



JBoss Admin Development Guide

JBoss 3.2.6

Copyright © 2004 JBoss, Inc.

Table of Contents

Forward	ix
About Open Source	x
About JBoss	xi
1. JBoss: A Full J2EE Implementation with JMX	xi
2. What this Book Covers	xii
1. Installing and Building the JBoss Server	1
1.1. Getting the Binary Files	1
1.1.1. Prerequisites	1
1.1.2. Installing the Binary Package	2
1.1.2.1. Directory Structure	2
1.1.3. The Default Server Configuration File Set	3
1.2. Basic Installation Testing	7
1.3. Booting from a Network Server	9
1.4. Building the Server from Source Code	12
1.4.1. Accessing the JBoss CVS Repositories at SourceForge	12
1.4.2. Understanding CVS	12
1.4.3. Anonymous CVS Access	12
1.4.4. Obtaining a CVS Client	13
1.4.5. Building the JBoss 3.2.6 Distribution Using the Source Code	13
1.4.6. Building the JBoss 3.2.6 Distribution Using the CVS Source Code	13
1.4.7. An Overview of the JBoss CVS Source Tree	14
1.4.8. Using the JBossTest unit testsuite	15
2. The JBoss JMX Microkernel	18
2.1. JMX	18
2.1.1. An Introduction to JMX	19
2.1.1.1. Instrumentation Level	20
2.1.1.2. Agent Level	20
2.1.1.3. Distributed Services Level	21
2.1.1.4. JMX Component Overview	21
2.2. JBoss JMX Implementation Architecture	24
2.2.1. The JBoss ClassLoader Architecture	24
2.2.2. Class Loading and Types in Java	24
2.2.2.1. ClassCastExceptions - I'm Not Your Type	25
2.2.2.2. IllegalAccessException - Doing what you should not	29
2.2.2.3. LinkageErrors - Making Sure You Are Who You Say You Are	30
2.2.2.4. Inside the JBoss Class Loading Architecture	36
2.2.3. JBoss XMBeans	42
2.2.3.1. Descriptors	43
2.2.3.2. The Management Class	45
2.2.3.3. The Constructors	45
2.2.3.4. The Attributes	46
2.2.3.5. The Operations	47
2.2.3.6. Notifications	48
2.3. Connecting to the JMX Server	50
2.3.1. Inspecting the Server - the JMX Console Web Application	50
2.3.1.1. Securing the JMX Console	52
2.3.2. Connecting to JMX Using RMI	54
2.3.3. Command Line Access to JMX	59

2.3.3.1. Connecting twiddle to a Remote Server	60
2.3.3.2. Sample twiddle Command Usage	60
2.3.4. Connecting to JMX Using Any Protocol	63
2.4. Using JMX as a Microkernel	63
2.4.1. The Startup Process	63
2.4.2. JBoss MBean Services	65
2.4.2.1. The SARDeployer MBean	65
2.4.2.2. The Service Life Cycle Interface	69
2.4.2.3. The ServiceController MBean	69
2.4.2.4. Specifying Service Dependencies	71
2.4.2.5. Identifying Unsatisfied Dependencies	72
2.4.2.6. Hot Deployment of Components, the URLDeploymentScanner	72
2.4.3. Writing JBoss MBean Services	74
2.4.3.1. A Standard MBean Example	74
2.4.3.2. XMBBean Examples	79
2.4.4. Deployment Ordering and Dependencies	94
2.5. JBoss Deployer Architecture	103
2.5.1. Deployers and ClassLoaders	106
2.6. Exposing MBean Events via SNMP	108
2.6.1. The SNMP Adaptor Service	108
2.6.2. The Event to Trap Service	109
2.7. Remote Access to Services, Detached Invokers	109
2.7.1. A Detached Invoker Example, the MBeanServer Invoker Adaptor Service	113
2.7.2. Detached Invoker Reference	117
2.7.2.1. The JRMPInvoker - RMI/JRMP Transport	117
2.7.2.2. The PooledInvoker - RMI/Socket Transport	118
2.7.2.3. The IIOPInvoker - RMI/IIOP Transport	119
2.7.2.4. The JRMPProxyFactory Service - Building Dynamic JRMP Proxies	119
2.7.2.5. The HttpInvoker - RMI/HTTP Transport	119
2.7.2.6. The HA JRMPInvoker - Clustered RMI/JRMP Transport	120
2.7.2.7. The HA HttpInvoker - Clustered RMI/HTTP Transport	120
2.7.2.8. HttpProxyFactory - Building Dynamic HTTP Proxies	120
2.7.2.9. Steps to Expose Any RMI Interface via HTTP	121
3. Naming on JBoss	123
3.1. An Overview of JNDI	123
3.1.1. The JNDI API	123
3.1.1.1. Names	123
3.1.1.2. Contexts	124
3.1.2. J2EE and JNDI - The Application Component Environment	125
3.1.2.1. ENC Usage Conventions	126
3.2. The JBossNS Architecture	140
3.2.1. The Naming InitialContext Factories	142
3.2.1.1. Naming Discovery in Clustered Environments	143
3.2.1.2. The HTTP InitialContext Factory Implementation	144
3.2.1.3. The Login InitialContext Factory Implementation	144
3.2.2. Accessing JNDI over HTTP	145
3.2.3. Accessing JNDI over HTTPS	148
3.2.4. Securing Access to JNDI over HTTP	151
3.2.5. Securing Access to JNDI with a Read-Only Unsecured Context	152
3.2.6. Additional Naming MBeans	153
3.2.6.1. org.jboss.naming.ExternalContext MBean	153
3.2.6.2. The org.jboss.naming.NamingAlias MBean	155

3.2.6.3. The org.jboss.naming.JNDIView MBean	155
4. Transactions on JBoss	159
4.1. Transaction/JTA Overview	159
4.1.1. Pessimistic and optimistic locking	160
4.1.2. The components of a distributed transaction	160
4.1.3. The two-phase XA protocol	161
4.1.4. Heuristic exceptions	161
4.1.5. Transaction IDs and branches	162
4.2. JBoss Transaction Internals	162
4.2.1. Adapting a Transaction Manager to JBoss	162
4.2.2. The Default Transaction Manager	163
4.2.2.1. org.jboss.tm.XidFactory	163
4.2.3. UserTransaction Support	164
5. EJBs on JBoss	165
5.1. The EJB Client Side View	165
5.1.1. Specifying the EJB Proxy Configuration	168
5.2. The EJB Server Side View	172
5.2.1. Detached Invokers - The Transport Middlemen	172
5.2.2. The HA JRMPInvoker - Clustered RMI/JRMP Transport	176
5.2.3. The HA HttpInvoker - Clustered RMI/HTTP Transport	176
5.3. The EJB Container	177
5.3.1. EJBDeployer MBean	177
5.3.1.1. Verifying EJB deployments	178
5.3.1.2. Deploying EJBs Into Containers	178
5.3.1.3. Container configuration information	179
5.3.2. Container Plug-in Framework	189
5.3.2.1. org.jboss.ejb.ContainerPlugin	190
5.3.2.2. org.jboss.ejb.Interceptor	190
5.3.2.3. org.jboss.ejb.InstancePool	191
5.3.2.4. org.jboss.ejb.InstanceCache	192
5.3.2.5. org.jboss.ejb.EntityPersistenceManager	193
5.3.2.6. org.jboss.ejb.StatefulSessionPersistenceManager	198
5.4. Entity Bean Locking and Deadlock Detection	199
5.4.1. Why JBoss Needs Locking	199
5.4.2. Entity Bean Lifecycle	199
5.4.3. Default Locking Behavior	200
5.4.4. Pluggable Interceptors and Locking Policy	200
5.4.5. Deadlock	201
5.4.5.1. Dedlock Detection	202
5.4.5.2. Catching ApplicationDeadlockException	203
5.4.5.3. Viewing Lock Information	203
5.4.6. Advanced Configurations and Optimizations	203
5.4.6.1. Short-lived Transactions	204
5.4.6.2. Ordered Access	204
5.4.6.3. Read-Only Beans	204
5.4.6.4. Explicitly Defining Read-Only Methods	204
5.4.6.5. Instance Per Transaction Policy	205
5.4.7. Running Within a Cluster	205
5.4.8. Troubleshooting	206
5.4.8.1. Locking Behavior Not Working	206
5.4.8.2. IllegalStateException	206
5.4.8.3. Hangs and Transaction Timeouts	206

6. Messaging on JBoss	207
6.1. JMS Examples	207
6.1.1. A Point-To-Point Example	207
6.1.2. A Pub-Sub Example	209
6.1.3. A Pub-Sub With Durable Topic Example	214
6.1.4. A Point-To-Point With MDB Example	216
6.2. JBoss Messaging Overview	222
6.2.1. Invocation Layer	222
6.2.1.1. RMI IL (deprecated)	222
6.2.1.2. OIL IL (deprecated)	222
6.2.1.3. UIL IL (deprecated)	222
6.2.1.4. UIL2 IL	222
6.2.1.5. JVM IL	223
6.2.1.6. HTTP IL	223
6.2.2. Security Manager	223
6.2.3. Destination Manager	223
6.2.4. Message Cache	223
6.2.5. State Manager	223
6.2.6. Persistence Manager	223
6.2.6.1. File PM	224
6.2.6.2. Rolling Logged PM	224
6.2.6.3. JDBC2 PM	224
6.2.7. Destinations	224
6.2.7.1. Queues	224
6.2.7.2. Topics	224
6.3. JBoss Messaging Configuration and MBeans	224
6.3.1. org.jboss.mq.il.jvm.JVMServerILService	225
6.3.2. org.jboss.mq.il.rmi.RMIServerILService (deprecated)	226
6.3.3. org.jboss.mq.il.oil.OILServerILService (deprecated)	226
6.3.4. org.jboss.mq.il.uil.UILServerILService (deprecated)	227
6.3.5. org.jboss.mq.il.uil2.UILServerILService	228
6.3.5.1. Configuring ILs for SSL	229
6.3.5.2. JMS client properties for the UIL2 transport	229
6.3.6. org.jboss.mq.il.http.HTTPServerILService	230
6.3.7. org.jboss.mq.server.jmx.Invoker	231
6.3.8. org.jboss.mq.server.jmx.InterceptorLoader	231
6.3.9. org.jboss.mq.sm.file.DynamicStateManager	231
6.3.10. org.jboss.mq.security.SecurityManager	232
6.3.11. org.jboss.mq.server.jmx.DestinationManager	234
6.3.12. org.jboss.mq.server.MessageCache	235
6.3.13. org.jboss.mq.pm.file.CacheStore	236
6.3.14. org.jboss.mq.pm.file.PersistenceManager	236
6.3.15. org.jboss.mq.pm.rollinglogged.PersistenceManager	236
6.3.16. org.jboss.mq.pm.jdbc2.PersistenceManager	236
6.3.17. Destination MBeans	238
6.3.17.1. org.jboss.mq.server.jmx.Queue	238
6.3.17.2. org.jboss.mq.server.jmx.Topic	239
6.3.18. Administration Via JMX	241
6.3.18.1. Creating Queues At Runtime	241
6.3.18.2. Creating Topics At Runtime	241
6.3.18.3. Managing a JBossMQ User IDs at Runtime	241
6.4. Specifying the MDB JMS Provider	241

6.4.1. org.jboss.jms.jndi.JMSProviderLoader MBean	242
6.4.2. org.jboss.jms.asf.ServerSessionPoolLoader MBean	244
6.4.3. Integrating non-JBoss JMS Providers	244
7. Connectors on JBoss	245
7.1. JCA Overview	245
7.2. An Overview of the JBossCX Architecture	247
7.2.1. BaseConnectionManager2 MBean	248
7.2.2. RARDeployment MBean	249
7.2.3. JBossManagedConnectionPool MBean	250
7.2.4. CachedConnectionManager MBean	251
7.2.5. A Sample Skeleton JCA Resource Adaptor	251
7.3. Configuring JCA Adaptors	257
7.3.1. Configuring JDBC DataSources	257
7.3.2. Configuring Generic JCA Adaptors	264
7.3.3. Sample Configurations	267
8. Security on JBoss	268
8.1. J2EE Declarative Security Overview	268
8.1.1. Security References	270
8.1.2. Security Identity	271
8.1.3. Security roles	272
8.1.4. EJB method permissions	273
8.1.5. Web Content Security Constraints	275
8.1.6. Enabling Declarative Security in JBoss	276
8.2. An Introduction to JAAS	276
8.2.1. What is JAAS?	276
8.2.1.1. The JAAS Core Classes	277
8.3. The JBoss Security Model	280
8.3.1. Enabling Declarative Security in JBoss Revisited	283
8.4. The JBoss Security Extension Architecture	286
8.4.1. How the JaasSecurityManager Uses JAAS	288
8.4.2. The JaasSecurityManagerService MBean	291
8.4.3. The JaasSecurityDomain MBean	293
8.4.4. An XML JAAS Login Configuration MBean	294
8.4.5. The JAAS Login Configuration Management MBean	296
8.4.6. Using and Writing JBossSX Login Modules	296
8.4.6.1. org.jboss.security.auth.spi.IdentityLoginModule	296
8.4.6.2. org.jboss.security.auth.spi.UsersRolesLoginModule	297
8.4.6.3. org.jboss.security.auth.spi.LdapLoginModule	299
8.4.6.4. org.jboss.security.auth.spi.DatabaseServerLoginModule	302
8.4.6.5. BaseCertLoginModule	304
8.4.6.6. org.jboss.security.auth.spi.ProxyLoginModule	306
8.4.6.7. org.jboss.security.auth.spi.RunAsLoginModule	307
8.4.6.8. org.jboss.security.ClientLoginModule	307
8.4.7. Writing Custom Login Modules	308
8.4.7.1. Support for the Subject Usage Pattern	309
8.4.7.2. A Custom LoginModule Example	312
8.4.8. The DynamicLoginConfig service	315
8.5. The Secure Remote Password (SRP) Protocol	316
8.5.1. Providing Password Information for SRP	319
8.5.2. Inside of the SRP algorithm	321
8.5.2.1. An SRP example	323
8.6. Running JBoss with a Java 2 security manager	325

8.7. Using SSL with JBoss using JSSE	328
8.8. Configuring JBoss for use Behind a Firewall	332
8.9. How to Secure the JBoss Server	333
8.9.1. The jmx-console.war	333
8.9.2. The web-console.war	334
8.9.3. The http-invoker.sar	334
8.9.4. The jmx-invoker-adaptor-server.sar	334
9. Integrating Servlet Containers	335
9.1. The AbstractWebContainer Class	335
9.1.1. The AbstractWebContainer Contract	337
9.1.2. Creating an AbstractWebContainer Subclass	341
9.1.2.1. Use the Thread Context Class Loader	341
9.1.2.2. Integrate Logging Using log4j	341
9.1.2.3. Delegate web container authentication and authorization to JBossSX	342
9.2. JBoss/Tomcat-5 bundle notes	343
9.2.1. The Tomcat server.xml file	343
9.2.1.1. Connector	344
9.2.1.2. Engine	346
9.2.1.3. Host	346
9.2.1.4. DefaultContext	346
9.2.1.5. Logger	347
9.2.1.6. Valve	347
9.2.2. Using SSL with the JBoss/Tomcat bundle	347
9.2.3. Setting up Virtual Hosts	350
9.2.4. Serving Static Content	351
9.2.5. Using Apache with the Tomcat	351
9.2.6. Using Clustering	352
10. MBean Services Miscellany	354
10.1. System Properties Management	354
10.2. Property Editor Management	355
10.3. Services Binding Management	355
10.3.1. Running Two JBoss Instances	357
10.4. Scheduling Tasks	361
10.4.1. org.jboss.varia.scheduler.Scheduler	361
10.5. The JBoss Logging Framework	364
10.5.1. org.jboss.logging.Log4jService	365
10.6. RMI Dynamic Class Loading	365
10.6.1. org.jboss.web.WebService	365
11. The CMP Engine	366
11.1. Getting Started	366
11.1.1. Example Code	366
11.1.2. Tests	368
11.1.3. Read-ahead	370
11.2. The jbosscmp-jdbc Structure	370
11.3. Entity Beans	372
11.3.1. Entity Mapping	374
11.4. CMP-Fields	376
11.4.1. CMP-Field Abstract Accessors	376
11.4.2. CMP-Field Declaration	377
11.4.3. CMP-Field Column Mapping	377
11.4.4. Read-only Fields	378
11.4.5. Auditing Entity Access	379

11.4.6. Dependent Value Classes (DVCs)	380
11.5. Container Managed Relationships	383
11.5.1. CMR-Field Abstract Accessors	384
11.5.2. Relationship Declaration	384
11.5.3. Relationship Mapping	385
11.5.3.1. Relationship Role Mapping	387
11.5.3.2. Foreign Key Mapping	389
11.5.3.3. Relation-table Mapping	389
11.6. Queries	391
11.6.1. Finder and ejbSelect Declaration	392
11.6.2. EJB-QL Declaration	392
11.6.3. Overriding the EJB-QL to SQL Mapping	393
11.6.4. JBossQL	394
11.6.5. DynamicQL	395
11.6.6. DeclaredSQL	396
11.6.6.1. Parameters	399
11.6.7. EJBQL 2.1 and SQL92 queries	400
11.6.8. BMP Custom Finders	400
11.7. Optimized Loading	401
11.7.1. Loading Scenario	401
11.7.2. Load Groups	403
11.7.3. Read-ahead	403
11.7.3.1. on-find	404
11.7.3.2. on-load	408
11.7.3.3. none	409
11.8. Loading Process	410
11.8.1. Commit Options	410
11.8.2. Eager-loading Process	410
11.8.3. Lazy loading Process	412
11.8.3.1. Relationships	412
11.8.4. Lazy loading result sets	415
11.9. Transactions	415
11.10. Optimistic Locking	417
11.11. Entity Commands and Primary Key Generation	421
11.11.1. Existing Entity Commands	422
11.12. Defaults	424
11.12.1. A sample jbossCMP-jdbc.xml defaults declaration	426
11.13. Datasource Customization	427
11.13.1. Function Mapping	430
11.13.2. Type Mapping	430
11.13.3. User Type Mappings	431
A. The JBoss Group and Our LGPL License	433
A.1. About The JBoss Group	433
A.2. The GNU Lesser General Public License (LGPL)	433
B. Book Example Installation	441

Forward

— Marc Fluery

If you are reading this foreword, first of all I want to thank you for buying our products. This is one of the ways in which you can support the development effort and ensure that JBoss continues to thrive and deliver the most technologically advanced web application server possible. The time this book was written corresponds to an interesting point in the evolution of Open Source. There are many projects out there and once the initial excitement has faded, the will to continue requires some professional dedication. JBoss seeks to define the forefront of "Professional Open Source" through commercial activities that subsidize the development of the free core product.

JBoss' modules are growing fast. The JMX base allows us to integrate all these disparate modules together using the MBeanServer of JMX as the basic abstraction for their life cycle and management. In this book, we cover the configuration and administration of all our MBeans. We also provide a comprehensive snapshot of the state of JBoss server modules, documented in a professional fashion by one of our very best developers. From the basic architecture, to the advanced modules like JBossSX for security and our CMP engine, you will find the information you need "to get the job done." In addition, we provide a wealth of information on all the modules you will want to understand better and eventually master as you progress in your day-to-day usage of JBoss.

JBoss has achieved a reputation for technical savvy and excellence. I would like this reputation to evolve a bit. Don't get me wrong, I am extremely proud of the group of people gathered around JBoss for the past 3+ years, but I want to make the circle bigger. I want to include all of you reading this book. Think of JBoss, not only as a great application server, but also as a community that thrives by the addition of new minds. We are not simply interested in gaining users; we are interested in giving you the tools and the knowledge necessary to master our product to the point of becoming a contributor. Understanding JBoss' configuration and architecture is a necessary step, not only for your day job using JBoss in development and production, but also an initiation into the joy of technology, as experienced in Open Source.

We hope this book will fulfill its potential to bring as many of you as possible to a strong enough understanding of the modules' functionality to dream up new tools and new functionalities, maybe even new modules. When you reach that point, make sure to come online, where you will find a thriving community of committed professionals sharing a passion for good technology. At www.jboss.org [<http://www.jboss.org>], you can also find additional information, forums, and the latest binaries.

Again thank you for buying our documentation. We hope to see you around. In the meantime, learn, get the job done and, most of all, enjoy,

About Open Source

The basic idea behind open source is very simple: When programmers can read, redistribute, and modify the source code for a piece of software, the software evolves. People improve it, people adapt it, people fix bugs. And this can happen at a speed that, if one is used to the slow pace of conventional software development, seems astonishing. Open Source is an often-misunderstood term relating to free software. The Open Source Initiative (OSI) web site provides a number of resources that define the various aspects of Open Source including an Open Source Definition at: <http://www.opensource.org/docs/definition.html>. The following quote from the OSI home page summarizes the key aspects as they relate to JBoss nicely:

We in the open source community have learned that this rapid evolutionary process produces better software than the traditional closed model, in which only a very few programmers can see the source and everybody else must blindly use an opaque block of bits.

Open Source Initiative exists to make this case to the commercial world.

Open source software is an idea whose time has finally come. For twenty years it has been building momentum in the technical cultures that built the Internet and the World Wide Web. Now it's breaking out into the commercial world, and that's changing all the rules. Are you ready?

—The Open Source Initiative

About JBoss

JBoss, Inc., headed by Marc Fleury, is composed of over 100 developers worldwide who are working to deliver a full range of J2EE tools, making JBoss the premier Enterprise Java application server for the Java 2 Enterprise Edition platform.

JBoss is an Open Source, standards-compliant, J2EE application server implemented in 100% Pure Java. The JBoss/Server and complement of products are delivered under a public license. With upwards of 100,000 downloads per month, JBoss is the most downloaded J2EE based server in the industry.

JBoss, one of the leading Java Open Source groups, integrates and develops these services for a full J2EE-based implementation. JBoss provides JBossServer, the basic EJB container, and Java Management Extension (JMX) infrastructure. It also provides JBossMQ, for JMS messaging, JBossTX, for JTA transactions, JBossCMP for CMP persistence, JBossSX for JAAS based security, and JBossCX for JCA connectivity. Support for web components, such as servlets and JSP pages, is provided by an abstract integration layer. Implementations of the integration service are provided for third party servlet engines like Tomcat and Jetty. JBoss enables you to mix and match these components through JMX by replacing any component you want with a JMX compliant implementation for the same APIs. JBoss doesn't even impose the JBoss components. Now that is modularity.

1. JBoss: A Full J2EE Implementation with JMX

Our goal is to provide the full Open Source J2EE stack. We have met our goal, and the reason for our success lies on JMX. JMX, or Java Management Extension, is the best weapon we have found for integration of software. JMX provides a common spine that allows one to integrate modules, containers, and plug-ins. Figure 1 illustrates how JMX is used a bus through which the components of the JBoss architecture interact.




Figure 1. The JBoss JMX integration bus and the standard JBossXX components

2. What this Book Covers

The primary focus of this book is the presentation of the standard JBoss 3.2.x architecture components from both the perspective of their configuration and architecture. As a user of a standard JBoss distribution you will be given an understanding of how to configure the standard components. Note that this book is not an introduction to J2EE or how to use J2EE in applications. It focuses on the internal details of the JBoss server architecture and how our implementation of a given J2EE container can be configured and extended.

As a JBoss developer, you will be given a good understanding of the architecture and integration of the standard components to enable you to extend or replace the standard components for your infrastructure needs. We also show you how to obtain the JBoss source code, along with how to build and debug the JBoss server.

Installing and Building the JBoss Server

JBoss, a free J2EE-based application server, is the most widely used Open Source application server on the market. The highly flexible and easy-to-use server architecture has made JBoss the ideal choice for users just starting out with J2EE, as well as senior architects looking for a customizable middleware platform. The server binary and source code distributions are available from the SourceForge repository. (<http://sourceforge.net/projects/jboss>). The ready availability of the source code allows you to debug the server, learn its inner workings and create customized versions for your personal or business use.

This chapter is a step-by-step tutorial about how to install and configure JBoss 3.2. You will learn how to:

- Obtain updated binaries from the JBoss SourceForge project site
- Install the binary
- Test the installation

You will also learn about:

- The installation directory structure
- Key configuration files an administrator may want to use to customize the JBoss installation
- Obtaining the source code for the 3.2.x release from the SourceForge CVS repository
- Building the server distribution.

1.1. Getting the Binary Files

The most recent release of JBoss is available from the SourceForge JBoss project files page, <http://sourceforge.net/projects/jboss>. You will also find previous releases as well as beta and release candidate versions of upcoming releases.

1.1.1. Prerequisites

Before installing and running the server, check your system to make sure you have a working JDK 1.3+ installation. The simplest way to do this is to execute the `java -version` command to ensure that the java executable is in your path, and that you are using Version 1.3 or higher. For example, running this command with a 1.3.1 JDK would produce version number like the following.

```
[nr@toki tmp]$ java -version
java version "1.3.1"
Java(TM) 2 Runtime Environment, Standard Edition (build 1.3.1-root_1.3.1_031103-17:49)
Java HotSpot(TM) Client VM (build 1.3.1_03-76, mixed mode)
```

It does not matter where on your system you install JBoss. Note, however, that installing JBoss into a directory that has a name containing spaces causes problems in some situations with Sun-based VMs. This is caused by bugs with file URLs not correctly escaping the spaces in the resulting URL. There is no requirement for root access to run JBoss on UNIX/Linux systems because none of the default ports are within the 0-1023 privileged port range.

1.1.2. Installing the Binary Package

After you have the binary archive you want to install, use the JDK jar tool (or any other ZIP extraction tool) to extract the `jboss-3.2.6.zip` archive contents into a location of your choice. The `jboss-3.2.6.tgz` archive is a gzipped tar file that requires a gnutar compatible tar which can handle the long pathnames in the archive. The default tar binaries on Solaris and OSX do not currently support the long pathnames. The extraction process will create a `jboss-3.2.6` directory. The following section explores the contents of this directory.

1.1.2.1. Directory Structure

As mentioned above, installing the JBoss distribution creates a `jboss-3.2.6` directory which contains server start scripts, JARs, server configuration sets and working directories. You do need to know your way around the distribution layout to locate JARs for compilation, updating configurations, deploying your code, etc. Figure 1.1 illustrates the installation directory of the JBoss server.

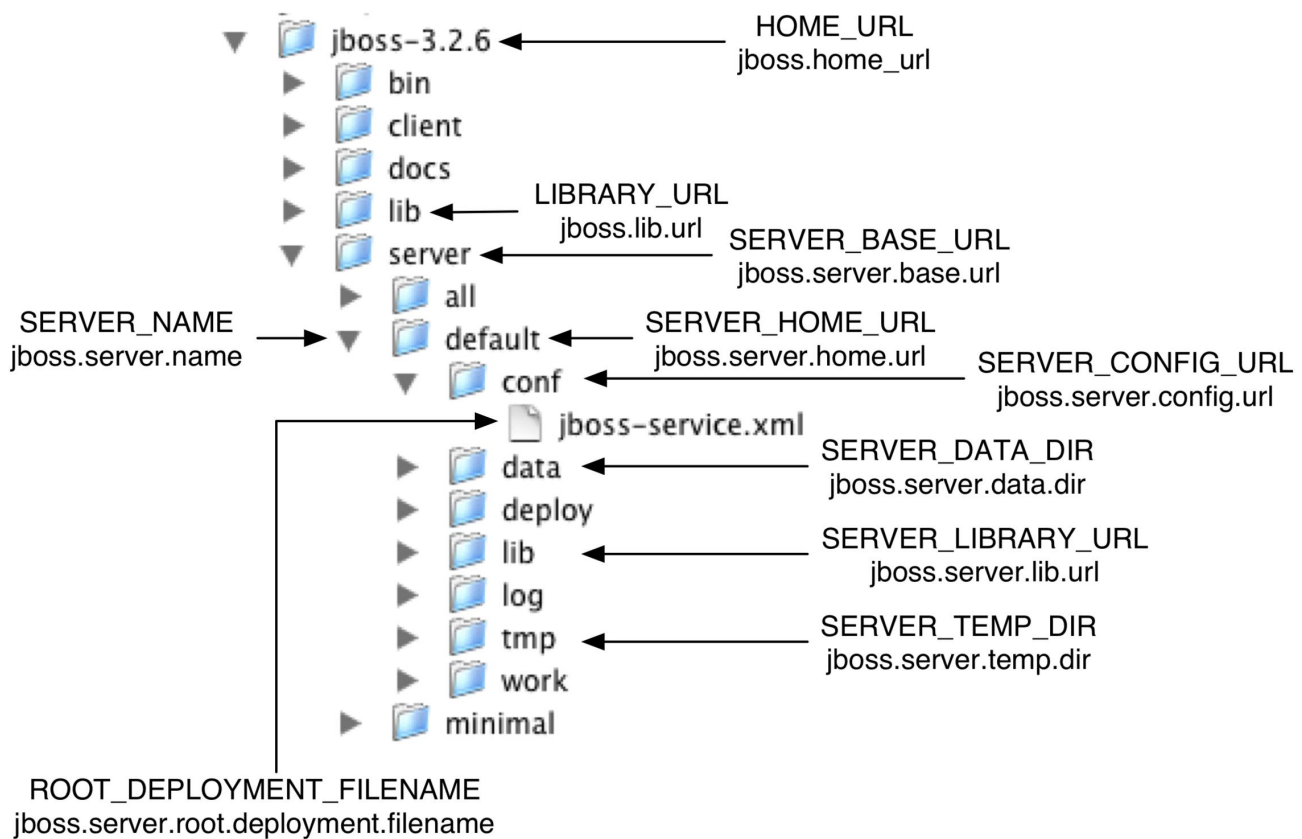


Figure 1.1. The view of the JBoss server installation directory structure with the default server configuration file set expanded and overridable locations identified

Throughout the book we will refer to the top-level `jboss-3.2.6` directory as the `JBOSS_DIST` directory. In Figure 1.1, the default server configuration file set is shown expanded. It contains a number of subdirectories: `conf`, `data`, `deploy`, `lib`, `log` and `tmp`. In a clean installation, only the `conf`, `deploy` and `lib` directories will exist. The purpose of each of these directories are discussed in Table 1.1. In this table, the *ServerConfig Property* column refers to the `org.jboss.system.server.ServerConfig` interface constant and its corresponding system property string. The *ServerConfig* constant names and corresponding system property name are displayed in the blue text in Figure 1.1. The `xxx_URL` names correspond to locations that can be specified using a URL to access remote locations, for example, HTTP URLs against a web server. You can use the properties listed in the following table to override the layout of a JBoss distribution

Table 1.1. The JBoss directory structure

Directory	Description	ServerConfig Property
bin	All the entry point JARs and start scripts included with the JBoss distribution are located in the <code>bin</code> directory.	
client	JARs required for clients are located in the <code>client</code> directory. A typical client requires the following: <ul style="list-style-type: none"> • <code>jbossall-client.jar</code> • <code>concurrent.jar</code> • <code>log4j.jar</code> • <code>jaas.jar</code>, <code>jnet.jar</code> (If not using JDK1.4+) • <code>jcrt.jar</code>, <code>jsse.jar</code> (for SSL if not using JDK1.4+) 	
server	The JBoss server configuration sets are located under the <code>server</code> directory. The default server configuration set is the <code>server/default</code> set. JBoss ships with minimal, default and all configuration sets. The sub-directories and key configuration files contained in the default configuration set are discussed in more detail in Section 1.1.3	<code>SERVER_BASE_DIR =</code> <code>"jboss.server.base.dir"</code> <code>SERVER_BASE_URL =</code> <code>"jboss.server.base.url"</code>
lib	The <code>lib</code> directory contains startup JARs used by JBoss. Do not place your own libraries in this directory.	<code>LIBRARY_URL =</code> <code>"jboss.lib.url"</code>
conf	The <code>conf</code> directory contains the bootstrap descriptor, <code>jboss-service.xml</code> by default, file for a given server configuration. This defines the core services that are fixed for the lifetime of the server.	<code>SERVER_CONFIG_URL =</code> <code>"jboss.server.config.url"</code>
data	The <code>data</code> directory is a location available for use by services that want to store content in the file system.	<code>SERVER_DATA_DIR =</code> <code>"jboss.server.data.dir"</code>
deploy	The <code>deploy</code> directory is the default location the hot deployment service looks to for dynamic deployment content. This may be overridden through the <code>URLDeploymentScanner</code> <code>URLs</code> attribute.	
lib	The <code>lib</code> directory is the default location referred to the by bootstrap descriptor. All JARs in this directory are loaded into the shared classpath.	<code>SERVER_LIBRARY_URL =</code> <code>"jboss.server.lib.url"</code>
log	The <code>log</code> directory is the default directory into which the bootstrap logging service places its logs. This may be overridden through the <code>conf/log4j.xml</code> configuration file.	none
tmp	The <code>tmp</code> directory is the location to which deployments are copied for local use.	<code>SERVER_TEMP_DIR=</code> <code>"jboss.server.temp.dir"</code>

1.1.3. The Default Server Configuration File Set

The `JBOSS_DIST/server` directory contains one or more configuration file sets. The default JBoss configuration file set is located in the `JBOSS_DIST/server/default` directory. JBoss allows you to add more than one configuration set so a server can easily be run using alternate configurations. Creating a new configuration file set typically starts with copying the default file set into a new directory name and then modifying the configuration files as desired. Figure 1.2 below shows the contents of the default configuration file set.

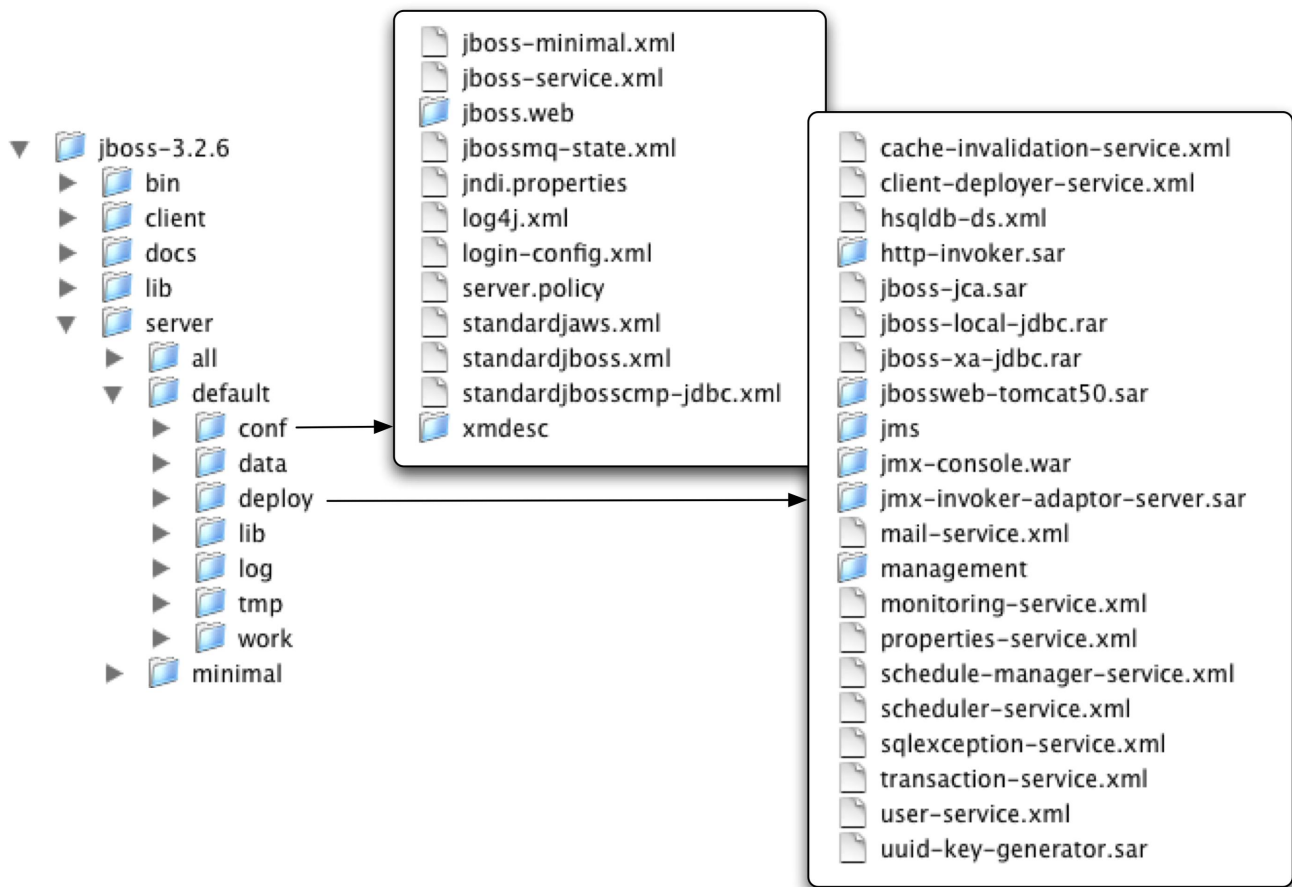


Figure 1.2. An expanded view of the default server configuration file set `conf` and `deploy` directories

conf/jboss-minimal.xml

This is a minimalist example of the `jboss-service.xml` configuration file. It is the `jboss-service.xml` file used in the `minimal` configuration file set.

conf/jboss-service.xml

The `jboss-service.xml` defines the core services configurations. The complete DTD and syntax of this file is described, along with the details on integrating custom services, in Section 2.4.2.

conf/jbossmq-state.xml

The `jbossmq-state.xml` is the JBossMQ configuration file that specifies the user to password mappings file, and the user to durable subscription. The format of this file is described in `org.jboss.mq.sm.file.DynamicStateManager`.

conf/jndi.properties

The `jndi.properties` file specifies the JNDI `InitialContext` properties that are used within the JBoss server when an `InitialContext` is created using the `no-arg` constructor.

conf/log4j.xml

The file configures the Apache log4j framework category priorities and appenders used by the JBoss server code. See the JBoss/Log4j book for details on configuring and using log4j with JBoss.

conf/login-config.xml

The `login-config.xml` file contains sample server side authentication configurations that are applicable when using JAAS based security. See Chapter 8 for additional details on the JBoss security framework and the format of this file.

conf/server.policy

The `server.policy` file is a place holder for Java2 security permissions. The default file simply grants permissions to all codebases.

conf/standardjaws.xml

The `standardjaws.xml` provides a default configuration file for the legacy EJB 1.1 JBossCMP engine. The `cmp` layer was rewritten in JBoss 3.0 to support EJB 2.0 and is fully documented for the 3.2 version in the chapter entitled The CMP Engine.

conf/standardjbosscmp-jdbc.xml

The `standardjbosscmp-jdbc.xml` provides a default configuration file for the JBoss 3.2 EJB 2.0 JBossCMP engine. See Chapter 11 for the details of this descriptor.

conf/standardjboss.xml

The `standardjboss.xml` file provides the default container configurations. Use of this file is covered in Chapter 5

conf/xmdesc/ClientTransaction-mbean.xml, JNDIView-xmbean.xml, TransactionManagerService-xmbean.xml

The `xmdesc` directory contains XBean descriptors for several services configured in the `jboss-service.xml` file.

deploy/cache-invalidation-service.xml

The `cache-invalidation-service.xml` is a service that allows for custom invalidation of the EJB caches via JMS notifications. It is disabled by default.

deploy/client-deployer-service.xml

The `client-deployer-service.xml` is a service that provides support for J2EE application clients. It manages the `java:comp/env` enterprise naming context for client applications based on the `application-client.xml` descriptor.

deploy/hsqldb-ds.xml

The `hsqldb-ds.xml` is the Hypersonic 1.7.1 embedded database service configuration file. It sets up the embedded database and related connection factories. The format of JCA datasource files is discussed in Section 7.3.1.

deploy/http-invoker.sar

The `http-invoker.sar` contains the detached invoker that supports RMI/HTTP. It also setups bindings of RMI/HTTP proxies for the JNDI naming service that allow the JBoss JNDI service to be accessed over http. This will be discussed in Section 2.7.2.5

deploy/jboss-jca.sar

The `jboss-jca.sar` is the application server implementation of the 1.0 JCA specification. It provides the connection management facilities for integrating resource adaptors into the JBoss server. The JCA layer is discussed in Chapter 7.

deploy/jboss-local-jdbc.rar

The `jboss-local-jdbc.rar` is a JCA resource adaptor that implements the JCA `ManagedConnectionFactory` interface for JDBC drivers that support the `DataSource` interface but not JCA.

deploy/jboss-xa.rar

The `jboss-xa.rar` is a JCA resource adaptor that implements the JCA `ManagedConnectionFactory` interface for JDBC drivers that support the `XADataSource` interface, but do not provide JCA adaptor.

deploy/jbossweb-tomcat50.sar

The `jbossweb-tomcat50.sar` directory is an unpacked MBean service archive for the configuration of the Tomcat 5 servlet engine. The SAR is unpacked rather than deployed as a JAR archive so that the tomcat configuration files can be easily edited. This service is discussed in Chapter 9.

deploy/jms/jbossmq-destinations-service.xml

The `jbossmq-destinations-service.xml` file configures a number of JMS queues and topics used by the JMS unit tests. Configuration of JMS destinations is discussed in Chapter 6.

deploy/jms/jbossmq-httpil.sar

The `jbossmq-httpil.sar` provides a JMS invocation layer that allows the use of JMS over http.

deploy/jms/jbossmq-service.xml

The `jbossmq-service.xml` file configures the core JBossMQ JMS service. These JMS services are discussed in Chapter 6.

deploy/jms/jms-ra.rar

The `jms-ra.rar` is a JCA resource adaptor that implements the JCA `ManagedConnectionFactory` interface for JMS connection factories.

deploy/jms/jms-ds.xml

The `jms-ds.xml` file configures the JBossMQ JMS provider for use with the `jms-ra.rar` JCA resource adaptor.

deploy/jms/jvm-il-service.xml

The `jvm-il-service.xml` configures the inter-vm JMS transport invocation layer. This transport layer is described in `org.jboss.mq.il.jvm.JVMServerILService`.

deploy/jms/oil-service.xml

The `oil-service.xml` configures the JMS optimized invocation layer. This transport layer is described in `org.jboss.mq.il.oil.OILServerILService`. (deprecated)

deploy/jms/oil2-service.xml

The `oil2-service.xml` configures the JMS version 2 optimized invocation layer. It is an experimental work that should not be used and will be dropped in the future.

deploy/jms/rmi-il-service.xml

The `rmi-il-service.xml` configures the JMS RMI based invocation layer. This is a slow transport layer that should not be used and will be dropped in the future.

deploy/jms/ui12-service.xml

The `ui12-service.xml` configures the JMS version 2 unified invocation layer. Its a custom socket based transport that is the fastest and most reliable and should be used for inter-vm messaging. This transport layer is described in `org.jboss.mq.il.ui12.UILServerILService`.

deploy/jmx-console.war

The `jmx-console.war` directory is an unpackaged web application archive that provides an HTML adaptor for the JMX `MBeanServer`. The WAR is unpackaged rather than deployed as a jar archive so that the `jmx-`

`console.war/WEB-INF/*.xml` descriptors may be edited to configure role based security easily. The `jmx-console` is discussed in Section 2.3.1

deploy/jmx-invoker-adaptor-server.sar

The `jmx-invoker-adaptor-server.sar` is an unpacked MBean service archive that exposes a subset of the `JMX MBeanServer` interface methods as an RMI interface to enable remote access to the JMX core functionality. This is similar to the legacy `jmx-rmi-adaptor.sar`, with the difference that the transport is handled by the detached invoker architecture. This service is discussed in Section 2.3.4.

deploy/mail-service.xml

The `mail-service.xml` file is an MBean service descriptor that provides JavaMail sessions for use inside of the JBoss server.

deploy/management/console-mgr.sar, web-console.war

The `console-mgr.sar` and `web-console.war` are an experimental web application/applet that provide a richer view of the JMX server management data than the JMX console. At this time they are still under development. You may view the console using the URL `http://localhost:8080/web-console/`.

deploy/monitoring-service.xml

The `monitoring-service.xml` file configures alert monitors like the console listener and email listener used by JMX notifications.

deploy/properties-service.xml

The `properties-service.xml` file is an MBean service descriptor that allows for customization of the `JavaBeans PropertyEditors` as well as the definition of system properties. This is discussed further in Section 10.1.

deploy/scheduler-service.xml, schedule-manager-service.xml

The `scheduler-service.xml`, `schedule-manager-service.xml` files are MBean service descriptors that provides a scheduling type of service. This is discussed further in Section 10.4.

deploy/sqlexception-service.xml

The `sqlexception-service.xml` file is an MBean service descriptor for handling vendor specific handling of `java.sql.SQLExceptions`. Its usage is discussed in Section 11.11.

deploy/transaction-service.xml

The `transaction-service.xml` service descriptor sets up the JBoss JTA transaction manager related services. This and related services are discussed in more detail in Chapter 4.

deploy/user-service.xml

The `user-service.xml` file is a template MBean service descriptor to which you may add your own custom MBean services. Its not necessary to use this file for this purpose however. Writing MBean services is discussed in Section 2.4.3.

deploy/uuid-key-generator.sar

The `uuid-key-generator.sar` service provides a UUID based key generation facility.

