy 494
ots type 400, 471
ots type kill strategy 496
ots type ping strategy 496
ots type start strategy 493
ots type stop strategy 495
Partition Services types 486
partition type 400, 476
partition type kill strategy 497
partition type ping strategy 496
partition type start strategy 493
partition type stop strategy 495
Ping action 396
process type 400, 432
process type kill strategy 496
process type ping strategy 495
process type start strategy 492
process type stop strategy 494
redundancy-group type 400, 427
redundancy-group type kill strategy 496
redundancy-group type ping strategy 495
redundancy-group type start strategy 493
redundancy-group type stop strategy 494
Start action 396
state-proxy type 400, 431
state-proxy type kill strategy 496
state-proxy type ping strategy 495
state-proxy type start strategy 493
state-proxy type stop strategy 494
Stop action 396
Templates 399
tibco type 400, 473
tibco type kill strategy 497
tibco type ping strategy 496
tibco type start strategy 493
tibco type stop strategy 495
time-rules element 499
type attribute 400, 422
types defined 422
Managed Resource
adding to a Configuration 399
Managed Resources 392
Managed Objects 392, 399
managed sign-on, VisiConnect Container 266
managed-object element
agent attribute 420

apache-process type 467
attributes 420
custom-executable type 457
custom-javascript type 449
data-directory attribute 421
deploy-data attribute 421
description attribute 421
display-name attribute 420
initial-desired-state attribute 420
initial-manage attribute 420
initial-monitor attribute 420
java-process type 439
name attribute 420
ordered-group type 400, 423
osagent type 465
ots type 471
partition type 476
process type 432
redundancy-group type 427
state-proxy type 431
tibco type 473
type attribute 420
vendor attribute 421
version attribute 421
managed-object sub-element 419
managed-objects element 416, 419
managed-object sub-element 419
managed-object sub-element attributes 420
management agent 375
starting 375, 384
Management Console
creating a Configuration 395
creating a Managed Object 399
Management Domain 392
Agent 392
Hub 392
Management Port 391
Agent 391
Hub 391
management port (Borland) 384
Management port ID
changing 391
Manifest
example 94
files, use of 94
MDB
connection recovery 176
dead queue 177
error recovery 176
fault tolerance 176
JMS provider clustering 175
queue configuration 177
rebind attempt 176
using with OpenJMS 232
member group-policy-level
ordered-groupo Managed Object 426
member start order
ordered-group Managed Object 426
member stop-order
ordered-group Managed Object 426
member sub-element
ordered-group element 426
partition-services element 484
redundancy-group element 430
state-proxy element 431
member sub-element attributes
ordered-group element 426
partition-services element 485
redundancy-group element 430
merge, iastool command 330
Message Driven Bean, using with OpenJMS 232
Message Driven Beans
ejb.mdb.init-size 361
ejb.mdb.local_transaction_optimization property 361
ejb.mdb.maxMessagesPerServerSession property 361
ejb.mdb.max-size 361
ejb.mdb.rebindAttemptCount property 362
ejb.mdb.rebindAttemptInterval property 362
ejb.mdb.unDeliverableQueue property 362
ejb.mdb.unDeliverableQueueConnectionFactory property 362
ejb.mdb.use_jms_threads property 361
ejb.mdb.wait_timeout 362
ejb.transactionManagerInstanceName property 359, 362
Message-Driven Beans 169
client view of 170
clustering 175
connecting to JMS connection factories 171
EJB 2.0 specification and 170
failover and fault tolerance 175
JMS and 169
transactions 177
metadata 81
Microsoft Internet Information Services web server. *See* IIS
migrate, iastool command 331
mode attribute
optimizeit sub-element 478
models
Configuration Templates 395
modes, OpenJMS 232
mo-ref attribute
main-root element 417
ordered-group member sub-element 426
partition-services member sub-element 485
redundancy-group member sub-element 430
state-proxy member sub-element 432
Mozilla Rhino project 407

N

name attribute 420
configuration-id element 416
name sub-element
script attribute 456
named sequence table, primary key generation 151
Naming service 9
namingservice
Managed Object type 400
newconfig, iastool command 332
nice-value attribute
platform-specific sub-element 438, 444, 448
null (none)
apache-process kill action strategy 496
apache-process start action strategy 493
apache-process stop action strategy 494
custom-executable kill action strategy 496
custom-executable start action strategy 493
custom-executable stop action strategy 494