1.2. Basic Installation Testing

Once you have installed the JBoss distribution, it is wise to perform a simple startup test to validate that there are no major problems with your Java VM/operating system combination. To test your installation, move to the `JBOSS_DIST/bin` directory and execute the `run.bat` or `run.sh` script as appropriate for your operating system. Your output should be similar to that shown below and contain no error or exception messages:

```
[orb@toki bin]$ sh run.sh
=====

JBoss Bootstrap Environment

JBOSS_HOME: /tmp/jboss-3.2.6

JAVA: /System/Library/Frameworks/JavaVM.framework/Home/bin/java

JAVA_OPTS: -Dprogram.name=run.sh

CLASSPATH: /tmp/jboss-3.2.6/bin/run.jar:/System/Library/Frameworks/JavaVM.framework/Home
//lib/tools.jar

=====

11:41:32,879 INFO [Server] Starting JBoss (MX MicroKernel)...
11:41:32,898 INFO [Server] Release ID: JBoss [WonderLand] 3.2.6 (build: CVSTag=JBoss_3_2_
6 date=200410140106)
11:41:32,902 INFO [Server] Home Dir: /private/tmp/jboss-3.2.6
11:41:33,039 INFO [Server] Home URL: file:/private/tmp/jboss-3.2.6/
11:41:33,043 INFO [Server] Library URL: file:/private/tmp/jboss-3.2.6/lib/
11:41:33,140 INFO [Server] Patch URL: null
11:41:33,147 INFO [Server] Server Name: default
11:41:33,237 INFO [Server] Server Home Dir: /private/tmp/jboss-3.2.6/server/default
11:41:33,250 INFO [Server] Server Home URL: file:/private/tmp/jboss-3.2.6/server/default/

11:41:33,256 INFO [Server] Server Data Dir: /private/tmp/jboss-3.2.6/server/default/data
11:41:33,260 INFO [Server] Server Temp Dir: /private/tmp/jboss-3.2.6/server/default/tmp
11:41:33,266 INFO [Server] Server Config URL: file:/private/tmp/jboss-3.2.6/server/default/conf/
11:41:33,270 INFO [Server] Server Library URL: file:/private/tmp/jboss-3.2.6/server/default/lib/
11:41:33,276 INFO [Server] Root Deployment Filename: jboss-service.xml
11:41:33,342 INFO [Server] Starting General Purpose Architecture (GPA)...
11:41:35,176 INFO [ServerInfo] Java version: 1.4.2_05,Apple Computer, Inc.
11:41:35,180 INFO [ServerInfo] Java VM: Java HotSpot(TM) Client VM 1.4.2-38,"Apple Computer, Inc."
11:41:35,190 INFO [ServerInfo] OS-System: Mac OS X 10.3.5,ppc
11:41:37,259 INFO [Server] Core system initialized
```

If your output is similar to this (accounting for installation directory differences), you should now be ready to use JBoss. To shutdown the server, simply issue a Ctrl-C sequence in the console in which JBoss was started. Alternatively, you can use the `shutdown.sh` command:

```
[nr@toki bin]$ ./shutdown.sh
A JMX client to shutdown (exit or halt) a remote JBoss server.

usage: shutdown [options] <operation>

options:
-h, --help Show this help message
-D<name>[=<value>] Set a system property
-- Stop processing options
-s, --server=<url> Specify the JNDI URL of the remote server
-n, --serverName=<url> Specify the JMX name of the ServerImpl
-a, --adapter=<name> Specify JNDI name of the RMI adapter to use
-u, --user=<name> Specify the username for authentication[not implemented yet]
-p, --password=<name> Specify the password for authentication[not implemented yet]

operations:
-S, --shutdown Shutdown the server (default)
-e, --exit=<code> Force the VM to exit with a status code
-H, --halt=<code> Force the VM to halt with a status code
```

Using `run.sh` without any arguments starts the server using the default server configuration file set. To start with an alternate configuration file set pass in the name of the directory under `JBOSS_DIST/server` you wish to

use as the value to the `-c` command line option. For example, to start with the minimal configuration file set you would specify:

```
[nr@toki bin]$ ./run.sh -c minimal
...
22:26:49,566 INFO [Server] JBoss (MX MicroKernel) [3.2.6RC2 (build: CVSTag=Branch_3_2 dat
e=200409270100)] Started in 11s:744ms
```

To view all of the supported command line options for the JBoss server bootstrap class issue `run -h` command, and the output will be:

```
usage: run.sh [options]

options:
  -h, --help                Show this help message
  -V, --version              Show version information
  --                        Stop processing options
  -D<name>[=<value>]        Set a system property
  -p, --patchdir=<dir>      Set the patch directory; Must be absolute
  -n, --netboot=<url>       Boot from net with the given url as base
  -c, --configuration=<name> Set the server configuration name
  -j, --jaxp=<type>         Set the JAXP impl type (ie. crimson)
  -L, --library=<filename>  Add an extra library to the loaders classpath
  -C, --classpath=<url>     Add an extra url to the loaders classpath
  -P, --properties=<url>    Load system properties from the given url
  -b, --host=<host or ip>   Bind address for all JBoss services
```

1.3. Booting from a Network Server

One very useful command line option is the `--netboot=url` option which causes JBoss to startup using the given URL as the base URL from which all libraries and configurations are loaded. Specifying the netboot option sets the `ServerConfig.HOME_URL` to the netboot option URL argument value. In the absence of any other overrides, all of the locations found in the standard `JBOSS_DIST` structure of will be resolved relative to the `HOME_URL` value. This means that if you make a JBoss distribution available from a web server you can boot JBoss using only the run scripts and `run.jar` file from the `JBOSS_DIST/bin` directory. Note that the web server must support the `PROPFIND` WebDAV command. JBoss includes a simple servlet filter that provides a minimal support for the `PROPFIND` command so that JBoss itself may be used as the netboot web server.

An example Ant build script that creates a custom netboot configuration fileset for booting the default configuration is available in the book `examples/src/main/org/jboss/chap1/build-netboot.xml` file. To test the netboot feature, run the `build-netboot.xml` script specifying the location of the `JBOSS_DIST` you want to use as the netboot webserver as shown here:

```
[nr@toki examples]$ ant -Djboss.dist=/tmp/jboss-3.2.6 -buildfile src/main/org/jboss/chap1/
build-netboot.xml
Buildfile: src/main/org/jboss/chap1/build-netboot.xml

netboot:
  [mkdir] Created dir: /tmp/jboss-3.2.6/server/netboot
  [copy] Copying 44 files to /tmp/jboss-3.2.6/server/netboot
  [unzip] Expanding: /tmp/jboss-3.2.6/docs/examples/netboot/netboot.war into /tmp/jboss-
3.2.6/server/netboot/deploy/netboot.war
  [copy] Copying 14 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war
  [copy] Copying 211 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server
/default
  [copy] Copied 1 empty directory to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war
/server/default
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/netboot/deploy/jbossweb-tomcat50.sar
/META-INF

zipdir:
  [move] Moving 10 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/d
```

```

efault/deploy/http-invoker.sarx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/http-invoker.sar
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/http-invoker.sarx

zipdir:
  [move] Moving 6 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jms/jbossmq-httpil.sarx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jms/jbossmq-httpil.sar
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jms/jbossmq-httpil.sarx

zipdir:
  [move] Moving 35 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jbossweb-tomcat50.sarx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jbossweb-tomcat50.sar
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jbossweb-tomcat50.sarx

zipdir:
  [move] Moving 22 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jmx-console.warx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jmx-console.war
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jmx-console.warx

zipdir:
  [move] Moving 5 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/management/console-mgr.sarx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/management/console-mgr.sar
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/management/console-mgr.sarx

zipdir:
  [move] Moving 53 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/management/web-console.warx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/management/web-console.war
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/management/web-console.warx

zipdir:
  [move] Moving 2 files to /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jmx-invoker-adaptor-server.sarx
  [zip] Building zip: /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jmx-invoker-adaptor-server.sar
  [delete] Deleting directory /tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/server/default/deploy/jmx-invoker-adaptor-server.sarx

BUILD SUCCESSFUL

```

You then startup the netboot server by specifying the netboot configuration as follows:

```

[nr@toki bin]$ ./run.sh -c netboot
=====

JBoss Bootstrap Environment

JBOSS_HOME: /tmp/jboss-3.2.6

JAVA: /System/Library/Frameworks/JavaVM.framework/Home/bin/java

JAVA_OPTS: -Dprogram.name=run.sh

CLASSPATH: /tmp/jboss-3.2.6/bin/run.jar:/System/Library/Frameworks/JavaVM.framework/Home

```

```
//lib/tools.jar

=====

13:53:38,479 INFO [Server] Starting JBoss (MX MicroKernel)...
13:53:38,488 INFO [Server] Release ID: JBoss [WonderLand] 3.2.6RC2 (build: CVSTag=Branch_3_2 date=200409270100)
13:53:38,503 INFO [Server] Home Dir: /private/tmp/jboss-3.2.6
13:53:38,552 INFO [Server] Home URL: file:/private/tmp/jboss-3.2.6/
13:53:38,555 INFO [Server] Library URL: file:/private/tmp/jboss-3.2.6/lib/
13:53:38,563 INFO [Server] Patch URL: null
13:53:38,568 INFO [Server] Server Name: netboot
13:53:38,598 INFO [Server] Server Home Dir: /private/tmp/jboss-3.2.6/server/netboot
13:53:38,644 INFO [Server] Server Home URL: file:/private/tmp/jboss-3.2.6/server/netboot/
13:53:38,648 INFO [Server] Server Data Dir: /private/tmp/jboss-3.2.6/server/netboot/data
13:53:38,653 INFO [Server] Server Temp Dir: /private/tmp/jboss-3.2.6/server/netboot/tmp
13:53:38,656 INFO [Server] Server Config URL: file:/private/tmp/jboss-3.2.6/server/netboot/conf/
13:53:38,660 INFO [Server] Server Library URL: file:/private/tmp/jboss-3.2.6/server/netboot/lib/
13:53:38,663 INFO [Server] Root Deployment Filename: jboss-service.xml
13:53:38,679 INFO [Server] Starting General Purpose Architecture (GPA)...
13:53:40,057 INFO [ServerInfo] Java version: 1.4.2_05,Apple Computer, Inc.
13:53:40,060 INFO [ServerInfo] Java VM: Java HotSpot(TM) Client VM 1.4.2-38,"Apple Computer, Inc."
13:53:40,131 INFO [ServerInfo] OS-System: Mac OS X 10.3.5,ppc
13:53:41,883 INFO [Server] Core system initialized
13:53:45,468 INFO [Log4jService$URLWatchTimerTask] Configuring from URL: resource:log4j.xml
13:53:47,754 INFO [NamingService] Started jnpPort=1099, rmiPort=1098, backlog=50, bindAddress=/192.168.168.110, Client SocketFactory=null, Server SocketFactory=org.jboss.net.sockets.DefaultSocketFactory@ad093076
13:53:53,702 INFO [Embedded] Catalina naming disabled
13:53:57,796 INFO [Http11Protocol] Initializing Coyote HTTP/1.1 on http-192.168.168.110-8080
13:53:58,102 INFO [Catalina] Initialization processed in 3578 ms
13:53:58,108 INFO [StandardService] Starting service jboss.web
13:53:58,185 INFO [StandardEngine] Starting Servlet Engine: Apache Tomcat/5.0.28
13:53:58,360 INFO [StandardHost] XML validation disabled
13:53:58,575 INFO [Catalina] Server startup in 469 ms
13:53:59,277 INFO [TomcatDeployer] deploy, ctxPath=/, warUrl=file:/private/tmp/jboss-3.2.6/server/netboot/deploy/jbossweb-tomcat50.sar/ROOT.war/
13:54:03,106 INFO [TomcatDeployer] deploy, ctxPath=/netboot, warUrl=file:/private/tmp/jboss-3.2.6/server/netboot/deploy/netboot.war/
13:54:04,111 INFO [Http11Protocol] Starting Coyote HTTP/1.1 on http-192.168.168.110-8080
13:54:04,797 INFO [ChannelSocket] JK2: ajp13 listening on /192.168.168.110:8009
13:54:04,883 INFO [JkMain] Jk running ID=0 time=1/303 config=null
13:54:04,898 INFO [Server] JBoss (MX MicroKernel) [3.2.6RC2 (build: CVSTag=Branch_3_2 date=200409270100)] Started in 25s:64ms
```

You can now startup any other instance of JBoss using just the run script and run.jar from the JBOSS_DIST/bin directory. For example:

```
[orb@rubik bin]$ sh run.sh --netboot=http://192.168.168.110:8080/netboot/
=====

JBoss Bootstrap Environment

JBOSS_HOME: /tmp/jboss-3.2.6

JAVA: /System/Library/Frameworks/JavaVM.framework/Home/bin/java

JAVA_OPTS: -Dprogram.name=run.sh

CLASSPATH: /tmp/jboss-3.2.6/bin/run.jar:/System/Library/Frameworks/JavaVM.framework/Home/lib/tools.jar

=====
```

```

13:55:40,847 INFO [Server] Starting JBoss (MX MicroKernel)...
13:55:40,867 INFO [Server] Release ID: JBoss [WonderLand] 3.2.6RC2 (build: CVSTag=Branch_
3_2 date=200409270100)
13:55:40,874 INFO [Server] Home Dir: /private/tmp/jboss-3.2.6
13:55:40,877 INFO [Server] Home URL: http://192.168.168.110:8080/netboot/
13:55:40,880 INFO [Server] Library URL: http://192.168.168.110:8080/netboot/lib/
13:55:40,928 INFO [Server] Patch URL: null
13:55:40,931 INFO [Server] Server Name: default
13:55:40,934 INFO [Server] Server Home Dir: /private/tmp/jboss-3.2.6/server/default
13:55:40,940 INFO [Server] Server Home URL: http://192.168.168.110:8080/netboot/server/de
fault/
13:55:40,943 INFO [Server] Server Data Dir: /private/tmp/jboss-3.2.6/server/default/data
13:55:40,946 INFO [Server] Server Temp Dir: /private/tmp/jboss-3.2.6/server/default/tmp
13:55:40,950 INFO [Server] Server Config URL: http://192.168.168.110:8080/netboot/server/
default/conf/
13:55:40,954 INFO [Server] Server Library URL: http://192.168.168.110:8080/netboot/server
/default/lib/
13:55:40,957 INFO [Server] Root Deployment Filename: jboss-service.xml
13:55:41,014 INFO [Server] Starting General Purpose Architecture (GPA)...
...

```

The custom netboot configuration fileset consists simply of the files needed to run the `bossweb-tomcat50.sar` web server and a `netboot.war` whose content is the `JBOSS_DIST/lib` and `JBOSS_DIST/server/default` files.

1.4. Building the Server from Source Code

Source code is available for every JBoss module, and you can build any version of JBoss from source by downloading the appropriate version of the code from SourceForge.

1.4.1. Accessing the JBoss CVS Repositories at SourceForge

The JBoss source is hosted at SourceForge, a great Open Source community service provided by VA Linux Systems. With over 88,000 Open Source projects and nearly 950,000 registered users, SourceForge.net is the largest Open Source hosting service available. Many of the top Open Source projects have moved their development to the `sourceforge.net` site. The services offered by SourceForge include hosting of project CVS repositories and a web interface for project management that includes bug tracking, release management, mailing lists and more. Best of all, these services are free to all Open Source developers. For additional details and to browse the plethora of projects, see the SourceForge home page. (<http://sourceforge.net/>).

1.4.2. Understanding CVS

CVS (Concurrent Versions System) is an Open Source version control system that is used pervasively throughout the Open Source community. CVS is a source control or revision control tool designed to keep track of source changes made by groups of developers who are working on the same files. CVS enables developers to stay in sync with each other as each individual chooses.

1.4.3. Anonymous CVS Access

The JBoss project's SourceForge CVS repository can be accessed through anonymous (pserver) CVS with the following instruction set. The module you want to check out must be specified as the `modulename`. When prompted for a password for `anonymous`, simply press the Enter key. The general syntax of the command line version of CVS for anonymous access to the JBoss repositories is:

```

cvs -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/jboss login
cvs -z3 -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/jboss co modulename

```


The first command logs into JBoss CVS repository as an anonymous user. This command only needs to be performed once for each machine on which you use CVS because the login information will be saved in your `HOME/.cvspass` file or equivalent for your system. The second command checks out a copy of the `modulename` source code into the directory from which you run the `cvs` command. To avoid having to type the long `cvs` command line each time, you can set up a `CVSROOT` environment variable with the value `:pserver:anonymous@cvs.jboss.sourceforge.net:/cvsroot/jboss` and then use the following abbreviated versions of the previous commands:

```
cvs login
cvs -z3 co modulename
```

The name of the JBoss module alias you use depends on the version of JBoss you want. For the 3.0 branch the module name is `jboss-3.0`, for the 3.2 branch it is `jboss-3.2`, and in general, for branch `x.y` the module name is `jboss-x.y`. To checkout the `HEAD` revision of `jboss` to obtain the latest code on the main branch you would use `jboss-head` as the module name. Releases of JBoss are tagged with the pattern `JBoss_X_Y_Z` where `x` is the major version, `y` is the minor version and `z` is the patch version. Release branches of JBoss are tagged with the pattern `Branch_X_Y`. Some checkout examples are:

```
cvs co -r Branch_3_0 jboss-3.0 # Checkout the current 3.0 branch code
cvs co -r JBoss_3_0_6 jboss-3.0 # Checkout the 3.0.6 release version code
cvs co -r Branch_3_2 jboss-3.2 # Checkout the current 3.2 branch code
cvs co -r JBoss_3_2_6 jboss-3.2 # Checkout the 3.2.6 release version code
cvs co jboss-head # Checkout the curent HEAD branch code
```

1.4.4. Obtaining a CVS Client

The command line version of the CVS program is freely available for nearly every platform, and is included by default on most Linux and UNIX distributions. A good port of CVS as well as numerous other UNIX programs for Win32 platforms is available from Cygwin at <http://sources.redhat.com/cygwin/>. The syntax of the command line version of CVS will be examined because this is common across all platforms.

For complete documentation on CVS, check out the CVS home page at <http://www.cvshome.org/>.

1.4.5. Building the JBoss 3.2.6 Distribution Using the Source Code

Every JBoss release includes a source archive that contains everything needed to build the release and is available from the files section of the JBoss project site here: <http://sourceforge.net/projects/jboss/>. The source directory structure matches that of the CVS source tree described below so once you have the source distribution you can build the release by following the instructions given in the next section, beginning with the instructions after the step to obtain the `jboss-3.2` source tree.

1.4.6. Building the JBoss 3.2.6 Distribution Using the CVS Source Code

This section will guide you through the task of building a JBoss distribution from the CVS source code. To start, create a directory into which you want to download the CVS source tree, and move into the newly created directory. This directory is referred to as the `CVS_WD` directory for CVS working directory. The example build in this book will check out code into a `/tmp/3.2.6` directory on a Linux system. Next, obtain the 3.2.6 version of the source code as shown here¹:

```
[nr@toki tmp]$ mkdir 3.2.6
[nr@toki tmp]$ cd 3.2.6
[nrrrr@toki 3.2.6]$ export CVSROOT=:pserver:anonymous@cvs.sourceforge.net:/cvsroot/jboss
```

```
[nr@toki 3.2.6]$ cvs co -r JBoss_3_2_6 jboss-3.2
cvs server: Updating tools
U tools/.classpath
U tools/.donotremove
U tools/.project
cvs server: Updating tools/apache
...
```

The resulting `jboss-3.2` directory structure contains all of the CVS modules required to build the server. To perform the build, `cd` to the `jboss-all/build` directory and execute the `build.sh` or `build.bat` file as appropriate for your OS. You will need to set the `JAVA_HOME` environment variable to the location of the JDK you wish to use for compilation.

```
[nr@toki 3.2.6]$ cd jboss-3.2/build/
[nr@toki build]$ export JAVA_HOME=/System/Library/Frameworks/JavaVM.framework/Home/
[nr@toki build]$ PATH=$JAVA_HOME/bin:$PATH
[nr@toki build]$ ./build.sh
Searching for build.xml ...
Buildfile: /tmp/3.2.6/jboss-3.2/build/build.xml

...

BUILD SUCCESSFUL
Total time: 2 minutes 41 seconds
```

Note that if you see an "Failed to launch JJTree" error do not have the `JAVA_HOME/bin` directory in your `PATH` required for the JavaCC JJTree Ant task.

The build process is driven by an Ant based configuration. The main Ant build script is the `build.xml` file located in the `jboss-3.2/build` directory. This script uses a number of custom Ant tasks masked as `buildmagic` constructs. The purpose of the main `build.xml` file is to compile the various module directories under `jboss-3.2` and then to integrate their output to produce the binary release. The binary release structure is found under the `jboss-3.2/build/output` directory. The example above used the `build.sh` script to kickoff the build process. This is just a wrapper that launches the ant binary included in the distribution. You can simply use `ant` if you have your environment setup to run Ant from the command line.

1.4.7. An Overview of the JBoss CVS Source Tree

The top-level directory structure under the `jboss-3.2` source tree is illustrated in Figure 1.3, the CVS source tree top-level directories. Table 1.2 gives the primary purpose of each of the top-level directories.

Table 1.2. Descriptions of the top-level directories of the JBoss CVS source tree.

Directory	Description
blocks	Not used
build	The main build directory from which the release builds are initiated
cache	The JBoss TreeCache module
cluster	The clustering support services source module.
common	A source module of common utility type code used by many of the other source mod-

¹There was a change in the module aliases used to obtain the complete JBoss source tree just prior to the 3.0.4 release. Now, instead of using `jboss-all` as the module alias for every branch, a branch specific module alias is defined. For the 3.0 branch this is `jboss-3.0`, for the 3.2 branch it is `jboss-3.2`, etc. To checkout the HEAD revision of `jboss` to obtain the latest code on the main branch you would use `jboss-head` as the module alias.

Directory	Description
	ules.
compatible	A backward compatibility module currently under development
connector	The JCA support and application server integration source module.
console	Admin apps for viewing the JMX MBeans
hibernate	The hibernate deployer service.
iiop	The RMI/IIOP transport service source module.
j2ee	The core J2EE interfaces and classes.
jaxrpc	The J2EE web services module.
jboss.net	A web services support source module that provides support for using SOAP to invoke operations on EJBs and MBeans.
jmx	The JBoss JMX implementation source module.
management	The JBoss JSR-77 source module.
messaging	The JBoss JMS implementation source module.
naming	The JBoss JNDI implementation source module.
security	The JBoss standard J2EE declarative security implementation based on JAAS.
server	The EJB 2.0 container implementation related source.
system	The JMX micro kernel based bootstrap services and standard deployment services source module.
testsuite	The JUnit unit test source module.
thirdparty	A module containing the third-party binary jars used by the JBoss modules.
tomcat	The Tomcat-5.0.x embedded service source module
tools	The jars used by the JBoss build process.
transaction	The JTA transaction manager
varia	Various utility services that have not or will not been integrated into one of the higher-level modules.

1.4.8. Using the JBossTest unit testsuite

More advanced testing of the JBoss installation and builds can be done using the JBoss testsuite. The JBossTest suite is a collection of client oriented unit tests of the JBoss server application. It is an Ant based package that uses the JUnit (<http://www.junit.org>) unit test framework. The JBossTest suite is used as a QA benchmark by the development team to help test new functionality and prevent introduction of bugs. It is run on a nightly basis and the results are posted to the development mailing list for all to see.

The unit tests are run using Ant and the source for the tests are contained in the `jboss-3.2/testsuite` directory of the source tree. The structure of the testsuite CVS module is illustrated in Figure 1.3.



Figure 1.3. The testsuite CVS module directory structure

The two main source branches are `src/main` and `src/resources`. The `src/main` tree contains the Java source code for the unit tests. The `src/resources` tree contains resource files like deployment descriptors, JAR manifest, web content, etc. The root package of every unit test is `org.jboss.test`. The typical structure below each specific unit test subpackage (for example, `security`) consists of a test package that contains the unit test classes. The `test` subpackage is a required naming convention as this is the only directory searched for unit tests by the Ant build scripts. If the tests involves EJBs then the convention is to include an `interfaces` and `ejb` subpackage for these components. The unit tests themselves need to follow a naming convention for the class file. The unit test class must be named `xxxUnitTest.java`, where `xxx` is either the class being tested or the name of the functionality being tested.

To run the unit tests use the build scripts located in the `testsuite` directory. The key targets in the `build.xml` file include:

- **tests:** this target builds and runs all unit tests and generates HTML and text reports of the tests into the `testsuite/output/reports/html` and `testsuite/output/reports/text` directories respectively.
- **tests-standard-unit:** builds all unit tests and runs a subset of the key unit tests. This is useful for quick check of the server to test for gross problems.
- **test:** this target allows one to run all tests within a particular package. To run this target you need to specify a test property that specifies a package name using `-Dtest=package` command line. The package value is the name of the package below `org.jboss.test` you want to run unit tests for. So, for example, to run all unit tests in the `org.jboss.test.naming` package, you would use: `build.sh-Dtest=naming test`
- **one-test:** this target allows you to run a single unit test. To run this target you need to specify a test property that specifies the classname of the unit test using `-Dtest=classname` on the command line. So, for example, to run the `org.jboss.test.naming.test.ENCUnitTestCase`, you would use: `build.sh -Dtest=org.jboss.test.naming.test.ENCUnitTestCase one-test`
- **tests-report:** this target generates html and text reports of the tests into the `testsuite/output/reports/html` and `testsuite/output/reports/text` directories respectively using the current JUnit XML results in the `testsuite/output/reports` directory. This is useful for generating the nice html reports when you have run a subset of the tests by hand and want to generate a summary.

On completion of a test the `testsuite/output/reports` directory will contain one or more XML files that represent the individual JUnit test runs. The `tests-report` target collates these into an HTML report located in the `html` subdirectory along with a text report located in the `text` subdirectory. Figure 1.4 shows an example of the HTML report for a run of the test suite against the JBoss 3.2.6 release.

You can find the results of the testsuite in the JBoss distribution in under the `JBOSS_DIST/docs/tests` directory.

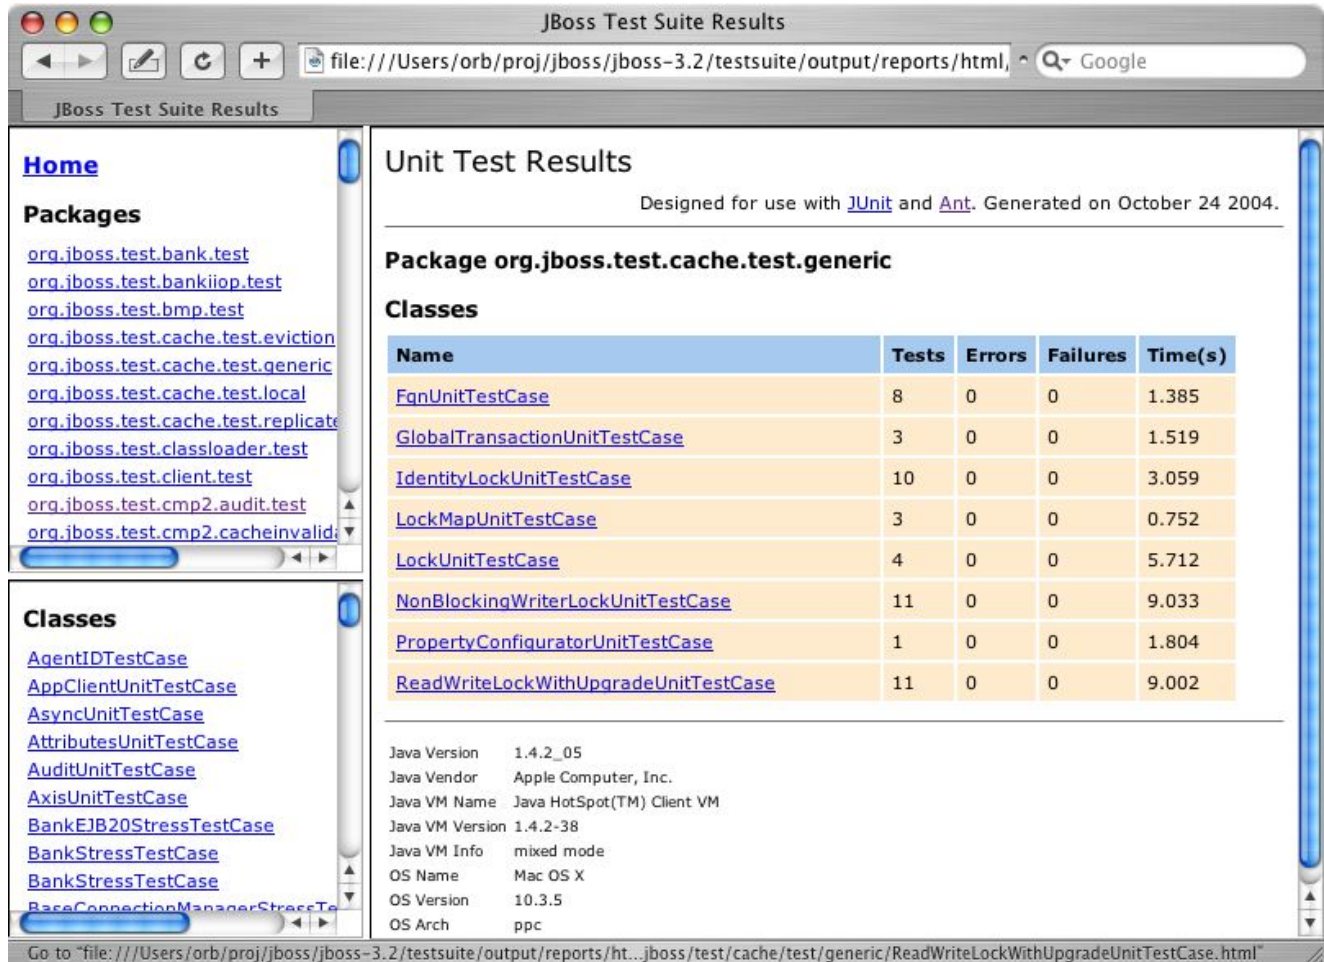


Figure 1.4. An example testsuite run report status HTML view as generated by the testsuite

Note that many tests require a running JBoss instance to deploy to. The build script does not start the JBoss instance. You will need to manually start the JBoss server in the `all` configuration before running the tests.

The JBoss JMX Microkernel

Modularly developed from the ground up, the JBoss server and container are completely implemented using component-based plug-ins. The modularization effort is supported by the use of JMX, the Java Management Extension API. Using JMX, industry-standard interfaces help manage both JBoss/Server components and the applications deployed on it. Ease of use is still the number one priority, and the JBoss Server version 3.x architecture sets a new standard for modular, plug-in design as well as ease of server and application management.

This high degree of modularity benefits the application developer in several ways. The already tight code can be further trimmed down to support applications that must have a small footprint. For example, if EJB passivation is unnecessary in your application, simply take the feature out of the server. If you later decide to deploy the same application under an Application Service Provider (ASP) model, simply enable the server's passivation feature for that web-based deployment. Another example is the freedom you have to drop your favorite object to relational database (O-R) mapping tool, such as TOPLink, directly into the container.

This chapter will introduce you to JMX and its role as the JBoss server component bus. You will also be introduced to the JBoss MBean service notion that adds life cycle operations to the basic JMX management component.

2.1. JMX

The success of the full Open Source J2EE stack lies with the use of JMX (Java Management Extension). JMX is the best tool for integration of software. It provides a common spine that allows the user to integrate modules, containers, and plug-ins. Figure 2.1 shows the role of JMX as an integration spine or bus into which components plug. Components are declared as MBean services that are then loaded into JBoss. The components may subsequently be administered using JMX.




Figure 2.1. The JBoss JMX integration bus and the standard JBoss components

2.1.1. An Introduction to JMX

Before looking at how JBoss uses JMX as its component bus, it would help to get a basic overview what JMX is by touching on some of its key aspects.

JMX components are defined by the Java Management Extensions Instrumentation and Agent Specification, v1.0, which is available from the JSR003 Web page at <http://jcp.org/aboutJava/communityprocess/final/jsr003/>. The material in this JMX overview section is derived from the JMX instrumentation specification, with a focus on the aspects most used by JBoss. A more comprehensive discussion of JMX and its application can be found in *JMX: Managing J2EE with Java Management Extensions* written by Juha Lindfors (Sams, 0672322889, 2002).

JMX is about providing a standard for managing and monitoring all varieties of software and hardware components from Java. Further, JMX aims to provide integration with the large number of existing management standards. Figure 2.2 shows examples of components found in a JMX environment, and illustrates the relationship between them as well as how they relate to the three levels of the JMX model. The three levels are:

- **Instrumentation**, which are the resources to manage
- **Agents**, which are the controllers of the instrumentation level objects
- **Distributed services**, the mechanism by which administration applications interact with agents and their managed objects

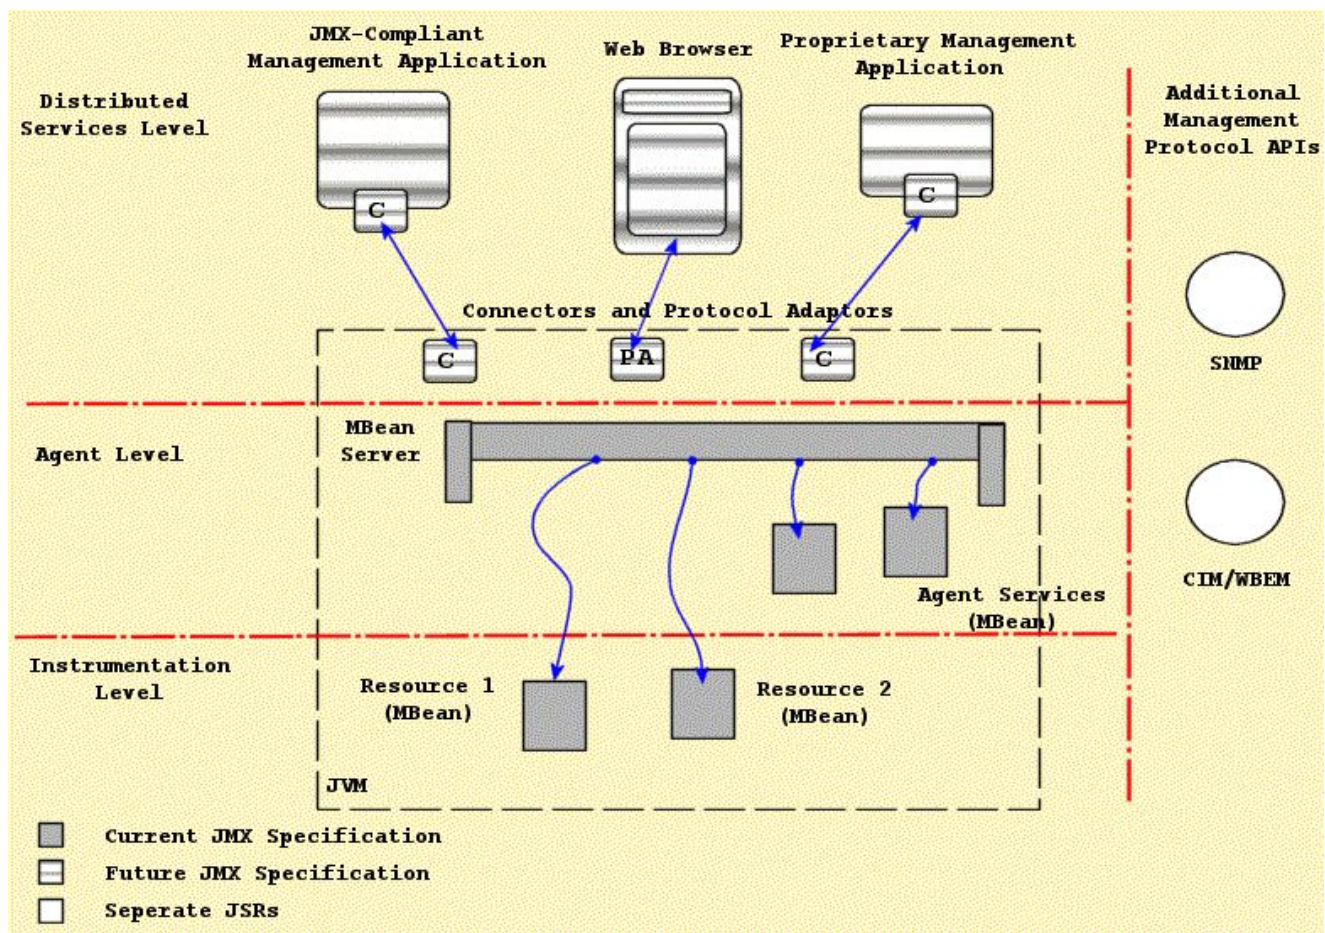


Figure 2.2. The Relationship between the components of the JMX architecture

2.1.1.1. Instrumentation Level

The instrumentation level defines the requirements for implementing JMX manageable resources. A JMX manageable resource can be virtually anything, including applications, service components, devices, and so on. The manageable resource exposes a Java object or wrapper that describes its manageable features, which makes the resource instrumented so that it can be managed by JMX-compliant applications.

The user provides the instrumentation of a given resource using one or more managed beans, or MBeans. There are four varieties of MBean implementations: standard, dynamic, model, and open. The differences between the various MBean types is discussed in Managed Beans or MBeans.

The instrumentation level also specifies a notification mechanism. The purpose of the notification mechanism is to allow MBeans to communicate changes with their environment. This is similar to the JavaBean property change notification mechanism, and can be used for attribute change notifications, state change notifications, and so on.

2.1.1.2. Agent Level

The agent level defines the requirements for implementing agents. Agents are responsible for controlling and exposing the managed resources that are registered with the agent. By default, management agents are located on the same hosts as their resources. This collocation is not a requirement.

The agent requirements make use of the instrumentation level to define a standard MBeanServer management agent, supporting services, and a communications connector. JBoss provides both an html adaptor as well as an

RMI adaptor.

The JMX agent can be located in the hardware that hosts the JMX manageable resources when a Java Virtual Machine (JVM) is available. This is how the JBoss server uses the MBeanServer. A JMX agent does not need to know which resources it will serve. JMX manageable resources may use any JMX agent that offers the services it requires.

Managers interact with an agent's MBeans through a protocol adaptor or connector, as described in the Section 2.1.1.3 in the next section. The agent does not need to know anything about the connectors or management applications that interact with the agent and its MBeans.

2.1.1.3. Distributed Services Level

The JMX specification notes that a complete definition of the distributed services level is beyond the scope of the initial version of the JMX specification. This was indicated by the component boxes with the horizontal lines in Figure 2.2. The general purpose of this level is to define the interfaces required for implementing JMX management applications or managers. The following points highlight the intended functionality of the distributed services level as discussed in the current JMX specification.

- Provide an interface for management applications to interact transparently with an agent and its JMX manageable resources through a connector
- Exposes a management view of a JMX agent and its MBeans by mapping their semantic meaning into the constructs of a data-rich protocol (for example HTML or SNMP)
- Distributes management information from high-level management platforms to numerous JMX agents
- Consolidates management information coming from numerous JMX agents into logical views that are relevant to the end user's business operations
- Provides security

It is intended that the distributed services level components will allow for cooperative management of networks of agents and their resources. These components can be expanded to provide a complete management application.

2.1.1.4. JMX Component Overview

This section offers an overview of the instrumentation and agent level components. The instrumentation level components include the following:

- MBeans (standard, dynamic, open, and model MBeans)
- Notification model elements
- MBean metadata classes

The agent level components include:

- MBean server
- Agent services

2.1.1.4.1. Managed Beans or MBeans

An MBean is a Java object that implements one of the standard MBean interfaces and follows the associated design patterns. The MBean for a resource exposes all necessary information and operations that a management application needs to control the resource.

The scope of the management interface of an MBean includes the following:

- Attribute values that may be accessed by name
- Operations or functions that may be invoked
- Notifications or events that may be emitted
- The constructors for the MBean's Java class

JMX defines four types of MBeans to support different instrumentation needs:

- **Standard MBeans:** These use a simple JavaBean style naming convention and a statically defined management interface. This is the most common type of MBean used by JBoss.
- **Dynamic MBeans:** These must implement the `javax.management.DynamicMBean` interface, and they expose their management interface at runtime when the component is instantiated for the greatest flexibility. JBoss makes use of Dynamic MBeans in circumstances where the components to be managed are not known until runtime.
- **Open MBeans:** These are an extension of dynamic MBeans. Open MBeans rely on basic, self-describing, user-friendly data types for universal manageability.
- **Model MBeans:** These are also an extension of dynamic MBeans. Model MBeans must implement the `javax.management.modelmbean.ModelMBean` interface. Model MBeans simplify the instrumentation of resources by providing default behavior. Although JBoss does not use any Model MBeans for its core services as of the 3.2.0 release, there is a Model MBean implementation known as an XMBean.

We will present an example of a Standard and a Model MBean in the section that discusses extending JBoss with your own custom services.

2.1.1.4.2. Notification Model

JMX Notifications are an extension of the Java event model. Both the MBean server and MBeans can send notifications to provide information. The JMX specification defines the `javax.management` package `Notification` event object, `NotificationBroadcaster` event sender, and `NotificationListener` event receiver interfaces. The specification also defines the operations on the MBean server that allow for the registration of notification listeners.

2.1.1.4.3. MBean Metadata Classes

There is a collection of metadata classes that describe the management interface of an MBean. Users can obtain a common metadata view of any of the four MBean types by querying the MBean server with which the MBeans are registered. The metadata classes cover an MBean's attributes, operations, notifications, and constructors. For each of these, the metadata includes a name, a description, and its particular characteristics. For example, one characteristic of an attribute is whether it is readable, writable, or both. The metadata for an operation contains the signature of its parameter and return types.

The different types of MBeans extend the metadata classes to be able to provide additional information as required. This common inheritance makes the standard information available regardless of the type of MBean. A management application that knows how to access the extended information of a particular type of MBean is able to do so.

2.1.1.4.4. MBean Server

A key component of the agent level is the managed bean server. Its functionality is exposed through an instance of the `javax.management.MBeanServer`. An MBean server is a registry for MBeans that makes the MBean management interface available for use by management applications. The MBean never directly exposes the MBean object itself; rather, its management interface is exposed through metadata and operations available in the MBean server interface. This provides a loose coupling between management applications and the MBeans they manage.

MBeans can be instantiated and registered with the MBeanServer by the following:

- Another MBean
- The agent itself
- A remote management application (through the distributed services)

When you register an MBean, you must assign it a unique object name. The object name then becomes the unique handle by which management applications identify the object on which to perform management operations. The operations available on MBeans through the MBean server include the following:

- Discovering the management interface of MBeans
- Reading and writing attribute values
- Invoking operations defined by MBeans
- Registering for notifications events
- Querying MBeans based on their object name or their attribute values

Protocol adaptors and connectors are required to access the MBeanServer from outside the agent's JVM. Each adaptor provides a view via its protocol of all MBeans registered in the MBean server the adaptor connects to. An example adaptor is an HTML adaptor that allows for the inspection and editing of MBeans using a Web browser. As was indicated in Figure 2.2, there are no protocol adaptors defined by the current JMX specification. Later versions of the specification will address the need for remote access protocols in standard ways.

A connector is an interface used by management applications to provide a common API for accessing the MBean server in a manner that is independent of the underlying communication protocol. Each connector type provides the same remote interface over a different protocol. This allows a remote management application to connect to an agent transparently through the network, regardless of the protocol. The specification of the remote management interface will be addressed in a future version of the JMX specification.

Adaptors and connectors make all MBean server operations available to a remote management application. For an agent to be manageable from outside of its JVM, it must include at least one protocol adaptor or connector. JBoss currently includes a custom HTML adaptor implementation and a custom JBoss RMI adaptor.

2.1.1.4.5. Agent Services

The JMX agent services are objects that support standard operations on the MBeans registered in the MBean server. The inclusion of supporting management services helps you build more powerful management solutions. Agent services are often themselves MBeans, which allow the agent and their functionality to be controlled through the MBean server. The JMX specification defines the following agent services:

- **A dynamic class loading MLet (management applet) service:** This allows for the retrieval and instantiation of new classes and native libraries from an arbitrary network location.
- **Monitor services:** These observe an MBean attribute's numerical or string value, and can notify other objects of several types of changes in the target.
- **Timer services:** These provide a scheduling mechanism based on a one-time alarm-clock notification or on a repeated, periodic notification.
- **The relation service:** This service defines associations between MBeans and enforces consistency on the relationships.

Any JMX-compliant implementation will provide all of these agent services. However, JBoss does not rely on any of these standard agent services.

2.2. JBoss JMX Implementation Architecture

2.2.1. The JBoss ClassLoader Architecture

JBoss 3.x employs a class loading architecture that facilitates sharing of classes across deployment units and hot deployment of services and applications. Before discussing the JBoss specific class loading model, we need to understand the nature of Java's type system and how class loaders fit in.

2.2.2. Class Loading and Types in Java

Class loading is a fundamental part of all server architectures. Arbitrary services and their supporting classes must be loaded into the server framework. This can be problematic due to the strongly typed nature of Java. Most developers know that the type of a class in Java is a function of the fully qualified name of the class. As of Java 1.2, the type is also a function of the `java.lang.ClassLoader` that is used to define that class. This additional qualification of type was added to ensure that environments in which classes may be loaded from arbitrary locations would be type-safe. A paper entitled *Java is not type-safe* by Vijay Saraswat in 1997 demonstrated that Java was not type-safe as intended. This could allow one to gain access to method and members of a class to which they should not have had access by fooling the Java VM into using an alternate implementation of a previously loaded class. Such circumvention of the type system was based on introducing class loaders that by-pass the normal delegation model. A class loader uses a delegation model to search for classes and resources. Each instance of `ClassLoader` has an associated parent class loader that is either explicitly set when it is created, or assigned by the VM if no parent was specified. When called upon to find a class, a class loader will typically delegate the search for the class to its parent class loader before attempting to find the class or resource itself. The VM has a root class loader, called the bootstrap class loader, does not have a parent but may serve as the parent of a `ClassLoader` instance.

To address the type-safety issue, the type system was strengthened to include a class's defining `ClassLoader` in addition to the name of the class to fully define the type. The original paper in which the solution was described is *Dynamic Class Loading in the Java Virtual Machine*, by Sheng Liang and Gilad Bracha, and can be obtained from <http://java.sun.com/people/sl/papers/oopsla98.ps.gz>. The ramifications of this change in a dynamic envir-

onment like an application server, and especially JBoss with its support for hot deployment are that class cast exceptions, linkage errors and illegal access errors can show up in ways not seen in more static class loading contexts. Let's take a look at the meaning of each of these exceptions and how they can happen.

2.2.2.1. ClassCastExceptions - I'm Not Your Type

A `java.lang.ClassCastException` results whenever an attempt is made to cast an instance to an incompatible type. A simple example is trying to obtain a `String` from a `List` into which a `URL` was placed:

```
ArrayList array = new ArrayList();
array.add(new URL("file:/tmp"));
String url = (String) array.get(0);

java.lang.ClassCastException: java.net.URL
at org.jboss.chap2.ex0.ExCCEa.main(ExlCCE.java:16)
```

The `ClassCastException` tells you that the attempt to cast the array element to a `String` failed because the actual type was `URL`. This trivial case is not what we are interested in however. Consider the case of a JAR being loaded by different class loaders. Although the classes loaded through each class loader are identical in terms of the bytecode, they are completely different types as viewed by the Java type system. An example of this is illustrated by the code shown in Example 2.1.

Example 2.1. The `ExCCEc` class used to demonstrate `ClassCastException` due to duplicate class loaders

```
package org.jboss.chap2.ex0;

import java.io.File;
import java.net.URL;
import java.net.URLClassLoader;
import java.lang.reflect.Method;

import org.apache.log4j.Logger;

import org.jboss.util.ChapterExRepository;
import org.jboss.util.Debug;

/**
 * An example of a ClassCastException that
 * results from classes loaded through
 * different class loaders.
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class ExCCEc
{
    public static void main(String[] args) throws Exception
    {
        ChapterExRepository.init(ExCCEc.class);

        String chapDir = System.getProperty("chapter.dir");
        Logger ucl0Log = Logger.getLogger("UCL0");
        File jar0 = new File(chapDir+"j0.jar");
        ucl0Log.info("jar0 path: "+jar0.toString());
        URL[] cp0 = {jar0.toURL()};
        URLClassLoader ucl0 = new URLClassLoader(cp0);
        Thread.currentThread().setContextClassLoader(ucl0);
        Class objClass = ucl0.loadClass("org.jboss.chap2.ex0.ExObj");
        StringBuffer buffer = new
            StringBuffer("ExObj Info");
        Debug.displayClassInfo(objClass, buffer, false);
        ucl0Log.info(buffer.toString());
        Object value = objClass.newInstance();

        File jar1 = new File(chapDir+"j0.jar");
```

```

        Logger ucllLog = Logger.getLogger("UCL1");
        ucllLog.info("jar1 path: "+jar1.toString());
        URL[] cp1 = {jar1.toURL()};
        URLClassLoader ucl1 = new URLClassLoader(cp1);
        Thread.currentThread().setContextClassLoader(ucl1);
        Class ctxClass2 = ucl1.loadClass("org.jboss.chap2.ex0.ExCtx");
        buffer.setLength(0);
        buffer.append("ExCtx Info");
        Debug.displayClassInfo(ctxClass2, buffer, false);
        ucllLog.info(buffer.toString());
        Object ctx2 = ctxClass2.newInstance();

        try {
            Class[] types = {Object.class};
            Method useValue =
                ctxClass2.getMethod("useValue", types);
            Object[] margs = {value};
            useValue.invoke(ctx2, margs);
        } catch (Exception e) {
            ucllLog.error("Failed to invoke ExCtx.useValue", e);
            throw e;
        }
    }
}

```

Example 2.2. The ExCtx, ExObj, and ExObj2 classes used by the examples

```

package org.jboss.chap2.ex0;

import java.io.IOException;
import org.apache.log4j.Logger;
import org.jboss.util.Debug;

/**
 * A classes used to demonstrate various class
 * loading issues
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class ExCtx
{
    ExObj value;

    public ExCtx()
        throws IOException
    {
        value = new ExObj();
        Logger log = Logger.getLogger(ExCtx.class);
        StringBuffer buffer = new StringBuffer("ctor.ExObj");
        Debug.displayClassInfo(value.getClass(), buffer, false);
        log.info(buffer.toString());
        ExObj2 obj2 = value.ivar;
        buffer.setLength(0);
        buffer = new StringBuffer("ctor.ExObj.ivar");
        Debug.displayClassInfo(obj2.getClass(), buffer, false);
        log.info(buffer.toString());
    }

    public Object getValue()
    {
        return value;
    }

    public void useValue(Object obj)
        throws Exception
    {
        Logger log = Logger.getLogger(ExCtx.class);
    }
}

```

```

        StringBuffer buffer = new
            StringBuffer("useValue2.arg class");
        Debug.displayClassInfo(obj.getClass(), buffer, false);
        log.info(buffer.toString());
        buffer.setLength(0);
        buffer.append("useValue2.ExObj class");
        Debug.displayClassInfo(ExObj.class, buffer, false);
        log.info(buffer.toString());
        ExObj ex = (ExObj) obj;
    }

    void pkgUseValue(Object obj)
        throws Exception
    {
        Logger log = Logger.getLogger(ExCtx.class);
        log.info("In pkgUseValue");
    }
}

```

Example 2.3. The ExObj and ExObj2 classes used in the examples

```

package org.jboss.chap2.ex0;

import java.io.Serializable;

/**
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class ExObj
    implements Serializable
{
    public ExObj2 ivar = new ExObj2();
}

-----
package org.jboss.chap2.ex0;

import java.io.Serializable;

/**
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class ExObj2
    implements Serializable
{
}

```

The ExCCEc.main method uses reflection to isolate the classes that are being loaded by the class loaders uc10 and uc11 from the application class loader. Both are setup to load classes from the output/chap2/j0.jar, the contents of which are:

```

[nr@toki examples]$ jar -tf output/chap2/j0.jar

org/jboss/chap2/ex0/ExCtx.class
org/jboss/chap2/ex0/ExObj.class
org/jboss/chap2/ex0/ExObj2.class

```

We will run an example that demonstrates how a class cast exection can occur and then look at the specific is-

sue with the example. See Appendix B for instructions on installing the examples accompanying the book, and then run the example from within the examples directory using the following command:

```
[nr@toki examples]$ ant -Dchap=chap2 -Dex=0c run-example
...
[java] [ERROR,UCL1] Failed to invoke ExCtx.useValue
[java] java.lang.reflect.InvocationTargetException
[java] at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
[java] at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39)
[java] at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl
.java:25)
[java] at java.lang.reflect.Method.invoke(Method.java:324)
[java] at org.jboss.chap2.ex0.ExCCEc.main(ExCCEc.java:58)
[java] Caused by: java.lang.ClassCastException
[java] at org.jboss.chap2.ex0.ExCtx.useValue(ExCtx.java:44)
[java] ... 5 more
```

Only the exception is shown here. The full output can be found in the `logs/chap2-ex0c.log` file. At line 55 of `ExCCEc.java` we are invoking `ExcCCEctx.useValue(Object)` on the instance loaded and created in lines 37-48 using `uc11`. The `ExObj` passed in is the one loaded and created in lines 25-35 via `uc10`. The exception results when the `ExCtx.useValue` code attempts to cast the argument passed in to a `ExObj`. To understand why this fails consider the debugging output from the `chap2-ex0c.log` file shown in Example 2.4.

Example 2.4. The `chap2-ex0c.log` debugging output for the `ExObj` classes seen

```
[INFO,UCL0] ExObj Info
org.jboss.chap2.ex0.ExObj(113fe2).ClassLoader=java.net.URLClassLoader@6e3914
..java.net.URLClassLoader@6e3914
...file:/C:/Scott/JBoss/Books/AdminDevel/education/books/admin-devel/examples/output/
chap2/j0.jar
++++CodeSource:
  (file:/C:/Scott/JBoss/Books/AdminDevel/education/books/admin-devel/examples/output/
  chap2/j0.jar <no certificates>)
Implemented Interfaces:
++interface java.io.Serializable(7934ad)
++++ClassLoader: null
++++Null CodeSource

[INFO,ExCtx] useValue2.ExObj class
org.jboss.chap2.ex0.ExObj(415de6).ClassLoader=java.net.URLClassLoader@30e280
..java.net.URLClassLoader@30e280
...file:/C:/Scott/JBoss/Books/AdminDevel/education/books/admin-devel/examples/output/
chap2/j0.jar
++++CodeSource:
  (file:/C:/Scott/JBoss/Books/AdminDevel/education/books/admin-devel/examples/output/
  chap2/j0.jar <no certificates>)
Implemented Interfaces:
++interface java.io.Serializable(7934ad)
++++ClassLoader: null
++++Null CodeSource
```

The first output prefixed with `[INFO,UCL0]` shows that the `ExObj` class loaded at line `ExCCEc.java:31` has a hash code of `113fe2` and an associated `URLClassLoader` instance with a hash code of `6e3914`, which corresponds to `uc10`. This is the class used to create the instance passed to the `ExCtx.useValue` method. The second output prefixed with `[INFO,ExCtx]` shows that the `ExObj` class as seen in the context of the `ExCtx.useValue` method has a hash code of `415de6` and a `URLClassLoader` instance with an associated hash code of `30e280`, which corresponds to `uc11`. So even though the `ExObj` classes are the same in terms of actual bytecode since it comes from the same `j0.jar`, the classes are different as seen by both the `ExObj` class hash codes, and the associated `URLClassLoader` instances. Hence, attempting to cast an instance of `ExObj` from one scope to the other results in the `ClassCastException`.

This type of error is common when one redeploys an application to which other applications are holding references to classes from the redeployed application. For example, a standalone WAR accessing an EJB. If you are redeploying an application, all dependent applications must flush their class references. Typically this requires that the dependent applications themselves be redeployed.

An alternate means of allowing independent deployments to interact in the presence of redeployment would be to isolate the deployments by configuring the EJB layer to use the standard call-by-value semantics rather than the call-by-reference JBoss will default to for components collocated in the same VM. An example of how to enable call-by-value semantics is presented in Chapter 5

2.2.2.2. `IllegalAccessException` - Doing what you should not

A `java.lang.IllegalAccessException` is thrown when one attempts to access a method or member that visibility qualifiers do not allow. Typical examples are attempting to access private or protected methods or instance variables. Another common example is accessing package protected methods or members from a class that appears to be in the correct package, but is really not due to caller and callee classes being loaded by different class loaders. An example of this is illustrated by the code shown in Example 2.6.

Example 2.5. The `ExIAEd` class used to demonstrate `IllegalAccessException` due to duplicate class loaders

```
package org.jboss.chap2.ex0;

import java.io.File;
import java.net.URL;
import java.net.URLClassLoader;
import java.lang.reflect.Method;

import org.apache.log4j.Logger;

import org.jboss.util.ChapterExRepository;
import org.jboss.util.Debug;

/**
 * An example of IllegalAccessExceptions due to
 * classes loaded by two class loaders.
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class ExIAEd
{
    public static void main(String[] args) throws Exception
    {
        ChapterExRepository.init(ExIAEd.class);

        String chapDir = System.getProperty("chapter.dir");
        Logger ucl0Log = Logger.getLogger("UCL0");
        File jar0 = new File(chapDir+"j0.jar");
        ucl0Log.info("jar0 path: "+jar0.toString());
        URL[] cp0 = {jar0.toURL()};
        URLClassLoader ucl0 = new URLClassLoader(cp0);
        Thread.currentThread().setContextClassLoader(ucl0);

        StringBuffer buffer = new
            StringBuffer("ExIAEd Info");
        Debug.displayClassInfo(ExIAEd.class, buffer, false);
        ucl0Log.info(buffer.toString());

        Class ctxClass1 = ucl0.loadClass("org.jboss.chap2.ex0.ExCtx");
        buffer.setLength(0);
        buffer.append("ExCtx Info");
        Debug.displayClassInfo(ctxClass1, buffer, false);
        ucl0Log.info(buffer.toString());
    }
}
```

```

    Object ctx0 = ctxClass1.newInstance();

    try {
        Class[] types = {Object.class};
        Method useValue =
            ctxClass1.getDeclaredMethod("pkgUseValue", types);
        Object[] margs = {null};
        useValue.invoke(ctx0, margs);
    } catch (Exception e) {
        ucl0Log.error("Failed to invoke ExCtx.pkgUseValue", e);
    }
}
}

```

The `ExIAEd.main` method uses reflection to load the `ExCtx` class via the `ucl0` class loader while the `ExIEAd` class was loaded by the application class loader. We will run this example to demonstrate how the `IllegalAccessException` can occur and then look at the specific issue with the example. Run the example using the following command:

```

[orb@toki examples]$ ant -Dchap=chap2 -Dex=0d run-example
Buildfile: build.xml
...
[java] [ERROR,UCL0] Failed to invoke ExCtx.pkgUseValue
[java] java.lang.IllegalAccessException: Class org.jboss.chap2.ex0.ExIAEd
    can not access a member of class org.jboss.chap2.ex0.ExCtx with modifiers " "
[java] at sun.reflect.Reflection.ensureMemberAccess(Reflection.java:57)
[java] at java.lang.reflect.Method.invoke(Method.java:317)
[java] at org.jboss.chap2.ex0.ExIAEd.main(ExIAEd.java:48)

```

The truncated output shown here illustrates the `IllegalAccessException`. The full output can be found in the `logs/chap2-ex0d.log` file. At line 48 of `ExIAEd.java` the `ExCtx.pkgUseValue(Object)` method is invoked via reflection. The `pkgUseValue` method has package protected access and even though both the invoking class `ExIAEd` and the `ExCtx` class whose method is being invoked reside in the `org.jboss.chap2.ex0` package, the invocation is seen to be invalid due to the fact that the two classes are loaded by different class loaders. This can be seen by looking at the debugging output from the `chap2-ex0d.log` file.

```

[INFO,UCL0] ExIAEd Info
org.jboss.chap2.ex0.ExIAEd(65855a).ClassLoader=sun.misc.Launcher$AppClassLoader@3f52a5
..sun.misc.Launcher$AppClassLoader@3f52a5
...
[INFO,UCL0] ExCtx Info
org.jboss.chap2.ex0.ExCtx(70eed6).ClassLoader=java.net.URLClassLoader@113fe2
..java.net.URLClassLoader@113fe2
...

```

The `ExIAEd` class is seen to have been loaded via the default application class loader instance `sun.misc.Launcher$AppClassLoader@3f52a5`, while the `ExCtx` class was loaded by the `java.net.URLClassLoader@113fe2` instance. Because the classes are loaded by different class loaders, access to the package protected method is seen to be a security violation. So, not only is type a function of both the fully qualified class name and class loader, the package scope is as well.

An example of how this can happen in practise is to include the same classes in two different SAR deployments. If classes in the deployment have a package protected relationship, users of the SAR service may end up loading one class from SAR class loading at one point, and then load another class from the second SAR at a later time. If the two classes in question have a protected access relationship an `IllegalAccessException` will result. The solution is to either include the classes in a separate jar that is referenced by the SARs, or to combine the SARs into a single deployment. This can either be a single SAR, or an EAR the includes both SARs.

2.2.2.3. LinkageErrors - Making Sure You Are Who You Say You Are

To address the type-safety problems of the early Java VMs, the notion of loading constraints were added to the 1.2 Java language spec. Loading constraints validate type expectations in the context of class loader scopes to ensure that a class x is consistently the same class when multiple class loaders are involved. This is important because Java allows for user defined class loaders. Linkage errors are essentially an extension of the class cast exception that is enforced by the VM when classes are loaded and used.

To understand what loading constraints are and how they ensure type-safety we will first introduce the nomenclature of the Liang and Bracha paper along with an example from this paper. There are two type of class loaders, initiating and defining. An initiating class loader is one that a `ClassLoader.loadClass` method has been invoked on to initiate the loading of the named class. A defining class loader is the loader that calls one of the `ClassLoader.defineClass` methods to convert the class byte code into a `Class` instance. The most complete expression of a class is given by $\langle C, L_d \rangle^{L_i}$, where C is the fully qualified class name, L_d is the defining class loader, and L_i is the initiating class loader. In a context where the initiating class loader is not important the type may be represented by $\langle C, L_d \rangle$, while when the defining class loader is not important, the type may be represented by C^{L_i} . In the latter case, there is still a defining class loader, its just not important what the identity of the defining class loader is. Also, a type is completely defined by $\langle C, L_d \rangle$. The only time the initiating loader is relevant is when a loading constraint is being validated. Now consider the classes shown in Example 2.6.

Example 2.6. Classes demonstrating the need for loading constraints

```
class <C,L1> {
    void f() {
        <Spooferd, L1>L1x = <Delegated, L2>L2
        x.secret_value = 1; // Should not be allowed
    }
}
```

```
class <Delegated,L2> {
    static <Spooferd, L2>L3 g() {...}
}
```

```
class <Spooferd, L1> {
    public int secret_value;
}
```

```
class <Spooferd, L2> {
    private int secret_value;
}
```

The class C is defined by L_1 and so L_1 is used to initiate loading of the classes `Spooferd` and `Delegated` referenced in the `C.f()` method. The `Spooferd` class is defined by L_1 , but `Delegated` is defined by L_2 because L_1 delegates to L_2 . Since `Delegated` is defined by L_2 , L_2 will be used to initiate loading of `Spooferd` in the context of the `Delegated.g()` method. In this example both L_1 and L_2 define different versions of `Spooferd` as indicated by the two versions shown at the end of Example 2.6. Since `C.f()` believes x is an instance of $\langle \text{Spooferd}, L_1 \rangle$ it is able to access the private field `secret_value` of $\langle \text{Spooferd}, L_2 \rangle$ returned by `Delegated.g()` due to the 1.1 and earlier Java VM's failure to take into account that a class type is determined by both the fully qualified name of the class and the defining class loader.

Java 1.2 and beyond addresses this problem by generating loader constraints to validate type consistency when the types being used are coming from different defining class loaders. For the Example 2.6 example, the VM generates a constraint $\text{Spooferd}^{L_1} = \text{Spooferd}^{L_2}$ when the first line of method `C.f()` is verified to indicate that the type `Spooferd` must be the same regardless of whether the load of `Spooferd` is initiated by L_1 or L_2 . It does not matter if L_1 or L_2 , or even some other class loader defines `Spooferd`. All that matters is that there is only one

Spoofer class defined regardless of whether L1 or L2 was used to initiate the loading. If L1 or L2 have already defined separate versions of Spoofer when this check is made a `LinkageError` will be generated immediately. Otherwise, the constraint will be recorded and when `Delegated.g()` is executed, any attempt to load a duplicate version of Spoofer will result in a `LinkageError`.

Now let's take a look at how a `LinkageError` can occur with a concrete example. Example 2.7 gives the example main class along with the custom class loader used.

Example 2.7. A concrete example of a `LinkageError`

```

1: package org.jboss.chap2.ex0;
2: import java.io.File;
3: import java.net.URL;
4:
5: import org.apache.log4j.Logger;
6: import org.jboss.util.ChapterExRepository;
7: import org.jboss.util.Debug;
8:
9: /**
10:  * An example of a LinkageError due to classes being defined by more
11:  * than one class loader in a non-standard class loading environment.
12:  *
13:  * @author Scott.Stark@jboss.org
14:  * @version $Revision: 1.11 $
15:  */
16: public class ExLE
17: {
18:     public static void main(String[] args)
19:         throws Exception
20:     {
21:         ChapterExRepository.init(ExLE.class);
22:
23:         String chapDir = System.getProperty("chapter.dir");
24:         Logger ucl0Log = Logger.getLogger("UCL0");
25:         File jar0 = new File(chapDir+"/j0.jar");
26:         ucl0Log.info("jar0 path: "+jar0.toString());
27:         URL[] cp0 = {jar0.toURL()};
28:         ExURLClassLoader ucl0 = new ExURLClassLoader(cp0);
29:         Thread.currentThread().setContextClassLoader(ucl0);
30:         Class ctxClass1 = ucl0.loadClass("org.jboss.chap2.ex0.ExCtx");
31:         Class obj2Class1 = ucl0.loadClass("org.jboss.chap2.ex0.ExObj2");
32:         StringBuffer buffer = new StringBuffer("ExCtx Info");
33:         Debug.displayClassInfo(ctxClass1, buffer, false);
34:         ucl0Log.info(buffer.toString());
35:         buffer.setLength(0);
36:         buffer.append("ExObj2 Info, UCL0");
37:         Debug.displayClassInfo(obj2Class1, buffer, false);
38:         ucl0Log.info(buffer.toString());
39:
40:         File jar1 = new File(chapDir+"/j1.jar");
41:         Logger ucl1Log = Logger.getLogger("UCL1");
42:         ucl1Log.info("jar1 path: "+jar1.toString());
43:         URL[] cp1 = {jar1.toURL()};
44:         ExURLClassLoader ucl1 = new ExURLClassLoader(cp1);
45:         Class obj2Class2 = ucl1.loadClass("org.jboss.chap2.ex0.ExObj2");
46:         buffer.setLength(0);
47:         buffer.append("ExObj2 Info, UCL1");
48:         Debug.displayClassInfo(obj2Class2, buffer, false);
49:         ucl1Log.info(buffer.toString());
50:
51:         ucl0.setDelegate(ucl1);
52:         try {
53:             ucl0Log.info("Try ExCtx.newInstance()");
54:             Object ctx0 = ctxClass1.newInstance();
55:             ucl0Log.info("ExCtx.ctor succeeded, ctx0: "+ctx0);
56:         } catch(Throwable e) {
57:             ucl0Log.error("ExCtx.ctor failed", e);
58:         }

```

```

59:     }
60: }

```

```

package org.jboss.chap2.ex0;

import java.net.URLClassLoader;
import java.net.URL;

import org.apache.log4j.Logger;

/**
 * A custom class loader that overrides the standard parent delegation
 * model
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class Ex0URLClassLoader extends URLClassLoader
{
    private static Logger log = Logger.getLogger(Ex0URLClassLoader.class);
    private Ex0URLClassLoader delegate;

    public Ex0URLClassLoader(URL[] urls)
    {
        super(urls);
    }

    void setDelegate(Ex0URLClassLoader delegate)
    {
        this.delegate = delegate;
    }

    protected synchronized Class loadClass(String name, boolean resolve)
        throws ClassNotFoundException
    {
        Class clazz = null;
        if (delegate != null) {
            log.debug(Integer.toHexString(hashCode()) +
                "; Asking delegate to loadClass: " + name);
            clazz = delegate.loadClass(name, resolve);
            log.debug(Integer.toHexString(hashCode()) +
                "; Delegate returned: "+clazz);
        } else {
            log.debug(Integer.toHexString(hashCode()) +
                "; Asking super to loadClass: "+name);
            clazz = super.loadClass(name, resolve);
            log.debug(Integer.toHexString(hashCode()) +
                "; Super returned: "+clazz);
        }
        return clazz;
    }

    protected Class findClass(String name)
        throws ClassNotFoundException
    {
        Class clazz = null;
        log.debug(Integer.toHexString(hashCode()) +
            "; Asking super to findClass: "+name);
        clazz = super.findClass(name);
        log.debug(Integer.toHexString(hashCode()) +
            "; Super returned: "+clazz);
        return clazz;
    }
}

```

The key component in this example is the `URLClassLoader` subclass `Ex0URLClassLoader`. This class loader implementation overrides the default parent delegation model to allow the `uc10` and `uc11` instances to both load the `ExObj2` class and then setup a delegation relationship from `uc10` to `uc11`. At lines 30 and 31, the `uc10`

`ExURLClassLoader` is used to load the `ExCtx` and `ExObj2` classes. At line 45 of `ExLE.main` the `uc11` `ExURLClassLoader` is used to load the `ExObj2` class again. At this point both the `uc10` and `uc11` class loaders have defined the `ExObj2` class. A delegation relationship from `uc10` to `uc11` is then setup at line 51 via the `uc10.setDelegate(uc11)` method call. Finally, at line 54 of `ExLE.main` an instance of `ExCtx` is created using the class loaded via `uc10`. The `ExCtx` class is the same as presented in Example 2.2, and the constructor was:

```
public ExCtx()
    throws IOException
{
    value = new ExObj();
    Logger log = Logger.getLogger(ExCtx.class);
    StringBuffer buffer = new StringBuffer("ctor.ExObj");
    Debug.displayClassInfo(value.getClass(), buffer, false);
    log.info(buffer.toString());
    ExObj2 obj2 = value.ivar;
    buffer.setLength(0);
    buffer = new StringBuffer("ctor.ExObj.ivar");
    Debug.displayClassInfo(obj2.getClass(), buffer, false);
    log.info(buffer.toString());
}
```

Now, since the `ExCtx` class was defined by the `uc10` class loader, and at the time the `ExCtx` constructor is executed, `uc10` delegates to `uc11`, line 24 of the `ExCtx` constructor involves the following expression which has been rewritten in terms of the complete type expressions:

$$\langle \text{ExObj2}, \text{uc10} \rangle^{\text{uc10}} \text{obj2} = \langle \text{ExObj}, \text{uc11} \rangle^{\text{uc10}} \text{value} * \text{ivar}$$

This generates a loading constraint of $\text{ExObj2}^{\text{uc10}} = \text{ExObj2}^{\text{uc11}}$ since the `ExObj2` type must be consistent across the `uc10` and `uc11` class loader instances. Because we have loaded `ExObj2` using both `uc10` and `uc11` prior to setting up the delegation relationship, the constraint will be violated and should generate a `LinkageError` when run. Run the example using the following command:

```
[nr@toki examples]$ ant -Dchap=chap2 -Dex=0e run-example
Buildfile: build.xml
...
[java] java.lang.LinkageError: loader constraints violated when linking org/jboss/chap2/ex0/ExObj2 cl
[java] at org.jboss.chap2.ex0.ExCtx.<init>(ExCtx.java:24)
[java] at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method)
[java] at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:39)
[java] at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:27)
[java] at java.lang.reflect.Constructor.newInstance(Constructor.java:274)
[java] at java.lang.Class.newInstance0(Class.java:308)
[java] at java.lang.Class.newInstance(Class.java:261)
[java] at org.jboss.chap2.ex0.ExLE.main(ExLE.java:53)
```

As expected, a `LinkageError` is thrown while validating the loader constraints required by line 24 of the `ExCtx` constructor.

2.2.2.3.1. Debugging Class Loading Issues

Debugging class loading issues comes down to finding out where a class was loaded from. A useful tool for this is the code snippet shown in Example 2.8 taken from the `org.jboss.util.Debug` class of the book examples.

Example 2.8. Obtaining debugging information for a Class

```
Class clazz =...;
StringBuffer results = new StringBuffer();

ClassLoader cl = clazz.getClassLoader();
results.append("\n" + clazz.getName() + "(" +
    Integer.toHexString(clazz.hashCode()) + ").ClassLoader=" + cl);
```

```

ClassLoader parent = cl;

while (parent != null) {
    results.append("\n.." + parent);
    URL[] urls = getClassLoaderURLs(parent);

    int length = urls != null ? urls.length : 0;
    for(int u = 0; u < length; u++) {
        results.append("\n..." + urls[u]);
    }

    if (showParentClassLoaders == false) {
        break;
    }
    if (parent != null) {
        parent = parent.getParent();
    }
}

CodeSource clazzCS = clazz.getProtectionDomain().getCodeSource();
if (clazzCS != null) {
    results.append("\n++++CodeSource: " + clazzCS);
} else {
    results.append("\n++++Null CodeSource");
}

```

The key items are shown in bold. The first is that every `Class` object knows its defining `ClassLoader` and this is available via the `getClassLoader()` method. The defines the scope in which the `Class` type is known as we have just seen in the previous sections on class cast exceptions, illegal access exceptions and linkage errors. From the `ClassLoader` you can view the hierarchy of class loaders that make up the parent delegation chain. If the class loader is a `URLClassLoader` you can also see the URLs used for class and resource loading.

The defining `ClassLoader` of a `Class` cannot tell you from what location that `Class` was loaded. To determine this you must obtain the `java.security.ProtectionDomain` and then the `java.security.CodeSource`. It is the `CodeSource` that has the URL location from which the class originated. Note that not every `Class` has a `CodeSource`. If a class is loaded by the bootstrap class loader then its `CodeSource` will be null. This will be the case for all classes in the `java.*` and `javax.*` packages, for example.

Beyond that it may be useful to view the details of classes being loaded into the JBoss server. You can enable verbose logging of the JBoss class loading layer using a Log4j configuration fragment like that shown in Example 2.9.

Example 2.9. An example log4j.xml configuration fragment for enabling verbose class loading logging

```

<appender name="UCL" class="org.apache.log4j.FileAppender">
    <param name="File" value="{jboss.server.home.dir}/log/ucl.log"/>
    <param name="Append" value="false"/>
    <layout class="org.apache.log4j.PatternLayout">
        <param name="ConversionPattern" value="[%r,%c{1},%t] %m%n"/>
    </layout>
</appender>
<category name="org.jboss.mx.loading" additivity="false">
    <priority value="TRACE" class="org.jboss.logging.XLevel"/>
    <appender-ref ref="UCL"/>
</category>

```

This places the output from the classes in the `org.jboss.mx.loading` package into the `ucl.log` file of the server configurations log directory. Although it may not be meaningful if you have not looked at the class loading code, it is vital information needed for submitting bug reports or questions regarding class loading problems. If

you have a class loading problem that appears to be a bug, submit it to the JBoss project on SourceForge and include this log file as an attachment. If the log file is too big, compress it and mail it to scott.stark@jboss.org.

2.2.2.4. Inside the JBoss Class Loading Architecture

Now that we have the role of class loaders in the Java type system defined, let's take a look at the JBoss class loading architecture. Figure 2.3.




Figure 2.3. The core JBoss class loading components

The central component is the `org.jboss.mx.loading.UnifiedClassLoader3` (UCL) class loader. This is an extension of the standard `java.net.URLClassLoader` that overrides the standard parent delegation model to use a shared repository of classes and resources. This shared repository is the `org.jboss.mx.loading.UnifiedLoaderRepository3`. Every UCL is associated with a single `UnifiedLoaderRepository3`, and a `UnifiedLoaderRepository3` typically has many UCLs. A UCL may have multiple URLs associated with it for class and resource loading. Deployers use the top-level deployment's UCL as a shared class loader and all deployment archives are assigned to this class loader. We will talk about the JBoss deployers and their interaction with the class loading system in more detail latter in Section 2.4.2.

When a UCL is asked to load a class, it first looks to the repository cache it is associated with to see if the class has already been loaded. Only if the class does not exist in the repository will it be loaded into the repository by the UCL. By default, there is a single `UnifiedLoaderRepository3` shared across all UCL instances. This means the UCLs form a single flat class loader namespace. The complete sequence of steps that occur when a `UnifiedClassLoader3.loadClass(String, boolean)` method is called is:

1. Check the `UnifiedLoaderRepository3` classes cache associated with the `UnifiedClassLoader3`. If the class is found in the cache it is returned.
2. Else, ask the `UnifiedClassLoader3` if it can load the class. This is essentially a call to the superclass `URLClassLoader.loadClass(String, boolean)` method to see if the class is among the URLs associated with the class loader, or visible to the parent class loader. If the class is found it is placed into the repository classes cache and returned.
3. Else, the repository is queried for all UCLs that are capable of providing the class based on the repository package name to UCL map. When a UCL is added to a repository an association between the package names available in the URLs associated with the UCL is made, and a mapping from package names to the

UCLs with classes in the package is updated. This allows for a quick determination of which UCLs are capable of loading the class. The UCLs are then queried for the requested class in the order in which the UCLs were added to the repository. If a UCL is found that can load the class it is returned, else a `java.lang.ClassNotFoundException` is thrown.

2.2.2.4.1. Viewing Classes in the Loader Repository

Another useful source of information on classes is the `UnifiedLoaderRepository` itself. This is an MBean that contains operations to display class and package information. The default repository is located under a standard JMX name of `JMImplementation:name=Default,service=LoaderRepository`, and its MBean can be accessed via the JMX console by following its link from the front page. The JMX console view of this MBean is shown in Figure 2.4.

The screenshot shows the 'MBean Inspector' window in a web browser. The address bar shows the URL: `http://localhost:8080/jmx-console/HtmlAdaptor?action=...`. The page title is 'JMX MBean View'. On the right is the JBoss logo. The main content area displays the MBean details for `JMImplementation:service=LoaderRepository,name=Default`. Below this is a table of attributes.

Name	Domain	JMImplementation
	service	LoaderRepository
	name	Default
Java Class	org.jboss.mx.loading.UnifiedLoaderRepository3	
Description	Management Bean.	

Below the table are links: [Back to Agent View](#) and [Refresh MBean View](#).

Attribute Name (Access) Type Description	Attribute Value
CacheSize (R) int MBean Attribute.	2010
	file:/private/tmp/jboss-3.2.3/server/default/lib/log4 file:/private/tmp/jboss-3.2.3/server/default/lib/jbos file:/private/tmp/jboss-3.2.3/server/default/lib/hsq

Figure 2.4. The default class `LoaderRepository` MBean view in the JMX console

Two useful operations you will find here are `getPackageClassLoaders(String)` and `displayClassInfo(String)`. The `getPackageClassLoaders` operation returns a set of class loaders that have been indexed to contain classes or resources for the given package name. The package name must have a trailing period. If you type in the package name `org.jboss.ejb.`, the following representation is displayed:

```
[org.jboss.mx.loading.UnifiedClassLoader3@7dac02{ url=file:/private/tmp/jboss-3.2.6/server/default/tmp
```

This is the string representation of the set. It shows one `UnifiedClassLoader3` instance with a primary URL pointing to the `default/conf/jboss-service.xml` descriptor. This is the second class loader added to the repository (shown by the `addedOrder=2`) and it is the class loader that owns all of the JARs in the `lib` directory of the server configuration (e.g., `server/default/lib`). If you enter the package name `org.jboss.jmx.adaptor.html.`, then the following set will be displayed:

```
[org.jboss.mx.loading.UnifiedClassLoader3@7dac02{ url=file:/private/tmp/jboss-3.2.6/server/default/tmp
```

This time there are two `UnifiedClassLoader3` instances, one for the `default/deploy/jmx-console.war` and one for the `default/deploy/jmx-console2.war`.

To view the information for a given class, use the `displayClassInfo` operation, passing in the fully qualified name of the class to view. For example, if we use `org.jboss.jmx.adaptor.html.HtmlAdaptorServlet` which is from the package we just looked at, the following description is displayed:

```
org.jboss.jmx.adaptor.html.HtmlAdaptorServlet Information
Repository cache version:
org.jboss.jmx.adaptor.html.HtmlAdaptorServlet(26f678).ClassLoader=org.jboss.mx.loading.Uni
fiedClassLoader3@30cd4a{ url=file:/private/tmp/jboss-3.2.6/server/default/deploy/jmx-conso
le.war/ ,addedOrder=32}
..org.jboss.mx.loading.UnifiedClassLoader3@30cd4a{ url=file:/private/tmp/jboss-3.2.6/serve
r/default/deploy/jmx-console.war/ ,addedOrder=32}
...file:/private/tmp/jboss-3.2.6/server/default/deploy/jmx-console.war/
...file:/private/tmp/jboss-3.2.6/server/default/deploy/jmx-console.war/WEB-INF/classes/
..org.jboss.system.server.NoAnnotationURLClassLoader@e48e1b
..sun.misc.Launcher$AppClassLoader@33056f
...file:/private/tmp/jboss-3.2.6/bin/run.jar
...file:/System/Library/Frameworks/JavaVM.framework/Versions/1.4.2/Home/lib/tools.jar
..sun.misc.Launcher$ExtClassLoader@94af67
...file:/System/Library/Java/Extensions/CoreAudio.jar
...file:/System/Library/Java/Extensions/j3daudio.jar
...file:/System/Library/Java/Extensions/j3dcore.jar
...file:/System/Library/Java/Extensions/j3dsupport.jar
...file:/System/Library/Java/Extensions/j3dutils.jar
...file:/System/Library/Java/Extensions/jai_codec.jar
...file:/System/Library/Java/Extensions/jai_core.jar
...file:/System/Library/Java/Extensions/libJ3D.jnilib
...file:/System/Library/Java/Extensions/libJ3DAudio.jnilib
...file:/System/Library/Java/Extensions/libJ3DUtils.jnilib
...file:/System/Library/Java/Extensions/libmllib_jai.jnilib
...file:/System/Library/Java/Extensions/mlibwrapper_jai.jar
...file:/System/Library/Java/Extensions/MRJTToolkit.jar
...file:/System/Library/Java/Extensions/QTJava.zip
...file:/System/Library/Java/Extensions/QTJSupport.jar
...file:/System/Library/Java/Extensions/vecmath.jar
...file:/System/Library/Frameworks/JavaVM.framework/Versions/1.4.2/Home/lib/ext/apple_pro
vider.jar
...file:/System/Library/Frameworks/JavaVM.framework/Versions/1.4.2/Home/lib/ext/ldapsec.j
ar
...file:/System/Library/Frameworks/JavaVM.framework/Versions/1.4.2/Home/lib/ext/localedat
a.jar
...file:/System/Library/Frameworks/JavaVM.framework/Versions/1.4.2/Home/lib/ext/sunjce_pr
ovider.jar
+++CodeSource: (file:/private/tmp/jboss-3.2.6/server/default/deploy/jmx-console.war/WEB-I
NF/classes/ )
Implemented Interfaces:

### Instance0 found in UCL: org.jboss.mx.loading.UnifiedClassLoader3@30cd4a{ url=file:/pri
vate/tmp/jboss-3.2.6/server/default/deploy/jmx-console.war/ ,addedOrder=32}

### Instance1 found in UCL: org.jboss.mx.loading.UnifiedClassLoader3@492aff{ url=file:/pri
vate/tmp/jboss-3.2.6/server/default/deploy/jmx-console2.war/ ,addedOrder=34}

### Instance2 via UCL: org.jboss.mx.loading.UnifiedClassLoader3@30cd4a{ url=file:/private/
```

```
tmp/jboss-3.2.6/server/default/deploy/jmx-console.war/ ,addedOrder=32}
```

The information is a dump of the information for the Class instance in the loader repository if one has been loaded, followed by the class loaders that are seen to have the class file available. If a class is seen to have more than one class loader associated with it, then there is the potential for class loading related errors.

2.2.2.4.2. Scoping Classes

If you need to deploy multiple versions of an application the default 3.x class loading model would require that each application be deployed in a separate JBoss server. Sometimes this is desirable as you have more control over security and resource monitoring, but it can be difficult to manage multiple server instances. An alternative mechanism exists that allows multiple versions of an application to be deployed using deployment based scoping.

With deployment based scoping, each deployment creates its own class loader repository in the form of a `HeirarchicalLoaderRepository3` that looks first to the `UnifiedClassLoader3` instances of the deployment units included in the EAR before delegating to the default `UnifiedLoaderRepository3`. To enable an EAR specific loader repository, you need to create a `META-INF/jboss-app.xml` descriptor as shown in Example 2.10.

Example 2.10. An example `jboss-app.xml` descriptor for enabled scoped class loading at the EAR level.

```
<jboss-app>
  <loader-repository>some.dot.com:loader=webtest.ear</loader-repository>
</jboss-app>
```

The value of the `loader-repository` element is the JMX object name to assign to the repository created for the EAR. This must be unique and valid JMX `ObjectName`, but the actual name is not important.

2.2.2.4.3. The Complete Class Loading Model

The previous discussion of the core class loading components introduced the custom `UnifiedClassLoader3` and `UnifiedLoaderRepository3` classes that form a shared class loading space. The complete class loading picture must also include the parent class loader used by `UnifiedClassLoader3`s as well as class loaders introduced for scoping and other speciality class loading purposes. Figure 2.5 shows an outline of the class hierarchy that would exist for an EAR deployment containing EJBs and WARs.

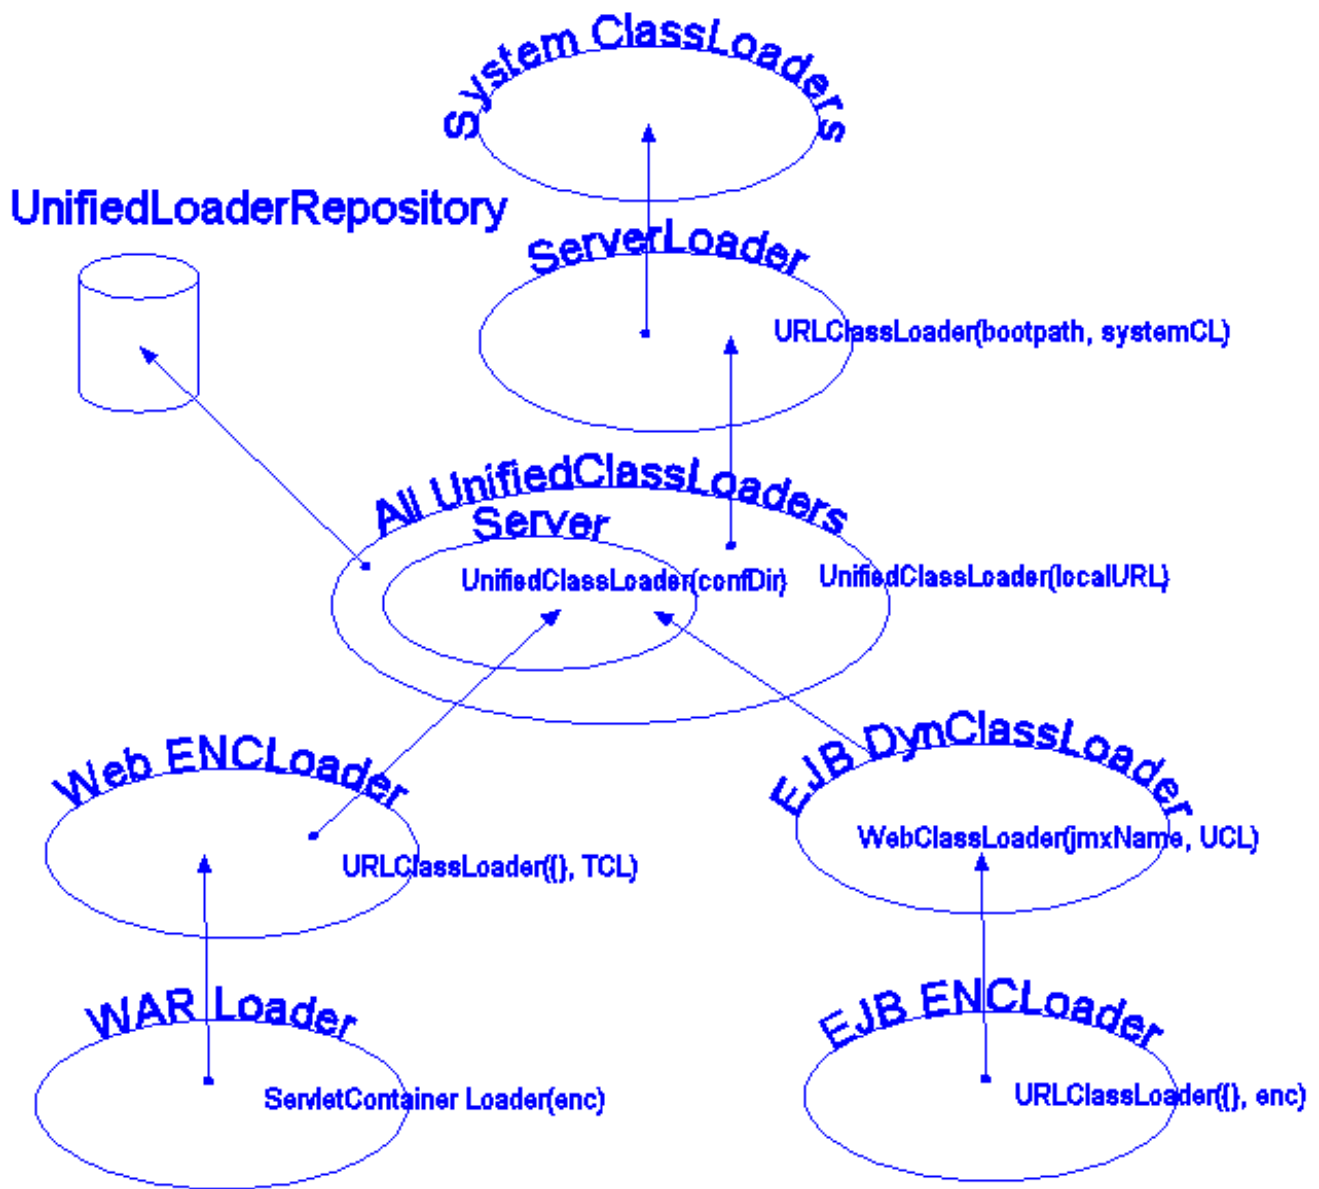


Figure 2.5. A complete class loader view

The following points apply to this figure:

- **System ClassLoaders:** The System ClassLoaders node refers to either the thread context class loader (TCL) of the VM main thread or of the thread of the application that is loading the JBoss server if it is embedded.
- **ServerLoader:** The ServerLoader node refers to the a `URLClassLoader` that delegates to the System ClassLoaders and contains the following boot URLs:
 - All URLs referenced via the `jboss.boot.library.list` system property. These are path specifications relative to the `libraryURL` defined by the `jboss.lib.url` property. If there is no `jboss.lib.url` property specified, it defaults to `jboss.home.url + /lib/`. If there is no `jboss.boot.library` property specified, it defaults to `jaxp.jar`, `log4j-boot.jar`, `jboss-common.jar`, and `jboss-system.jar`.
 - The JAXP JAR which is either `crimson.jar` or `xerces.jar` depending on the `-j` option to the Main

entry point. The default is `crimson.jar`.

- The JBoss JMX jar and GNU regex jar, `jboss-jmx.jar` and `gnu-regexp.jar`.
- Oswego concurrency classes JAR, `concurrent.jar`
- Any JARs specified as libraries via `-L` command line options
- Any other JARs or directories specified via `-c` command line options
- **Server:** The `Server` node represent a collection of UCLs created by the `org.jboss.system.server.Server` interface implementation. The default implementation creates UCLs for the `patchDir` entries as well as the `server.conf` directory. The last UCL created is set as the JBoss main thread context class loader. This will be combined into a single UCL now that multiple URLs per UCL are supported.
- **All UnifiedClassLoader3s:** The *All UnifiedClassLoader3* node represents the UCLs created by deployers. This covers EARs, jars, WARs, SARs and directories seen by the deployment scanner as well as JARs referenced by their manifests and any nested deployment units they may contain. This is a flat namespace and there should not be multiple instances of a class in different deployment JARs. If there are, only the first loaded will be used and the results may not be as expected. There is a mechanism for scoping visibility based on EAR deployment units that we discussed in Section 2.2.2.4.2. Use this mechanism if you need to deploy multiple versions of a class in a given JBoss server.
- **EJB DynClassLoader:** The `EJB DynClassLoader` node is a subclass of `URLClassLoader` that is used to provide RMI dynamic class loading via the simple HTTP WebService. It specifies an empty `URL[]` and delegates to the `TCL` as its parent class loader. If the WebService is configured to allow system level classes to be loaded, all classes in the `UnifiedLoaderRepository3` as well as the system classpath are available via HTTP.
- **EJB ENCLoader:** The *EJB ENCLoader* node is a `URLClassLoader` that exists only to provide a unique context for an EJB deployment's `java:comp` JNDI context. It specifies an empty `URL[]` and delegates to the `EJB DynClassLoader` as its parent class loader.
- **Web ENCLoader:** The *Web ENCLoader* node is a `URLClassLoader` that exists only to provide a unique context for a web deployment's `java:comp` JNDI context. It specifies an empty `URL[]` and delegates to the `TCL` as its parent class loader.
- **WAR Loader:** The *WAR Loader* is a servlet container specific classloader that delegates to the `Web ENCLoader` as its parent class loader. The default behavior is to load from its parent class loader and then the `WEB-INF/classes` and `lib` directories. If the servlet 2.3 class loading model is enabled it will first load from the its `WEB-INF` directories and then the parent class loader.

In its current form there are some advantages and disadvantages to the JBoss class loading architecture. Advantages include:

- Classes do not need to be replicated across deployment units in order to have access to them.
- Many future possibilities including novel partitioning of the repositories into domains, dependency and conflict detection, etc.

Disadvantages include:

- Existing deployments may need to be repackaged to avoid duplicate classes. Duplication of classes in a

loader repository can lead to class cast exceptions and linkage errors depending on how the classes are loaded.

- Deployments that depend on different versions of a given class need to be isolated in separate EARs and a unique `HeirarchicalLoaderRepository3` defined using a `jboss-app.xml` descriptor.

2.2.3. JBoss XMBeans

XMBeans are the JBoss JMX implementation version of the JMX model MBean. XMBeans have the richness of the dynamic MBean metadata without the tedious programming required by a direct implementation of the `DynamicMBean` interface. The JBoss model MBean implementation allows one to specify the management interface of a component through a XML descriptor, hence the X in XMBean. In addition to providing a simple mechanism for describing the metadata required for a dynamic MBean, XMBeans also allow for the specification of attribute persistence, caching behavior, and even advanced customizations like the MBean implementation interceptors. The high level elements of the `jboss_xmbean_1_0.dtd` for the XMBean descriptor is given in Figure 2.6.