Index

custom-javascript kill action strategy 496
custom-javascript start action strategy 493
custom-javascript stop action strategy 494
ordered-group start action strategy 493
ordered-group stop action strategy 494
osagent kill action strategy 496
osagent start action strategy 493
osagent stop action strategy 494
ots kill action strategy 496
ots start action strategy 493
ots stop action strategy 495
partition kill action strategy 497
partition start action strategy 493
partition stop action strategy 495
process kill action strategy 496
process start action strategy 492
process stop action strategy 494
redundancy-group start action strategy 493
redundancy-group stop action strategy 494
state-proxy start action strategy 493
state-proxy stop action strategy 494
tibco kill action strategy 497
tibco start action strategy 493
tibco stop action strategy 495

O

objects
factory 411
in BMSH 411
pre-instantiated 411
user-instantiated 411
one-phase commit, VisiConnect 266
OpenJMS 225
changing datasource 228
configuring for 2PC optimization 229
configuring JNDI objects 226
connection modes 228
creating tables 229
modes 232
running 232
specifying partition level properties 230
using MDB with 232
optimistic concurrency 120
SelectForUpdate 121
SelectForUpdateNoWAIT 121
UpdateAllFields 121
UpdateModifiedFields 121
VerifyAllFields 122
VerifyModifiedFields 121
optimisticConcurrencyBehavior, table properties 136
optimization, 2PC optimization for OpenJMS 229
Optimizeit
Partition 26
optimizeit sub-element
enable attribute 478
jdkpath attribute 478
mode attribute 478
partition-process element 477
sthome attribute 478
xmlpath attribute 478
options sub-element
profiler element 482
strace element 481
VBJ-process element 446
Ordered Group Managed Object 400, 401, 422
ordered-group
Managed Object type 400, 422, 423
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 493
Managed Object type stop strategy 494
ordered-group element
cannot-start-policy attribute 425
desired-range-max attribute 425
desired-range-min attribute 425
fail-policy attribute 425
member sub-element 426
member sub-element attributes 426
ordered-group type 423
require attribute 425
ordered-group Managed Object
member start order 426
member stop-order 426
ordered-group member sub-element
group-policy-level attribute 426
mo-ref attribute 426
start attribute 426
stop attribute 426
stop-policy attribute 426
ordered-group type
control-overrides element 424
kill action strategy 496
ordered-group element 423
ping action strategy 495
start action strategy 493
stop action strategy 494
time-rules attributes 500
time-rules element 423
osagent
Managed Object 422
Managed Object type 400, 422, 465
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 493
Managed Object type stop strategy 494
process start action strategy 493
osagent and web components 32
osagent element
host sub-element 467
logdir sub-element 467
osagent type 465, 466
port sub-element 467
osagent Managed Object 400, 404, 422
osagent type
control-overrides element 465
kill action strategy 496
osagent element 465, 466
ping action strategy 495
process element 465
start action strategy 493
stop action strategy 494
time-rules element 465
OTS
Managed Object 422

ots
Managed Object type 400, 422, 471
Managed Object type kill strategy 496
Managed Object type ping strategy 496
Managed Object type start strategy 493
Managed Object type stop strategy 495
process start action strategy 493
OTS Managed Object 400, 404, 422
ots Managed Object type
process element attributes 471
process element sub-elements 471
ots type
control-overrides element 472
kill action strategy 496
ping action strategy 496
process element 471
start action strategy 493
stop action strategy 495
time-rules element 472