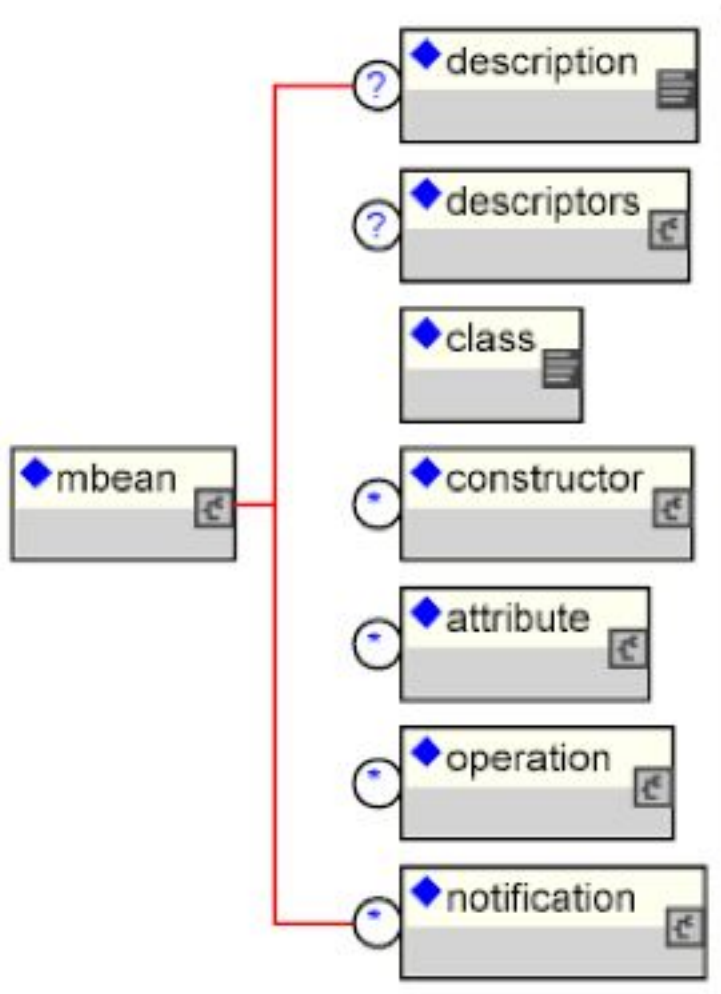


Figure 2.6. The JBoss 1.0 XMBean DTD Overview (jboss_xmbean_1_0.dtd)

The `mbean` element is the root element of the document containing the required elements for describing the management interface of one MBean (constructors, attributes, operations and notifications). It also includes an

optional description element, which can be used to describe the purpose of the MBean, as well as an optional descriptors element which allows for persistence policy specification, attribute caching, etc.

2.2.3.1. Descriptors

The descriptors element contains all the descriptors for a containing element, as subelements. The descriptors suggested in the JMX specification as well as those used by JBoss have predefined elements and attributes, whereas custom descriptors have a generic descriptor element with name and value attributes as show in Figure 2.7.

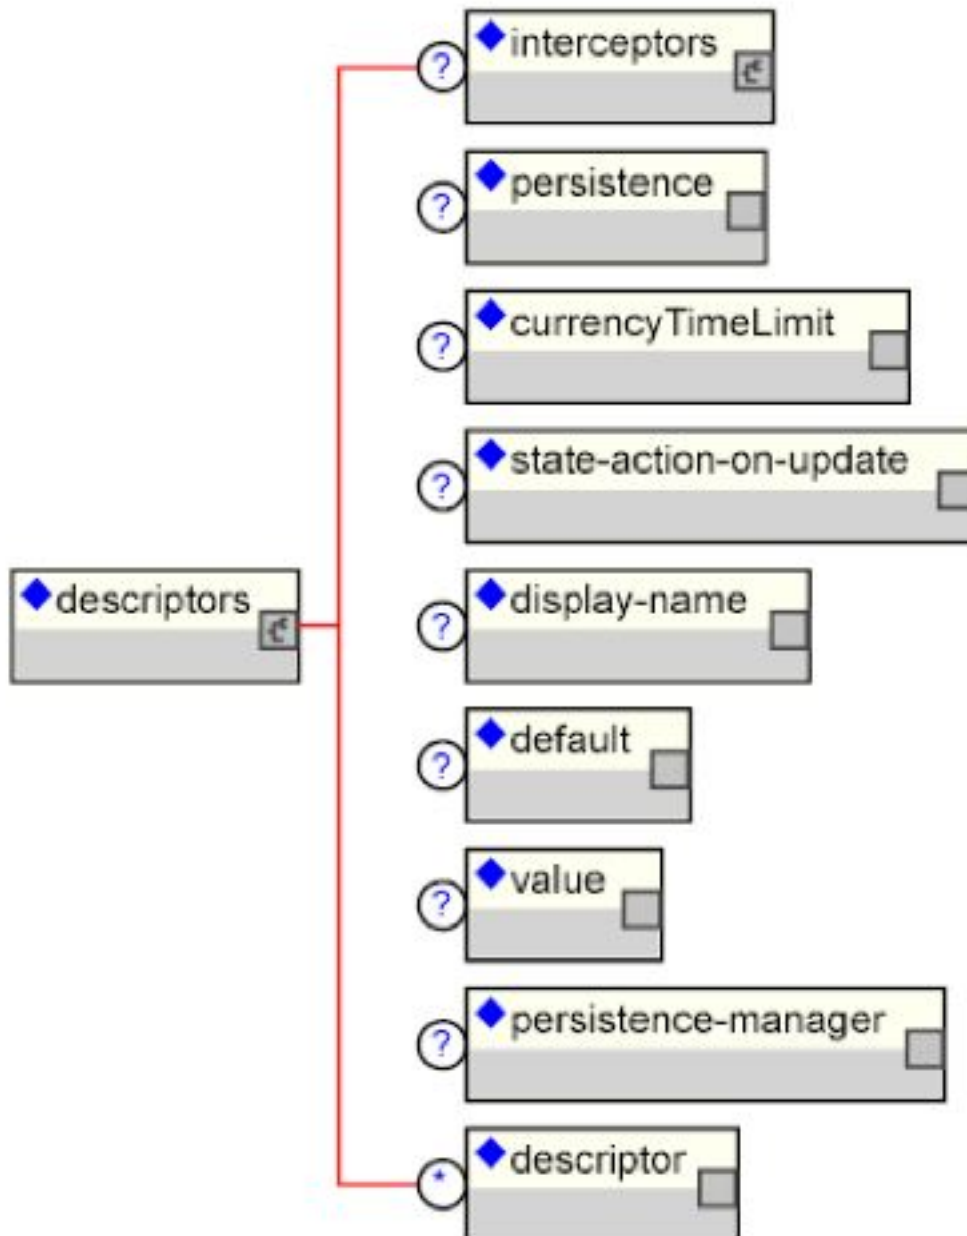


Figure 2.7. The descriptors element content model

The key descriptors child elements include:

- **interceptors:** The `interceptors` element specifies a customized stack of interceptors that will be used in

place of the default stack. Currently this is only used when specified at the MBean level, but it could define a custom attribute or operation level interceptor stack in the future. The content of the `interceptors` element specifies a custom interceptor stack. If no `interceptors` element is specified the standard `ModelMBean` interceptors will be used. The standard interceptors are:

- `org.jboss.mx.interceptor.PersistenceInterceptor`
- `org.jboss.mx.interceptor.MBeanAttributeInterceptor`
- `org.jboss.mx.interceptor.ObjectReferenceInterceptor`

When specifying a custom interceptor stack you would typically include the standard interceptors along with your own unless you are replacing the corresponding standard interceptor.

Each interceptor element content value specifies the fully qualified class name of the interceptor implementation, and the class must implement the `org.jboss.mx.interceptor.Interceptor` interface. The interceptor class must also have either a no-arg constructor, or a constructor that accepts a `(javax.management.MBeanInfo, org.jboss.mx.server.MBeanInvoker)` pair.

The interceptor elements may have any number of attributes that correspond to JavaBean style properties on the interceptor class implementation. For each `interceptor` element attribute specified, the interceptor class is queried for a matching setter method. The attribute value is converted to the true type of the interceptor class property using the `java.beans.PropertyEditor` associated with the type. It is an error to specify an attribute for which there is no setter or `PropertyEditor`.

- **persistence:** The `persistence` element allows the specification of the `persistPolicy`, `persistPeriod`, `persistLocation`, and `persistName` persistence attributes suggested by the JMX specification. The persistence element attributes are:
 - **persistPolicy:** The `persistPolicy` attribute defines when attributes should be persisted and its value must be one of
 - **Never:** attribute values are transient values that are never persisted
 - **OnUpdate:** attribute values are persisted whenever they are updated
 - **OnTimer:** attribute values are persisted based on the time given by the `persistPeriod`.
 - **NoMoreOftenThan:** attribute values are persisted when updated unless but no more often than the `persistPeriod`.
 - **persistPeriod:** The `persistPeriod` attribute gives the update frequency in milliseconds if the `persistPolicy` attribute is `NoMoreOftenThan` OR `OnTimer`.
 - **persistLocation:** The `persistLocation` attribute specifies the location of the persistence store. Its form depends on the JMX persistence implementation. Currently this should refer to a directory into which the attributes will be serialized if using the default JBoss persistence manager.
 - **persistName:** The `persistName` attribute can be used in conjunction with the `persistLocation` attribute to further qualify the persistent store location. For a directory `persistLocation` the `persistName` specifies the file to which the attributes are stored within the directory.
- **currencyTimeLimit:** The `currencyTimeLimit` element specifies the time in seconds that a cached value of

an attribute remains valid. Its value attribute gives the time in seconds. A 0 value indicates that an attribute value should always be retrieved from the MBean and never cached. A -1 value indicates that a cache value is always valid.

- **state-action-on-update:** The `state-action-on-update` element specifies the what happens to an MBean when one of its attributes is updated. The action is given by the `value` attribute. Its value attribute defines what happens to the mbean lifecycle state when one of its attributes is update. It must be one of: **keep-running**, **restart**, **reconfigure**, **restantiate**. However, note that this descriptor is not currently used.
- **display-name:** The `display-name` element specifies the human friendly name of an item.
- **default:** The `default` element specifies a default value to use when a field has not been set. Note that this value is not written to the MBean on startup as is the case with the `jboss-service.xml` attribute element content value. Rather, the default value is used only if there is no attribute accessor defined, and there is no value element defined.
- **value:** The `value` element specifies a management attribute's current value. Unlike the `default` element, the `value` element is written through to the MBean on startup provided there is a setter method available.
- **persistence-manager:** The `persistence-manager` element gives the name of a class to use as the persistence manager. The `value` attribute specifies the class name that supplies the `org.jboss.mx.persistence.PersistenceManager` interface implementation. The only implementation currently supplied by JBoss is the `org.jboss.mx.persistence.ObjectStreamPersistenceManager` which serializes the `ModelMBeanInfo` content to a file using Java serialization.
- **descriptor:** The `descriptor` element specifies an arbitrary descriptor not known to JBoss. Its `name` attribute specifies the type of the descriptor and its `value` attribute specifies the descriptor value. The `descriptor` element allows for the attachment of arbitrary management metadata.

Note that any of the constructor, attribute, operation or notification elements may have a `descriptors` element to specify the specification defined descriptors as well as arbitrary extension descriptor settings.

2.2.3.2. The Management Class

The `class` element specifies the fully qualified name of the managed object whose management interface is described by the XMBean descriptor.

2.2.3.3. The Constructors

The `constructor` element(s) specifies the constructors available for creating an instance of the managed object. The constructor element and its content model are shown in Figure 2.8.

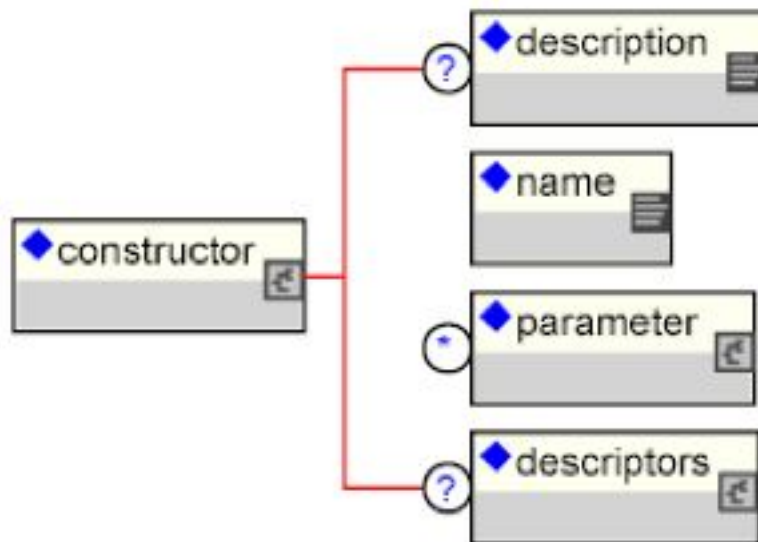


Figure 2.8. The XMBean constructor element and its content model

The key child elements are:

- **description:** A description of the constructor.
- **name:** The name of the constructor, which must be the same as the implementation class.
- **parameter:** The parameter element describes a constructor parameter. The parameter element has the following attributes:
 - **description:** An optional description of the parameter.
 - **name:** The required variable name of the parameter.
 - **type:** The required fully qualified class name of the parameter type.
- **descriptors:** Any descriptors to associate with the constructor metadata.

2.2.3.4. The Attributes

The `attribute` element(s) specifies the management attributes exposed by the MBean. The attribute element and its content model are shown in Figure 2.9.

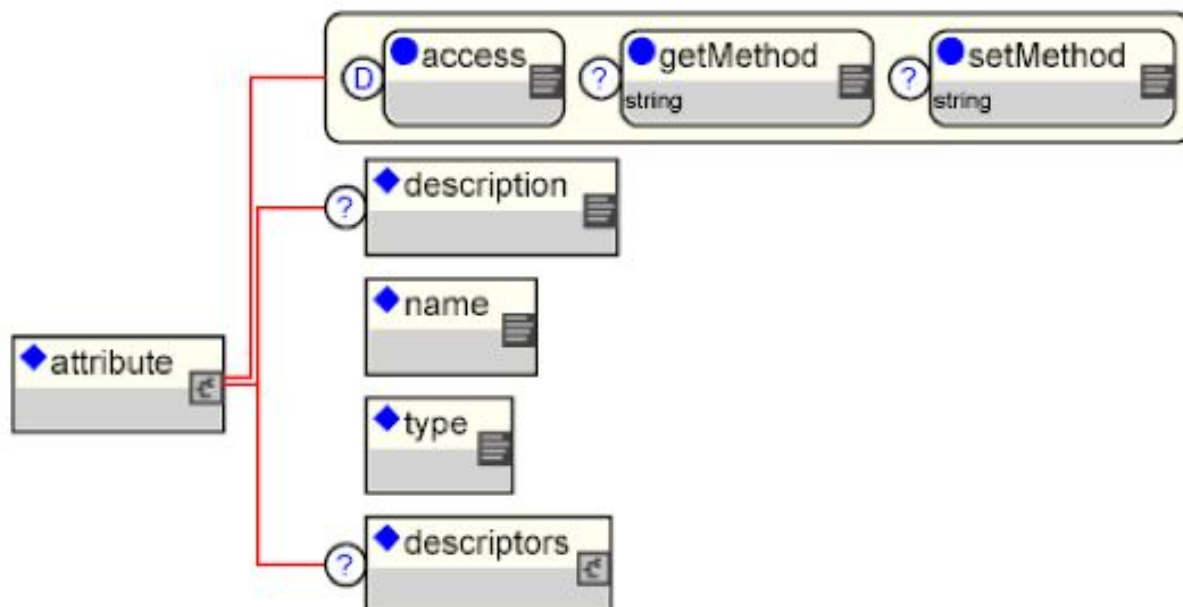


Figure 2.9. The XMBean attribute element and its content model

The `attribute` element supported attributes include:

- **access:** The optional `access` attribute defines the read/write access modes of an attribute. It must be one of:
 - **read-only:** The attribute may only be read.
 - **write-only:** The attribute may only be written.
 - **read-write:** The attribute is both readable and writable. This is the implied default.
- **getMethod:** The `getMethod` attribute defines the name of the method which reads the named attribute. This must be specified if the managed attribute should be obtained from the MBean instance.
- **setMethod:** The `setMethod` attribute defines the name of the method which writes the named attribute. This must be specified if the managed attribute should be obtained from the MBean instance.

The key child elements of the attribute element include:

- **description:** A description of the attribute.
- **name:** The name of the attribute as would be used in the `MBeanServer.getAttribute()` operation.
- **type:** The fully qualified class name of the attribute type.
- **descriptors:** Any additional descriptors that affect the attribute persistence, caching, default value, etc.

2.2.3.5. The Operations

The management operations exposed by the XMBean are specified via one or more operation elements. The

operation element and its content model are shown in Figure 2.10.

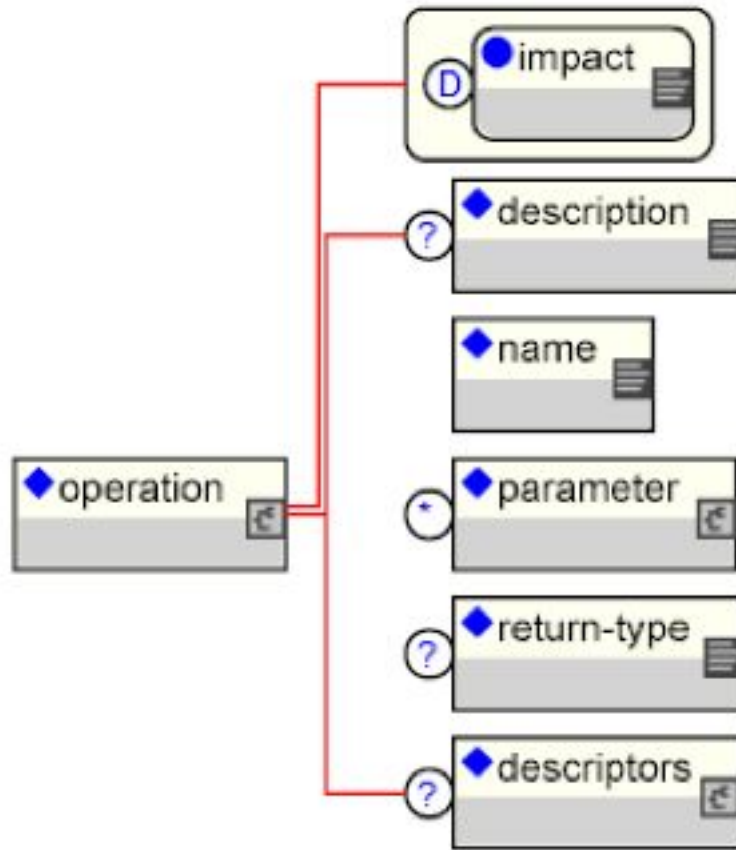


Figure 2.10. The XMBean operation element and its content model

The impact attribute defines the impact of executing the operation and must be one of:

- **ACTION:** The operation changes the state of the MBean component (write operation)
- **INFO:** The operation should not alter the state of the MBean component (read operation).
- **ACTION_INFO:** The operation behaves like a read/write operation.

The child elements are:

- **description:** This element specifies a human readable description of the operation.
- **name:** This element contains the operation's name
- **parameter:** This element describes the operation's signature.
- **return-type:** This element contains a fully qualified class name of the return type from this operation. If not specified, it defaults to void.
- **descriptors:** Any descriptors to associate with the operation metadata.

2.2.3.6. Notifications

The `notification` element(s) describes the management notifications that may be emitted by the XMBean.

The notification element and its content model is shown in Figure 2.11.

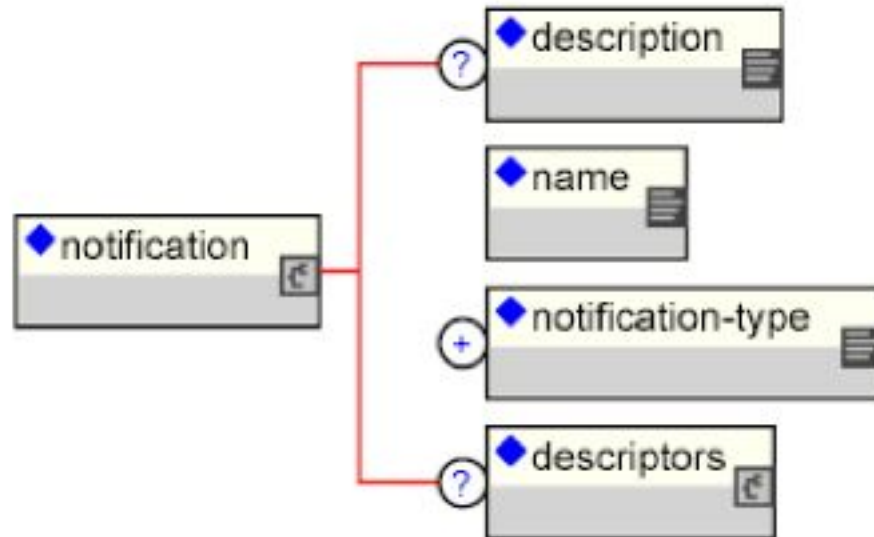


Figure 2.11. The XMBean notification element and content model

The child elements are:

- **description:** This element gives a human readable description of the notification.
- **name:** This element contains the fully qualified name of the notification class.
- **notification-type:** This element contains the dot-separated notification type string.
- **descriptors:** Any descriptors to associate with the notification metadata.

For a reference of the complete DTD content model see the expanded view of the complete provided in Figure 2.12. We will work through examples of creating an XMBeans when we discuss the JBoss MBean services notion. See Section 2.4.3.2 for these examples.

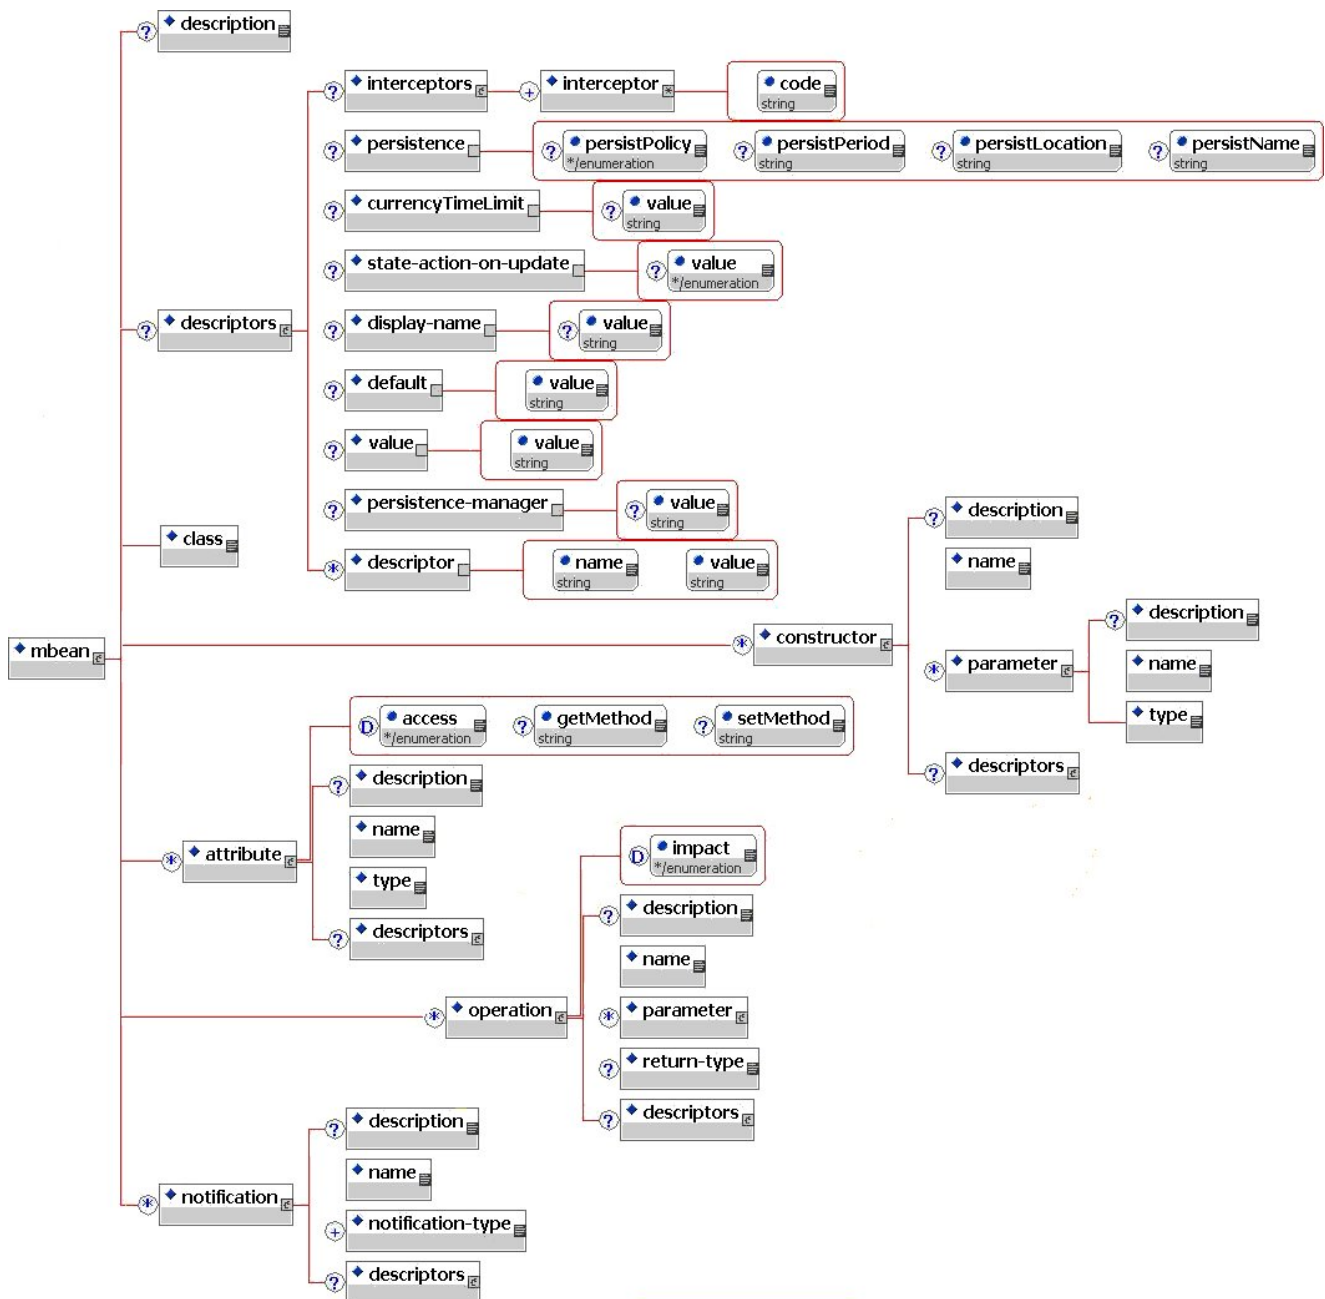


Figure 2.12. An expanded view of the `jboss_xmbean_1_0` DTD

2.3. Connecting to the JMX Server

JBoss includes adaptors that allow access to the JMX MBeanServer from outside of the JBoss server VM. The current adaptors include HTML, an RMI interface, and an EJB.

2.3.1. Inspecting the Server - the JMX Console Web Application

JBoss comes with its own implementation of a JMX HTML adaptor that allows one to view the server's MBeans using a standard web browser. The default URL for the console web application is `http://localhost:8080/jmx-console/`. If you browse this location you will see something similar to that presented in Figure 2.13.

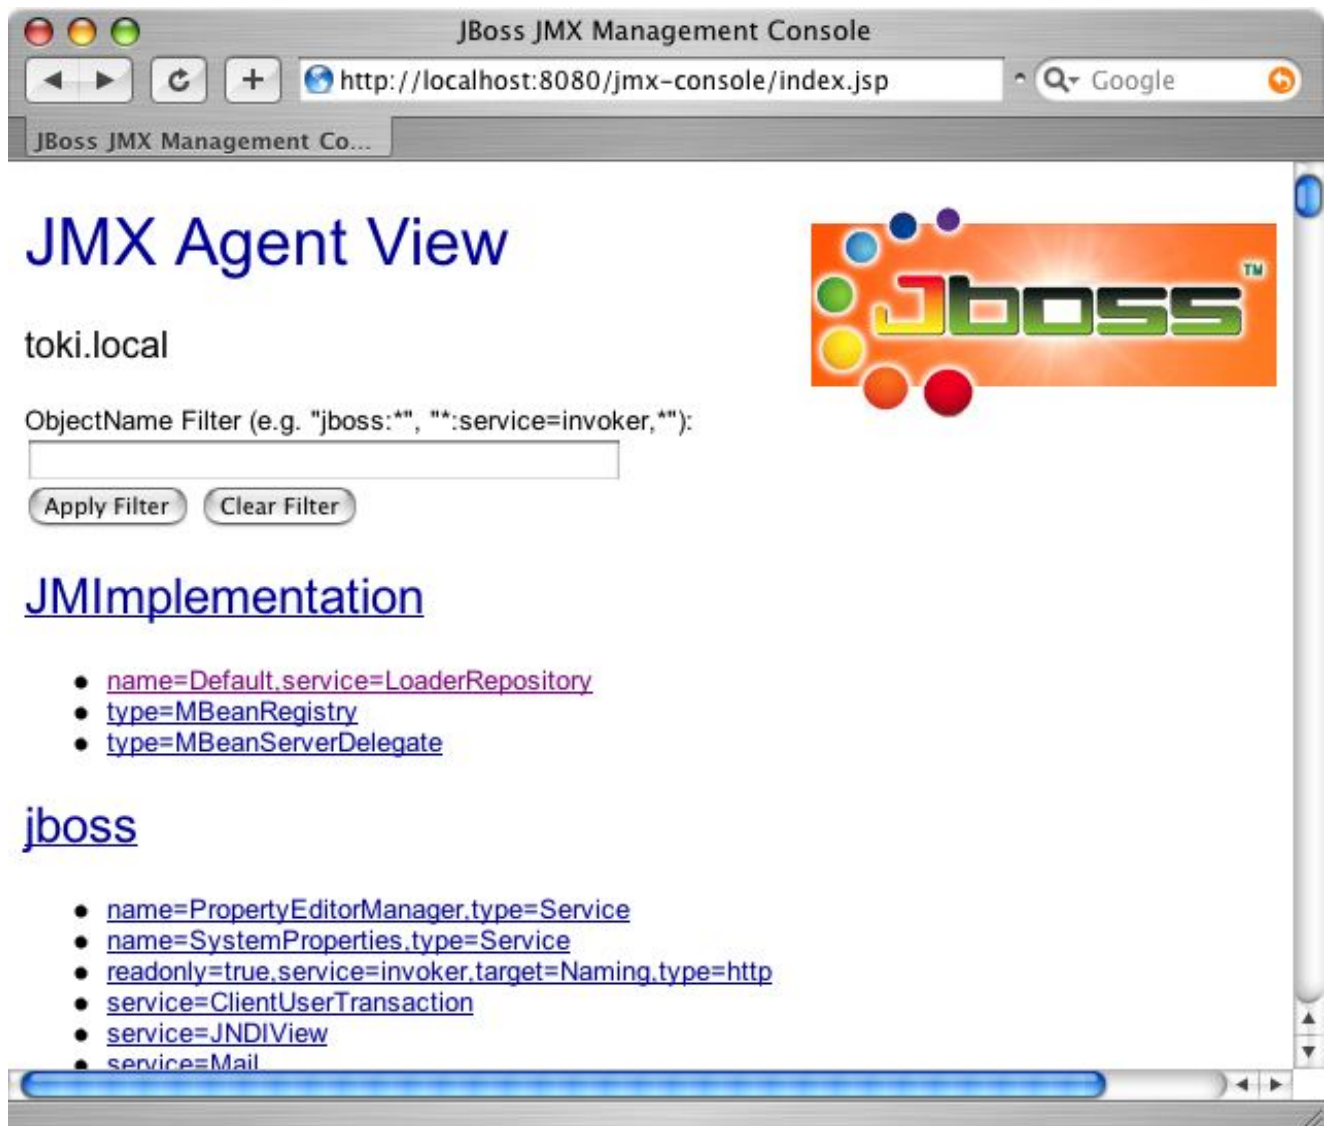
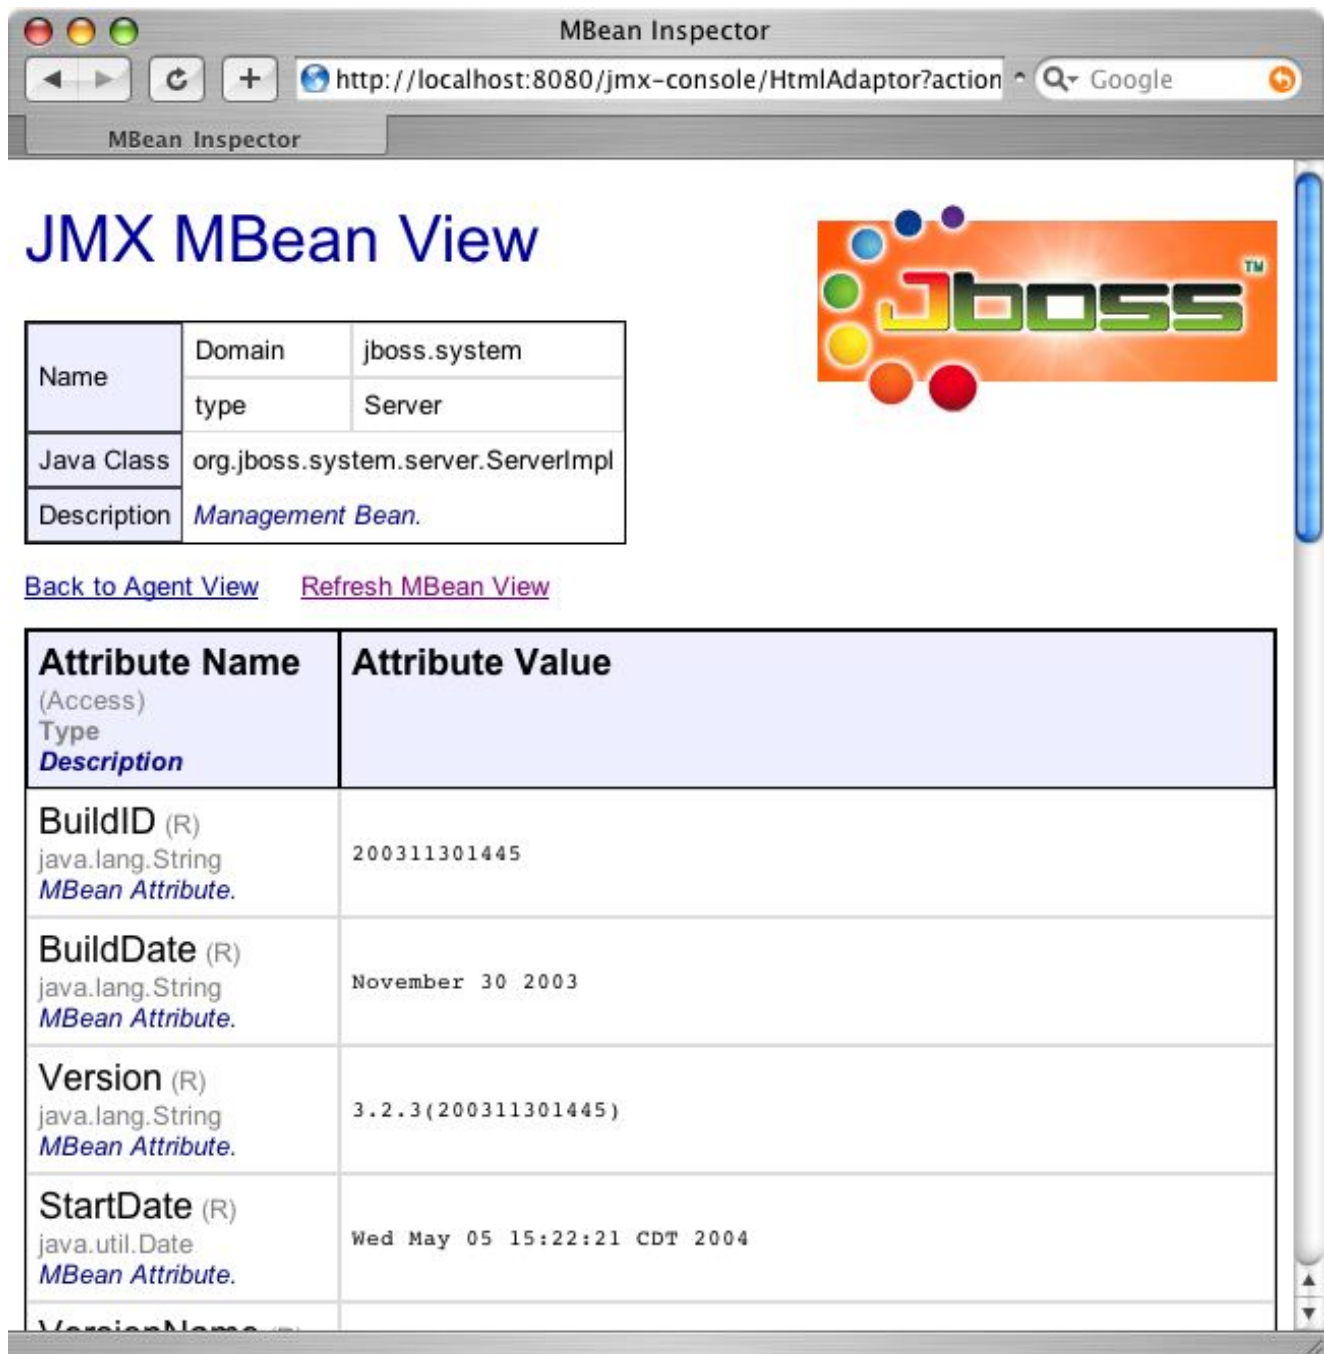


Figure 2.13. The JBoss JMX console web application agent view

The top view is called the agent view and it provides a listing of all MBeans registered with the `MBeanServer` sorted by the domain portion of the MBean's `ObjectName`. Under each domain are the MBeans under that domain. When you select one of the MBeans you will be taken to the MBean view. This allows one to view and edit an MBean's attributes as well as invoke operations. As an example, Figure 2.14 shows the MBean view for the `jboss.system:type=Server` MBean.



JMX MBean View

Name	Domain	jboss.system
	type	Server
Java Class	org.jboss.system.server.ServerImpl	
Description	<i>Management Bean.</i>	

[Back to Agent View](#) [Refresh MBean View](#)

Attribute Name (Access) Type Description	Attribute Value
BuildID (R) java.lang.String <i>MBean Attribute.</i>	200311301445
BuildDate (R) java.lang.String <i>MBean Attribute.</i>	November 30 2003
Version (R) java.lang.String <i>MBean Attribute.</i>	3.2.3(200311301445)
StartDate (R) java.util.Date <i>MBean Attribute.</i>	Wed May 05 15:22:21 CDT 2004
VersionName (R)	

Figure 2.14. The MBean view for the "jboss.system:type=Server" MBean

The source code for the JMX console web application is located in the `varia` module under the `src/main/org/jboss/jmx` directory. Its web pages are located under `varia/src/resources/jmx`. The application is a simple MVC servlet with JSP views that utilize the `MBeanServer`.

2.3.1.1. Securing the JMX Console

Since the JMX console web application is just a standard servlet, it may be secured using standard J2EE role based security. The `jmx-console.war` that is deployed as an unpacked WAR that includes template settings for quickly enabling simple username and password based access restrictions. If you look at the `jmx-console.war` in the `server/default/deploy` directory you will find the `web.xml` and `jboss-web.xml` descriptors in the `WEB-INF` directory and a `jmx-console-roles.properties` and `jmx-console-users.properties` file under `WEB-INF/classes`.

By uncommenting the security sections of the `web.xml` and `jboss-web.xml` descriptors as shown in Example 2.11, you enable HTTP basic authentication that restricts access to the `jmx-console` application to the user `admin` with password `admin`. The username and password are determined by the `admin=admin` line in the `jmx-console-users.properties` file.

Example 2.11. The `jmx-console.war` `web.xml` descriptors with the security elements uncommented.

```
<?xml version="1.0"?>
<!DOCTYPE web-app PUBLIC
    "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
    "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>
    <!-- ... -->

    <!-- A security constraint that restricts access to the HTML JMX console
         to users with the role JBossAdmin. Edit the roles to what you want and
         uncomment the WEB-INF/jboss-web.xml/security-domain element to enable
         secured access to the HTML JMX console.
    -->
    <security-constraint>
        <web-resource-collection>
            <web-resource-name>HtmlAdaptor</web-resource-name>
            <description> An example security config that only allows users with
                the role JBossAdmin to access the HTML JMX console web
                application </description>
            <url-pattern>/*</url-pattern>
            <http-method>GET</http-method>
            <http-method>POST</http-method>
        </web-resource-collection>
        <auth-constraint>
            <role-name>JBossAdmin</role-name>
        </auth-constraint>
    </security-constraint>
    <login-config>
        <auth-method>BASIC</auth-method>
        <realm-name>JBoss JMX Console</realm-name>
    </login-config>
    <security-role>
        <role-name>JBossAdmin</role-name>
    </security-role>
</web-app>
```

Example 2.12. The `jmx-console.war` `jboss-web.xml` descriptors with the security elements uncommented.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE jboss-web
    PUBLIC "-//JBoss//DTD Web Application 2.3//EN"
    "http://www.jboss.org/j2ee/dtd/jboss-web_3_0.dtd">
<jboss-web>
    <!--
        Uncomment the security-domain to enable security. You will
        need to edit the htmladaptor login configuration to setup the
        login modules used to authentication users.
    -->
    <security-domain>java:/jaas/jmx-console</security-domain>
</jboss-web>
```

Make these changes and then when you try to access the JMX Console URL. You will see a dialog similar to that shown in Figure 2.15.



Figure 2.15. The jmx-console basic HTTP login dialog.

Its generally a bad idea to use the properties files for securing access to the JMX console application. To see how to properly configure the security settings of web applications see Chapter 8.

2.3.2. Connecting to JMX Using RMI

JBoss supplies an RMI interface for connecting to the JMX MBeanServer. This interface is `org.jboss.jmx.adaptor.rmi.RMIAdaptor`, and it is shown in Example 2.13.

Example 2.13. The RMIAdaptor interface

```
/*
 * JBoss, the OpenSource J2EE webOS
 *
 * Distributable under LGPL license.
 * See terms of license at gnu.org.
 */
package org.jboss.jmx.adaptor.rmi;

import javax.management.Attribute;
import javax.management.AttributeList;
import javax.management.ObjectName;
import javax.management.QueryExp;
import javax.management.ObjectInstance;
import javax.management.NotificationFilter;
import javax.management.NotificationListener;
import javax.management.MBeanInfo;

import javax.management.AttributeNotFoundException;
import javax.management.InstanceAlreadyExistsException;
import javax.management.InstanceNotFoundException;
import javax.management.IntrospectionException;
import javax.management.InvalidAttributeValueException;
import javax.management.ListenerNotFoundException;
```

```

import javax.management.MBeanException;
import javax.management.MBeanRegistrationException;
import javax.management.NotCompliantMBeanException;
import javax.management.OperationNotSupportedException;
import javax.management.ReflectionException;

public interface RMIAdaptor
    extends java.rmi.Remote
{
    public ObjectInstance createMBean(String pClassName,
                                     ObjectName pName)
        throws ReflectionException,
               InstanceAlreadyExistsException,
               MBeanRegistrationException,
               MBeanException,
               NotCompliantMBeanException,
               RemoteException;

    public ObjectInstance createMBean(String pClassName,
                                     ObjectName pName,
                                     ObjectName pLoaderName)
        throws ReflectionException,
               InstanceAlreadyExistsException,
               MBeanRegistrationException,
               MBeanException,
               NotCompliantMBeanException,
               InstanceNotFoundException,
               RemoteException;

    public ObjectInstance createMBean(String pClassName,
                                     ObjectName pName,
                                     Object[] pParams,
                                     String[] pSignature)
        throws ReflectionException,
               InstanceAlreadyExistsException,
               MBeanRegistrationException,
               MBeanException,
               NotCompliantMBeanException,
               RemoteException;

    public ObjectInstance createMBean(String pClassName,
                                     ObjectName pName,
                                     ObjectName pLoaderName,
                                     Object[] pParams,
                                     String[] pSignature)
        throws ReflectionException,
               InstanceAlreadyExistsException,
               MBeanRegistrationException,
               MBeanException,
               NotCompliantMBeanException,
               InstanceNotFoundException,
               RemoteException;

    public void unregisterMBean(ObjectName pName)
        throws InstanceNotFoundException,
               MBeanRegistrationException,
               RemoteException;

    public ObjectInstance getObjectInstance(ObjectName pName)
        throws InstanceNotFoundException,
               RemoteException;

    public Set queryMBeans(ObjectName pName, QueryExp pQuery)
        throws RemoteException;

    public Set queryNames(ObjectName pName, QueryExp pQuery)
        throws RemoteException;

    public boolean isRegistered(ObjectName pName)
        throws RemoteException;

```

```

public boolean isInstanceOf(ObjectName pName, String pClassName)
    throws InstanceNotFoundException,
           RemoteException;

public Integer getMBeanCount()
    throws RemoteException;

public Object getAttribute(ObjectName pName, String pAttribute)
    throws MBeanException,
           AttributeNotFoundException,
           InstanceNotFoundException,
           ReflectionException,
           RemoteException;

public AttributeList getAttributes(ObjectName pName,
                                   String[] pAttributes)
    throws InstanceNotFoundException,
           ReflectionException,
           RemoteException;

public void setAttribute(ObjectName pName, Attribute pAttribute)
    throws InstanceNotFoundException,
           AttributeNotFoundException,
           InvalidAttributeValueException,
           MBeanException,
           ReflectionException,
           RemoteException;

public AttributeList setAttributes(ObjectName pName,
                                   AttributeList pAttributes)
    throws InstanceNotFoundException,
           ReflectionException,
           RemoteException;

public Object invoke(ObjectName pName, String pActionName,
                    Object[] pParams, String[] pSignature)
    throws InstanceNotFoundException,
           MBeanException,
           ReflectionException,
           RemoteException;

public String getDefaultDomain()
    throws RemoteException;

public void addNotificationListener(ObjectName pName,
                                   ObjectName pListener,
                                   NotificationFilter pFilter,
                                   Object pHandback)
    throws InstanceNotFoundException,
           RemoteException;

public void removeNotificationListener(ObjectName pName,
                                       ObjectName pListener)
    throws InstanceNotFoundException,
           ListenerNotFoundException,
           RemoteException;

public MBeanInfo getMBeanInfo(ObjectName pName)
    throws InstanceNotFoundException,
           IntrospectionException,
           ReflectionException,
           RemoteException;
}

```

The `RMIAdaptor` interface was bound into JNDI by the `org.jboss.jmx.adaptor.rmi.RMIAdaptorService` MBean, but as of 3.2.2 this service has been removed from the `dist` deploy directory by default. It can still be

found in the `docs/examples/jmx` directory, but it has been deprecated in favor of the invoker adaptor service. This service also supports the `RMIAdaptor` interface and its configuration also provides a binding of this interface in the default location of `jmx/rmi/RMIAdaptor` for backwards compatibility with existing clients. The `RMIAdaptorService` still has utility for remote clients that need to receive JMX notifications. The invoker adaptor service does not yet support this capability so if this is required, the `jmx-invoker-adaptor-server.sar` must be replaced with the `jmx-rmi-adaptor.sar` from the examples directory.

The `RMIAdaptorService` is deployed as the `jmx-rmi-adaptor.sar` package, and supports the following attributes:

- **JndiName:** The JNDI name under which the `RMIAdaptor` interface will be bound. The default name is `jmx/rmi/RMIAdaptor`.
- **RMIObjectPort:** The server side listening port number for the exported RMI object. This defaults to 0 meaning choose an anonymous available port.
- **ServerAddress:** The server interface name or IP address to bind the export RMI listening port to. This defaults to an empty value meaning to bind on all available interfaces.
- **Backlog:** The RMI object server socket backlog of client connection requests that will be accepted before a connection error occurs.

Example 2.14 shows a client that makes use of the `RMIAdaptor` interface to query the `MBeanInfo` for the `JNDIView` MBean. It also invokes the MBean's `list(boolean)` method and displays the result.

Example 2.14. A JMX client that uses the `RMIAdaptor`

```
public class JMXBrowser
{
    /**
     * @param args the command line arguments
     */
    public static void main(String[] args)
        throws Exception
    {
        InitialContext ic = new InitialContext();
        RMIAdaptor server = (RMIAdaptor) ic.lookup("jmx/rmi/RMIAdaptor");

        // Get the MBeanInfo for the JNDIView MBean
        ObjectName name = new ObjectName("jboss:service=JNDIView");
        MBeanInfo info = server.getMBeanInfo(name);
        System.out.println("JNDIView Class: "+info.getClassName());

        MBeanOperationInfo[] opInfo = info.getOperations();
        System.out.println("JNDIView Operations: ");

        for (int o = 0; o < opInfo.length; o++) {
            MBeanOperationInfo op = opInfo[o];

            String returnType = op.getReturnType();
            String opName = op.getName();

            System.out.print(" "+returnType+" "+opName+"(");

            MBeanParameterInfo[] params = op.getSignature();
            for (int p = 0; p < params.length; p++) {
                MBeanParameterInfo paramInfo = params[p];

                String pname = paramInfo.getName();
                String type = paramInfo.getType();

                if (pname.equals(type)) {
```

```

        System.out.print(type);
    } else {
        System.out.print(type+" "+name);
    }

    if (p < params.length-1) {
        System.out.print(',');
    }
}
System.out.println("");
}

// Invoke the list(boolean) op
String[] sig = {"boolean"};
Object[] opArgs = {Boolean.TRUE};
Object result = server.invoke(name,
                               "list", opArgs, sig);
System.out.println("JNDIView.list(true) output:\n"+result);
}
}

```

To test the client access using the RMIAdaptor, run the following:

```

[orb@toki examples]$ ant -Dchap=chap2 -Dex=4 run-example
Buildfile: build.xml

...

run-example4:
[java] JNDIView Class: org.jboss.mx.modelmbean.XMBean
[java] JNDIView Operations:
[java] + java.lang.String list(boolean jboss:service=JNDIView)
[java] + java.lang.String listXML()
[java] + void create()
[java] + void start()
[java] + void stop()
[java] + void destroy()
[java] + void jbossInternalLifecycle(java.lang.String jboss:service=JNDIView)
[java] + java.lang.String getName()
[java] + int getState()
[java] + java.lang.String getStateString()
[java] JNDIView.list(true) output:
[java] <h1>java: Namespace</h1>
[java] <pre>
[java] +- XAConnectionFactory (class: org.jboss.mq.SpyXAConnectionFactory)
[java] +- DefaultDS (class: org.jboss.resource.adapter.jdbc.WrapperDataSource)
[java] +- SecurityProxyFactory (class: org.jboss.security.SubjectSecurityProxyFacto
ry)
[java] +- DefaultJMSProvider (class: org.jboss.jms.jndi.JNDIProviderAdapter)
[java] +- comp (class: javax.naming.Context)
[java] +- JmsXA (class: org.jboss.resource.adapter.jms.JmsConnectionFactoryImpl)
[java] +- ConnectionFactory (class: org.jboss.mq.SpyConnectionFactory)
[java] +- jaas (class: javax.naming.Context)
[java] | +- JmsXARealm (class: org.jboss.security.plugins.SecurityDomainContext)
[java] | +- jbossmq (class: org.jboss.security.plugins.SecurityDomainContext)
[java] | +- HsqlDbRealm (class: org.jboss.security.plugins.SecurityDomainContext)
[java] +- timedCacheFactory (class: javax.naming.Context)
[java] Failed to lookup: timedCacheFactory, errmsg=null
[java] +- TransactionPropagationContextExporter (class: org.jboss.tm.TransactionPro
pagationContextFactory)
[java] +- Mail (class: javax.mail.Session)
[java] +- StdJMSPool (class: org.jboss.jms.asf.StdServerSessionPoolFactory)
[java] +- TransactionPropagationContextImporter (class: org.jboss.tm.TransactionPro
pagationContextImporter)
[java] +- TransactionManager (class: org.jboss.tm.TxManager)
[java] </pre>
[java] <h1>Global JNDI Namespace</h1>
[java] <pre>
[java] +- HAILConnectionFactory[link -> ConnectionFactory] (class: javax.naming.

```

```

    LinkRef)
[java] +- jmx (class: org.jnp.interfaces.NamingContext)
[java] | +- invoker (class: org.jnp.interfaces.NamingContext)
[java] | | +- RMIAdaptor (proxy: $Proxy26 implements interface org.jboss.jmx.ad
    aptor.rmi.RMIAdaptor,interface org.jboss.jmx.adaptor.rmi.RMIAdaptorExt)
[java] | +- rmi (class: org.jnp.interfaces.NamingContext)
[java] | | +- RMIAdaptor[link -> jmx/invoker/RMIAdaptor] (class: javax.namin
g.LinkRef)
[java] +- HTTPXAConnectionFactory (class: org.jboss.mq.SpyXAConnectionFactory)
[java] +- ConnectionFactory (class: org.jboss.mq.SpyConnectionFactory)
[java] +- UserTransactionSessionFactory (proxy: $Proxy13 implements interface org.j
    boss.tm.usertx.interfaces.UserTransactionSessionFactory)
[java] +- HTTPConnectionFactory (class: org.jboss.mq.SpyConnectionFactory)
[java] +- XAConnectionFactory (class: org.jboss.mq.SpyXAConnectionFactory)
[java] +- invokers (class: org.jnp.interfaces.NamingContext)
[java] | +- 0.0.0.0 (class: org.jnp.interfaces.NamingContext)
[java] | | +- pooled (class: org.jboss.invocation.pooled.interfaces.PooledInvok
erProxy)
[java] +- UserTransaction (class: org.jboss.tm.usertx.client.ClientUserTransaction)
[java] +- UILXAConnectionFactory[link -> XAConnectionFactory] (class: javax.nami
ng.LinkRef)
[java] +- HAILXAConnectionFactory[link -> XAConnectionFactory] (class: javax.nam
ing.LinkRef)
[java] +- UIL2XAConnectionFactory[link -> XAConnectionFactory] (class: javax.nam
ing.LinkRef)
[java] +- queue (class: org.jnp.interfaces.NamingContext)
[java] | +- A (class: org.jboss.mq.SpyQueue)
[java] | +- testQueue (class: org.jboss.mq.SpyQueue)
[java] | +- ex (class: org.jboss.mq.SpyQueue)
[java] | +- DLQ (class: org.jboss.mq.SpyQueue)
[java] | +- D (class: org.jboss.mq.SpyQueue)
[java] | +- C (class: org.jboss.mq.SpyQueue)
[java] | +- B (class: org.jboss.mq.SpyQueue)
[java] +- topic (class: org.jnp.interfaces.NamingContext)
[java] | +- testDurableTopic (class: org.jboss.mq.SpyTopic)
[java] | +- testTopic (class: org.jboss.mq.SpyTopic)
[java] | +- securedTopic (class: org.jboss.mq.SpyTopic)
[java] +- console (class: org.jnp.interfaces.NamingContext)
[java] | +- PluginManager (proxy: $Proxy27 implements interface org.jboss.console
    .manager.PluginManagerMBean)
[java] +- UIL2ConnectionFactory[link -> ConnectionFactory] (class: javax.naming.
LinkRef)
[java] +- UILConnectionFactory[link -> ConnectionFactory] (class: javax.naming.L
inkRef)
[java] +- UUIDKeyGeneratorFactory (class: org.jboss.ejb.plugins.keygenerator.uuid.U
UIDKeyGeneratorFactory)
[java] </pre>

```

2.3.3. Command Line Access to JMX

JBoss provides a simple command line tool that allows for interaction with a remote JMX server instance. This tool is called `twiddle` (for twiddling bits via JMX) and is located in the `bin` directory of the distribution. Twiddle is a command execution tool, not a general command shell. It is run using either the `twiddle.sh` or `twiddle.bat` scripts, and passing in a `-h(--help)` argument provides the basic syntax, and `--help-commands` shows what you can do with the tool:

```

[nr@toki bin]$ ./twiddle.sh -h
A JMX client to 'twiddle' with a remote JBoss server.

usage: twiddle.sh [options] <command> [command_arguments]

options:
-h, --help Show this help message
--help-commands Show a list of commands
-H=<command> Show command specific help
-c=command.properties Specify the command.properties file to use
-D<name>[=<value>] Set a system property

```

```
-- Stop processing options
-s, --server=<url> The JNDI URL of the remote server
-a, --adapter=<name> The JNDI name of the RMI adapter to use
[nr@toki bin]$ ./twiddle.sh --help-commands
twiddle.sh commands:
get Get the values of one or more MBean attributes
invoke Invoke an operation on an MBean
unregister Unregister one or more MBeans
create Create an MBean
serverinfo Get information about the MBean server
query Query the server for a list of matching MBeans
info Get the metadata for an MBean
```

2.3.3.1. Connecting twiddle to a Remote Server

By default the twiddle command will connect to the localhost at port 1099 to lookup the default `jmx/rmi/RMIAdaptor` binding of the `RMIAdaptor` service as the connector for communicating with the JMX server. To connect to a different server/port combination you can use the `-s` (`--server`) option:

```
[nr@rubik bin]$ ./twiddle.sh -s toki serverinfo -d jboss
[nr@rubik bin]$ ./twiddle.sh -s toki:1099 serverinfo -d jboss
```

To connect using a different `RMIAdaptor` binding use the `-a` (`--adapter`) option:

```
[nr@rubik bin]$ ./twiddle.sh -s toki -a jmx/rmi/RMIAdaptor serverinfo -d jboss
[nr@rubik bin]$ ./twiddle.sh -s toki --adapter=jmx/rmi/RMIAdaptor serverinfo -d jboss
```

2.3.3.2. Sample twiddle Command Usage

To access basic information about a server, use the `serverinfo` command. This currently supports:

```
[nr@toki bin]$ ./twiddle.sh -H serverinfo
Get information about the MBean server

usage: serverinfo [options]

options:
-d, --domain Get the default domain
-c, --count Get the MBean count
-l, --list List the MBeans
-- Stop processing options

[nr@rubik bin]$ ./twiddle.sh --server=toki serverinfo --count
385
[nr@rubik bin]$ ./twiddle.sh --server=toki serverinfo --domain
jboss
```

To query the server for the name of MBeans matching a pattern, use the `query` command. This currently supports:

```
[nr@rubik bin]$ ./twiddle.sh -H query
Query the server for a list of matching MBeans

usage: query [options] <query>
options:
  -c, --count      Display the matching MBean count
  --              Stop processing options

Examples:
  query all mbeans: query '*:*'
  query all mbeans in the jboss.j2ee domain: query 'jboss.j2ee:*'
[nr@rubik bin]$ ./twiddle.sh -s toki query 'jboss:service=invoker,*'
jboss:readonly=true,service=invoker,target=Naming,type=http
jboss:service=invoker,type=jrmp
```



```
jboss:service=invoker,type=httpHA
jboss:service=invoker,type=local
jboss:service=invoker,socketType=SSL,type=jrmp
jboss:service=invoker,type=pooled
jboss:service=invoker,type=http
jboss:service=invoker,target=Naming,type=http
```

To get the attributes of an MBean, use the get command:

```
[nr@toki bin]$ ./twiddle.sh -H get
Get the values of one or more MBean attributes

usage: get [options] <name> [<attr>+]
  If no attribute names are given all readable attributes are retrieved
options:
  --noprfix      Do not display attribute name prefixes
  --             Stop processing options
[nr@toki bin]$ ./twiddle.sh get jboss:service=invoker,type=jrmp RMIObjectPort StateString
RMIObjectPort=4444
StateString=Started
[nr@toki bin]$ ./twiddle.sh get jboss:service=invoker,type=jrmp
ServerAddress=0.0.0.0
StateString=Started
State=3
EnableClassCaching=false
SecurityDomain=null
RMIServerSocketFactory=null
Backlog=200
RMIObjectPort=4444
Name=JRMPInvoker
RMIClientSocketFactory=null
```

To query the MBeanInfo for an MBean, use the info command:

```
[nr@toki bin]$ ./twiddle.sh -H info
Get the metadata for an MBean

usage: info <mbean-name>
  Use '*' to query for all attributes
[nr@toki bin]$ ./twiddle.sh info jboss:service=invoker,type=jrmp
Description: Management Bean.
+++ Attributes:
Name: ServerAddress
Type: java.lang.String
Access: rw
Name: StateString
Type: java.lang.String
Access: r-
Name: State
Type: int
Access: r-
Name: EnableClassCaching
Type: boolean
Access: rw
Name: SecurityDomain
Type: java.lang.String
Access: rw
Name: RMIServerSocketFactory
Type: java.lang.String
Access: rw
Name: Backlog
Type: int
Access: rw
Name: RMIObjectPort
Type: int
Access: rw
Name: Name
Type: java.lang.String
Access: r-
```

```

Name: RMIClientSocketFactory
Type: java.lang.String
Access: rw
+++ Operations:
void start()
void jbossInternalLifecycle(java.lang.String java.lang.String)
void destroy()
void create()
void stop()

```

To invoke an operation on an MBean, use the invoker command:

```

[nr@toki bin]$ ./twiddle.sh -H invoke
Invoke an operation on an MBean

usage: invoke [options] <query> <operation> (<arg>)*

options:
  -q, --query-type[=<type>]    Treat object name as a query
  --                          Stop processing options

query type:
  f[first]    Only invoke on the first matching name [default]
  a[all]      Invoke on all matching names
[nr@toki bin]$
<hl>java: Namespace</hl>
<pre>
+- XAConnectionFactory (class: org.jboss.mq.SpyXAConnectionFactory)
+- DefaultDS (class: org.jboss.resource.adapter.jdbc.WrapperDataSource)
+- SecurityProxyFactory (class: org.jboss.security.SubjectSecurityProxyFactory)
+- DefaultJMSProvider (class: org.jboss.jms.jndi.JNDIProviderAdapter)
+- comp (class: javax.naming.Context)
+- JmsXA (class: org.jboss.resource.adapter.jms.JmsConnectionFactoryImpl)
+- ConnectionFactory (class: org.jboss.mq.SpyConnectionFactory)
+- jaas (class: javax.naming.Context)
|   +- JmsXARealm (class: org.jboss.security.plugins.SecurityDomainContext)
|   +- jbossmq (class: org.jboss.security.plugins.SecurityDomainContext)
|   +- HsqlDbRealm (class: org.jboss.security.plugins.SecurityDomainContext)
+- timedCacheFactory (class: javax.naming.Context)
Failed to lookup: timedCacheFactory, errmsg=null
+- TransactionPropagationContextExporter (class: org.jboss.tm.TransactionPropagationContextFactory)
+- Mail (class: javax.mail.Session)
+- StdJMSPool (class: org.jboss.jms.asf.StdServerSessionPoolFactory)
+- TransactionPropagationContextImporter (class: org.jboss.tm.TransactionPropagationContextImporter)
+- TransactionManager (class: org.jboss.tm.TxManager)
</pre>
<hl>Global JNDI Namespace</hl>
<pre>
+- HAILConnectionFactory[link -> ConnectionFactory] (class: javax.naming.LinkRef)
+- jmx (class: org.jnp.interfaces.NamingContext)
|   +- invoker (class: org.jnp.interfaces.NamingContext)
|   |   +- RMIAdaptor (proxy: $Proxy26 implements interface org.jboss.jmx.adaptor.rmi.RMIAdaptor, interface org.jboss.jmx.adaptor.rmi.RMIAdaptorExt)
|   |   |   +- rmi (class: org.jnp.interfaces.NamingContext)
|   |   |   |   +- RMIAdaptor[link -> jmx/invoke/RMIAdaptor] (class: javax.naming.LinkRef)
+- HTTPXAConnectionFactory (class: org.jboss.mq.SpyXAConnectionFactory)
+- ConnectionFactory (class: org.jboss.mq.SpyConnectionFactory)
+- UserTransactionSessionFactory (proxy: $Proxy13 implements interface org.jboss.tm.user.tx.interfaces.UserTransactionSessionFactory)
+- HTTPConnectionFactory (class: org.jboss.mq.SpyConnectionFactory)
+- XAConnectionFactory (class: org.jboss.mq.SpyXAConnectionFactory)
+- invokers (class: org.jnp.interfaces.NamingContext)
|   +- 0.0.0.0 (class: org.jnp.interfaces.NamingContext)
|   |   +- pooled (class: org.jboss.invocation.pooled.interfaces.PooledInvokerProxy)
+- UserTransaction (class: org.jboss.tm.usertx.client.ClientUserTransaction)
+- UILXAConnectionFactory[link -> XAConnectionFactory] (class: javax.naming.LinkRef)
+- HAILXAConnectionFactory[link -> XAConnectionFactory] (class: javax.naming.LinkRef)
+- UIL2XAConnectionFactory[link -> XAConnectionFactory] (class: javax.naming.LinkRef)

```

```

+- queue (class: org.jnp.interfaces.NamingContext)
|
| +- A (class: org.jboss.mq.SpyQueue)
| +- testQueue (class: org.jboss.mq.SpyQueue)
| +- ex (class: org.jboss.mq.SpyQueue)
| +- DLQ (class: org.jboss.mq.SpyQueue)
| +- D (class: org.jboss.mq.SpyQueue)
| +- C (class: org.jboss.mq.SpyQueue)
| +- B (class: org.jboss.mq.SpyQueue)
+- topic (class: org.jnp.interfaces.NamingContext)
|
| +- testDurableTopic (class: org.jboss.mq.SpyTopic)
| +- testTopic (class: org.jboss.mq.SpyTopic)
| +- securedTopic (class: org.jboss.mq.SpyTopic)
+- console (class: org.jnp.interfaces.NamingContext)
|
| +- PluginManager (proxy: $Proxy27 implements interface org.jboss.console.manager.PluginManagerMBean)
+- UIL2ConnectionFactory[link -> ConnectionFactory] (class: javax.naming.LinkRef)
+- UILConnectionFactory[link -> ConnectionFactory] (class: javax.naming.LinkRef)
+- UUIDKeyGeneratorFactory (class: org.jboss.ejb.plugins.keygenerator.uuid.UUIDKeyGeneratorFactory)
</pre>

```

2.3.4. Connecting to JMX Using Any Protocol

With the detached invokers and a somewhat generalized proxy factory capability, you can really talk to the JMX server using the `InvokerAdaptorService` and a proxy factory service to expose an `RMIAaptor` or similar interface over your protocol of choice. We will introduce the detached invoker notion along with proxy factories in Section 2.7. See Section 2.7.1 for an example of an invoker service that allows one to access the MBean server using to the `RMIAaptor` interface over any protocol for which a proxy factory service exists.

2.4. Using JMX as a Microkernel

When JBoss starts up, one of the first steps performed is to create an MBean server instance (`javax.management.MBeanServer`). The JMX MBean server in the JBoss architecture plays the role of a microkernel. All other manageable MBean components are plugged into JBoss by registering with the MBean server. The kernel in that sense is only an framework, and not a source of actual functionality. The functionality is provided by MBeans, and in fact all major JBoss components are manageable MBeans interconnected through the MBean server.

2.4.1. The Startup Process

In this section we will describe the JBoss server startup process. A summary of the steps that occur during the JBoss server startup sequence is:

1. The run start script initiates the boot sequence using the `org.jboss.Main.main(String[])` method entry point.
2. The `Main.main` method creates a thread group named `jboss` and then starts a thread belonging to this thread group. This thread invokes the `Main.boot` method.
3. The `Main.boot` method processes the `Main.main` arguments and then creates an `org.jboss.system.server.ServerLoader` using the system properties along with any additional properties specified as arguments.
4. The XML parser libraries, `jboss-jmx.jar`, `concurrent.jar` and extra libraries and classpaths given as arguments are registered with the `ServerLoader`.

5. The JBoss server instance is created using the `ServerLoader.load(ClassLoader)` method with the current thread context class loader passed in as the `ClassLoader` argument. The returned server instance is an implementation of the `org.jboss.system.server.Server` interface. The creation of the server instance entails:
 - Creating a `java.net.URLClassLoader` with the URLs of the jars and directories registered with the `ServerLoader`. This `URLClassLoader` uses the `ClassLoader` passed in as its parent and it is pushed as the thread context class loader.
 - The class name of the implementation of the `Server` interface to use is determined by the `jboss.server.type` property. This defaults to `org.jboss.system.server.ServerImpl`.
 - The `Server` implementation class is loaded using the `URLClassLoader` created in step 6 and instantiated using its no-arg constructor. The thread context class loader present on entry into the `ServerLoader.load` method is restored and the server instance is returned.
6. The server instance is initialized with the properties passed to the `ServerLoader` constructor using the `Server.init(Properties)` method.
7. The server instance is then started using the `Server.start()` method. The default implementation performs the following steps:
 - Set the thread context class loader to the class loader used to load the `ServerImpl` class.
 - Create an `MBeanServer` under the `jboss` domain using the `MBeanServerFactory.createMBeanServer(String)` method.
 - Register the `ServerImpl` and `ServerConfigImpl` MBeans with the MBean server.
 - Initialize the unified class loader repository to contain all JARs in the optional patch directory as well as the server configuration file `conf` directory, for example, `server/default/conf`. For each JAR and directory an `org.jboss.mx.loading.UnifiedClassLoader` is created and registered with the unified repository. One of these `UnifiedClassLoader` is then set as the thread context class loader. This effectively makes all `UnifiedClassLoaders` available through the thread context class loader.
 - The `org.jboss.system.ServiceController` MBean is created. The `ServiceController` manages the JBoss MBean services life cycle. We will discuss the JBoss MBean services notion in detail in Section 2.4.2.
 - The `org.jboss.deployment.MainDeployer` is created and started. The `MainDeployer` manages deployment dependencies and directing deployments to the correct deployer.
 - The `org.jboss.deployment.JARDeployer` is created and started. The `JARDeployer` handles the deployment of JARs that are simple library JARs.
 - The `org.jboss.deployment.SARDeployer` is created and started. The `SARDeployer` handles the deployment of JBoss MBean services.
 - The `MainDeployer` is invoked to deploy the services defined in the `conf/jboss-service.xml` of the current server file set.
 - Restore the thread context class loader.

The JBoss server starts out as nothing more than a container for the JMX MBean server, and then loads its per-

sonality based on the services defined in the `jboss-service.xml` MBean configuration file from the named configuration set passed to the server on the command line. Because MBeans define the functionality of a JBoss server instance, it is important to understand how the core JBoss MBeans are written, and how you should integrate your existing services into JBoss using MBeans. This is the topic of the next section.

2.4.2. JBoss MBean Services

As we have seen, JBoss relies on JMX to load in the MBean services that make up a given server instance's personality. All of the bundled functionality provided with the standard JBoss distribution is based on MBeans. The best way to add services to the JBoss server is to write your own JMX MBeans.

There are two classes of MBeans: those that are independent of JBoss services, and those that are dependent on JBoss services. MBeans that are independent of JBoss services are the trivial case. They can be written per the JMX specification and added to a JBoss server by adding an `mbean` tag to the `deploy/user-service.xml` file. Writing an MBean that relies on a JBoss service such as naming requires you to follow the JBoss service pattern. The JBoss MBean service pattern consists of a set of life cycle operations that provide state change notifications. The notifications inform an MBean service when it can create, start, stop, and destroy itself. The management of the MBean service life cycle is the responsibility of three JBoss MBeans: `SARDeployer`, `ServiceConfigurator` and `ServiceController`.

2.4.2.1. The SARDeployer MBean

JBoss manages the deployment of its MBean services via a custom MBean that loads an XML variation of the standard JMX MLet configuration file. This custom MBean is implemented in the `org.jboss.deployment.SARDeployer` class. The `SARDeployer` MBean is loaded when JBoss starts up as part of the bootstrap process. The SAR acronym stands for *service archive*.

The `SARDeployer` handles services archives. A service archive can be either a jar that ends with a `.sar` suffix and contains a `META-INF/jboss-service.xml` descriptor, or a standalone XML descriptor with a naming pattern that matches `*-service.xml`. The DTD for the service descriptor is given in Figure 2.16.

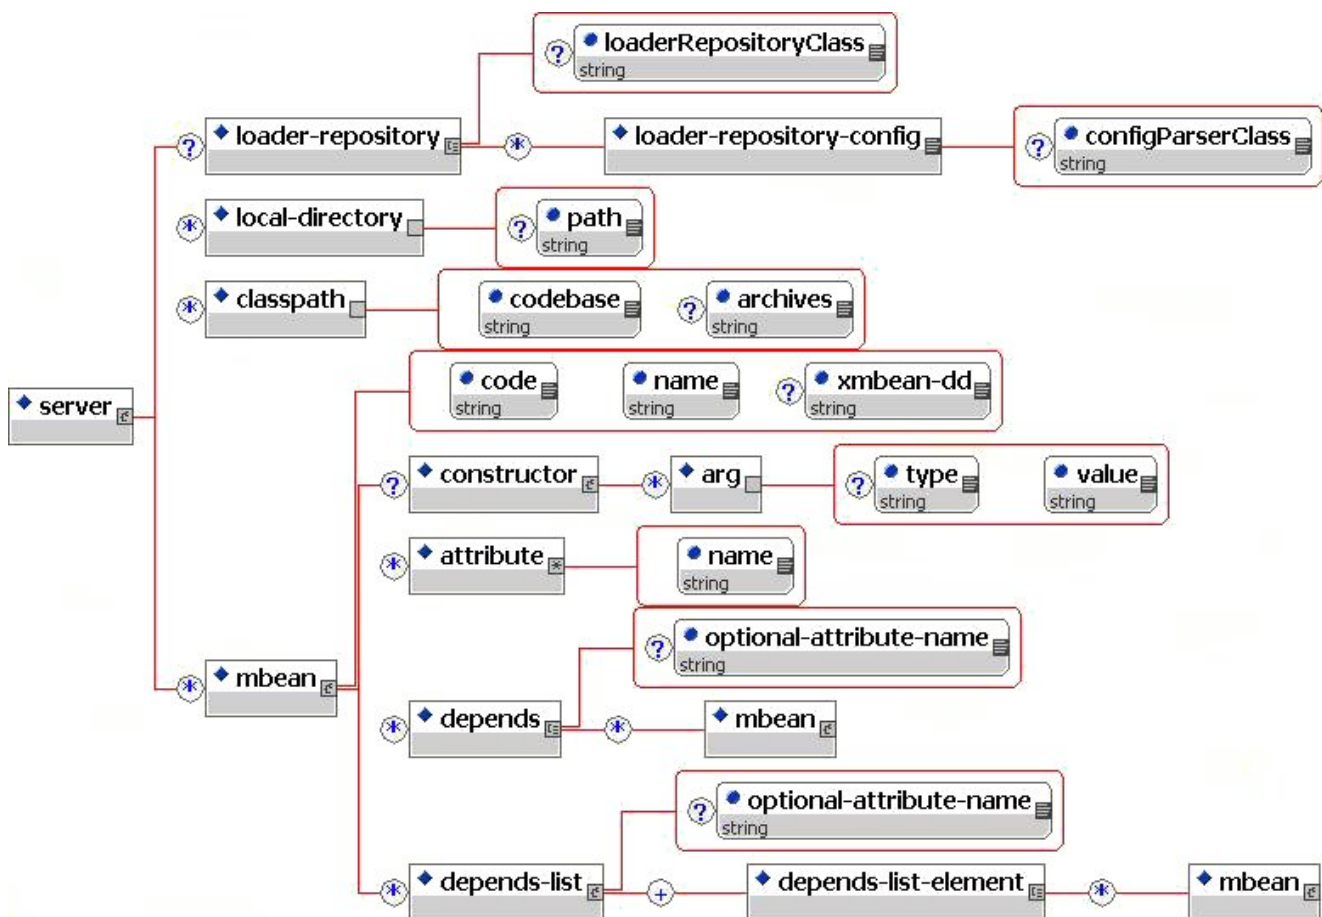


Figure 2.16. The DTD for the MBean service descriptor parsed by the SARDeployer

The elements of the DTD are:

- **server/loader-repository:** This element specifies the name of the `UnifiedLoaderRepository` MBean to use for the SAR to provide SAR level scoping of classes deployed in the sar. It is a unique JMX Object-Name string. It may also specify an arbitrary configuration by including a `loader-repository-config` element. The optional `loaderRepositoryClass` attribute specifies the fully qualified name of the loader repository implementation class. It defaults to `org.jboss.mx.loading.HeirachicalLoaderRepository3`.
- **server/loader-repository/loader-repository-config:** This optional element specifies an arbitrary configuration that may be used to configure the `loadRepositoryClass`. The optional `configParserClass` attribute gives the fully qualified name of the `org.jboss.mx.loading.LoaderRepositoryFactory.LoaderRepositoryConfigParser` implementation to use to parse the `loader-repository-config` content.
- **server/local-directory:** This element specifies a path within the deployment archive that should be copied to the `server/<config>/db` directory for use by the MBean. The `path` attribute is the name of an entry within the deployment archive.
- **server/classpath:** This element specifies one or more external JARs that should be deployed with the MBean(s). The optional `archives` attribute specifies a comma separated list of the JAR names to load, or the `*` wild card to signify that all jars should be loaded. The wild card only works with file URLs, and http URLs if the web server supports the WEBDAV protocol. The `codebase` attribute specifies the URL from which the JARs specified in the `archive` attribute should be loaded. If the `codebase` is a path rather than a URL string, the full URL is built by treating the `codebase` value as a path relative to the JBoss distribution

`server/<config>` directory. The order of JARs specified in the archives as well as the ordering across multiple classpath element is used as the classpath ordering of the JARs. Therefore, if you have patches or inconsistent versions of classes that require a certain ordering, use this feature to ensure the correct ordering. Both the `codebase` and `archives` attributes values may reference a system property using a pattern `${x}` to refer to replacement of the `x` system property.

- **server/mbean:** This element specifies an MBean service. The required `code` attribute gives the fully qualified name of the MBean implementation class. The required `name` attribute gives the JMX `ObjectName` of the MBean. The optional `xmbean-dd` attribute specifies the path to the XMBean resource if this MBean service uses the JBoss XMBean descriptor to define a Model MBean management interface.
- **server/mbean/attribute:** Each attribute element specifies a name/value pair of the attribute of the MBean. The name of the attribute is given by the `name` attribute, and the attribute element body gives the value. The body may be a text representation of the value, or an arbitrary element and child elements if the type of the MBean attribute is `org.w3c.dom.Element`. For text values, the text is converted to the attribute type using the JavaBean `java.beans.PropertyEditor` mechanism.

The text value of an attribute may reference a system property `x` by using the pattern `${x}`. In this case the value of the attribute will be the result of `System.getProperty("x")`, or null if no such property exists.

- **server/mbean/depends** and **server/mbean/depends-list:** these elements specify a dependency from the MBean using the element to the MBean(s) named by the `depends` or `depends-list` elements. Section 2.4.2.4. Note that the dependency value can be another `mbean` element which defines a nested `mbean`.

When the `SARDeployer` is asked to deploy a service performs several steps. Figure 2.17 is a sequence diagram that shows the init through start phases of a service.

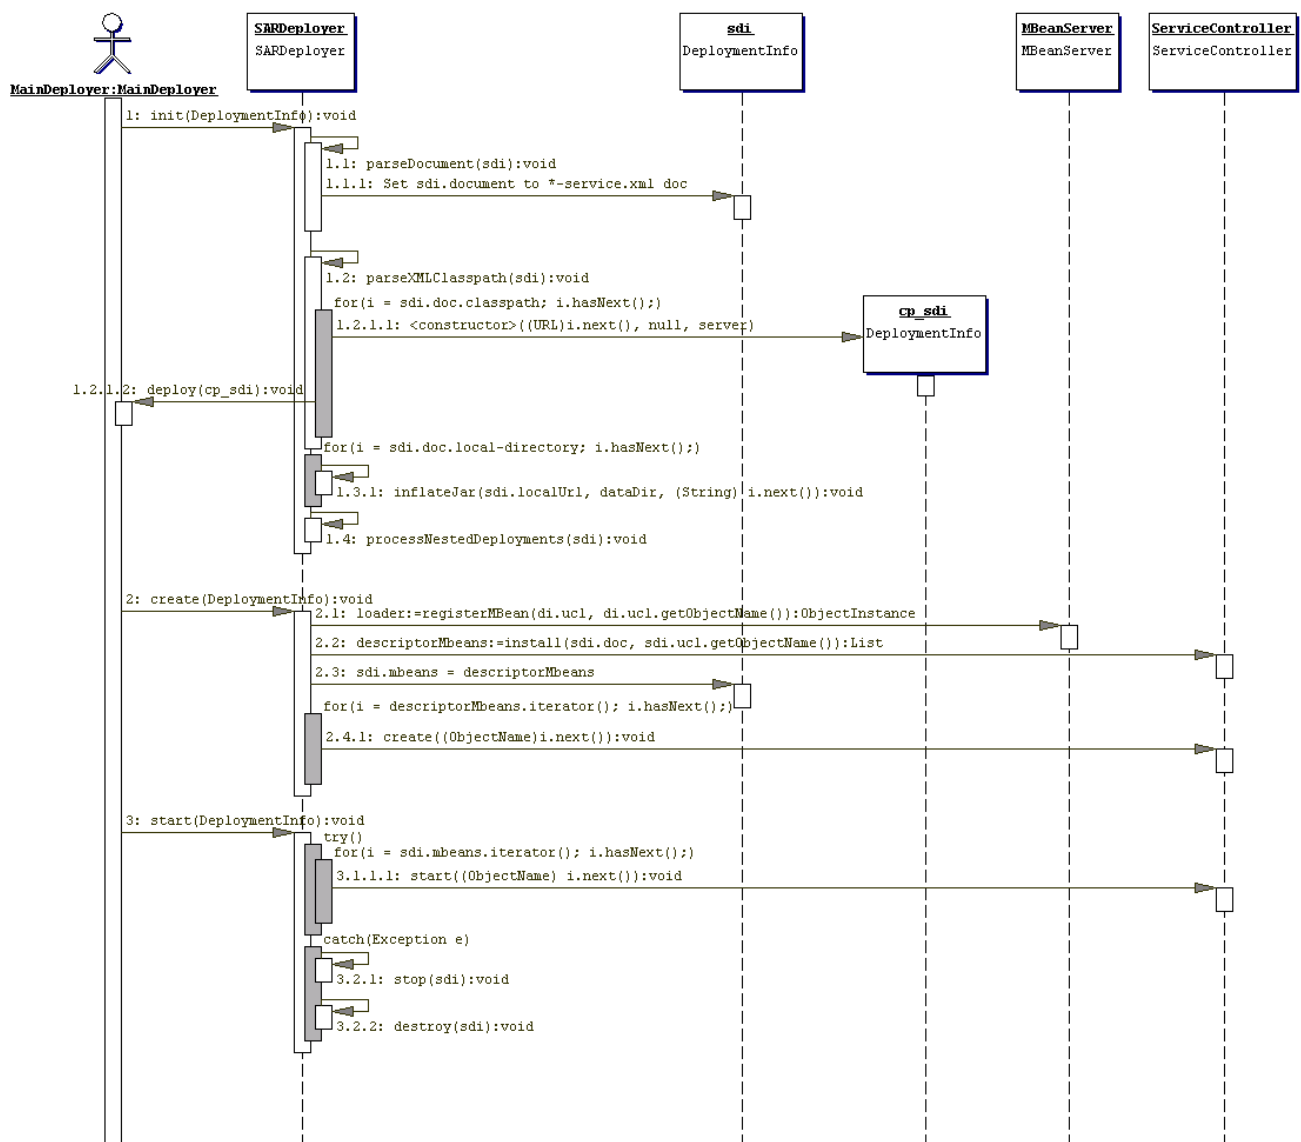


Figure 2.17. A sequence diagram highlighting the main activities performed by the SARDeployer to start a JBoss MBean service

In Figure 2.17 the following is illustrated:

- Methods prefixed with 1.1 correspond to the load and parse of the XML service descriptor.
- Methods prefixed with 1.2 correspond to processing each classpath element in the service descriptor to create an independent deployment that makes the jar or directory available through a UnifiedClassLoader registered with the unified loader repository.
- Methods prefixed with 1.3 correspond to processing each `local-directory` element in the service descriptor. This does a copy of the SAR elements specified in the `path` attribute to the `server/<config>/db` directory.
- Method 1.4. Process each deployable unit nested in the service a child deployment is created and added to the service deployment info subdeployment list.
- Method 2.1. The `UnifiedClassLoader` of the SAR deployment unit is registered with the MBean Server so that it can be used for loading of the SAR MBeans.

- Method 2.2. For each MBean element in the descriptor, create an instance and initialize its attributes with the values given in the service descriptor. This is done by calling the `ServiceController.install` method.
- Method 2.4.1. For each MBean instance created, obtain its JMX `ObjectName` and ask the `ServiceController` to handle the create step of the service life cycle. The `ServiceController` handles the dependencies of the MBean service. Only if the service's dependencies are satisfied is the service create method invoked.
- Methods prefixed with 3.1 correspond to the start of each MBean service defined in the service descriptor. For each MBean instance created, obtain its JMX `ObjectName` and ask the `ServiceController` to handle the start step of the service life cycle. The `ServiceController` handles the dependencies of the MBean service. Only if the service's dependencies are satisfied is the service start method invoked.

2.4.2.2. The Service Life Cycle Interface

The JMX specification does not define any type of life cycle or dependency management for MBeans. The JBoss `ServiceController` MBean introduces this notion. A JBoss MBean is an extension of the JMX MBean in that an MBean is expected to decouple creation from the life cycle of its service duties. This is necessary to implement any type of dependency management. For example, if you are writing an MBean that needs a JNDI naming service to be able to function, your MBean needs to be told when its dependencies are satisfied. This ranges from difficult to impossible to do if the only life cycle event is the MBean constructor. Therefore, JBoss introduces a service life cycle interface that describes the events a service can use to manage its behavior. The following listing shows the `org.jboss.system.Service` interface:

```
package org.jboss.system;
public interface Service
{
    public void create() throws Exception;
    public void start() throws Exception;
    public void stop();
    public void destroy();
}
```

The `ServiceController` MBean invokes the methods of the `Service` interface at the appropriate times of the service life cycle. We'll discuss the methods in more detail in the `ServiceController` section.

Note that there is a J2EE management specification request (JSR 77, <http://jcp.org/jsr/detail/77.jsp>) that introduces a state management notion that includes a start/stop lifecycle notion. When this standard is finalized JBoss will likely support an extension of the JSR 77 based service lifecycle implementation. As of the 3.2.0 release we do support JSR77 management objects and most of the statistics, but the lifecycle operations are not supported.

2.4.2.3. The ServiceController MBean

JBoss manages dependencies between MBeans via the `org.jboss.system.ServiceController` custom MBean. The `SARDeployer` delegates to the `ServiceController` when initializing, creating, starting, stopping and destroying MBean services. Figure 2.18 shows a sequence diagram that highlights interaction between the `SARDeployer` and `ServiceController`.

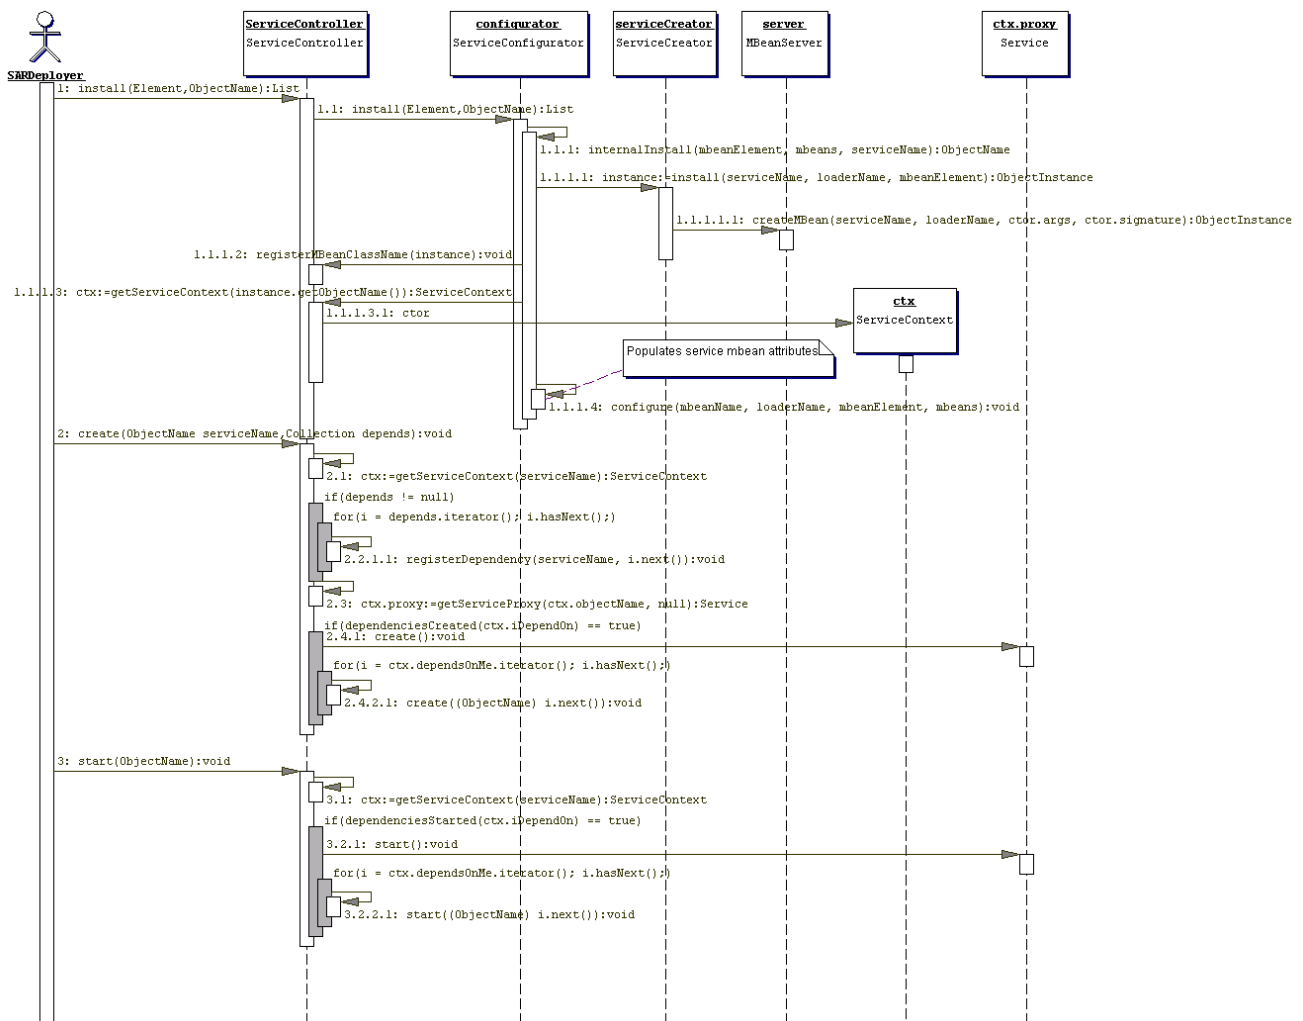


Figure 2.18. The interaction between the SARDeployer and ServiceController to start a service

The ServiceController MBean has four key methods for the management of the service life cycle: create, start, stop and destroy.

2.4.2.3.1. The create(ObjectName) method

The create(ObjectName) method is called whenever an event occurs that affects the named services state. This could be triggered by an explicit invocation by the SARDeployer, a notification of a new class, or another service reaching its created state.

When a service's create method is called, all services on which the service depends have also had their create method invoked. This gives an MBean an opportunity to check that required MBeans or resources exist. A service cannot utilize other MBean services at this point, as most JBoss MBean services do not become fully functional until they have been started via their start method. Because of this, service implementations often do not implement create in favor of just the start method because that is the first point at which the service can be fully functional.

2.4.2.3.2. The start(ObjectName) method

The start(ObjectName) method is called whenever an event occurs that affects the named services state. This could be triggered by an explicit invocation by the SARDeployer, a notification of a new class, or another service reaching its started state.

When a service's `start` method is called, all services on which the service depends have also had their `start` method invoked. Receipt of a `start` method invocation signals a service to become fully operational since all services upon which the service depends have been created and started.

2.4.2.3.3. The `stop(ObjectName)` method

The `stop(ObjectName)` method is called whenever an event occurs that affects the named services state. This could be triggered by an explicit invocation by the `SARDeployer`, notification of a class removal, or a service on which other services depend reaching its stopped state.

2.4.2.3.4. The `destroy(ObjectName)` method

The `destroy(ObjectName)` method is called whenever an event occurs that affects the named services state. This could be triggered by an explicit invocation by the `SARDeployer`, notification of a class removal, or a service on which other services depend reaching its destroyed state.

Service implementations often do not implement `destroy` in favor of simply implementing the `stop` method, or neither `stop` nor `destroy` if the service has no state or resources that need cleanup.

2.4.2.4. Specifying Service Dependencies

To specify that an MBean service depends on other MBean services you need to declare the dependencies in the `mbean` element of the service descriptor. This is done using the `depends` and `depends-list` elements. One difference between the two elements relates to the `optional-attribute-name` attribute usage. If you track the `ObjectNames` of dependencies using single valued attributes you should use the `depends` element. If you track the `ObjectNames` of dependencies using `java.util.List` compatible attributes you would use the `depends-list` element. If you only want to specify a dependency and don't care to have the associated service `ObjectName` bound to an attribute of your MBean then use whatever element is easiest. The following listing shows example service descriptor fragments that illustrate the usage of the dependency related elements.