P

parameters element
escalate-stop attribute 489
kill-first-ping-interval attribute 490
kill-ping-interval attribute 490
kill-retry-interval attribute 490
kill-timeout attribute 490
local-restart attribute 489
ping-interval attribute 489
ping-policy attribute 489
start-first-ping-delay attribute 489
start-ping-interval attribute 489
start-timeout attribute 489
stop-ping-interval attribute 490
stop-retry-interval attribute 489
stop-timeout attribute 489
parameters escalate-stop attribute
control-overrides element 489
parameters kill-first-ping-interval attribute
control-overrides element 490
parameters kill-ping-interval attribute
control-overrides element 490
parameters kill-retry-interval attribute
control-overrides element 490
parameters kill-timeout attribute
control-overrides element 490
parameters local-restart attribute
control-overrides element 489
parameters ping-interval attribute
control-overrides element 489
parameters ping-policy attribute
control-overrides element 489
parameters start-first-ping-delay attribute
control-overrides element 489
parameters start-ping-interval attribute
control-overrides element 489
parameters start-timeout attribute
control-overrides element 489
parameters stop-ping-interval attribute
control-overrides element 490
parameters stop-retry-interval attribute
control-overrides element 489
parameters stop-timeout attribute
control-overrides element 489
parameters sub-element
control-overides element 452, 461
control-overrides element 487
parameters sub-element attributes
control-overrides element 489
Partition
Borland web container ENV variables 31
Managed Object 423
server.xml 29, 35
services 8
Tomcat configuration files 35
web container env variables 31
web container service 29
web services 67, 68
partition
Managed Object type 400, 423, 476
Managed Object type kill strategy 497
Managed Object type ping strategy 496
Managed Object type start strategy 493
Managed Object type stop strategy 495
process start action strategy 493
Partition Lifecycle Interceptors 8
deploying 276
interceptor class 274
interceptor class example 275
module-borland.xml DTD 273
Partition Managed Object 400, 404, 423, 476
Partition properties 17
advanced JDK options 20
Advanced tab 21
advanced VisiBroker properties 20
debugging 19
General tab 17
JDK tab 20
JMX Agent tab 19
Log Settings tab 21
management 21
Partition Settings tab 19
Security tab 21
Statistics tab 19
time rules 21
VisiBroker tab 20
partition properties, specifying for OpenJMS 230
Partition Service Managed Object 400, 404, 484
Partition Services
Borland web container 29
Java Session Service (JSS) 51
Managed Object types 486
VisiConnect 277
partition type
jmx element 476
kill action strategy 497
partition-process element 476
partition-services element 476
ping action strategy 496
start action strategy 493
stop action strategy 495
partition-process element
jpda sub-element 477
jpda sub-element attributes 483
optimizeit sub-element 477
partition type 476
partition-process Managed Object type
process element attributes 477
process element sub-elements 477
Partitions 8, 13
advanced deployment 16

Index

Class Loading tab 24
cloning 15
configuring 17
configuring in XML 345
configuring the Connection Pool 20
configuring the Dispatcher Pool 20
creating 13
deleting 15
deploying JAR files 320
deploying modules to 15
dumping stack traces 26
editing VisiBroker properties 20
editing VM settings 20
General tab 22
general use 13
hosting modules 17
JDBC Pool States tab 24
JDK properties 20
JPDA debugging 19
lifecycle interceptors 8
Logs tab 24
Management settings 21
naming 17
Optimizeit 26
options 17
partition.xml reference 345
performance tuning 24
Properties tab 23
setting action strategies 21
setting heap size 20
setting location 17
setting thread stack size 20
Status tab 24
stub generator options 16
undeploying JAR files 340
verifying deployments 16
viewing 22
viewing classloader information 24
viewing configuration.xml definition 23
viewing general information 22
viewing properties 23
viewing statistics 24
VisiBroker settings 20
XML tab 23
partitions (Borland), starting in JBuilder 384
partition-services element
member sub-element 484
member sub-element attributes 485
partition type 476
partition-services member sub-element
mo-ref attribute 485
partition.xml reference 345
partner-url attribute
tibco-data element 475
password credential storage 271
password information
protecting 344
running from a script file 344
patch, iastool command 333
path attribute
stderr element 437
stdin element 437
stdout element 437
path sub-element
apache-data element 470
java-process element 443
profiler element 483
strace element 481
VBJ-process element 448
pessimistic concurrency 120
Ping action 396
ping action strategy
apache-process Managed Object type 495
apache-process type 495
custom-executable Managed Object type 495
custom-executable type 495
custom-javascript Managed Object type 495
custom-javascript type 495
ordered-group Managed Object type 495
ordered-group type 495
osagent Managed Object type 495
osagent type 495
ots Managed Object type 496
ots type 496
partition Managed Object type 496
partition type 496
process Managed Object type 495
process type 495
redundancy-group Managed Object type 495
redundancy-group type 495
state-proxy Managed Object type 495
state-proxy type 495
tibco Managed Object type 496
tibco type 496
ping element
strategy attribute 492, 495
ping sub-element
control-overides element 450, 459
control-overrides element 488
ping sub-element attribute
control-overrides element 492
ping sub-element attributes
control-overrrides element 453, 462
ping sub-element strategy attribute
control-overrides element 492
ping, iastool command 333
ping-apache-process
apache-process ping action strategy 495
ping-custom-executable
custom-executable ping action strategy 495
ping-interval attribute
parameters element 489
ping-partition
partition ping action strategy 496
ping-policy attribute
parameters element 489
ping-tibco-server
tibco ping action strategy 496
platform checks 508
platform-specific (UNIX) sub-element
java-process element 444
process element 438
VBJ-process element 448
platform-specific (Windows) sub-element

java-process element 444
process element 437
VBJ-process element 448
platform-specific sub-element
group-name attribute 438, 444, 448
nice-value attribute 438, 444, 448
root-directory attribute 438, 444, 448
show-gui attribute 437, 444, 448
start-minimized attribute 437, 444, 448
stderr-mode attribute 438, 444, 448
stdout-mode attribute 438, 444, 448
umask attribute 438, 444, 448
user-name attribute 438, 444, 448
window-title attribute 437, 444, 448
plugin, IIOP 35
port
Management 391
port sub-element
osagent element 467
ports
changing Borland management 384
VisiBroker Smart Agent 375
precompiling JSPs 317
pre-instantiated objects
in BMSH 411
primary keys 149
automatic generation 151
generating 149, 150, 151
key cache size 152
named sequence table 151
privileged port, Apache web server 28
process
Managed Object type 400, 422, 432
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 492
Managed Object type stop strategy 494
scheduled tasks for 405
process element
(UNIX) platform-specific sub-element 438
(Windows) platform-specific sub-element 437
arguments sub-element 434
command attribute 434
custom-executable type 458, 464
data-id attribute 434
directory attribute 434
osagent type 465
ots Managed Object type 471
ots type 471
partition-process Managed Object type 477
process type 432
stderr sub-element 436
stdin sub-element 436
stdout sub-element 436
tibco Managed Object type 476
tibco type 473
visiserver Managed Object type 466
process element attributes
custom-executable type 464
process element sub-elements
custom-executable type 464
Process Managed Object 400, 403, 422
process type
control-overrides element 433
kill action strategy 496
ping action strategy 495
process type 432
start action strategy 492
stop action strategy 494
time-rules element 433
process() method and ReqProcessor IDL 63
Profiler
Partition 26
profiler element
bootclasspath sub-element 482
classpath sub-element 481
options sub-element 482
path sub-element 483
profiler sub-element
jdkpath attribute 479
sthome attribute 479
xmlpath attribute 479
properties
Configuration-wide 505
configuration.xml 416
container-level 353
defining Configuration-wide 505
deprecated 507
EJB 353, 357
EJB common 357
EJB customization 356
EJB security 364
ejb.classload_policy 354
ejb.collect.display_detail_statistics 356
ejb.collect.display_statistics 356
ejb.collect.statistics 356
ejb.collect.stats_gather_frequency 356
ejb.copy_arguments 353
ejb.default_transaction_attribute 357
ejb.findByPrimaryKeyBehavior 360
ejb.finder.no_custom_marshal 355
ejb.interop.marshal_handle_as_ior 355
ejb.jdb.pstore_location 355
ejb.jsec.doInstanceBasedAC 363
ejb.jss.pstore_location 355
ejb.logging.doFullExceptionLogging 355
ejb.logging.verbose 355
ejb.maxBeansInCache 358
ejb.maxBeansInPool 358
ejb.maxBeansInTransactions 358
ejb.mdb.init-size 361
ejb.mdb.local_transaction_optimization 361
ejb.mdb.maxMessagesPerServerSession 361
ejb.mdb.max-size 361
ejb.mdb.rebindAttemptCount 362
ejb.mdb.rebindAttemptInterval 362
ejb.mdb.threadMax 356
ejb.mdb.threadMaxIdle 356
ejb.mdb.threadMin 356
ejb.mdb.unDeliverableQueue 362
ejb.mdb.unDeliverableQueueConnection Factory 362
ejb.mdb.use_jms_threads 361
ejb.mdb.wait_timeout 362
ejb.module_preload 354
ejb.no_sleep 354
ejb.security.transportType 364
ejb.security.trustInClient 364
ejb.sfsb.aggressive_passivation 355
ejb.sfsb.factory_name 355
ejb.sfsb.instance_max 363
ejb.sfsb.instance_max_timeout 363