```
<mbean code="org.jboss.mq.server.jmx.Topic"
  name="jms.topic:service=Topic,name=testTopic">
  <!-- Declare a dependency on the "jboss.mq:service=DestinationManager" and
    bind this name to the DestinationManager attribute -->
  <depends optional-attribute-name="DestinationManager">
    jboss.mq:service=DestinationManager
  </depends>

  <!-- Declare a dependency on the "jboss.mq:service=SecurityManager" and
    bind this name to the SecurityManager attribute -->
  <depends optional-attribute-name="SecurityManager">
    jboss.mq:service=SecurityManager
  </depends>

  <!-- ... -->

  <!-- Declare a dependency on the
    "jboss.mq:service=CacheManager" without
    any binding of the name to an attribute-->
  <depends>jboss.mq:service=CacheManager</depends>
</mbean>

<mbean code="org.jboss.mq.server.jmx.TopicMgr"
  name="jboss.mq.destination:service=TopicMgr">
  <!-- Declare a dependency on the given topic destination mbeans and
    bind these names to the Topics attribute -->
  <depends-list optional-attribute-name="Topics">
    <depends-list-element>jms.topic:service=Topic,name=A</depends-list-element>
    <depends-list-element>jms.topic:service=Topic,name=B</depends-list-element>
    <depends-list-element>jms.topic:service=Topic,name=C</depends-list-element>
  </depends-list>
```

```
</mbean>
```

Another difference between the `depends` and `depends-list` elements is that the value of the `depends` element may be a complete MBean service configuration rather than just the `ObjectName` of the service. Example 2.15 shows an example from the `hsqldb-service.xml` descriptor. In this listing the `org.jboss.resource.connectionmanager.RARDeployment` service configuration is defined using a nested `mbean` element as the `depends` element value. This indicates that the `org.jboss.resource.connectionmanager.LocalTxConnectionManager` MBean depends on this service. The `org.jboss.jca:service=LocalTxDS,name=hsqldbDS` `ObjectName` will be bound to the `ManagedConnectionFactoryName` attribute of the `LocalTxConnectionManager` class.

Example 2.15. An example of using the `depends` element to specify the complete configuration of a depended on service.

```
<mbean code="org.jboss.resource.connectionmanager.LocalTxConnectionManager"
  name="jboss.jca:service=LocalTxCM,name=hsqldbDS">
  <depends optional-attribute-name="ManagedConnectionFactoryName">
    <!--embedded mbean-->
    <mbean code="org.jboss.resource.connectionmanager.RARDeployment"
      name="jboss.jca:service=LocalTxDS,name=hsqldbDS">
      <attribute name="JndiName">DefaultDS</attribute>
      <attribute name="ManagedConnectionFactoryProperties">
        <properties>
          <config-property name="ConnectionURL"
            type="java.lang.String">
            jdbc:hsqldb:hsqldb://localhost:1476
          </config-property>
          <config-property name="DriverClass" type="java.lang.String">
            org.hsqldb.jdbcDriver
          </config-property>
          <config-property name="UserName" type="java.lang.String">
            sa
          </config-property>
          <config-property name="Password" type="java.lang.String"/>
        </properties>
      </attribute>
    <!-- ... -->
    </mbean>
  </depends>
  <!-- ... -->
</mbean>
```

2.4.2.5. Identifying Unsatisfied Dependencies

The `ServiceController` MBean supports two operations that help with debugging what MBeans are not running due to unsatisfied dependencies. The first operation is `listIncompletelyDeployed`. This returns a `java.util.List` of `org.jboss.system.ServiceContext` objects for the MBean services that are not in the `RUNNING` state.

The second operation is `listWaitingMBeans`. This operation returns a `java.util.List` of the JMX `ObjectNames` of MBean services that cannot be deployed because the class specified by the `code` attribute is not available.

2.4.2.6. Hot Deployment of Components, the `URLDeploymentScanner`

The `URLDeploymentScanner` MBean service provides the JBoss hot deployment capability. This service watches one or more URLs for deployable archives and deploys the archives as they appear or change. It also

undeploys previously deployed applications if the archive from which the application was deployed is removed. The configurable attributes include:

- **URLs:** A comma separated list of URL strings for the locations that should be watched for changes. Strings that do not correspond to valid URLs are treated as file paths. Relative file paths are resolved against the server home URL, for example, `JBOSS_DIST/server/default` for the default config file set. If a URL represents a file then the file is deployed and watched for subsequent updates or removal. If a URL ends in `/` to represent a directory, then the contents of the directory are treated as a collection of deployables and scanned for content that are to be watched for updates or removal. The requirement that a URL end in a `/` to identify a directory follows the RFC2518 convention and allows discrimination between collections and directories that are simply unpacked archives.

The default value for the URLs attribute is `deploy/` which means that any SARs, EARs, JARs, WARs, RARs, etc. dropped into the `server/<name>/deploy` directory will be automatically deployed and watched for updates.

Example URLs include:

- `"deploy/"` scans `${jboss.server.url}/deploy/`, which is local or remote depending on the URL used to boot the server
- `"${jboss.server.home.dir}/deploy/"` scans `${jboss.server.home.dir}/deploy`, which is always local
- `"file:/var/opt/myapp.ear"` deploy myapp.ear from a local location
- `"file:/var/opt/apps/"` scans the specified directory
- `"http://www.test.com/netboot/myapp.ear"` deploys myapp.ear from a remote location
- `"http://www.test.com/netboot/apps/"` scans the specified remote location using WebDAV. This will only work if the remote http server supports the WebDAV PROPFIND command.
- **ScanPeriod:** The time in milliseconds between runs of the scanner thread. The default is 5000 (5 seconds).
- **URLComparator:** The class name of a `java.util.Comparator` implementation used to specify a deployment ordering for deployments found in a scanned directory. The implementation must be able to compare two `java.net.URL` objects passed to its `compare` method. The default setting is the `org.jboss.deployment.DeploymentSorter` class which orders based on the deployment URL suffix. The ordering of suffixes is: "sar", "service.xml", "rar", "jar", "war", "wsr", "ear", "zip".

An alternate implementation is the `org.jboss.deployment.scanner.PrefixDeploymentSorter` class. This orders the URLs based on numeric prefixes. The prefix digits are converted to an int (ignoring leading zeroes), smaller prefixes are ordered ahead of larger numbers. Deployments that do not start with any digits will be deployed after all numbered deployments. Deployments with the same prefix value are further sorted by the `DeploymentSorter` logic.

- **Filter:** The class name of a `java.io.FileFilter` implementation that is used to filter the contents of scanned directories. Any file not accepted by this filter will not be deployed. The default is `org.jboss.deployment.scanner.DeploymentFilter` which is an implementation that rejects the following patterns:

`"#*", "%*", ".*", ".*_", "$*", "*#", "*$", "*%", ".*.BAK", ".*.old", ".*.orig", ".*.rej", ".*.bak", ".*.v", ".*~", ".make.state", ".nse_depinfo", "CVS", "CVS.admin", "RCS", "RCSLOG", "SCCS", "TAGS", "core", "tags"`

- **RecursiveSearch:** This property indicates whether or not deploy subdirectories are seen to be holding de-

ployable content. If this is false, deploy subdirectories that do not contain ``.'` in their name are seen to be unpacked jars with nested subdeployments. If true, then deploy subdirectories are just groupings of deployable content. The difference between the two views shows is related to the depth first deployment model JBoss supports. The false setting which treats directories as unpacked jars with nested content triggers the deployment of the nested content as soon as the jar directory is deployed. The true setting simply ignores the directory and adds its content to the list of deployables and calculates the order based on the previous filter logic. The default is true. However, note that the jboss-3.2.1 release shipped with a default configuration with this set to false.

- **Deployer:** The JMX ObjectName string of the MBean that implements the `org.jboss.deployment.Deployer` interface operations. The default setting is to use the `MainDeployer` created by the bootstrap startup process.

2.4.3. Writing JBoss MBean Services

Writing a custom MBean service that integrates into the JBoss server requires the use of the `org.jboss.system.Service` interface pattern if the custom service is dependent on other services. When a custom MBean depends on other MBean services you cannot perform any service dependent initialization in any of the `javax.management.MBeanRegistration` interface methods since JMX has no dependency notion. Instead, you must manage dependency state using the `Service` interface `create` and/or `start` methods. You can do this using any one of the following approaches:

- Add any of the `Service` methods that you want called on your MBean to your MBean interface. This allows your MBean implementation to avoid dependencies on JBoss specific interfaces.
- Have your MBean interface extend the `org.jboss.system.Service` interface.
- Have your MBean interface extend the `org.jboss.system.ServiceMBean` interface. This is a subinterface of `org.jboss.system.Service` that adds `String getName()`, `int getState()`, and `String getStateString()` methods.

Which approach you choose depends on if you want to be associated with JBoss specific code. If you don't, then you would use the first approach. If you don't care about dependencies on JBoss classes, the simplest approach is to have your MBean interface extend from `org.jboss.system.ServiceMBean` and your MBean implementation class extend from the abstract `org.jboss.system.ServiceMBeanSupport` class. This class implements the `org.jboss.system.ServiceMBean` interface. `ServiceMBeanSupport` provides implementations of the `create`, `start`, `stop`, and `destroy` methods that integrate logging and JBoss service state management tracking. Each method delegates any subclass specific work to `createService`, `startService`, `stopService`, and `destroyService` methods respectively. When subclassing `ServiceMBeanSupport`, you would override one or more of the `createService`, `startService`, `stopService`, and `destroyService` methods as required.

2.4.3.1. A Standard MBean Example

This section develops a simple MBean that binds a `HashMap` into the JBoss JNDI namespace at a location determined by its `JndiName` attribute to demonstrate what is required to create a custom MBean. Because the MBean uses JNDI, it depends on the JBoss naming service MBean and must use the JBoss MBean service pattern to be notified when the naming service is available.

The MBean you develop is called `JNDIMap`. Version one of the `JNDIMapMBean` interface and `JNDIMap` implementation class, which is based on the service interface method pattern, is given in Example 2.16. This version of the interface makes use of the first approach in that it incorporates the `Service` interface methods needed to start up correctly, but does not do so by using a JBoss-specific interface. The interface includes the `Service.start` method, which will be informed when all required services have been started, and the `stop` method, which will clean up the service.

Example 2.16. JNDIMapMBean interface and implementation based on the service interface method pattern

```

package org.jboss.chap2.ex1;

// The JNDIMap MBean interface
import javax.naming.NamingException;

public interface JNDIMapMBean
{
    public String getJndiName();
    public void setJndiName(String jndiName) throws NamingException;
    public void start() throws Exception;
    public void stop() throws Exception;
}

```

```

package org.jboss.chap2.ex1;

// The JNDIMap MBean implementation
import java.util.HashMap;
import javax.naming.InitialContext;
import javax.naming.Name;
import javax.naming.NamingException;
import org.jboss.naming.NonSerializableFactory;

public class JNDIMap implements JNDIMapMBean
{
    private String jndiName;
    private HashMap contextMap = new HashMap();
    private boolean started;

    public String getJndiName()
    {
        return jndiName;
    }
    public void setJndiName(String jndiName) throws NamingException
    {
        String oldName = this.jndiName;
        this.jndiName = jndiName;
        if (started) {
            unbind(oldName);
            try {
                rebind();
            } catch (Exception e) {
                NamingException ne = new NamingException("Failed to update jndiName");
                ne.setRootCause(e);
                throw ne;
            }
        }
    }

    public void start() throws Exception
    {
        started = true;
        rebind();
    }

    public void stop()
    {
        started = false;
        unbind(jndiName);
    }

    private void rebind() throws NamingException
    {
        InitialContext rootCtx = new InitialContext();
        Name fullName = rootCtx.getNameParser("").parse(jndiName);
        System.out.println("fullName="+fullName);
    }
}

```

```

        NonSerializableFactory.rebind(fullName, contextMap, true);
    }

    private void unbind(String jndiName)
    {
        try {
            InitialContext rootCtx = new InitialContext();
            rootCtx.unbind(jndiName);
            NonSerializableFactory.unbind(jndiName);
        } catch (NamingException e) {
            e.printStackTrace();
        }
    }
}

```

```

package org.jboss.chap2.ex1;

// The JNDIMap MBean interface
import javax.naming.NamingException;

public interface JNDIMapMBean
{
    public String getJndiName();
    public void setJndiName(String jndiName) throws NamingException;
    public void start() throws Exception;
    public void stop() throws Exception;
}

```

```

package org.jboss.chap2.ex1;
// The JNDIMap MBean implementation
import java.util.HashMap;
import javax.naming.InitialContext;
import javax.naming.Name;
import javax.naming.NamingException;
import org.jboss.naming.NonSerializableFactory;

public class JNDIMap implements JNDIMapMBean
{
    private String jndiName;
    private HashMap contextMap = new HashMap();
    private boolean started;

    public String getJndiName()
    {
        return jndiName;
    }

    public void setJndiName(String jndiName) throws NamingException
    {
        String oldName = this.jndiName;
        this.jndiName = jndiName;
        if (started) {
            unbind(oldName);
            try {
                rebind();
            } catch (Exception e) {
                NamingException ne = new NamingException("Failed to update jndiName");
                ne.setRootCause(e);
                throw ne;
            }
        }
    }

    public void start() throws Exception
    {
        started = true;
        rebind();
    }
}

```



```

public void stop()
{
    started = false;
    unbind(jndiName);
}

private void rebind() throws NamingException
{
    InitialContext rootCtx = new InitialContext();
    Name fullName = rootCtx.getNameParser("").parse(jndiName);
    System.out.println("fullName="+fullName);
    NonSerializableFactory.rebind(fullName, contextMap, true);
}

private void unbind(String jndiName)
{
    try {
        InitialContext rootCtx = new InitialContext();
        rootCtx.unbind(jndiName);
        NonSerializableFactory.unbind(jndiName);
    } catch (NamingException e) {
        e.printStackTrace();
    }
}
}

```

Version two of the JNDIMapMBean interface and JNDIMap implementation class, which is based on the ServiceMBean interface and ServiceMBeanSupport class, is given in Example 2.16. In this version, the implementation class extends the ServiceMBeanSupport class and overrides the startService method and the stopService method. JNDIMapMBean also implements the abstract getName to return a descriptive name for the MBean. The JNDIMapMBean interface extends the org.jboss.system.ServiceMBean interface and only declares the setter and getter methods for the JndiName attribute because it inherits the Service life cycle methods from ServiceMBean. This is the third approach mentioned at the start of the Section 2.4.2. The implementation differences between Example 2.16 and Example 2.17 are highlighted in bold in Example 2.17.

Example 2.17. JNDIMap MBean interface and implementation based on the ServiceMBean interface and ServiceMBeanSupport class

```

package org.jboss.chap2.ex2;

// The JNDIMap MBean interface
import javax.naming.NamingException;

public interface JNDIMapMBean extends org.jboss.system.ServiceMBean
{
    public String getJndiName();
    public void setJndiName(String jndiName) throws NamingException;
}

package org.jboss.chap2.ex2;
// The JNDIMap MBean implementation
import java.util.HashMap;
import javax.naming.InitialContext;
import javax.naming.Name;
import javax.naming.NamingException;
import org.jboss.naming.NonSerializableFactory;

public class JNDIMap extends org.jboss.system.ServiceMBeanSupport
    implements JNDIMapMBean
{
    private String jndiName;
    private HashMap contextMap = new HashMap();

    public String getJndiName()

```

```

{
    return jndiName;
}

public void setJndiName(String jndiName)
    throws NamingException
{
    String oldName = this.jndiName;
    this.jndiName = jndiName;
    if (super.getState() == STARTED) {
        unbind(oldName);
        try {
            rebind();
        } catch (Exception e) {
            NamingException ne = new NamingException("Failed to update jndiName");
            ne.setRootCause(e);
            throw ne;
        }
    }
}

public void startService() throws Exception
{
    rebind();
}

public void stopService()
{
    unbind(jndiName);
}

private void rebind() throws NamingException
{
    InitialContext rootCtx = new InitialContext();
    Name fullName = rootCtx.getNameParser("").parse(jndiName);
    log.info("fullName="+fullName);
    NonSerializableFactory.rebind(fullName, contextMap, true);
}

private void unbind(String jndiName)
{
    try {
        InitialContext rootCtx = new InitialContext();
        rootCtx.unbind(jndiName);
        NonSerializableFactory.unbind(jndiName);
    } catch (NamingException e) {
        log.error("Failed to unbind map", e);
    }
}
}

```

The source code for these MBeans along with the service descriptors is located in the `examples/src/main/org/jboss/chap2/{ex1,ex2}` directories.

The example 1 service descriptor is shown below along with a sample client usage code fragment. The JNDIMap MBean binds a HashMap object under the `inmemory/maps/MapTest` JNDI name and the client code fragment demonstrates retrieving the HashMap object from the `inmemory/maps/MapTest` location.

```

<!-- The SAR META-INF/jboss-service.xml descriptor -->
<server>
    <mbean code="org.jboss.chap2.ex1.JNDIMap"
          name="chap2.ex1:service=JNDIMap">
        <attribute name="JndiName">inmemory/maps/MapTest</attribute>
        <depends>jboss:service=Naming</depends>
    </mbean>
</server>

```

```
// Sample lookup code
InitialContext ctx = new InitialContext();
HashMap map = (HashMap) ctx.lookup("inmemory/maps/MapTest");
```

2.4.3.2. XMBean Examples

In this section we will develop a variation of the JNDIMap MBean introduced in the preceding section that exposes its management metadata using the JBoss XMBean framework. Our core managed component will be exactly the same core code from the JNDIMap class, but this will not implement any specific management related interface. We will illustrate the following capabilities not possible with a Standard MBean:

- The ability to add rich descriptions to attribute and operations
- The ability to expose notification information
- The ability to add persistence of attributes
- The ability to add custom interceptors for security and remote access through a typed interface

2.4.3.2.1. Version 1, The Annotated JNDIMap XMBean

Let's start with a simple XMBean variation of the standard MBean version of the JNDIMap that adds the descriptive information about the attributes and operations and their arguments. The following listing shows the jboss-service.xml descriptor and the jndimap-xmbean1.xml XMBean descriptor. The source can be found in the src/main/org/jboss/chap2/xmbean directory of the book examples.

```
<?xml version='1.0' encoding='UTF-8' ?>
<!DOCTYPE server PUBLIC
    "-//JBoss//DTD MBean Service 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jboss-service_3_2.dtd">

<server>
  <mbean code="org.jboss.chap2.xmbean.JNDIMap"
    name="chap2.xmbean:service=JNDIMap"
    xmbean-dd="META-INF/jndimap-xmbean.xml">
    <attribute name="JndiName">inmemory/maps/MapTest</attribute>
    <depends>jboss:service=Naming</depends>
  </mbean>
</server>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE mbean PUBLIC
    "-//JBoss//DTD JBOSS XMBean 1.0//EN"
    "http://www.jboss.org/j2ee/dtd/jboss_xmbean_1_0.dtd">

<mbean>
  <description>The JNDIMap XMBean Example Version 1</description>
  <descriptors>
    <persistence persistPolicy="Never" persistPeriod="10"
      persistLocation="data/JNDIMap.data" persistName="JNDIMap"/>
    <currencyTimeLimit value="10"/>
    <state-action-on-update value="keep-running"/>
  </descriptors>
  <class>org.jboss.test.jmx.xmbean.JNDIMap</class>
  <constructor>
    <description>The default constructor</description>
    <name>JNDIMap</name>
  </constructor>
  <!-- Attributes -->
  <attribute access="read-write" getMethod="getJndiName" setMethod="setJndiName">
    <description>
      The location in JNDI where the Map we manage will be bound
    </description>
    <name>JndiName</name>
```

```

    <type>java.lang.String</type>
    <descriptors>
        <default value="inmemory/maps/MapTest" />
    </descriptors>
</attribute>
<attribute access="read-write" getMethod="getInitialValues"
    setMethod="setInitialValues">
    <description>The array of initial values that will be placed into the
        map associated with the service. The array is a collection of
        key,value pairs with elements[0,2,4,...2n] being the keys and
        elements [1,3,5,...,2n+1] the associated values. The
        "[Ljava.lang.String;" type signature is the VM representation of the
        java.lang.String[] type. </description>
    <name>InitialValues</name>
    <type>[Ljava.lang.String;</type>
    <descriptors>
        <default value="key0,value0" />
    </descriptors>
</attribute>
<!-- Operations -->
<operation>
    <description>The start lifecycle operation</description>
    <name>start</name>
</operation>
<operation>
    <description>The stop lifecycle operation</description>
    <name>stop</name>
</operation>
<operation impact="ACTION">
    <description>Put a value into the map</description>
    <name>put</name>
    <parameter>
        <description>The key the value will be store under</description>
        <name>key</name>
        <type>java.lang.Object</type>
    </parameter>
    <parameter>
        <description>The value to place into the map</description>
        <name>value</name>
        <type>java.lang.Object</type>
    </parameter>
</operation>
<operation impact="INFO">
    <description>Get a value from the map</description>
    <name>get</name>
    <parameter>
        <description>The key to lookup in the map</description>
        <name>get</name>
        <type>java.lang.Object</type>
    </parameter>
    <return-type>java.lang.Object</return-type>
</operation>
<!-- Notifications -->
<notification>
    <description>The notification sent whenever a value is get into the map
        managed by the service</description>
    <name>javax.management.Notification</name>
    <notification-type>org.jboss.chap2.xmlbean.JNDIMap.get</notification-type>
</notification>
<notification>
    <description>The notification sent whenever a value is put into the map
        managed by the service</description>
    <name>javax.management.Notification</name>
    <notification-type>org.jboss.chap2.xmlbean.JNDIMap.put</notification-type>
</notification>
</mbean>

```

As noted previously, the 3.2.2 release replaced the binding of the RMIAdaptor interface with the invoker adaptor service and this service does not yet support remoting of JMX notifications. Therefore, we need to create

a config that uses the `RMIAdaptorService`. There is a config target that sets up a rmi-adaptor configuration with the `jmx-rmi-adaptor.sar` installed. Build this setup using:

```
[nr@toki]$ ant -Dchap=chap2 config
...

[echo] Preparing rmi-adaptor configuration fileset
[copy] Copying 214 files to /tmp/jboss-3.2.6/server/rmi-adaptor
[copy] Copied 2 empty directories to /tmp/jboss-3.2.6/server/rmi-adaptor
[copy] Copying 2 files to /tmp/jboss-3.2.6/server/rmi-adaptor/deploy/jmx-rmi-adaptor.sar
[delete] Deleting directory /tmp/jboss-3.2.6/server/rmi-adaptor/deploy/jmx-invoker-adaptor-server.sar
[delete] Deleting directory /tmp/jboss-3.2.6/server/rmi-adaptor/deploy/management
```

Now, run the rmi-adaptor configuration, and then build, deploy and test the `XMBean` as follows:

```
[nr@toki examples]$ ant -Dchap=chap2 -Dex=xmbean1 -Djboss.deploy.conf=rmi-adaptor run-example
...
run-examplexmbean1:
[copy] Copying 1 file to /tmp/jboss-3.2.6/server/rmi-adaptor/deploy
[java] JNDIMap Class: org.jboss.mx.modelmbean.XMBean
[java] JNDIMap Operations:
[java] + void start()
[java] + void stop()
[java] + void put(java.lang.Object chap2.xmbean:service=JNDIMap,java.lang.Object chap2.xmbean:service=JNDIMap)
[java] + java.lang.Object get(java.lang.Object chap2.xmbean:service=JNDIMap)
[java] + java.lang.String getJndiName()
[java] + void setJndiName(java.lang.String chap2.xmbean:service=JNDIMap)
[java] + [Ljava.lang.String; getInitialValues()
[java] + void setInitialValues([Ljava.lang.String; chap2.xmbean:service=JNDIMap)
[java] handleNotification, event: null
[java] key=key0, value=value0
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:service=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=3,timeStamp=1098631527823,message=null,userData=null]
[java] JNDIMap.put(key1, value1) successful
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:service=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.get,sequenceNumber=4,timeStamp=1098631527940,message=null,userData=null]
[java] JNDIMap.get(key0): null
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:service=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.get,sequenceNumber=5,timeStamp=1098631527985,message=null,userData=null]
[java] JNDIMap.get(key1): value1
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:service=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=6,timeStamp=1098631527999,message=null,userData=null]
```

The functionality is largely the same as the Standard MBean with the notable exception of the JMX notifications. A Standard MBean has no way of declaring that it will emit notifications. An XMBean may declare the notifications it emits using notification elements as is shown in the version 1 descriptor. We see the notifications from the get and put operations on the test client console output. Note that there is also an `jmx.attribute.change` notification emitted when the `InitialValues` attribute was changed. This is a standard feature of `ModelMBeans` owing to the fact that the `ModelMBean` interface extends the `ModelMBeanNotificationBroadcaster` which supports `AttributeChangeNotificationListeners`.

The other major difference between the Standard and XMBean versions of `JNDIMap` is the descriptive metadata. Look at the `chap2.xmbean:service=JNDIMap` in the JMX Console, and you will see the attributes section as shown in Figure 2.19.

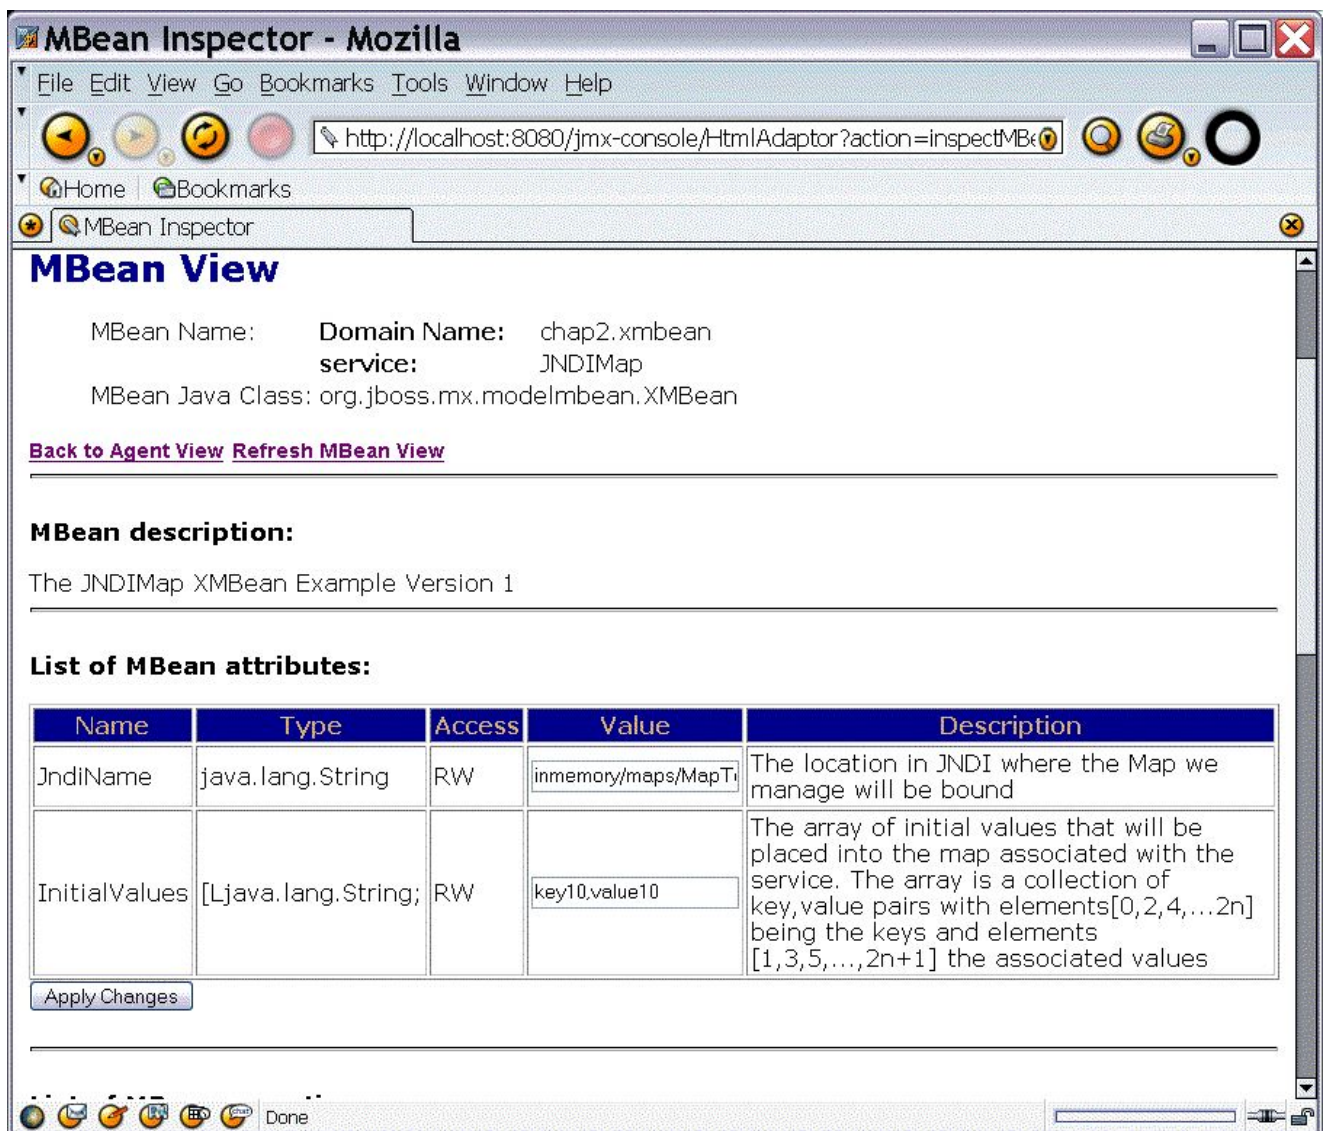


Figure 2.19. The Version 1 JNDIMapXMBean jmx-console view

Notice that the JMX Console now displays the full attribute description as specified in the XMBean descriptor rather than `MBean Attribute` text seen in standard MBean implementations. Scroll down to the operations and you will also see that these now also have nice descriptions of their function and parameters.

2.4.3.2.2. Version 2, Adding Persistence to the JNDIMap XMBean

In version 2 of the XMBean we add support for persistence of the XMBean attributes. The updated XMBean deployment descriptor is given below. The changes with respect to the version 1 descriptor of `Lxxx1` are shown in bold.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE mbean PUBLIC
    "-//JBoss//DTD JBOSS XMBean 1.0//EN"
    "http://www.jboss.org/j2ee/dtd/jboss_xmbean_1_0.dtd">
<mbean>
  <description>The JNDIMap XMBean Example Version 2</description>
  <descriptors>
    <persistence persistPolicy="OnUpdate" persistPeriod="10"
      persistLocation="{jboss.server.data.dir}" persistName="JNDIMap.ser"/>
    <currencyTimeLimit value="10"/>
    <state-action-on-update value="keep-running"/>
    <persistence-manager value="org.jboss.mx.persistence.ObjectStreamPersistenceManager"/>
  </descriptors>
</mbean>
```

```

</descriptors>    <class>org.jboss.test.jmx.xmlbean.JNDIMap</class>
<constructor>
    <description>The default constructor</description>
    <name>JNDIMap</name>
</constructor>
<!-- Attributes -->
<attribute access="read-write" getMethod="getJndiName" setMethod="setJndiName">
    <description>
        The location in JNDI where the Map we manage will be bound
    </description>
    <name>JndiName</name>
    <type>java.lang.String</type>
    <descriptors>
        <default value="inmemory/maps/MapTest"/>
    </descriptors>
</attribute>
<attribute access="read-write" getMethod="getInitialValues"
    setMethod="setInitialValues">
    <description>The array of initial values that will be placed into the
        map associated with the service. The array is a collection of
        key,value pairs with elements[0,2,4,...2n] being the keys and
        elements [1,3,5,...,2n+1] the associated values</description>
    <name>InitialValues</name>
    <type>[Ljava.lang.String;</type>
    <descriptors>
        <default value="key0,value0"/>
    </descriptors>
</attribute>
<!-- Operations -->
<operation>
    <description>The start lifecycle operation</description>
    <name>start</name>
</operation>
<operation>
    <description>The stop lifecycle operation</description>
    <name>stop</name>
</operation>
<operation impact="ACTION">
    <description>Put a value into the nap</description>
    <name>put</name>
    <parameter>
        <description>The key the value will be store under</description>
        <name>key</name>
        <type>java.lang.Object</type>
    </parameter>
    <parameter>
        <description>The value to place into the map</description>
        <name>value</name>
        <type>java.lang.Object</type>
    </parameter>
</operation>
<operation impact="INFO">
    <description>Get a value from the map</description>
    <name>get</name>
    <parameter>
        <description>The key to lookup in the map</description>
        <name>get</name>
        <type>java.lang.Object</type>
    </parameter>
    <return-type>java.lang.Object</return-type>
</operation>
<!-- Notifications -->
<notification>
    <description>The notification sent whenever a value is get into the map
        managed by the service</description>
    <name>javax.management.Notification</name>
    <notification-type>org.jboss.chap2.xmlbean.JNDIMap.get</notification-type>
</notification>
<notification>
    <description>The notification sent whenever a value is put into the map
        managed by the service</description>

```

```

<name>javax.management.Notification</name>
<notification-type>org.jboss.chap2.xmbean.JNDIMap.put</notification-type>
</notification>
</mbean>

```

Build, deploy and test the version 2 XMBean as follows:

```

[examples]$ ant -Dchap=chap2 -Dex=xmbean2 -Djboss.deploy.conf=rmi-adaptor run-example
...
run-examplexmbean2:
  [java] JNDIMap Class: org.jboss.mx.modelmbean.XMBean
  [java] JNDIMap Operations:
  [java]   + void start()
  [java]   + void stop()
  [java]   + void put(java.lang.Object chap2.xmbean:service=JNDIMap,java.lang.Object cha
p2.xmbean:service=JNDIMap)
  [java]   + java.lang.Object get(java.lang.Object chap2.xmbean:service=JNDIMap)
  [java]   + java.lang.String getJndiName()
  [java]   + void setJndiName(java.lang.String chap2.xmbean:service=JNDIMap)
  [java]   + [Ljava.lang.String; getInitialValues()
  [java]   + void setInitialValues([Ljava.lang.String; chap2.xmbean:service=JNDIMap)
  [java] handleNotification, event: null
  [java] key=key10, value=value10
  [java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=7,timeStamp=10986326
93716,message=null,userData=null]
  [java] JNDIMap.put(key1, value1) successful
  [java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.get,sequenceNumber=8,timeStamp=10986326
93857,message=null,userData=null]
  [java] JNDIMap.get(key0): null
  [java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.get,sequenceNumber=9,timeStamp=10986326
93896,message=null,userData=null]
  [java] JNDIMap.get(key1): value1
  [java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=10,timeStamp=1098632
693925,message=null,userData=null]

```

There is nothing manifestly different about this version of the XMBean at this point because we have done nothing to test that changes to attribute value are actually persisted. Perform this test by running example xmbean2a several times:

```

[nr@toki examples] ant -Dchap=chap2 -Dex=xmbean2a -Djboss.deploy.conf=rmi-adaptor \
run-example
...
  [java] InitialValues.length=2
  [java] key=key10, value=value10

```

```

[nr@toki examples] ant -Dchap=chap2 -Dex=xmbean2a -Djboss.deploy.conf=rmi-adaptor \
run-example
...
  [java] InitialValues.length=4
  [java] key=key10, value=value10
  [java] key=key2, value=value2

```

```

[nr@toki examples] ant -Dchap=chap2 -Dex=xmbean2a -Djboss.deploy.conf=rmi-adaptor \
run-example
...
  [java] InitialValues.length=6
  [java] key=key10, value=value10
  [java] key=key2, value=value2
  [java] key=key3, value=value3

```


The `org.jboss.chap2.xmbean.TestXMBeanRestart` used in this example obtains the current `InitialValues` attribute setting, and then adds another key/value pair to it. The client code is shown below.

```
package org.jboss.chap2.xmbean;

import javax.management.Attribute;
import javax.management.ObjectName;
import javax.naming.InitialContext;

import org.jboss.jmx.adaptor.rmi.RMIAdaptor;

/**
 * A client that demonstrates the persistence of the xmbean
 * attributes. Every time it is run it looks up the InitialValues
 * attribute, prints it out and then adds a new key/value to the
 * list.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class TestXMBeanRestart
{
    /**
     * @param args the command line arguments
     */
    public static void main(String[] args) throws Exception
    {
        InitialContext ic = new InitialContext();
        RMIAdaptor server = (RMIAdaptor) ic.lookup("jmx/rmi/RMIAdaptor");

        // Get the InitialValues attribute
        ObjectName name = new ObjectName("chap2.xmbean:service=JNDIMap");
        String[] initialValues = (String[])
            server.getAttribute(name, "InitialValues");
        System.out.println("InitialValues.length="+initialValues.length);
        int length = initialValues.length;
        for (int n = 0; n < length; n += 2) {
            String key = initialValues[n];
            String value = initialValues[n+1];

            System.out.println("key="+key+", value="+value);
        }
        // Add a new key/value pair
        String[] newInitialValues = new String[length+2];
        System.arraycopy(initialValues, 0, newInitialValues,
            0, length);
        newInitialValues[length] = "key"+(length/2+1);
        newInitialValues[length+1] = "value"+(length/2+1);

        Attribute ivalues = new
            Attribute("InitialValues", newInitialValues);
        server.setAttribute(name, ivalues);
    }
}
```

At this point you may even shutdown the JBoss server, restart it and then rerun the initial example 2 to see if the changes are persisted across server restarts:

```
[examples]$ ant -Dchap=chap2 -Dex=xmbean2 -Djboss.deploy.conf=rmi-adaptor run-example
...

run-examplexmbean2:
[java] JNDIMap Class: org.jboss.mx.modelmbean.XMBean
[java] JNDIMap Operations:
[java] + void start()
[java] + void stop()
[java] + void put(java.lang.Object chap2.xmbean:service=JNDIMap,java.lang.Object cha
p2.xmbean:service=JNDIMap)
[java] + java.lang.Object get(java.lang.Object chap2.xmbean:service=JNDIMap)
```

```

[java] + java.lang.String getJndiName()
[java] + void setJndiName(java.lang.String chap2.xmbean:service=JNDIMap)
[java] + [Ljava.lang.String; getInitialValues()
[java] + void setInitialValues([Ljava.lang.String; chap2.xmbean:service=JNDIMap)
[java] handleNotification, event: null
[java] key=key10, value=value10
[java] key=key2, value=value2
[java] key=key3, value=value3
[java] key=key4, value=value4
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=3,timeStamp=10986336
64712,message=null,userData=null]
[java] JNDIMap.put(key1, value1) successful
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.get,sequenceNumber=4,timeStamp=10986336
64821,message=null,userData=null]
[java] JNDIMap.get(key0): null
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.get,sequenceNumber=5,timeStamp=10986336
64860,message=null,userData=null]
[java] JNDIMap.get(key1): value1
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=6,timeStamp=10986336
64877,message=null,userData=null]
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=7,timeStamp=10986336
64895,message=null,userData=null]
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=8,timeStamp=10986336
64899,message=null,userData=null]
[java] handleNotification, event: javax.management.Notification[source=chap2.xmbean:s
ervice=JNDIMap,type=org.jboss.chap2.xmbean.JNDIMap.put,sequenceNumber=9,timeStamp=10986336
65614,message=null,userData=null]

```

You see that the last InitialValues attribute setting is in fact visible.

2.4.3.2.3. Version 3, Adding Security and Remote Access to the JNDIMap XMBean

The last example version of the JNDIMap XMBean will demonstrate customization of the server interceptor stack as well as exposing a subset of the XMBean management interface via a typed proxy to a remote client using RMI/JRMP. On the server side we will add a simple security interceptor that only allows access to attributes or operations by a user specified in the interceptor configuration. We will also use another custom interceptor to implement the MBean detached invoker pattern described in Section 2.7. By implementing this pattern in an invoker rather than the XMBean, we demonstrate how to introduce a remote access aspect without having to modify the existing JNDIMap implementation.

We will use the JRMPProxyFactory service to expose the ClientInterface to remote clients.

```

public interface ClientInterface
{
    public String[] getInitialValues();
    public void setInitialValues(String[] keyValuePairs);
    public Object get(Object key);
    public void put(Object key, Object value);
}

```

Our test client will obtain the ClientInterface proxy from JNDI and interact with the XMBean through RMI style calls instead of the RMIAdaptor and MBean Server style used previously.

```

package org.jboss.chap2.xmbean;

import javax.naming.InitialContext;

```

```

import org.jboss.security.SecurityAssociation;
import org.jboss.security.SimplePrincipal;

/**
 * A client that accesses an XMBean through its RMI interface
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class TestXMBean3
{
    /**
     * @param args the command line arguments
     */
    public static void main(String[] args) throws Exception
    {
        InitialContext ic = new InitialContext();
        ClientInterface xmbean = (ClientInterface)
            ic.lookup("secure-xmbean/ClientInterface");

        // This call should fail because we have not set a security context
        try {
            String[] tmp = xmbean.getInitialValues();
            throw new IllegalStateException("Was able to call getInitialValues");
        } catch (Exception e) {
            System.out.println("Called to getInitialValues failed as expected: "
                + e.getMessage());
        }

        // Set a security context using the SecurityAssociation
        SecurityAssociation.setPrincipal(new SimplePrincipal("admin"));

        // Get the InitialValues attribute
        String[] initialValues = xmbean.getInitialValues();
        for(int n = 0; n < initialValues.length; n += 2) {
            String key = initialValues[n];
            String value = initialValues[n+1];

            System.out.println("key="+key+", value="+value);
        }

        // Invoke the put(Object, Object) op
        xmbean.put("key1", "value1");
        System.out.println("JNDIMap.put(key1, value1) successful");
        Object result0 = xmbean.get("key0");
        System.out.println("JNDIMap.get(key0): "+result0);
        Object result1 = xmbean.get("key1");
        System.out.println("JNDIMap.get(key1): "+result1);

        // Change the InitialValues
        initialValues[0] += ".1";
        initialValues[1] += ".2";
        xmbean.setInitialValues(initialValues);

        initialValues = xmbean.getInitialValues();
        for(int n = 0; n < initialValues.length; n += 2) {
            String key = initialValues[n];
            String value = initialValues[n+1];

            System.out.println("key="+key+", value="+value);
        }
    }
}

```

The deployment descriptor is shown below:

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE mbean PUBLIC
    "-//JBoss//DTD JBOSS XMBean 1.0//EN"

```

```

"http://www.jboss.org/j2ee/dtd/jboss_xmbean_1_0.dtd"
[<!-- ATTLIST interceptor adminName CDATA #IMPLIED-->]
<mbean>
  <description>The JNDIMap XMBean Example Version 3</description>
  <descriptors>
    <interceptors>
      <interceptor code="org.jboss.chap2.xmbean.ServerSecurityInterceptor"
        adminName="admin"/>
      <interceptor code="org.jboss.chap2.xmbean.InvokerInterceptor"/>
      <interceptor code="org.jboss.mx.interceptor.PersistenceInterceptor2"/>
      <interceptor code="org.jboss.mx.interceptor.ModelMBeanInterceptor"/>
      <interceptor code="org.jboss.mx.interceptor.ObjectReferenceInterceptor"/>
    </interceptors>
    <persistence persistPolicy="Never"/>
    <currencyTimeLimit value="10"/>
    <state-action-on-update value="keep-running"/>
  </descriptors>
  <class>org.jboss.test.jmx.xmbean.JNDIMap</class>
  <constructor>
    <description>The default constructor</description>
    <name>JNDIMap</name>
  </constructor>
  <!-- Attributes -->
  <attribute access="read-write" getMethod="getJndiName" setMethod="setJndiName">
    <description>
      The location in JNDI where the Map we manage will be bound
    </description>
    <name>JndiName</name>
    <type>java.lang.String</type>
    <descriptors>
      <default value="inmemory/maps/MapTest"/>
    </descriptors>
  </attribute>
  <attribute access="read-write" getMethod="getInitialValues"
    setMethod="setInitialValues">
    <description>The array of initial values that will be placed into the
      map associated with the service. The array is a collection of
      key,value pairs with elements[0,2,4,...,2n] being the keys and
      elements [1,3,5,...,2n+1] the associated values</description>
    <name>InitialValues</name>
    <type>[Ljava.lang.String;</type>
    <descriptors>
      <default value="key0,value0"/>
    </descriptors>
  </attribute>
  <!-- Operations -->
  <operation>
    <description>The start lifecycle operation</description>
    <name>start</name>
  </operation>
  <operation>
    <description>The stop lifecycle operation</description>
    <name>stop</name>
  </operation>
  <operation impact="ACTION">
    <description>Put a value into the nap</description>
    <name>put</name>
    <parameter>
      <description>The key the value will be store under</description>
      <name>key</name>
      <type>java.lang.Object</type>
    </parameter>
    <parameter>
      <description>The value to place into the map</description>
      <name>value</name>
      <type>java.lang.Object</type>
    </parameter>
  </operation>
  <operation impact="INFO">
    <description>Get a value from the map</description>
    <name>get</name>

```

```

    <parameter>
      <description>The key to lookup in the map</description>
      <name>get</name>
      <type>java.lang.Object</type>
    </parameter>
    <return-type>java.lang.Object</return-type>
  </operation>
</mbean>

```

The addition over the previous versions of the `JNDIMap XMBean` is the interceptors element shown in bold in the listing. This defines the interceptor stack through which all MBean attribute access and operations pass. The first two interceptors, `org.jboss.chap2.xmbean.ServerSecurityInterceptor` and `org.jboss.chap2.xmbean.InvokerInterceptor` are the example custom interceptors. The remaining three interceptors are the standard ModelMBean interceptors. Because we have a persistence policy of `Never`, we could in fact remove the standard `org.jboss.mx.interceptor.PersistenceInterceptor2`. The JMX interceptors are an ordered chain of filters. The standard base class of an interceptor is shown below.

```

package org.jboss.mx.interceptor;

import javax.management.MBeanInfo;
import org.jboss.mx.server.MBeanInvoker;

/**
 * Base class for all interceptors.
 *
 * @see org.jboss.mx.interceptor.StandardMBeanInterceptor
 * @see org.jboss.mx.interceptor.LogInterceptor
 *
 * @author <a href="mailto:juha@jboss.org">Juha Lindfors</a>.
 * @version $Revision: 1.11 $
 */
public class AbstractInterceptor implements Interceptor
{
    // Attributes -----
    protected Interceptor next = null;
    protected String name = null;
    protected MBeanInfo info;
    protected MBeanInvoker invoker;

    // Constructors -----
    public AbstractInterceptor()
    {
        this(null);
    }
    public AbstractInterceptor(String name)
    {
        this.name = name;
    }
    public AbstractInterceptor(MBeanInfo info,
                              MBeanInvoker invoker)
    {
        this.name = getClass().getName();
        this.info = info;
        this.invoker = invoker;
    }

    // Public -----
    public Object invoke(Invocation invocation)
        throws InvocationException
    {
        return getNext().invoke(invocation);
    }

    public Interceptor getNext()
    {
        return next;
    }
}

```

```

    public Interceptor setNext(Interceptor interceptor)
    {
        this.next = interceptor;
        return interceptor;
    }
}

```

The custom interceptors for the version 3 XMBean example are the `ServerSecurityInterceptor` and the `InvokerInterceptor`. The `ServerSecurityInterceptor` intercepts invoke operations and validates that the invocation context include an admin principal.