Index

ejb.sfsb.keep_alive_timeout 355
ejb.sfsb.passivation_timeout 363
ejb.system_classpath_first 354
ejb.trace_container 354
ejb.transactionCommitMode 359
ejb.transactionManagerInstanceName 359, 362
ejb.use_java_serialization 353
ejb.useDynamicStubs 354
ejb.usePKHashCodeAndEquals 354
ejb.xml_validation 354
ejb.xml_verification 354
Entity Beans 358
if statement conditions in a Configuration 509
Java Session Service 54
JSS 54, 364
jss.backingStoreType 366
jss.debug 365
jss.factoryName 365
jss.maxIdle 365
jss.passWord 366
jss.pstore 365
jss.softCommit 365
jss.userName 366
jss.workingDir 365
JTS 367
jts.allow_unrecoverable_completion 367
jts.default_max_timeout 368
jts.default_timeout 368
jts.no_global_tids 367
jts.no_local_tids 367
jts.timeout_enable 367
jts.timeout_interval 367
jts.trace 368
jts.transaction_debug_timeout 368
MDBs 361
referencing Configuration-wide 506
referencing pre-defined Agent 507
Session Service 54
Stateful Session Beans 363
using Configuration-wide 506
using if statements in a Configuration 509
using pre-defined Agent 507
using switch statements in a Configuration 509
properties element 416, 505
properties sub-element 505
property element
value attribute 509
property expressions 509
Providers
web services 68, 69, 70
web services examples 72
proxy Apache directory 29
pservice, iastool command 334

R

RA managed sign-on 266
realm information
protecting 344
running from a script file 344
redirector, IIS/IIOP configuration 48
Redundancy Group Managed Object 400, 401, 422
redundancy-group
Managed Object type 400, 422, 427
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 493
Managed Object type stop strategy 494
redundancy-group element
clear-member-start-failures attribute 429
count-member-problem-as-stopped attribute 429
count-member-unknown-as-stopped attribute 429
desired-range-max attribute 429
desired-range-min attribute 429
excess-running-ok attribute 429
member sub-element 430
member sub-element attributes 430
require attribute 429
redundancy-group member sub-element
mo-ref attribute 430
stop-policy attribute 430
redundancy-group type
control-overrides element 428
kill action strategy 496
ping action strategy 495
start action strategy 493
stop action strategy 494
time-rules element 428
References, links 91
RegExp object
in BMSH 411
regular expression processor
for JavaScript 413
remote debugging, Borland servers 387
remote interface, obtain reference to 76
removestubs, iastool command 335
ReqProcessor IDL 62
process()method 63
ReqProcessor Interface Definition Language (IDL) 61
require attribute
ordered-group element 425
redundancy-group element 429
resource adapters 272
VisiConnect 277
res-ref-name 89
res-ref-names 90
restart, iastool command 336
revision attribute
configuration element 415
Rhino project 407
root-directory attribute
platform-specific sub-element 438, 444, 448
run element
arguments sub-element 456
run element sub-elements
jscript element 456
run sub-element
jscript element 455
script attribute 456
run sub-element attributes
jscript element 456
run-apache-process
apache-process start action strategy 493
run-custom-executable
custom-executable kill action strategy 496

custom-executable stop action strategy 494
process start action strategy 493
run-process
osagent start action strategy 493
ots start action strategy 493
process start action strategy 492
tibco start action strategy 493

S

scheduled tasks
about 405
Scheduler Service 255
1PC optimization 257
clustering support 260
configuring database to persist data 256
partition service properties 257
relevant Quartz properties 259
using 255
scheduling Configuration tasks 396
script attribute
jscript element name sub-element 456
jscript element run sub-element 456
script file
-file option 344
passing a file to iastool 344
piping a file to iastool 344
running iastool utilities from 344
scriptable objects
with BMSH 412
scripting
in BMSH 408
scripts
for BMSH 409
in BMSH 407
SCU
starting 391
SCU executable 391
stopping 397
SCU process 389
agent.shutdown.policy property 397, 398
restarting 398
stopping 397
search path feature
in BMSH 410
Security
ejb.security.transportType property 364
ejb.security.trustInClient property 364
security
enabling for Tibco 224
JMS 212
policy ra.xml processing 271
server-config.wsdd file, web services 71, 72
server-name attribute
tibco-data element 475
server-side stub file, generating 324
ServerTrace, Partition 26
server-url attribute
tibco-data element 475
server.xml 29
IIOP connector configuration 35
server.xml file 35
Service Broker, web services 67
Service Provider, web services 67
Service Requestor, web services 67
servlet 30
session bean
remote interface, reference to 76
remove instances of 79
transaction attributes 161
session management 51
web component clustering 55
Session Service 9
automatic storage 58
programmatic storage 58
properties 54, 364
storage implementation 58
web components 58
shared-data-storage attribute
tibco-data element 475
Sheduler Service, configuring 255
show-gui attribute
platform-specific sub-element 437, 444, 448
shutdown-apache-process
apache-process stop action strategy 494
small-icon attribute
configuration element 415
Smart Agent 32, 375
and web components 32
Managed Object 422
Smart Agent Managed Object 400, 422
smart session handling, IIOP connector 55, 57
SOAP
Web services 71
web services 67
Software updates 4
SQL types mapped to Java types 113, 148
square brackets 3
stack trace, generating 321
stack traces, Partition 26
Standard Partition, creating 13
Start action 396
start action strategy
apache-process Managed Object type 493
apache-process type 493
custom-executable Managed Object type 493
custom-executable type 493
custom-javascript Managed Object type 493
custom-javascript type 493
ordered-group Managed Object type 493
ordered-group type 493
osagent Managed Object type 493
osagent type 493
ots Managed Object type 493
ots type 493
partition Managed Object type 493
partition type 493
process Managed Object type 492
process type 492
redundancy-group Managed Object type 493
redundancy-group type 493
start-proxy type 493
state-proxy Managed Object type 493
tibco Managed Object type 493
tibco type 493
start attribute
ordered-group member sub-element 426
start element
strategy attribute 490, 492
start sub-element
control-overides element 450, 459
control-overrides element 487