```

package org.jboss.chap2.xmbean;

import java.security.Principal;

import org.jboss.logging.Logger;
import org.jboss.mx.interceptor.AbstractInterceptor;
import org.jboss.mx.interceptor.Invocation;
import org.jboss.mx.interceptor.InvocationException;
import org.jboss.security.SimplePrincipal;

/**
 * A simple security interceptor example that restricts access to a
 * single principal
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class ServerSecurityInterceptor extends AbstractInterceptor
{
    private static Logger log = Logger.getLogger(ServerSecurityInterceptor.class);
    private SimplePrincipal admin = new SimplePrincipal("admin");

    public String getAdminName()
    {
        return admin.getName();
    }
    public void setAdminName(String name)
    {
        admin = new SimplePrincipal(name);
    }

    public Object invoke(Invocation invocation)
        throws InvocationException
    {
        String opName = invocation.getName();

        // If this is not the invoke(Invocation) op just pass it along
        if (opName.equals("invoke") == false) {
            return getNext().invoke(invocation);
        }

        Object[] args = invocation.getArgs();
        org.jboss.invocation.Invocation invokeInfo =
            (org.jboss.invocation.Invocation) args[0];
        Principal caller = invokeInfo.getPrincipal();
        log.info("invoke, opName="+opName+", caller="+caller);

        // Only the admin caller is allowed access
        if (caller == null || caller.equals(admin) == false) {
            throw new InvocationException(new SecurityException("Caller=" +
                caller +
                " is not allowed access"));
        }
        return getNext().invoke(invocation);
    }
}

```

```

    }
}

```

The `InvokerInterceptor` implements the detached invoker pattern. This is discussed in detail in [Remote Access to Services, Detached Invokers](#).

```

package org.jboss.chap2.xmbean;

import java.lang.reflect.Method;
import java.util.HashMap;
import javax.management.Descriptor;
import javax.management.MBeanInfo;

import org.jboss.logging.Logger;
import org.jboss.mx.interceptor.AbstractInterceptor;
import org.jboss.mx.interceptor.Invocation;
import org.jboss.mx.interceptor.InvocationException;
import org.jboss.mx.server.MBeanInvoker;
import org.jboss.invocation.MarshalledInvocation;

/** An interceptor that handles the
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class InvokerInterceptor
    extends AbstractInterceptor
{
    private static Logger log = Logger.getLogger(InvokerInterceptor.class);
    private Class exposedInterface = ClientInterface.class;
    private HashMap methodMap = new HashMap();
    private HashMap invokeMap = new HashMap();

    public InvokerInterceptor(MBeanInfo info,
                             MBeanInvoker invoker)
    {
        super(info, invoker);
        try {
            Descriptor[] descriptors = invoker.getDescriptors();
            Object resource = invoker.getResource();
            Class[] getInitialValuesSig = {};
            Method getInitialValues =
                exposedInterface.getDeclaredMethod("getInitialValues",
                                                    getInitialValuesSig);
            Long hash = new Long(MarshalledInvocation.calculateHash(getInitialValues));
            InvocationInfo invokeInfo =
                new InvocationInfo("InitialValues",
                                   Invocation.ATTRIBUTE,
                                   Invocation.READ, getInitialValuesSig,
                                   descriptors, resource);
            methodMap.put(hash, getInitialValues);
            invokeMap.put(getInitialValues, invokeInfo);
            log.debug("getInitialValues hash:"+hash);

            Class[] setInitialValuesSig = {String[].class};
            Method setInitialValues =
                exposedInterface.getDeclaredMethod("setInitialValues",
                                                    setInitialValuesSig);
            hash = new Long(MarshalledInvocation.calculateHash(setInitialValues));
            invokeInfo = new InvocationInfo("InitialValues",
                                           Invocation.ATTRIBUTE,
                                           Invocation.WRITE,
                                           setInitialValuesSig,
                                           descriptors, resource);
            methodMap.put(hash, setInitialValues);
            invokeMap.put(setInitialValues, invokeInfo);
            log.debug("setInitialValues hash:"+hash);
        }
    }
}

```

```

        Class[] getSig = {Object.class};
        Method get = exposedInterface.getDeclaredMethod("get",
                                                         getSig);

        hash = new Long(MarshalledInvocation.calculateHash(get));
        invokeInfo = new InvocationInfo("get",
                                         Invocation.OPERATION,
                                         Invocation.READ, getSig,
                                         descriptors, resource);

        methodMap.put(hash, get);
        invokeMap.put(get, invokeInfo);
        log.debug("get hash:"+hash);

        Class[] putSig = {Object.class, Object.class};
        Method put = exposedInterface.getDeclaredMethod("put",
                                                         putSig);

        hash = new Long(MarshalledInvocation.calculateHash(put));
        invokeInfo = new InvocationInfo("put",
                                         Invocation.OPERATION,
                                         Invocation.WRITE, putSig,
                                         descriptors, resource);

        methodMap.put(hash, put);
        invokeMap.put(put, invokeInfo);
        log.debug("put hash:"+hash);
    } catch (Exception e) {
        log.error("Failed to init InvokerInterceptor", e);
    }
}

public Object invoke(Invocation invocation)
    throws InvocationException
{
    String opName = invocation.getName();
    Object[] args = invocation.getArgs();
    Object returnValue = null;
    if (opName.equals("invoke") == true) {
        org.jboss.invocation.Invocation invokeInfo =
            (org.jboss.invocation.Invocation) args[0];
        // Set the method hash to Method mapping
        if (invokeInfo instanceof MarshalledInvocation) {
            MarshalledInvocation mi = (MarshalledInvocation) invokeInfo;
            mi.setMethodMap(methodMap);
        }

        // Invoke the exposedInterface method via reflection if
        // this is an invoke
        Method method = invokeInfo.getMethod();
        Object[] methodArgs = invokeInfo.getArguments();
        InvocationInfo info = (InvocationInfo) invokeMap.get(method);
        Invocation methodInvocation = info.getInvocation(methodArgs);
        returnValue = getNext().invoke(methodInvocation);
    } else {
        returnValue = getNext().invoke(invocation);
    }
    return returnValue;
}

/**
 * A class that holds the ClientInterface method info needed to build
 * the JMX Invocation to pass down the interceptor stack.
 */
private class InvocationInfo
{
    private int type;
    private int impact;
    private String name;
    private String[] signature;
    private Descriptor[] descriptors;
    private Object resource;

    InvocationInfo(String name, int type, int impact,

```



```

        Class[] signature, Descriptor[] descriptors,
        Object resource)
    {
        this.name = name;
        this.type = type;
        this.impact = impact;
        this.descriptors = descriptors;
        this.resource = resource;
        this.signature = new String[signature.length];
        for(int s = 0; s < signature.length; s++) {
            this.signature[s] = signature[s].getName();
        }
    }

    Invocation getInvocation(Object[] args)
    {
        return new Invocation(name, type, impact, args, signature,
                               descriptors, resource);
    }
}

```

The deployment descriptor should include the interceptor stack.

```

<?xml version='1.0' encoding='UTF-8' ?>
<server>
  <mbean code="org.jboss.chap2.xmbean.JNDIMap"
    name="chap2.xmbean:service=JNDIMap,version=3"
    xmbean-dd="META-INF/jndimap-xmbean3.xml">
    <attribute name="JndiName">inmemory/maps/MapTest</attribute>
    <depends>jboss:service=Naming</depends>
  </mbean>
  <!-- The JRMP invoker proxy configuration for
        the naming service -->
  <mbean code="org.jboss.invocation.jrmp.server.JRMPProxyFactory"
    name="jboss.test:service=proxyFactory,type=jrmp,target=JNDIMap">
    <!-- Use the standard JRMPInvoker from
        conf/jboss-service.xml -->
    <attribute name="InvokerName">jboss:service=invoker,type=jrmp</attribute>
    <attribute name="TargetName">chap2.xmbean:service=JNDIMap,version=3</attribute>
    <attribute name="JndiName">secure-xmbean/ClientInterface</attribute>
    <attribute name="ExportedInterface">
      org.jboss.chap2.xmbean.ClientInterface
    </attribute>
    <attribute name="ClientInterceptors">
      <interceptors>
        <interceptor>org.jboss.proxy.ClientMethodInterceptor</interceptor>
        <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
        <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
      </interceptors>
    </attribute>
    <depends>jboss:service=invoker,type=jrmp</depends>
    <depends>chap2.xmbean:service=JNDIMap,version=3</depends>
  </mbean>
</server>

```

```

[nr@toki examples] ant -Dchap=chap2 -Dex=xmbean3 config
...
config:
  [echo] Preparing rmi-adaptor configuration fileset
  [copy] Copying 60 files to /tmp/jboss-3.2.6/server/rmi-adaptor
  [delete] Deleting directory /tmp/jboss-3.2.6/server/rmi-adaptor/deploy/jmx-invoker-adap
tor-server.sar
  [delete] Deleting directory /tmp/jboss-3.2.6/server/rmi-adaptor/deploy/management

```

```

[nr@toki examples]$ ant -Dchap=chap2 -Dex=xmbean3 run-example
...
run-examplexmbean3:

```

```
[java] Called to getInitialValues failed as expected: Caller=null is not allowed access
[java] key=key0, value=value0
[java] JNDIMap.put(key1, value1) successful
[java] JNDIMap.get(key0): null
[java] JNDIMap.get(key1): value1
[java] key=key0.1, value=value0.2
```

```
[nr@toki examples]$ ant -Dchap=chap2 -Dex=xmbean3 run-example
...
run-examplexmbean3:
[java] Called to getInitialValues failed as expected: Caller=null is not allowed access
[java] key=key0.1, value=value0.2
[java] JNDIMap.put(key1, value1) successful
[java] JNDIMap.get(key0): null
[java] JNDIMap.get(key1): value1
[java] key=key0.1.1, value=value0.2.2
```

2.4.4. Deployment Ordering and Dependencies

We have seen how to manage dependencies using the service descriptor depends and depends-list tags. The deployment ordering supported by the deployment scanners provides a coarse-grained dependency management in that there is an order to deployments. If dependencies are consistent with the deployment packages then this is a simpler mechanism than having to enumerate the explicit MBean-MBean dependencies. By writing your own filters you can change the coarse grained ordering performed by the deployment scanner.

When a component archive is deployed, its nested deployment units are processed in a depth first ordering. Structuring of components into an archive hierarchy is yet another way to manage deployment ordering.

Typically you will need to explicitly state your MBean dependencies because your packaging structure does not happen to resolve the dependencies. Let's consider an example component deployment that consists of an MBean that uses an EJB. Here is the structure of the example EAR.

```
output/chap2/chap2-ex3.ear
+- META-INF/MANIFEST.MF
+- META-INF/jboss-app.xml
+- chap2-ex3.jar (archive) [EJB jar]
| +- META-INF/MANIFEST.MF
| +- META-INF/ejb-jar.xml
| +- org/jboss/chap2/ex3/EchoBean.class
| +- org/jboss/chap2/ex3/EchoLocal.class
| +- org/jboss/chap2/ex3/EchoLocalHome.class
+- chap2-ex3.sar (archive) [MBean sar]
| +- META-INF/MANIFEST.MF
| +- META-INF/jboss-service.xml
| +- org/jboss/chap2/ex3/EjbMBeanAdaptor.class
+- META-INF/application.xml
```

The EAR contains a chap2-ex3.jar and chap2-ex3.sar. The chap2-ex3.jar is the EJB archive and the chap2-ex3.sar is the MBean service archive. We have implemented the service as a Dynamic MBean to provide an illustration of their use. .

```
package org.jboss.chap2.ex3;

import java.lang.reflect.Method;
import javax.ejb.CreateException;
import javax.management.Attribute;
import javax.management.AttributeList;
import javax.management.AttributeNotFoundException;
import javax.management.DynamicMBean;
import javax.management.InvalidAttributeValueException;
```

```

import javax.management.JMRuntimeException;
import javax.management.MBeanAttributeInfo;
import javax.management.MBeanConstructorInfo;
import javax.management.MBeanInfo;
import javax.management.MBeanNotificationInfo;
import javax.management.MBeanOperationInfo;
import javax.management.MBeanException;
import javax.management.MBeanServer;
import javax.management.ObjectName;
import javax.management.ReflectionException;
import javax.naming.InitialContext;
import javax.naming.NamingException;

import org.jboss.system.ServiceMBeanSupport;

/**
 * An example of a DynamicMBean that exposes select attributes and
 * operations of an EJB as an MBean.
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.11 $
 */
public class EjbMBeanAdaptor extends ServiceMBeanSupport
    implements DynamicMBean
{
    private String helloPrefix;
    private String ejbJndiName;
    private EchoLocalHome home;

    /** These are the mbean attributes we expose
     */
    private MBeanAttributeInfo[] attributes = {
        new MBeanAttributeInfo("HelloPrefix", "java.lang.String",
            "The prefix message to append to the session echo reply",
            true, // isReadable
            true, // isWritable
            false), // isIs
        new MBeanAttributeInfo("EjbJndiName", "java.lang.String",
            "The JNDI name of the session bean local home",
            true, // isReadable
            true, // isWritable
            false) // isIs
    };

    /**
     * These are the mbean operations we expose
     */
    private MBeanOperationInfo[] operations;

    /**
     * We override this method to setup our echo operation info. It
     * could also be done in a ctor.
     */
    public ObjectName preRegister(MBeanServer server,
        ObjectName name)
        throws Exception
    {
        log.info("preRegister notification seen");

        operations = new MBeanOperationInfo[5];

        Class thisClass = getClass();
        Class[] parameterTypes = {String.class};
        Method echoMethod =
            thisClass.getMethod("echo", parameterTypes);
        String desc = "The echo op invokes the session bean echo method and"
            + " returns its value prefixed with the helloPrefix attribute value";
        operations[0] = new MBeanOperationInfo(desc, echoMethod);

        // Add the Service interface operations from our super class
        parameterTypes = new Class[0];
        Method createMethod =

```

```

        thisClass.getMethod("create", parameterTypes);
        operations[1] = new MBeanOperationInfo("The
            JBoss Service.create", createMethod);
        Method startMethod =
            thisClass.getMethod("start", parameterTypes);
        operations[2] = new MBeanOperationInfo("The
            JBoss Service.start", startMethod);
        Method stopMethod =
            thisClass.getMethod("stop", parameterTypes);
        operations[3] = new MBeanOperationInfo("The
            JBoss Service.stop", startMethod);
        Method destroyMethod =
            thisClass.getMethod("destroy", parameterTypes);
        operations[4] = new MBeanOperationInfo("The
            JBoss Service.destroy", startMethod);
        return name;
    }

    // --- Begin ServiceMBeanSupport overrides
    protected void createService() throws Exception
    {
        log.info("Notified of create state");
    }

    protected void startService() throws Exception
    {
        log.info("Notified of start state");
        InitialContext ctx = new InitialContext();
        home = (EchoLocalHome) ctx.lookup(ejbJndiName);
    }

    protected void stopService()
    {
        log.info("Notified of stop state");
    }

    // --- End ServiceMBeanSupport overrides

    public String getHelloPrefix()
    {
        return helloPrefix;
    }
    public void setHelloPrefix(String helloPrefix)
    {
        this.helloPrefix = helloPrefix;
    }

    public String getEjbJndiName()
    {
        return ejbJndiName;
    }
    public void setEjbJndiName(String ejbJndiName)
    {
        this.ejbJndiName = ejbJndiName;
    }

    public String echo(String arg)
        throws CreateException, NamingException
    {
        log.debug("Lookup EchoLocalHome@"+ejbJndiName);
        EchoLocal bean = home.create();
        String echo = helloPrefix + bean.echo(arg);
        return echo;
    }

    // --- Begin DynamicMBean interface methods
    /**
     * Returns the management interface that describes this dynamic
     * resource. It is the responsibility of the implementation to
     * make sure the description is accurate.

```

```

*
* @return the management interface descriptor.
*/
public MBeanInfo getMBeanInfo()
{
    String classname = getClass().getName();
    String description = "This is an MBean that uses a session bean in the"
        + " implementation of its echo operation.";
    MBeanConstructorInfo[] constructors = null;
    MBeanNotificationInfo[] notifications = null;
    MBeanInfo mbeanInfo = new MBeanInfo(classname,
                                         description, attributes,
                                         constructors, operations,
                                         notifications);

    // Log when this is called so we know when in the
    lifecycle this is used
    Throwable trace = new Throwable("getMBeanInfo trace");
    log.info("Don't panic, just a stack
        trace", trace);
    return mbeanInfo;
}

/**
 * Returns the value of the attribute with the name matching the
 * passed string.
 *
 * @param attribute the name of the attribute.
 * @return the value of the attribute.
 * @exception AttributeNotFoundException when there is no such
 * attribute.
 * @exception MBeanException wraps any error thrown by the
 * resource when
 * getting the attribute.
 * @exception ReflectionException wraps any error invoking the
 * resource.
 */
public Object getAttribute(String attribute)
    throws AttributeNotFoundException,
           MBeanException,
           ReflectionException
{
    Object value = null;
    if (attribute.equals("HelloPrefix")) {
        value = getHelloPrefix();
    } else if (attribute.equals("EjbJndiName")) {
        value = getEjbJndiName();
    } else {
        throw new AttributeNotFoundException("Unknown
            attribute(" + attribute + ") requested");
    }
    return value;
}

/**
 * Returns the values of the attributes with names matching the
 * passed string array.
 *
 * @param attributes the names of the attribute.
 * @return an {@link AttributeList AttributeList} of name
 * and value pairs.
 */
public AttributeList getAttributes(String[] attributes)
{
    AttributeList values = new AttributeList();
    for (int a = 0; a < attributes.length; a++) {
        String name = attributes[a];
        try {
            Object value = getAttribute(name);
            Attribute attr = new Attribute(name, value);
            values.add(attr);
        } catch (Exception e) {

```

```

        log.error("Failed to find attribute: "+name, e);
    }
}
return values;
}

/**
 * Sets the value of an attribute. The attribute and new value
 * are passed in the name value pair {@link Attribute
 * Attribute}.
 *
 * @see javax.management.Attribute
 *
 * @param attribute the name and new value of the attribute.
 * @exception AttributeNotFoundException when there is no such
 * attribute.
 * @exception InvalidAttributeValueException when the new value
 * cannot be converted to the type of the attribute.
 * @exception MBeanException wraps any error thrown by the
 * resource when setting the new value.
 * @exception ReflectionException wraps any error invoking the
 * resource.
 */
public void setAttribute(Attribute attribute)
    throws AttributeNotFoundException,
        InvalidAttributeValueException,
        MBeanException,
        ReflectionException
{
    String name = attribute.getName();
    if (name.equals("HelloPrefix")) {
        String value = attribute.getValue().toString();
        setHelloPrefix(value);
    } else if (name.equals("EjbJndiName")) {
        String value = attribute.getValue().toString();
        setEjbJndiName(value);
    } else {
        throw new AttributeNotFoundException("Unknown attribute("+name+") requested");
    }
}

/**
 * Sets the values of the attributes passed as an
 * {@link AttributeList AttributeList} of name and new
 * value pairs.
 *
 * @param attributes the name and new value pairs.
 * @return an {@link AttributeList AttributeList} of name and
 * value pairs that were actually set.
 */
public AttributeList setAttributes(AttributeList attributes)
{
    AttributeList setAttributes = new AttributeList();
    for(int a = 0; a < attributes.size(); a++) {
        Attribute attr = (Attribute) attributes.get(a);
        try {
            setAttribute(attr);
            setAttributes.add(attr);
        } catch (Exception ignore) {
        }
    }
    return setAttributes;
}

/**
 * Invokes a resource operation.
 *
 * @param actionName the name of the operation to perform.
 * @param params the parameters to pass to the operation.
 * @param signature the signatures of the parameters.
 * @return the result of the operation.

```

```

* @exception MBeanException wraps any error thrown by the
* resource when performing the operation.
* @exception ReflectionException wraps any error invoking the
* resource.
*/
public Object invoke(String actionName, Object[] params,
                    String[] signature)
    throws MBeanException,
           ReflectionException
{
    Object rtnValue = null;
    log.debug("Begin invoke, actionName="+actionName);
    try {
        if (actionName.equals("echo")) {
            String arg = (String) params[0];
            rtnValue = echo(arg);
            log.debug("Result: "+rtnValue);
        } else if (actionName.equals("create")) {
            super.create();
        } else if (actionName.equals("start")) {
            super.start();
        } else if (actionName.equals("stop")) {
            super.stop();
        } else if (actionName.equals("destroy")) {
            super.destroy();
        } else {
            throw new JMRuntimeException("Invalid state,
            don't know about op="+actionName);
        }
    } catch (Exception e) {
        throw new ReflectionException(e, "echo failed");
    }

    log.debug("End invoke, actionName="+actionName);
    return rtnValue;
}

// --- End DynamicMBean interface methods
}

```

Believe it or not, this is a very trivial MBean. The vast majority of the code is there to provide the MBean metadata and handle the callbacks from the MBean Server. This is required because a Dynamic MBean is free to expose whatever management interface it wants. A Dynamic MBean can in fact change its management interface at runtime simply by returning a different metadata value from the `getMBeanInfo` method. Of course, clients may not be too happy with such a dynamic object, but the MBean Server will do nothing to prevent a Dynamic MBean from changing its interface.

There are two points to this example. First, demonstrate how an MBean can depend on an EJB for some of its functionality and second, how to create MBeans with dynamic management interfaces. If we were to write a standard MBean with a static interface for this example it would look like the following.

```

public interface EjbMBeanAdaptorMBean
{
    public String getHelloPrefix();
    public void setHelloPrefix(String prefix);
    public String getEjbJndiName();
    public void setEjbJndiName(String jndiName);
    public String echo(String arg) throws CreateException, NamingException;
    public void create() throws Exception;
    public void start() throws Exception;
    public void stop();
    public void destroy();
}

```

Moving to lines 67-83, this is where the MBean operation metadata is constructed. The `echo(String)`, `create()`, `start()`, `stop()` and `destroy()` operations are defined by obtaining the corresponding `java.lang.reflect.Method` object and adding a description. Let's go through the code and discuss where this interface implementation exists and how the MBean uses the EJB. Beginning with lines 40-51, the two `MBeanAttributeInfo` instances created define the attributes of the MBean. These attributes correspond to the `getHelloPrefix/setHelloPrefix` and `getEjbJndiName/setEjbJndiName` of the static interface. One thing to note in terms of why one might want to use a Dynamic MBean is that you have the ability to associate descriptive text with the attribute metadata. This is not something you can do with a static interface.

Lines 88-103 correspond to the JBoss service life cycle callbacks. Since we are subclassing the `ServiceMBeanSupport` utility class, we override the `createService`, `startService`, and `stopService` template callbacks rather than the `create`, `start`, and `stop` methods of the service interface. Note that we cannot attempt to lookup the `EchoLocalHome` interface of the EJB we make use of until the `startService` method. Any attempt to access the home interface in an earlier life cycle method would result in the name not being found in JNDI because the EJB container had not gotten to the point of binding the home interfaces. Because of this dependency we will need to specify that the MBean service depends on the `EchoLocal` EJB container to ensure that the service is not started before the EJB container is started. We will see this dependency specification when we look at the service descriptor.

Lines 105-121 are the `HelloPrefix` and `EjbJndiName` attribute accessors implementations. These are invoked in response to `getAttribute/setAttribute` invocations made through the MBean Server.

Lines 123-130 correspond to the `echo(String)` operation implementation. This method invokes the `EchoLocal.echo(String)` EJB method. The local bean interface is created using the `EchoLocalHome` that was obtained in the `startService` method.

The remainder of the class makes up the Dynamic MBean interface implementation. Lines 133-152 correspond to the MBean metadata accessor callback. This method returns a description of the MBean management interface in the form of the `javax.management.MBeanInfo` object. This is made up of a description, the `MBeanAttributeInfo` and `MBeanOperationInfo` metadata created earlier, as well as constructor and notification information. This MBean does not need any special constructors or notifications so this information is null.

Lines 154-258 handle the attribute access requests. This is rather tedious and error prone code so a toolkit or infrastructure that helps generate these methods should be used. A Model MBean framework based on XML called XBeans is currently being investigated in JBoss. Other than this, no other Dynamic MBean frameworks currently exist.

Lines 260-310 correspond to the operation invocation dispatch entry point. Here the request operation action name is checked against those the MBean handles and the appropriate method is invoked.

The `jboss-service.xml` descriptor for the MBean is given below. The dependency on the EJB container MBean is highlighted in bold. The format of the EJB container MBean ObjectName is:

```
"jboss.j2ee:service=EJB,jndiName=" + <home-jndi-name>
```

where the `<home-jndi-name>` is the EJB home interface JNDI name.

```
<server>
  <mbean code="org.jboss.chap2.ex3.EjbMBeanAdaptor"
    name="jboss.book:service=EjbMBeanAdaptor">
    <attribute name="HelloPrefix">AdaptorPrefix</attribute>
    <attribute name="EjbJndiName">local/chap2.EchoBean</attribute>
    <depends>jboss.j2ee:service=EJB,jndiName=local/chap2.EchoBean</depends>
  </mbean>
</server>
```

Deploy the example ear by running:


```
[starksm@banshee examples]$ ant -Dchap=chap2 -Dex=3 run-example
...
run-example3:
    [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
```

On the server console there will be messages similar to the following:

```
14:57:12,906 INFO [EARDeployer] Init J2EE application: file:/private/tmp/jboss-3.2.6/server/default/deploy/chap2-ex3.ear
14:57:13,044 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.RawDynamicInvoker.preRegister(RawDynamicInvoker.java:187)
...
14:57:13,088 INFO [EjbMBeanAdaptor] preRegister notification seen
14:57:13,093 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.registry.BasicMBeanRegistry.registerMBean(BasicMBeanRegistry.java:207)
...
14:57:13,117 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.registry.BasicMBeanRegistry.registerMBean(BasicMBeanRegistry.java:235)
...
14:57:13,140 WARN [EjbMBeanAdaptor] Unexcepected error accessing MBeanInfo for null
java.lang.NullPointerException
    at org.jboss.system.ServiceMBeanSupport.postRegister(ServiceMBeanSupport.java:418)
    at org.jboss.mx.server.RawDynamicInvoker.postRegister(RawDynamicInvoker.java:226)
    at org.jboss.mx.server.registry.BasicMBeanRegistry.registerMBean(BasicMBeanRegistry.java:312)
...
14:57:13,203 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.MBeanServerImpl.getMBeanInfo(MBeanServerImpl.java:481)
...
14:57:13,232 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.MBeanServerImpl.getMBeanInfo(MBeanServerImpl.java:481)
...
14:57:13,420 INFO [EjbModule] Deploying Chap2EchoInfoBean
14:57:13,443 INFO [EjbModule] Deploying chap2.EchoBean
14:57:13,488 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.MBeanServerImpl.getMBeanInfo(MBeanServerImpl.java:481)
...
14:57:13,542 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.MBeanServerImpl.getMBeanInfo(MBeanServerImpl.java:481)
...
14:57:13,558 INFO [EjbMBeanAdaptor] Begin invoke, actionName=create
14:57:13,560 INFO [EjbMBeanAdaptor] Notified of create state
14:57:13,562 INFO [EjbMBeanAdaptor] End invoke, actionName=create
14:57:13,604 INFO [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
```

```
    at org.jboss.mx.server.MBeanServerImpl.getMBeanInfo(MBeanServerImpl.java:481)
    at org.jboss.mx.server.MBeanServerImpl.isInstanceOf(MBeanServerImpl.java:639)
...
14:57:13,621 INFO  [EjbMBeanAdaptor] Don't panic, just a stack trace
java.lang.Throwable: getMBeanInfo trace
    at org.jboss.chap2.ex3.EjbMBeanAdaptor.getMBeanInfo(EjbMBeanAdaptor.java:153)
    at org.jboss.mx.server.RawDynamicInvoker.getMBeanInfo(RawDynamicInvoker.java:172)
    at org.jboss.mx.server.MBeanServerImpl.getMBeanInfo(MBeanServerImpl.java:481)
    at org.jboss.mx.util.JMXInvocationHandler.<init>(JMXInvocationHandler.java:110)
    at org.jboss.mx.util.MBeanProxy.get(MBeanProxy.java:76)
    at org.jboss.mx.util.MBeanProxy.get(MBeanProxy.java:64)
14:57:13,641 INFO  [EjbMBeanAdaptor] Begin invoke, actionName=getState
14:57:13,942 INFO  [EjbMBeanAdaptor] Begin invoke, actionName=start
14:57:13,944 INFO  [EjbMBeanAdaptor] Notified of start state
14:57:13,951 INFO  [EjbMBeanAdaptor] Testing Echo
14:57:13,983 INFO  [EchoBean] echo, info=echo info, arg=, arg=startService
14:57:13,986 INFO  [EjbMBeanAdaptor] echo(startService) = startService
14:57:13,988 INFO  [EjbMBeanAdaptor] End invoke, actionName=start
14:57:13,991 INFO  [EJBDeployer] Deployed: file:/private/tmp/jboss-3.2.6/server/default/tmp
p/deploy/tmpl418chap2-ex3.ear-contents/chap2-ex3.jar
14:57:14,075 INFO  [EARDeployer] Started J2EE application: file:/private/tmp/jboss-3.2.6/s
erver/default/deploy/chap2-ex3.ear
```

The stack traces are not exceptions. They are traces coming from line 150 of the `EjbMBeanAdaptor` code to demonstrate that clients ask for the MBean interface when they want to discover the MBean's capabilities. Notice that the EJB container (lines with `[EjbModule]`) is started before the example MBean (lines with `[EjbMBeanAdaptor]`).

Now, let's invoke the echo method using the JMX console web application. Browse to <http://localhost:8080/jmx-console/HtmlAdaptor?action=inspectMBean&name=jboss.book%3AService%3DEjbMBeanAdaptor> and scroll down to the echo operation section. The view should be like that shown in Figure 2.20.

The screenshot shows the MBean Inspector web console. The browser address bar displays `http://localhost:8080/jmx-console/HtmlAdaptor?action=inspectM`. The console is divided into two main sections: Attribute configuration and Operation configuration.

Attribute Name (Access) Type Description	Attribute Value
HelloPrefix (RW) java.lang.String The prefix message to append to the session echo reply	<input type="text" value="AdaptorPrefix"/>
EjbJndiName (RW) java.lang.String The JNDI name of the session bean local home	<input type="text" value="local/chap2.EchoBean"/>

Operation Name Return Type Description	Parameters
echo java.lang.String The echo op invokes the session bean echo method and returns its value prefixed with the helloPrefix attribute value	arg0 java.lang.String MBean Operation Parameter. <input type="text"/> <input type="button" value="Invoke"/>

Figure 2.20. The EjbMBeanAdaptor MBean operations JMX console view

As shown, we have already entered an argument string of "-echo-arg" into the ParamValue text field. Press the Invoke button and a result string of "AdaptorPrefix-echo-arg" is displayed on the results page. The server console will show several stack traces from the various metadata queries issues by the JMX console and the MBean invoke method debugging lines:

```
10:51:48,671 INFO [EjbMBeanAdaptor] Begin invoke, actionName=echo
10:51:48,671 INFO [EjbMBeanAdaptor] Lookup EchoLocalHome@local/chap2.EchoBean
10:51:48,687 INFO [EchoBean] echo, info=echo info, arg=, arg=-echo-arg
10:51:48,687 INFO [EjbMBeanAdaptor] Result: AdaptorPrefix-echo-arg
10:51:48,687 INFO [EjbMBeanAdaptor] End invoke, actionName=echo
```

2.5. JBoss Deployer Architecture

JBoss has an extensible deployment architecture that allows one to incorporate components into the bare JBoss JMX microkernel. Figure 2.21 shows the classes in the deployment layer.

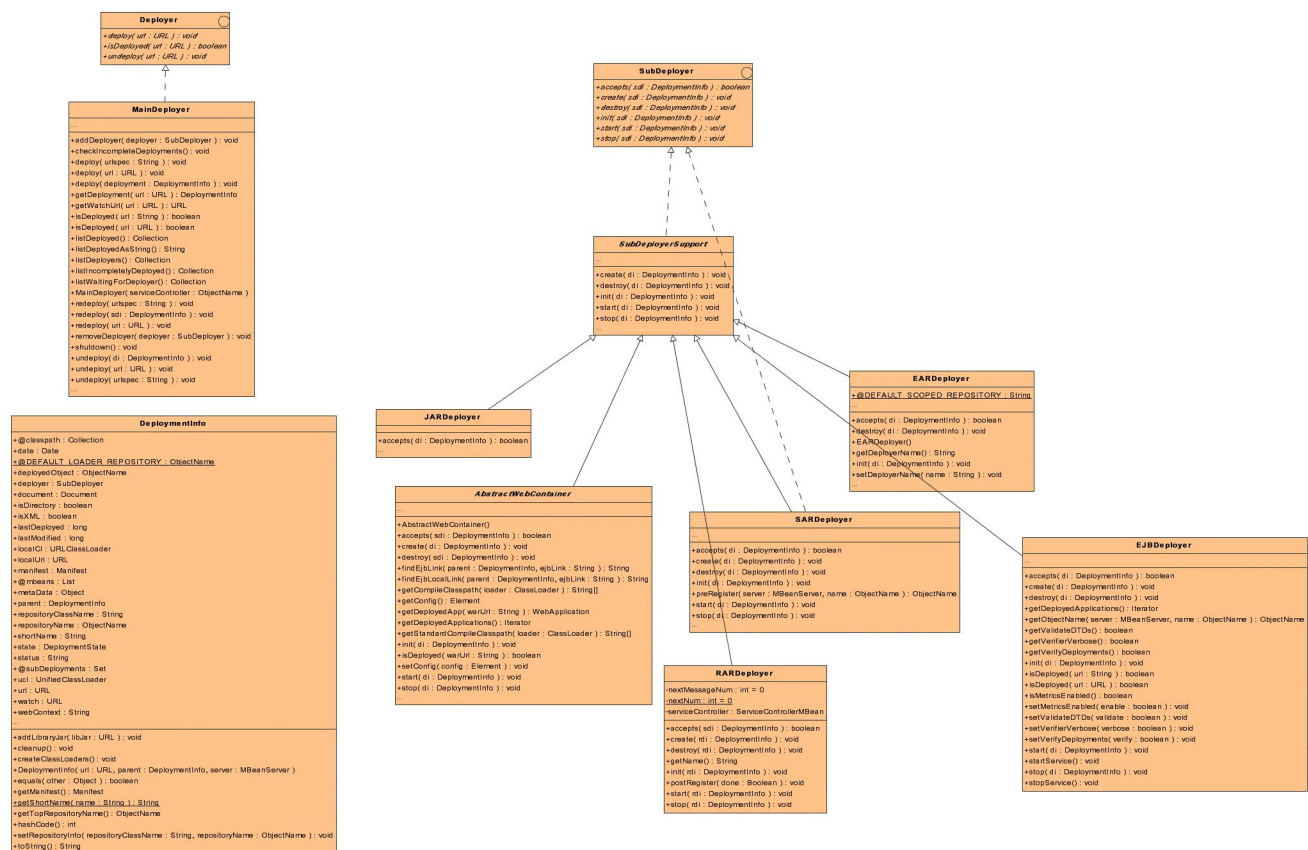


Figure 2.21. The deployment layer classes

The MainDeployer is the deployment entry point. Requests to deploy a component are sent to the MainDeployer and it determines if there is a subdeployer capable of handling the deployment, and if there is, it delegates the deployment to the subdeployer. We saw an example of this when we looked at how the MainDeployer used the SARDeployer to deploy MBean services. The current deployers included with JBoss are:

- **AbstractWebContainer:** This subdeployer handles web application archives (WARs). It accepts deployment archives and directories whose name ends with a "war" suffix. WARs must have a WEB-INF/web.xml descriptor and may have a WEB-INF/jboss-web.xml descriptor.
- **EARDeployer:** This subdeployer handles enterprise application archives (EARs). It accepts deployment archives and directories whose name ends with an "ear" suffix. EARs must have a META-INF/application.xml descriptor and may have a META-INF/jboss-app.xml descriptor.
- **EJBDeployer:** This subdeployer handles enterprise bean jars. It accepts deployment archives and directories whose name ends with a "jar" suffix. EJB jars must have a META-INF/ejb-jar.xml descriptor and may have a META-INF/jboss.xml descriptor.
- **JARDeployer:** This subdeployer handles library jar archives. The only restriction it places on an archive is that it cannot contain a WEB-INF directory.
- **RARDeployer:** This subdeployer handles JCA resource archives (RARs). It accepts deployment archives and directories whose name ends with a "rar" suffix. RARs must have a META-INF/ra.xml descriptor.
- **SARDeployer:** This subdeployer handles JBoss MBean service archives (SARs). It accepts deployment archives and directories whose name ends with a "sar" suffix, as well as standalone XML files that end with "service.xml". SARs that are jars must have a META-INF/jboss-service.xml descriptor.

The MainDeployer, JARDeployer and SARDeployer are hard coded deployers in the JBoss server core. The AbstractWebContainer, EARDeployer, EJBDeployer, and RARDeployer are MBean services that register themselves as deployers with the MainDeployer using the addDeployer(SubDeployer) operation. The SubDeployer interface is shown below.

```
public interface SubDeployer
{
    /**
     * The <code>accepts</code> method
     * is called by MainDeployer to
     * determine which deployer is suitable for a DeploymentInfo.
     *
     * @param sdi a <code>DeploymentInfo</code> value
     * @return a <code>boolean</code> value
     *
     * @jmx:managed-operation
     */
    boolean accepts(DeploymentInfo sdi);

    /**
     * The <code>init</code> method lets the deployer set
     * a few properties of the DeploymentInfo, such as the watch url.
     *
     * @param sdi a <code>DeploymentInfo</code> value
     * @throws DeploymentException if an error occurs
     *
     * @jmx:managed-operation
     */
    void init(DeploymentInfo sdi) throws DeploymentException;

    /**
     * Set up the components of the deployment that do not
     * refer to other components
     *
     * @param sdi a <code>DeploymentInfo</code> value
     * @throws DeploymentException Failed to deploy
     *
     * @jmx:managed-operation
     */
    void create(DeploymentInfo sdi) throws DeploymentException;

    /**
     * The <code>start</code> method sets up relationships
     * with other components.
     *
     * @param sdi a <code>DeploymentInfo</code> value
     * @throws DeploymentException if an error occurs
     *
     * @jmx:managed-operation
     */
    void start(DeploymentInfo sdi) throws DeploymentException;

    /**
     * The <code>stop</code> method removes relationships
     * between components.
     *
     * @param sdi a <code>DeploymentInfo</code> value
     * @throws DeploymentException if an error occurs
     *
     * @jmx:managed-operation
     */
    void stop(DeploymentInfo sdi) throws DeploymentException;

    /**
     * The <code>destroy</code> method
     * removes individual components
     *
     * @param sdi a <code>DeploymentInfo</code> value
     * @throws DeploymentException if an error occurs
     */
}
```