Index

start sub-element attribute
control-overrides element 490
start sub-element attributes
control-overrrides element 452, 461
start sub-element strategy attribute
control-overrides element 490
start, iastool command 337, 338
start-first-ping-delay attribute
parameters element 489
starting the BMSH interactive shell 408
start-minimized attribute
platform-specific sub-element 437, 444, 448
start-partition
partition start action strategy 493
start-ping-interval attribute
parameters element 489
start-timeout attribute
parameters element 489
State Proxy Managed Object 400, 402, 422
stateful service 55
Stateful Session Beans
ejb.jsec.doInstanceBasedAC property 363
ejb.sfsb.instance_max property 363
ejb.sfsb.instance_max_timeout property 363
ejb.sfsb.passivation_timeout property 363
Stateful Sessions
aggressive passivation 98
caching 97
passivation 97
secondary storage 99
simple passivation 97
stateful storage timeout 99
stateless service 55
stateless session bean, exposing as a web service 68
state-proxy
Managed Object type 400, 422, 431
Managed Object type kill strategy 496
Managed Object type ping strategy 495
Managed Object type start strategy 493
Managed Object type stop strategy 494
state-proxy element
member sub-element 431
state-proxy member sub-element
mo-ref attribute 432
state-proxy type
control-overrides element 431
kill action strategy 496
ping action strategy 495
start action strategy 493
stop action strategy 494
time-rules element 431
Statistics, viewing for Partitions 24
stderr element
append attribute 437
attributes 437
path attribute 437
stderr path sub-element
apache-data element 471
stderr sub-element
java-process element 443
process element 436
VBJ-process element 448
stderr-mode attribute
platform-specific sub-element 438, 444, 448
stdin element
attributes 437
path attribute 437
stdin path sub-element
apache-data element 471
stdin sub-element
java-process element 443
process element 436
VBJ-process element 448
stdout element
append attribute 437
attributes 437
path attribute 437
stdout path sub-element
apache-data element 471
stdout sub-element
java-process element 443
process element 436
VBJ-process element 448
stdout.log, generating a stack trace 321
stdout-mode attribute
platform-specific sub-element 438, 444, 448
sthome attribute
optimizeit sub-element 478
profiler sub-element 479
strace sub-element 479
Stop action 396
stop action strategy
apache-process Managed Object type 494
apache-process type 494
custom-executable Managed Object type 494
custom-executable type 494
custom-javascript Managed Object type 494
custom-javascript type 494
ordered-group Managed Object type 494
ordered-group type 494
osagent Managed Object type 494
osagent type 494
ots Managed Object type 495
ots type 495
partition Managed Object type 495
partition type 495
process Managed Object type 494
process type 494
redundancy-group Managed Object type 494
redundancy-group type 494
state-proxy Managed Object type 494
state-proxy type 494
tibco Managed Object type 495
tibco type 495
stop attribute
ordered-group member sub-element 426
stop element
strategy attribute 491, 494
stop sub-element
control-overides element 451, 460
control-overrides element 488
stop sub-element attribute
control-overrides element 491
stop sub-element attributes

control-overrrides element 453, 463
stop sub-element strategy attribute
control-overrides element 491
stop, iastool command 339
stop-appserver-launcher
process stop action strategy 494
stop-gui-process
process stop action strategy 494
stop-partition
partition stop action strategy 495
stop-ping-interval attribute
parameters element 490
stop-policy attribute
ordered-group member sub-element 426
redundancy-group member sub-element 430
stop-process-term-signal
ots stop action strategy 495
process stop action strategy 494
tibco stop action strategy 495
stop-process-windows-c-exit
osagent stop action strategy 494
process stop action strategy 494
stop-retry-interval attribute
parameters element 489
stop-timeout attribute
parameters element 489
strace element
bootclasspath sub-element 480
classpath sub-element 480
options sub-element 481
path sub-element 481
strace sub-element
jdkpath attribute 479
sthome attribute 479
xmlpath attribute 479
strategies for Managed Objects 405
strategy attribute
control-overrides kill sub-element 454, 463
control-overrides ping sub-element 453, 462
control-overrides start sub-element 452, 461
control-overrides stop sub-element 453, 463
kill element 491, 496
ping element 492, 495
start element 490, 492
stop element 491, 494
string arrays
for Java in BMSH 413
string-replacement properties 505
stub file, generating 324
stubs
deployment options 16
generating at deployment 16
Support, contacting 4
switch statement
in a Configuration property 509
symbols
ellipsis ... 3
square brackets [ ] 3
vertical bar | 3
system configuration information 325
system contracts, VisiConnect 263

T

table properties, optimisticConcurrencyBehavior 136
Technical Support, contacting 4
Templates
adding Configurations 395
adding Managed Objects 399
Configuration 395
Managed Objects 399
thread dump, generating 321
thread stack, setting for a Partition 20
tibco
Managed Object type 400, 423, 473
Managed Object type kill strategy 497
Managed Object type ping strategy 496
Managed Object type start strategy 493
Managed Object type stop strategy 495
process start action strategy 493
Tibco Admin Console 223
Tibco JMS Managed Object 400, 404, 423, 473
Tibco JMS provider
Managed Object 423
tibco Managed Object type
process element attributes 476
tibco type
kill action strategy 497
ping action strategy 496
process element 473
start action strategy 493
stop action strategy 495
tibco-data element 473
time-rules element 474
tibco-data element
jms-home attribute 475
partner-url attribute 475
server-name attribute 475
server-url attribute 475
shared-data-storage attribute 475
tibco type 473
time rules, Partition properties 21
Timer Service 255
time-rule attributes 500
time-rules
exceptions for usage 406
Managed Object, about 406
scheduled tasks, about 405
time-rules attributes
ordered-group type 500
time-rules element
apache-process type 468
custom-executable type 458
custom-javascript type 449
java-process type 439
Managed Object type 499
ordered-group type 423
osagent type 465
ots type 472
process type 433
redundancy-group type 428
state-proxy type 431
tibco type 474
time-rule sub-element 500
Tomcat-based web container 29
adding environment variables 31
configuration files 29
connecting to JSS 32
ENV variables 31
IIOP configuration 35
IIOP connector 35
JavaServer Pages 30

Index

server.xml 29, 35
servlets 30
trace dumping 26
transaction
bean-managed 160
characteristics of 153
client management of 80
commit protocol 154
container support for 154
container-managed 159, 160
continuing 165
declarative management of 159
definition of 153
distributed 155
two-phase commit 155
EJBException 164
enterprise bean management of 159
exceptions 164
application-level 164
continuing 165
handling of 165
rollback 165
system-level 164
flat 154
global and local 160
Java Transaction API 162
Java Transaction Service 154
management 153
VisiConnect 265
manager 9
VisiTransact 154
Mandatory attribute 161
nested 154
Never attribute 161
NotSupported attribute 161
properties 154
recovery 155
Required attribute 161
RequiresNew attribute 161
rollback 165
Supports attribute 161
transaction attributes 161
two-phase commit 155
understanding 153
transactions, VisiConnect 265
two-phase commit
best practices 155
completion flag 155
distributed transactions 155
transactions 155
tunneling databases 155
VisiTransact 154
when to use 155
type attribute 420, 422
apache-process 400, 422, 467
custom-executable 400, 422, 457
custom-javascript 400, 422, 449
java-process 400, 439
namingservice 400
of Managed Objects 400
ordered-group 400, 422, 423
osagent 400, 422, 465
ots 400, 422, 471
partition 400, 423, 476
Partition Services 486
process 400, 422, 432
redundancy-group 400, 422, 427
state-proxy 400, 422, 431
tibco 400, 423, 473
type mapping 113, 148

U

UDDI web services 67
umask attribute
platform-specific sub-element 438, 444, 448
uncompress, iastool command 339
undeploy, iastool command 340
unmanage, iastool command 341
UriMapFile.properties 41, 50, 66
Apache to CORBA connections 64
usage, iastool command 342
use-current-env attribute
env-vars element 443, 447
use-default-env attribute
env-vars element 443, 447
user-instantiated objects
in BMSH 411
user-name attribute
platform-specific sub-element 438, 444, 448
UserTransaction interface 80, 162
use-vbroker-env attribute
env-vars element 443, 447

V

vbj-java-options sub-element
VBJ-process element 445
VBJ-process element
(UNIX) platform-specific sub-element 448
(Windows) platform-specific sub-element 448
arguments sub-element 447
attributes 444
env-vars sub-element 447
java-process type 439, 444
library-path sub-element 447
options sub-element 446
path sub-element 448
stderr sub-element 448
stdin sub-element 448
stdout sub-element 448
vbj-java-options sub-element 445
vendor attribute 421
verify, iastool command 342
version attribute 421
configuration-id element 416
VisiBroker
configuring Partitions for 20
editing for Partitions 20
ORB, making available to JBuilder 375
Partition options 20
Smart Agent 375
VisiBroker Naming Service Managed Object 400
VisiClient 92
about 87

deployment descriptors 88
example 92
VisiClient Container, embedding into existing application 93
VisiConnect
component managed sign-on 266
connection management 264
description 270
managed sign-on 266
overview 277
Resource Adapter managed sign-on 266
security 266
using 277
VisiConnect Service, overview 277
VisiExchange component 27, 35, 61
VisiNaming process Managed Object 400, 404
visiserver Managed Object type
process element attributes 466
process element sub-elements 466
vm-type attribute
java-process element 440

W

WAR file 29, 31
containing web services 71
web services 71, 72
WEB-INF directory 31
WAR files, precompiling Java Server Pages 317
web application
WAR file 31
WEB-INF directory 31
Web Application Archive File (WAR file) 29, 31
web component connection, modifying 35
web components 27
and Smart Agent (osagent) 32
clustering 55, 58
Web Container 29
adding environment variables 31
configuration files 29
connecting to JSS 32
ENV variables 31
JavaServer Pages 30
server.xml 29
servlets 30
web container 9
IIOP configuration 35
IIOP connector 35
server.xml 35
Web module 31
web server
Apache 27
connecting to CORBA 61
directory structure 29
.htaccess files 29
IIOP configuration 39, 64
IIOP connector 35
Web Service Deployment Descriptor (WSDD) web services 71
Web service Providers 68, 70, 72
Web services 67
Apache ANT tool 73
Apache Axis 68, 69
Apache Axis Admin tool 74
Apache Axis samples 73
architecture 67
Axis Toolkit libraries 72
creating a WAR 72
deploy.wssd file 69
EJB provider 70, 72
examples 72, 73
Java2WSDL tool 74
overview 67
Partitions 68
provider examples 72
Providers 71
providers 69, 70
RPC provider 70
server-config.wsdd file 72
Service Broker 67
Service Providers 67
service providers 68, 70
Service Requestor 67
SOAP 67, 71
stateless session bean 68
tools 73
UDDI 67
WAR file 71
WSDD 71
WSDL2Java tool 74
XML 69, 71
xml 67
Web Services pack 27, 35, 61
Web Services, server-config.wsdd file 71
web-borland.xml 29, 31
WebClusters.properties file 39, 49, 64
WebClusters.properties, Apache to CORBA connections 64
webcontainer_id attribute 65
WEB-INF directory 31
web.xml 31
Windows 2000, IIS/IIOP redirector configuration 46
Windows 2003, IIS/IIOP redirector configuration 46
Windows XP, IIS/IIOP redirector configuration 46
window-title attribute
platform-specific sub-element 437, 444, 448
World Wide Web, Borland updated software 4
WSDL2Java tool, web services 74

X

XDOM object
in BMSH 411
XML
DTD 89, 90
VisiClient 88
grammar 89
Web services 71
web services 67, 69
xml
Managed Objects 419
xml files
access to in BMSH 411
xmlpath attribute