```
    *  
    * @jmx:managed-operation  
    */  
    void destroy(DeploymentInfo sdi) throws DeploymentException;  
}
```

The `DeploymentInfo` object is basically a data structure that encapsulates the complete state of a deployable component. When the `MainDeployer` receives a deployment request, it iterates through its registered subdeployers and invokes the `accepts(DeploymentInfo)` method on the subdeployer. The first subdeployer to return true is chosen and the deployment deployer and the `MainDeployer` will delegate the init, create, start, stop and destroy deployment life cycle operations to the subdeployer.

2.5.1. Deployers and ClassLoaders

Deployers are the mechanism by which components are brought into a JBoss server. Deployers are also the creators of the majority of UCL instances, and the primary creator is the `MainDeployer`. The `MainDeployer` creates the UCL for a deployment early on during its init method. The UCL is created by calling the `DeploymentInfo.createClassLoaders()` method. As of the 3.0.5RC1 release, only the topmost `DeploymentInfo` will actually create a UCL. All subdeployments will add their class paths to their parent `DeploymentInfo` UCL. Previously every subdeployment created a UCL for its deployment, and a separate UCL for every manifest or classpath reference. This could cause problems because classes ended up being loaded by more than one UCL and `IllegalAccessError`s and `LinkageError`s would result. Every deployment does have a standalone `URLClassLoader` that uses the deployment URL as its path. This is used to localize the loading of resources such as deployment descriptors. Figure 2.22 provides an illustration of the interaction between Deployers, `DeploymentInfos` and class loaders.

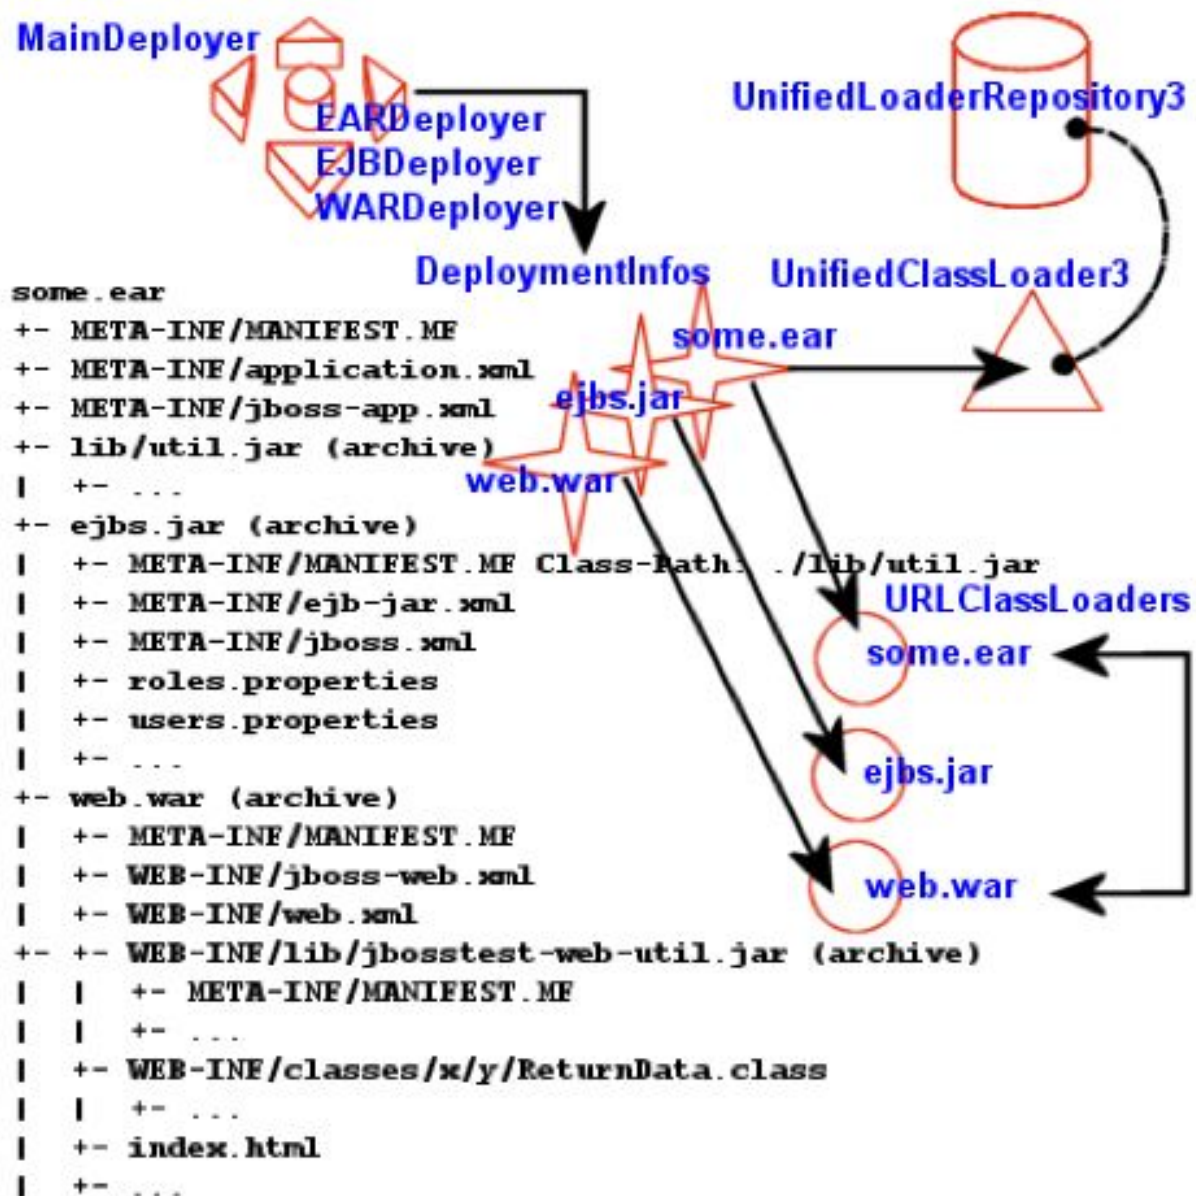


Figure 2.22. An illustration of the class loaders involved with an EAR deployment

The figure illustrates an EAR deployment with EJB and WAR subdeployments. The EJB deployment references the lib/util.jar utility jar via its manifest. The WAR includes classes in its WEB-INF/classes directory as well as the WEB-INF/lib/jbosstest-web-util.jar. Each deployment has a DeploymentInfo instance that has a URLClassLoader pointing to the deployment archive. The DeploymentInfo associated with some.ear is the only one to have a UCL created. The ejbs.jar and web.war DeploymentInfo s add their deployment archive to the some.ear UCL classpath, and share this UCL as their deployment UCL. The EJBDeployer also adds any manifest jars to the EAR UCL.

The WARDeployer behaves differently than other deployers in that it only adds its WAR archive to the DeploymentInfo UCL classpath. The loading of classes from the WAR WEB-INF/classes and WEB-INF/lib locations is handled by the servlet container class loader. The servlet container class loaders delegate to the WAR DeploymentInfo UCL as their parent class loader, but the server container class loader is not part of the JBoss class loader repository. Therefore, classes inside of a WAR are not visible to other components. Classes that need to be shared between web application components and other components such as EJBs, and MBeans need to be loaded into the shared class loader repository either by including the classes into a SAR or EJB deploy-

ment, or by referencing a jar containing the shared classes through a manifest Class-Path entry. In the case of a SAR, the SAR classpath element in the service deployment serves the same purpose as a jar manifest Class-Path.

2.6. Exposing MBean Events via SNMP

3.2.2 added an snmp-adaptor service that can be used to intercept JMX notifications emitted by MBeans, convert them to traps and send them to SNMP managers. In this respect the snmp-adaptor acts as a SNMP agent. Future versions may offer support for full agent get/set functionality that maps onto MBean attributes or operations.

It can be used to integrate JBoss with higher order system/network management platforms (e.g., HP Open-View), thus making the MBeans visible to those systems. The MBean developer can instrument the MBeans by producing notifications for any significant event (e.g. server coldstart). The adaptor can then be configured to intercept and map those notifications to SNMP traps. The adaptor uses the JoeSNMP package from OpenNMS as the SNMP engine.

SnmpAgentService is the main MBean that implements the SNMP agent. It is configured by means of three different configuration files:

- **managers.xml**: configures where to send traps
- **mbeans.xml**: configures the monitored MBeans/notifications types
- **notifications.xml**: specifies the exact mapping of each notification type to a corresponding SNMP trap

2.6.1. The SNMP Adaptor Service

The org.jboss.jmx.adaptor.snmp.agent.SnmpAgentService allows one to send V1 or V2 SNMP traps to one or more SNMP managers defined by their IP address, listening port number and expected SNMP version.

- **HeartBeatPeriod**: The period in seconds at which heartbeat notifications are generated.
- **ManagersResName**: Specifies the resource name of the file containing SNMP manager specifications. The content model for this file is shown in Figure 2.23.
- **MonitoredObjectsResName**: Specifies the resource name of the file that configures which JMX objects to monitor for events. The content model for this file is shown in Figure 2.25.
- **NotificationMapResName**: Specifies the resource name of the file containing the JMX notification to SNMP trap mappings. The content model for this file is shown in Figure 2.24.
- **TrapFactoryClassName**: The org.jboss.jmx.adaptor.snmp.agent.TrapFactory implementation class that takes care of translation of JMX Notifications into SNMP V1 and V2 traps.
- **TimerName**: Specifies the JMX ObjectName of the JMX timer service to use for heartbeat notifications.




Figure 2.23. The schema for the SNMP managers file

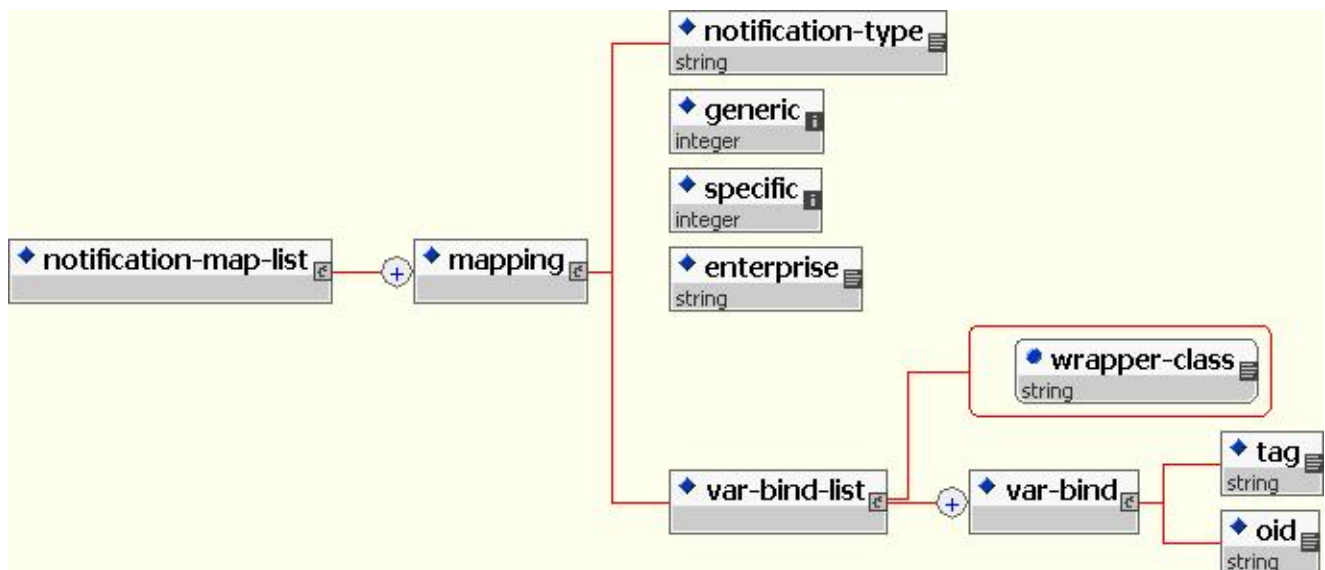


Figure 2.24. The schema for the notification to trap mapping file



Figure 2.25. The schema for the monitored objects file

2.6.2. The Event to Trap Service

`org.jboss.jmx.adaptor.snmp.trapd.TrapdService` is a simple MBean that acts as an SNMP Manager. It listens to a configurable port for incoming traps and logs them as DEBUG messages using the system logger. You can modify the log4j configuration to redirect the log output to a file. `SnmpAgentService` and `TrapdService` are not dependent with each other.

2.7. Remote Access to Services, Detached Invokers

In addition to the MBean services notion that allows for the ability to integrate arbitrary functionality, JBoss also has a detached invoker concept that allows MBean services to expose functional interfaces via arbitrary protocols for remote access by clients. This notion first showed up in 3.0 for the EJB container and it has been further generalized to any MBean service in 3.2. The notion of a detached invoker is that remoting and the protocol by which a service is accessed is a functional aspect or service from independent of the component. Thus, one can make a naming service available for use via RMI/JRMP, RMI/HTTP, RMI/SOAP, or any arbitrary custom transport.

Let's begin our discussion of the detached invoker architecture with an overview of the components involved. The main components in the detached invoker architecture are shown in Figure 2.26.

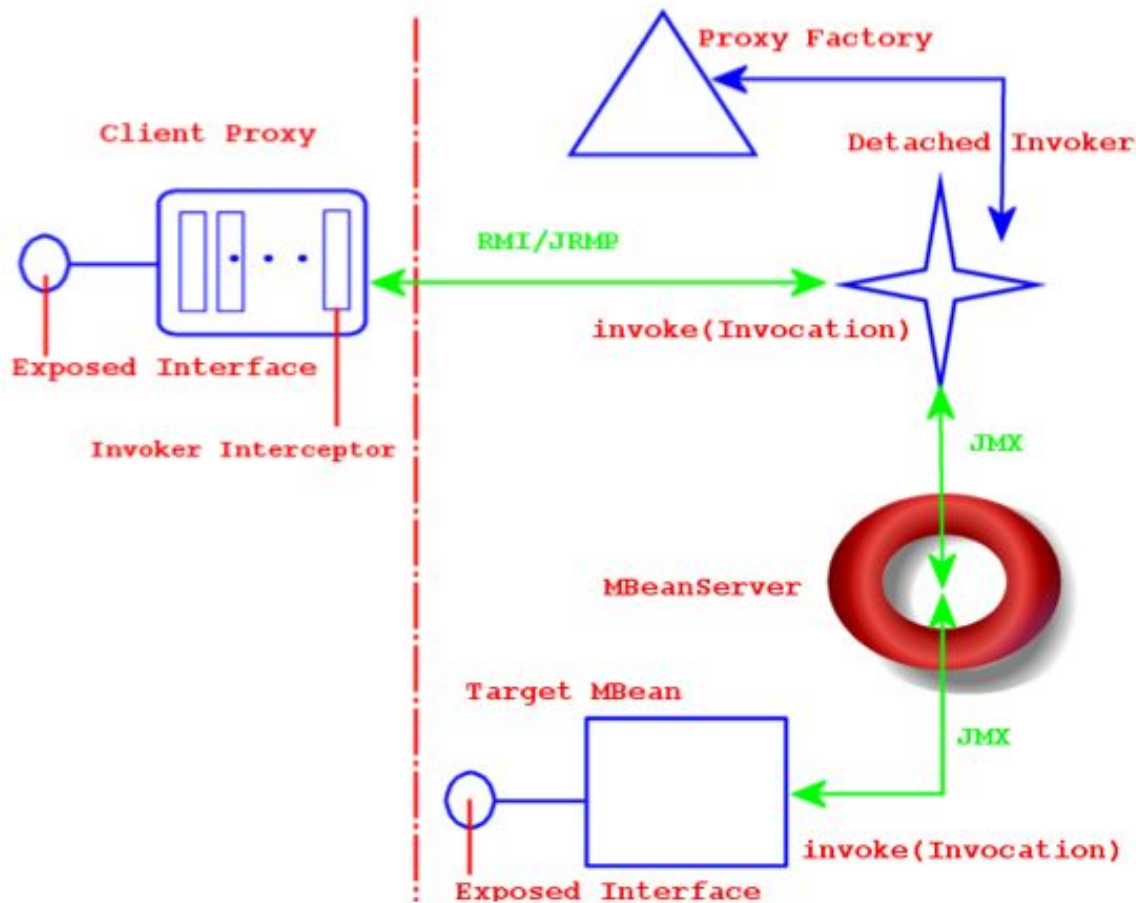


Figure 2.26. The main components in the detached invoker architecture

On the client side, there exist a client proxy which exposes the interface(s) of the MBean service. This is the same smart, compile-less dynamic proxy that we use for EJB home and remote interfaces. The only difference between the proxy for an arbitrary service and the EJB is the set of interfaces exposed as well as the client side interceptors found inside the proxy. The client interceptors are represented by the rectangles found inside of the client proxy. An interceptor is an assembly line type of pattern that allows for transformation of a method invocation and/or return values. A client obtains a proxy through some lookup mechanism, typically JNDI. Although RMI is indicated in Figure 2.26, the only real requirement on the exposed interface and its types is that they are serializable between the client server over JNDI as well as the transport layer.

The choice of the transport layer is determined by the last interceptor in the client proxy, which is referred to as the "Invoker Interceptor" in Figure 2.26. The invoker interceptor contains a reference to the transport specific

stub of the server side "Detached Invoker" MBean service. The invoker interceptor also handles the optimization of calls that occur within the same VM as the target MBean. When the invoker interceptor detects that this is the case the call is passed to a call-by-reference invoker that simply passes the invocation along to the target MBean.

The detached invoker service is responsible for making a generic invoke operation available via the transport the detached invoker handles. The Invoker interface illustrates the generic invoke operation.

```
package org.jboss.invocation;

import java.rmi.Remote;
import org.jboss.proxy.Interceptor;
import org.jboss.util.id.GUID;

public interface Invoker
    extends Remote
{
    GUID ID = new GUID();

    Object invoke(Invocation invocation) throws Exception;
}
```

The Invoker interface extends Remote to be compatible with RMI, but this does not mean that an invoker must expose an RMI service stub. The detached invoker service simply acts as a transport gateway that accepts invocations represented as the org.jboss.invocation.Invocation object over its specific transport, unmarshalls the invocation, forwards the invocation onto the destination MBean service, represented by the "Target MBean" in Figure 2.26, and marshalls the return value or exception resulting from the forwarded call back to the client.

The Invocation object is just a representation of a method invocation context. This includes the target MBean name, the method, the method arguments, a context of information associated with the proxy by the proxy factory, and an arbitrary map of data associated with the invocation by the client proxy interceptors. The following listing shows the key methods of the Invocation class.

```
package org.jboss.invocation;

import java.lang.reflect.Method;
import java.security.Principal;
import java.util.Map;
import java.util.HashMap;
import javax.transaction.Transaction;

public class Invocation
{
    /** The signature of the invoke() method */
    public static final String[] INVOKE_SIGNATURE =
        {"org.jboss.invocation.Invocation"};

    /**
     * Contextual information to the invocation that is not part of
     * the payload.
     */
    public Map transient_payload;

    /**
     * as_is classes that will not be marshalled by the invocation
     * (java.* and javax.* or anything in system classpath is OK)
     */
    public Map as_is_payload;

    /**
     * Payload will be marshalled for type hiding at the RMI layers.
     */
    public Map payload;
}
```

```

protected InvocationContext invocationContext;
protected Object[] args;
protected Object objectName;
protected Method method;

public Invocation()
{
    payload = new HashMap();
    as_is_payload = new HashMap();
    transient_payload = new HashMap();
}

public Invocation(Object id, Method m, Object[] args,
                  Transaction tx,
                  Principal identity, Object credential )
{
    this.payload = new HashMap();
    this.as_is_payload = new HashMap();
    this.transient_payload = new HashMap();

    setId(id);
    setMethod(m);
    setArguments(args);
    setTransaction(tx);
    setPrincipal(identity);
    setCredential(credential);
}

public void setValue(Object key, Object value)
{
    setValue(key, value, PayloadKey.PAYLOAD);
}

public void setValue(Object key, Object value, PayloadKey type)
{
    if(type == PayloadKey.TRANSIENT) {
        transient_payload.put(key,value);
    } else if(type == PayloadKey.AS_IS) {
        as_is_payload.put(key,value);
    } else if(type == PayloadKey.PAYLOAD) {
        payload.put(key,value);
    } else {
        throw new IllegalArgumentException("Unknown
            PayloadKey: " + type);
    }
}

public Object getValue(Object key)
{
    // find where it is
    Object rtn = payload.get(key);
    if (rtn != null) return rtn;

    rtn = as_is_payload.get(key);
    if (rtn != null) return rtn;

    rtn = transient_payload.get(key);
    return rtn;
}

public Object getPayloadValue(Object key)
{
    return payload.get(key);
}

// ... Convenience accessor methods deleted...
}

```

The configuration of the client proxy is done by the server side proxy factory MBean service, indicated by the

"Proxy Factory" component in Figure 2.26. The proxy factory preforms the following tasks:

- Create a dynamic proxy that implements the interface the target MBean wishes to expose.
- Associate the client proxy interceptors with the dynamic proxy handler.
- Associate the invocation context with the dynamic proxy. This includes the target MBean, detached invoker stub and the proxy JNDI name.
- Make the proxy available to clients by binding the proxy into JNDI.

The last component in Figure 2.26 is the "Target MBean" service that wishes to expose an interface for invocations to remote clients. The steps required for an MBean service to be accessible through a given interface are:

- Define a JMX operation matching the signature: `public Object invoke(org.jboss.invocation.Invocation) throws Exception`
- Create a `HashMap<Long, Method>` mapping from the exposed interface `java.lang.reflect.Method`s to the long hash representation using the `org.jboss.invocation.MarshalledInvocation.calculateHash` method.
- Implement the `invoke(Invocation)` JMX operation and use the interface method hash mapping to transform from the long hash representation of the invoked method to the `java.lang.reflect.Method` of the exposed interface. Reflection is used to perform the actual invocation on the object associated with the MBean service that actually implements the exposed interface.

2.7.1. A Detached Invoker Example, the MBeanServer Invoker Adaptor Service

In the section on connecting to the JMX server we mentioned that there was a service that allows one to access the `javax.management.MBeanServer` via any protocol using an invoker service. In this section we present the `org.jboss.jmx.connector.invoker.InvokerAdaptorService` and its configuration for access via RMI/JRMP as an example of the steps required to provide remote access to an MBean service.

The `InvokerAdaptorService` is a simple MBean service that only exists to fulfill the target MBean role in the detached invoker pattern.

Example 2.18. The `InvokerAdaptorService` MBean

```
package org.jboss.jmx.connector.invoker;
public interface InvokerAdaptorServiceMBean
    extends org.jboss.system.ServiceMBean
{
    Class getExportedInterface();
    void setExportedInterface(Class exportedInterface);

    Object invoke(org.jboss.invocation.Invocation invocation)
        throws Exception;
}

package org.jboss.jmx.connector.invoker;

import java.lang.reflect.InvocationTargetException;
import java.lang.reflect.Method;
import java.lang.reflect.UndeclaredThrowableException;
import java.util.Collections;
import java.util.HashMap;
```

```

import java.util.Map;

import javax.management.MBeanServer;
import javax.management.ObjectName;

import org.jboss.invocation.Invocation;
import org.jboss.invocation.MarshalledInvocation;
import org.jboss.mx.server.ServerConstants;
import org.jboss.system.ServiceMBeanSupport;
import org.jboss.system.Registry;

public class InvokerAdaptorService
    extends ServiceMBeanSupport
    implements InvokerAdaptorServiceMBean, ServerConstants
{
    private static ObjectName mbeanRegistry;

    static {
        try {
            mbeanRegistry = new ObjectName(MBEAN_REGISTRY);
        } catch (Exception e) {
            throw new RuntimeException(e.toString());
        }
    }

    private Map marshalledInvocationMapping = new HashMap();
    private Class exportedInterface;

    public Class getExportedInterface()
    {
        return exportedInterface;
    }

    public void setExportedInterface(Class exportedInterface)
    {
        this.exportedInterface = exportedInterface;
    }

    protected void startService()
        throws Exception
    {
        // Build the interface method map
        Method[] methods = exportedInterface.getMethods();
        HashMap tmpMap = new HashMap(methods.length);
        for (int m = 0; m < methods.length; m++) {
            Method method = methods[m];
            Long hash = new Long(MarshalledInvocation.calculateHash(method));
            tmpMap.put(hash, method);
        }

        marshalledInvocationMapping = Collections.unmodifiableMap(tmpMap);
        // Place our ObjectName hash into the Registry so invokers can
        // resolve it
        Registry.bind(new Integer(serviceName.hashCode()), serviceName);
    }

    protected void stopService()
        throws Exception
    {
        Registry.unbind(new Integer(serviceName.hashCode()));
    }

    public Object invoke(Invocation invocation)
        throws Exception
    {
        // Make sure we have the correct classloader before unmarshalling
        Thread thread = Thread.currentThread();
        ClassLoader oldCL = thread.getContextClassLoader();

        // Get the MBean this operation applies to

```

```

ClassLoader newCL = null;
ObjectName objectName = (ObjectName)
    invocation.getValue("JMX_OBJECT_NAME");
if (objectName != null) {
    // Obtain the ClassLoader associated with the MBean deployment
    newCL = (ClassLoader)
        server.invoke(mbeanRegistry, "getValue",
            new Object[] { objectName, CLASSLOADER },
            new String[] { ObjectName.class.getName(),
                "java.lang.String" });
}

if (newCL != null && newCL != oldCL) {
    thread.setContextClassLoader(newCL);
}

try {
    // Set the method hash to Method mapping
    if (invocation instanceof MarshalledInvocation) {
        MarshalledInvocation mi = (MarshalledInvocation) invocation;
        mi.setMethodMap(marshalledInvocationMapping);
    }

    // Invoke the MBeanServer method via reflection
    Method method = invocation.getMethod();
    Object[] args = invocation.getArguments();
    Object value = null;
    try {
        String name = method.getName();
        Class[] sig = method.getParameterTypes();
        Method mbeanServerMethod =
            MBeanServer.class.getMethod(name, sig);
        value = mbeanServerMethod.invoke(server, args);
    } catch (InvocationTargetException e) {
        Throwable t = e.getTargetException();
        if (t instanceof Exception) {
            throw (Exception) t;
        } else {
            throw new UndeclaredThrowableException(t, method.toString());
        }
    }

    return value;
} finally {
    if (newCL != null && newCL != oldCL) {
        thread.setContextClassLoader(oldCL);
    }
}
}

```

Let's go through the key details of this service. The `InvokerAdaptorServiceMBean` Standard MBean interface of the `InvokerAdaptorService` has a single `ExportedInterface` attribute and a single `invoke(Invocation)` operation. The `ExportedInterface` attribute allows customization of the type of interface the service exposes to clients. This has to be "compatible" with the `MBeanServer` class in terms of method name and signature as we will see. The `invoke(Invocation)` operation is the required entry point that target MBean services must expose to participate in the detached invoker pattern. This operation is invoked by the detached invoker services that have been configured to provide access to the `InvokerAdaptorService`.

Lines 54-64 of the `InvokerAdaptorService` build the `HashMap<Long, Method>` of the `ExportedInterface` Class using the `org.jboss.invocation.MarshalledInvocation.calculateHash(Method)` utility method. Because `java.lang.reflect.Method` instances are not serializable, a `MarshalledInvocation` version of the non-serializable `Invocation` class is used to marshall the invocation between the client and server. The `MarshalledInvocation` replaces the `Method` instances with their corresponding hash representation. On the server side, the `MarshalledIn-`

vocation must be told what the hash to Method mapping is.

Line 64 creates a mapping between the `InvokerAdaptorService` service name and its `hashCode` representation. This is used by detached invokers to determine what the target MBean `ObjectName` of an Invocation is. When the target MBean name is store in the Invocation, its store as its `hashCode` because `ObjectName` s are relatively expensive objects to create. The `org.jboss.system.Registry` is a global map like construct that invokers use to store the `hashCode` to `ObjectName` mappings in.

Lines 77-93 obtain the name of the MBean on which the MBeanServer operation is being performed and look-up the `ClassLoader` associated with the MBean's SAR deployment. This information is available via the `org.jboss.mx.server.registry.BasicMBeanRegistry`, a JBoss JMX implementation specific class. It is generally necessary for an MBean to establish the correct class loading context because the detached invoker protocol layer may not have access to the class loaders needed to unmarshall the types associated with an invocation.

Lines 101-105 install the `ExposedInterface` class method hash to method mapping if the invocation argument is of type `MarshaledInvocation`. The method mapping calculated previously at lines 54-62 is used here.

Lines 107-114 perform a second mapping from the `ExposedInterface` Method to the matching method of the MBeanServer class. The `InvokerServiceAdaptor` decouples the `ExposedInterface` from the MBeanServer class in that it allows an arbitrary interface. This is needed on one hand because the standard `java.lang.reflect.Proxy` class can only proxy interfaces. It also allows one to only expose a subset of the MBeanServer methods and add transport specific exceptions like `java.rmi.RemoteException` to the `ExposedInterface` method signatures.

Line 115 dispatches the MBeanServer method invocation to the MBeanServer instance to which the `InvokerAdaptorService` was deployed. The server instance variable is inherited from the `ServiceMBeanSupport` super-class.

Lines 117-124 handle any exceptions coming from the reflective invocation including the unwrapping of any declared exception thrown by the invocation.

Line 126 is the return of the successful MBeanServer method invocation result.

Note that the `InvokerAdaptorService` MBean does not deal directly with any transport specific details. There is the calculation of the method hash to Method mapping, but this is a transport independent detail.

Now let's take a look at how the `InvokerAdaptorService` may be used to expose the same `org.jboss.jmx.adaptor.rmi.RMIAdaptor` interface via RMI/JRMP as seen in [Connecting to JMX Using RMI](#). We will start by presenting the proxy factory and `InvokerAdaptorService` configurations found in the default setup in the `jmx-invoker-adaptor-service.sar` deployment. Example 2.19 shows the `jboss-service.xml` descriptor for this deployment.

Example 2.19. The default `jmx-invoker-adaptor-server.sar` `jboss-service.xml` deployment descriptor

```
<server>
  <!-- The JRMP invoker proxy configuration for the InvokerAdaptorService -->
  <mbean code="org.jboss.invocation.jrmp.server.JRMPProxyFactory"
    name="jboss.jmx:type=adaptor,name=Invoker,protocol=jrmp,service=proxyFactory">
    <!-- Use the standard JRMPInvoker from conf/jboss-service.xml -->
    <attribute name="InvokerName">jboss:service=invoker,type=jrmp</attribute>
    <!-- The target MBean is the InvokerAdaptorService configured below -->
    <attribute name="TargetName">jboss.jmx:type=adaptor,name=Invoker</attribute>
    <!-- Where to bind the RMIAdaptor proxy -->
    <attribute name="JndiName">jmx/invoker/RMIAdaptor</attribute>
    <!-- The RMI compabitle MBeanServer interface -->
    <attribute name="ExportedInterface">org.jboss.jmx.adaptor.rmi.RMIAdaptor</attribute>
    <attribute name="ClientInterceptors">
      <interceptors>
```



```

        <interceptor>org.jboss.proxy.ClientMethodInterceptor</interceptor>
        <interceptor>
            org.jboss.jmx.connector.invoker.client.InvokerAdaptorClientInterceptor </interceptor>
        <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
    </interceptors>
</attribute>
<depends>jboss:service=invoker,type=jrmp</depends>
</mbean>
<!-- This is the service that handles the RMIAdaptor invocations by routing
      them to the MBeanServer the service is deployed under. -->
<mbean code="org.jboss.jmx.connector.invoker.InvokerAdaptorService"
      name="jboss:jmx:type=adaptor,name=Invoker">
    <attribute name="ExportedInterface">org.jboss.jmx.adaptor.rmi.RMIAdaptor </attribute>
</mbean>
</server>

```

The first MBean, `org.jboss.invocation.jrmp.server.JRMPProxyFactory`, is the proxy factory MBean service that creates proxies for the RMI/JRMP protocol. The complete reference information on the `JRMPProxyFactory` may be found in section [The JRMPProxyFactory Service - Building Dynamic JRMP Proxies](#). The configuration of this service as shown in Example 2.19 states that the `JRMPInvoker` will be used as the detached invoker, the `InvokerAdaptorService` is the target mbean to which requests will be forwarded, that the proxy will expose the `RMIAdaptor` interface, the proxy will be bound into JNDI under the name `"jmx/invoker/RMIAdaptor"`, and the proxy will contain 3 interceptors: `ClientMethodInterceptor`, `InvokerAdaptorClientInterceptor`, `InvokerInterceptor`. The configuration of the `InvokerAdaptorService` simply sets the `RMIAdaptor` interface that the service is exposing.

The last piece of the configuration for exposing the `InvokerAdaptorService` via RMI/JRMP is the detached invoker. The detached invoker we will use is the standard RMI/JRMP invoker used by the EJB containers for home and remote invocations, and this is the `org.jboss.invocation.jrmp.server.JRMPInvoker` MBean service configured in the `conf/jboss-service.xml` descriptor. That we can use the same service instance emphasizes the detached nature of the invokers. The `JRMPInvoker` simply acts as the RMI/JRMP endpoint for all RMI/JRMP proxies regardless of the interface(s) the proxies expose or the service the proxies utilize.

2.7.2. Detached Invoker Reference

2.7.2.1. The JRMPInvoker - RMI/JRMP Transport

The `org.jboss.invocation.jrmp.server.JRMPInvoker` class is an MBean service that provides the RMI/JRMP implementation of the `Invoker` interface. The `JRMPInvoker` exports itself as an RMI server so that when it is used as the `Invoker` in a remote client, the `JRMPInvoker` stub is sent to the client instead and invocations use the RMI/JRMP protocol.

The `JRMPInvoker` MBean supports a number of attribute to configure the RMI/JRMP transport layer. Its configurable attributes are:

- **RMIObjectPort:** sets the RMI server socket listening port number. This is the port RMI clients will connect to when communicating through the proxy interface. The default setting in the `jboss-service.xml` descriptor is 4444, and if not specified, the attribute defaults to 0 to indicate an anonymous port should be used.
- **RMIClientSocketFactory:** specifies a fully qualified class name for the `java.rmi.server.RMIClientSocketFactory` interface to use during export of the proxy interface.
- **RMIServerSocketFactory:** specifies a fully qualified class name for the

java.rmi.server.RMIServerSocketFactory interface to use during export of the proxy interface.

- **ServerAddress:** specifies the interface address that will be used for the RMI server socket listening port. This can be either a DNS hostname or a dot-decimal Internet address. Since the RMIServerSocketFactory does not support a method that accepts an InetAddress object, this value is passed to the RMIServerSocketFactory implementation class using reflection. A check for the existence of a: public void setBindAddress(java.net.InetAddress addr) method is made, and if one exists, the RMIServerSocketAddr value is passed to the RMIServerSocketFactory implementation. If the RMIServerSocketFactory implementation does not support such a method, the ServerAddress value will be ignored.
- **SecurityDomain:** specifies the JNDI name of an org.jboss.security.SecurityDomain interface implementation to associate with the RMIServerSocketFactory implementation. The value will be passed to the RMIServerSocketFactory using reflection to locate a method with a signature of: public void setSecurityDomain(org.jboss.security.SecurityDomain d) If no such method exists the SecurityDomain value will be ignored.

2.7.2.2. The PooledInvoker - RMI/Socket Transport

The org.jboss.invocation.pooled.server.PooledInvoker is an MBean service that provides RMI over a custom socket transport implementation of the Invoker interface. The PooledInvoker exports itself as an RMI server so that when it is used as the Invoker in a remote client, the PooledInvoker stub is sent to the client instead and invocations use the a custom socket protocol.

The PooledInvoker MBean supports a number of attribute to configure the socket transport layer. Its configurable attributes are:

- **NumAcceptThreads:** The number of threads that exist for accepting client connections. The default is 1.
- **MaxPoolSize:** The number of server threads for processing client. The default is 300.
- **SocketTimeout:** The socket timeout value passed to the Socket.setSoTimeout() cmethod. The default is 60000.
- **ServerBindPort:** The port used for the server socket. A value of 0 indicates that an anonymous port should be chosen.
- **ClientConnectAddress:** The address that the client passes to the Socket(addr, port) constructor. This defaults to the server InetAddress.getLocalHost() value.
- **ClientConnectPort:** The port that the client passes to the Socket(addr, port) constructor. The default is the port of the server listening socket.
- **ClientMaxPoolSize:** The client side maximum number of threads. The default is 300.
- **Backlog:** The backlog associated with the server accept socket. The default is 200.
- **EnableTcpNoDelay:** A boolean flag indicating if client sockets will enable the TcpNoDelay flag on the socket. The default is false.
- **ServerBindAddress:** The address on which the server binds its listening socket. The default is an empty value which indicates the server should be bound on all interfaces.
- **TransactionManagerService:** The JMX ObjectName of the JTA transaction manager service.

2.7.2.3. The IIOPInvoker - RMI/IIOP Transport

The `org.jboss.invocation.iiop.IIOPInvoker` class is an MBean service that provides the RMI/IIOP implementation of the Invoker interface. The IIOPInvoker IIOP invoker that routes IIOP requests to CORBA servants are used by the `org.jboss.proxy.ejb.IORFactory` proxy factory to create RMI/IIOP proxies. However, rather than creating Java proxies (as the JRMP proxy factory does), this factory creates CORBA IORs. An `IORFactory` is associated to a given enterprise bean. It registers with the IIOP invoker two CORBA servants: an `EjbHomeCorbaServant` for the bean's EJBHome and an `EjbObjectCorbaServant` for the bean's EJBObjects.

The IIOPInvoker MBean has no configurable properties, since all properties are configured from the `conf/jacorb.properties` property file used by the JacORB CORBA service.

2.7.2.4. The JRMPProxyFactory Service - Building Dynamic JRMP Proxies

The `org.jboss.invocation.jrmp.server.JRMPProxyFactory` MBean service is a proxy factory that can expose any interface with RMI compatible semantics for access to remote clients using JRMP as the transport.

The JRMPProxyFactory supports the following attributes:

- **InvokerName:** The server side JRMPInvoker MBean service JMX ObjectName string that will handle the RMI/JRMP transport.
- **TargetName:** The server side MBean that exposes the `invoke(Invocation)` JMX operation for the exported interface. This is used as the destination service for any invocations done through the proxy.
- **JndiName:** The JNDI name under which the proxy will be bound.
- **ExportedInterface:** The fully qualified class name of the interface that the proxy implements. This is the typed view of the proxy that the client uses for invocations.
- **ClientInterceptors:** An XML fragment of interceptors/interceptor elements with each interceptor element body specifying the fully qualified class name of an `org.jboss.proxy.Interceptor` implementation to include in the proxy interceptor stack. The ordering of the interceptors/interceptor elements defines the order of the interceptors.

2.7.2.5. The HttpInvoker - RMI/HTTP Transport

The `org.jboss.invocation.http.server.HttpInvoker` MBean service provides the provides support for making invocations into the JMX bus over HTTP. Unlike the JRMPInvoker, the HttpInvoker is not an implementation of Invoker, but it does implement the `Invoker.invoke` method. The HttpInvoker is accessed indirectly by issuing an HTTP POST against the `org.jboss.invocation.http.servlet.InvokerServlet`. The HttpInvoker exports a client side proxy in the form of the `org.jboss.invocation.http.interfaces.HttpInvokerProxy` class, which is an implementation of Invoker, and is serializable. The HttpInvoker is a drop in replacement for the JRMPInvoker as the target of the bean-invoker and home-invoker EJB configuration elements. The HttpInvoker and InvokerServlet are deployed in the `http-inovker.sar` discussed in the JNDI chapter in the section entitled Accessing JNDI over HTTP

The HttpInvoker supports the following attributes:

- **InvokerURL:** This is either the http URL to the InvokerServlet mapping, or the name of a system property that will be resolved inside the client VM to obtain the http URL to the InvokerServlet. This value can itself be a reference to a system property resolved in the server if the value is of the form `${x}` where x is the

name of the system property. This allows the URL or client side system property to be set in one place and reused in the `HttpInvoker` config as well as the `InvokerServlet` config.

- **InvokerURLPrefix:** If there is no `invokerURL` set, then one will be constructed via the concatenation of `invokerURLPrefix` + the local host + `invokerURLSuffix`. An example prefix is "http://", and this is the default.
- **InvokerURLSuffix:** If there is no `invokerURL` set, then one will be constructed via the concatenation of `invokerURLPrefix` + the local host + `invokerURLSuffix`. An example suffix is ":8080/invoker/JMXInvokerServlet" and this is the default.
- **UseHostName:** A boolean flag if the `InetAddress.getHostName()` or `getHostAddress()` method should be used as the host component of `invokerURLPrefix` + host + `invokerURLSuffix`. If true `getHostName()` is used, false `getHostAddress()`.

2.7.2.6. The HA JRMPInvoker - Clustered RMI/JRMP Transport

The `org.jboss.proxy.generic.ProxyFactoryHA` service is an extension of the `ProxyFactoryHA` that is a cluster aware factory. The `ProxyFactoryHA` fully supports all of the attributes of the `JRMPProxyFactory`. This means that customized bindings of the port, interface and socket transport are available to clustered RMI/JRMP as well. In addition, the following cluster specific attributes are supported:

- **PartitionObjectName:** The JMX ObjectName of the cluster service to which the proxy is to be associated with.
- **LoadBalancePolicy:** The class name of the `org.jboss.ha.framework.interfaces.LoadBalancePolicy` interface implementation to associate with the proxy.

2.7.2.7. The HA HttpInvoker - Clustered RMI/HTTP Transport

The RMI/HTTP layer added in JBoss-3.0.2 has been extended to allow for software load balancing of the invocations in a clustered environment in JBoss-3.0.3. An HA capable extension of the HTTP invoker has been added that borrows much of its functionality from the HA-RMI/JRMP clustering.

To enable HA-RMI/HTTP you need to configure the invokers for the EJB container. This is done through either a `jboss.xml` descriptor, or the `standardjboss.xml` descriptor.

2.7.2.8. HttpProxyFactory - Building Dynamic HTTP Proxies

The `org.jboss.invocation.http.server.HttpProxyFactory` MBean service is a proxy factory that can expose any interface with RMI compatible semantics for access to remote clients using HTTP as the transport.

The `HttpProxyFactory` supports the following attributes:

- **InvokerName:** The server side MBean that exposes the invoke operation for the exported interface. The name is embedded into the `HttpInvokerProxy` context as the target to which the invocation should be forwarded by the `HttpInvoker`.
- **JndiName:** The JNDI name under which the `HttpInvokerProxy` will be bound. This is the name clients lookup to obtain the dynamic proxy that exposes the service interfaces and marshalls invocations over HTTP. This may be specified as an empty value to indicate that the proxy should not be bound into JNDI.
- **InvokerURL:** This is either the http URL to the `InvokerServlet` mapping, or the name of a system property

that will be resolved inside the client VM to obtain the http URL to the `InvokerServlet`. This value can itself be a reference to a system property resolved in the server if the value is of the form `${x}` where `x` is the name of the system property.

- **InvokerURLPrefix:** If there is no `invokerURL` set, then one will be constructed via the concatenation of `invokerURLPrefix` + the local host + `invokerURLSuffix`. An example prefix is `"http://"`, and this is the default.
- **InvokerURLSuffix:** If there is no `invokerURL` set, then one will be constructed via the concatenation of `invokerURLPrefix` + the local host + `invokerURLSuffix`. An example suffix is `":8080/invoker/JMXInvokerServlet"` and this is the default.
- **UseHostName:** A boolean flag indicating if the `InetAddress.getHostName()` or `getHostAddress()` method should be used as the host component of `invokerURLPrefix` + host + `invokerURLSuffix`. If true `getHostName()` is used, false `getHostAddress()`.
- **ExportedInterface:** The name of the RMI compatible interface that the `HttpInvokerProxy` implements.

2.7.2.9. Steps to Expose Any RMI Interface via HTTP

Using the `HttpProxyFactory` MBean and JMX, you can expose any interface for access using HTTP as the transport. The interface to expose does not have to be an RMI interface, but it does have to be compatible with RMI in that all method parameters and return values are serializable. There is also no support for converting RMI interfaces used as method parameters or return values into their stubs.

The three steps to making your object invocable via HTTP are:

```
import java.lang.reflect.Method;
import java.util.HashMap;
import org.jboss.invocation.MarshalledInvocation;

HashMap marshalledInvocationMapping = new HashMap();

// Build the Naming interface method map
Method[] methods = SRPRemoteServerInterface.class.getMethods();
for(int m = 0; m < methods.length; m++) {
    Method method = methods[m];
    Long hash = new Long(MarshalledInvocation.calculateHash(method));
    marshalledInvocationMapping.put(hash, method);
}
```

```
import org.jboss.invocation.Invocation;
import org.jboss.invocation.MarshalledInvocation;

public Object invoke(Invocation invocation)
    throws Exception
{
    SRPRemoteServerInterface theServer = <the_actual_rmi_server_object>;
    // Set the method hash to Method mapping
    if (invocation instanceof MarshalledInvocation) {
        MarshalledInvocation mi = (MarshalledInvocation) invocation;
        mi.setMethodMap(marshalledInvocationMapping);
    }

    // Invoke the Naming method via reflection
    Method method = invocation.getMethod();
    Object[] args = invocation.getArguments();
    Object value = null;
    try {
        value = method.invoke(theServer, args);
    } catch (InvocationTargetException e) {
        Throwable t = e.getTargetException();
    }
}
```

```
        if (t instanceof Exception) {
            throw (Exception) e;
        } else {
            throw new UndeclaredThrowableException(t, method.toString());
        }
    }

    return value;
}
```

```
<!-- Expose the SRP service interface via HTTP -->
<mbean code="org.jboss.invocation.http.server.HttpProxyFactory"
      name="jboss.security.tests:service=SRP/HTTP">
    <attribute name="InvokerURL">http://localhost:8080/invoker/JMXInvokerServlet</attribute>
    <attribute name="InvokerName">jboss.security.tests:service=SRPService</attribute>
    <attribute name="ExportedInterface">org.jboss.security.srp.SRPRemoteServerInterface
    </attribute><attribute name="JndiName">srp-test-http/SRPServerInterface</attribute>
</mbean>
```

- Create a mapping of longs to the RMI interface methods using the `MarshaledInvocation.calculateHash` method. Here for example, is the procedure for an RMI `SRPRemoteServerInterface` interface:
- Either create or extend an existing MBean to support an invoke operation. Its signature is `Object invoke(Invocation invocation) throws Exception`, and the steps it performs are as shown here for the `SRPRemoteServerInterface` interface. Note that this uses the `marshalledInvocationMapping` from step 1 to map from the `Long` method hashes in the `MarshaledInvocation` to the `Method` for the interface.
- Create a configuration of the `HttpProxyFactory` MBean to make the RMI/HTTP proxy available through JNDI. For example:

Any client may now lookup the RMI interface from JNDI using the name specified in the `HttpProxyFactory` (e.g., `srp-test-http/SRPServerInterface`) and use the obtain proxy in exactly the same manner as the RMI/JRMP version.

Naming on JBoss

The JNDI Naming Service

This chapter discusses the JBoss JNDI based naming service and the role of JNDI in JBoss and J2EE. An introduction to the basic JNDI API and common usage conventions will also be discussed. The JBoss specific configuration of J2EE component naming environments defined by the standard deployment descriptors will also be addressed. The final topic is the configuration and architecture of the JBoss naming service.

The JBoss naming service is an plays a key role in J2EE because it provides a naming service that allows a user to map a name onto an object. This is a fundamental need in any programming environment because developers and administrators want to be able to refer to objects and services by recognizable names. A good example of a pervasive naming service is the Internet Domain Name System (DNS). The DNS service allows you to refer to hosts using logical names, rather than their numeric Internet addresses. JNDI serves a similar role in J2EE by enabling developers and administrators to create name-to-object bindings for use in J2EE components.

3.1. An Overview of JNDI

JNDI is a standard Java API that is bundled with JDK1.3 and higher. JNDI provides a common interface to a variety of existing naming services: DNS, LDAP, Active Directory, RMI registry, COS registry, NIS, and file systems. The JNDI API is divided logically into a client API that is used to access naming services, and a service provider interface (SPI) that allows the user to create JNDI implementations for naming services.

The SPI layer is an abstraction that naming service providers must implement to enable the core JNDI classes to expose the naming service using the common JNDI client interface. An implementation of JNDI for a naming service is referred to as a JNDI provider. JBoss naming is an example JNDI implementation, based on the SPI classes. Note that the JNDI SPI is not needed by J2EE component developers.

For a thorough introduction and tutorial on JNDI, which covers both the client and service provider APIs, see the Sun tutorial at <http://java.sun.com/products/jndi/tutorial/>.

3.1.1. The JNDI API

The main JNDI API package is the `javax.naming` package. It contains five interfaces, 10 classes, and several exceptions. There is one key class, `InitialContext`, and two key interfaces, `Context` and `Name`.

3.1.1.1. Names

The notion of a name is of fundamental importance in JNDI. The naming system determines the syntax that the name must follow. The syntax of the naming system allows the user to parse string representations of names into its components. A name is used with a naming system to locate objects. In the simplest sense, a naming system is just a collection of objects with unique names. To locate an object in a naming system you provide a name to the naming system, and the naming system returns the object store under the name.

As an example, consider the Unix file system's naming convention. Each file is named from its path relative to

the root of the file system, with each component in the path separated by the forward slash character ("/"). The file's path is ordered from left to right. The pathname `/usr/jboss/readme.txt`, for example, names a file `readme.txt` in the directory `jboss`, under the directory `usr`, located in the root of the file system. JBoss naming uses a UNIX-style namespace as its naming convention.

The `javax.naming.Name` interface represents a generic name as an ordered sequence of components. It can be a composite name (one that spans multiple namespaces), or a compound name (one that is used within a single hierarchical naming system). The components of a name are numbered. The indexes of a name with `N` components range from 0 up to, but not including, `N`. The most significant component is at index 0. An empty name has no components.

A composite name is a sequence of component names that span multiple namespaces. An example of a composite name would be the hostname+file commonly used with UNIX commands like `scp`. For example, this command copies `localfile.txt` to the file `remotefile.txt` in the `tmp` directory on host `ahost.someorg.org`:

```
scp localfile.txt ahost.someorg.org:/tmp/remotefile.txt
```

A compound name is derived from a hierarchical namespace. Each component in a compound name is an atomic name, meaning a string that cannot be parsed into smaller components. A file pathname in the UNIX file system is an example of a compound name. `ahost.someorg.org:/tmp/remotefile.txt` is a composite name that spans the DNS and UNIX file system namespaces. The components of the composite name are `ahost.someorg.org` and `/tmp/remotefile.txt`. A component is a string name from the namespace of a naming system. If the component comes from a hierarchical namespace, that component can be further parsed into its atomic parts by using the `javax.naming.CompoundName` class. The JNDI API provides the `javax.naming.CompositeName` class as the implementation of the `Name` interface for composite names.

3.1.1.2. Contexts

The `javax.naming.Context` interface is the primary interface for interacting with a naming service. The `Context` interface represents a set of name-to-object bindings. Every context has an associated naming convention that determines how the context parses string names into `javax.naming.Name` instances. To create a name to object binding you invoke the `bind` method of a `Context` and specify a name and an object as arguments. The object can later be retrieved using its name using the `Context` lookup method. A `Context` will typically provide operations for binding a name to an object, unbinding a name, and obtaining a listing of all name-to-object bindings. The object you bind into a `Context` can itself be of type `Context`. The `Context` object that is bound is referred to as a subcontext of the `Context` on which the `bind` method was invoked.

As an example, consider a file directory with a pathname `/usr`, which is a context in the UNIX file system. A file directory named relative to another file directory is a subcontext (commonly referred to as a subdirectory). A file directory with a pathname `/usr/jboss` names a `jboss` context that is a subcontext of `usr`. In another example, a DNS domain, such as `org`, is a context. A DNS domain named relative to another DNS domain is another example of a subcontext. In the DNS domain `jboss.org`, the DNS domain `jboss` is a subcontext of `org` because DNS names are parsed right to left.

3.1.1.2.1. Obtaining a Context using InitialContext

All naming service operations are performed on some implementation of the `Context` interface. Therefore, you need a way to obtain a `Context` for the naming service you are interested in using. The `javax.naming.InitialContext` class implements the `Context` interface, and provides the starting point for interacting with a naming service.

When you create an `InitialContext`, it is initialized with properties from the environment. JNDI determines each property's value by merging the values from the following two sources, in order.

- The first occurrence of the property from the constructor's environment parameter and (for appropriate properties) the applet parameters and system properties.
- All `jndi.properties` resource files found on the classpath.

For each property found in both of these two sources, the property's value is determined as follows. If the property is one of the standard JNDI properties that specify a list of JNDI factories, all of the values are concatenated into a single colon-separated list. For other properties, only the first value found is used. The preferred method of specifying the JNDI environment properties is through a `jndi.properties` file, which allows your code to externalize the JNDI provider specific information so that changing JNDI providers will not require changes to your code or recompilation.

The `Context` implementation used internally by the `InitialContext` class is determined at runtime. The default policy uses the environment property `java.naming.factory.initial`, which contains the class name of the `javax.naming.spi.InitialContextFactory` implementation. You obtain the name of the `InitialContextFactory` class from the naming service provider you are using.

Example 3.1 gives a sample `jndi.properties` file a client application would use to connect to a JBossNS service running on the local host at port 1099. The client application would need to have the `jndi.properties` file available on the application classpath. These are the properties that the JBossNS JNDI implementation requires. Other JNDI providers will have different properties and values.

Example 3.1. A sample `jndi.properties` file

```
### JBossNS properties
java.naming.factory.initial=org.jnp.interfaces.NamingContextFactory
java.naming.provider.url=jnp://localhost:1099
java.naming.factory.url.pkgs=org.jboss.naming:org.jnp.interfaces
```

3.1.2. J2EE and JNDI - The Application Component Environment

JNDI is a fundamental aspect of the J2EE specifications. One key usage is the isolation of J2EE component code from the environment in which the code is deployed. Use of the application component's environment allows the application component to be customized without the need to access or change the application component's source code. The application component environment is referred to as the ENC, the enterprise naming context. It is the responsibility of the application component container to make an ENC available to the container components in the form of JNDI Context. The ENC is utilized by the participants involved in the life cycle of a J2EE component in the following ways.

- Application component business logic should be coded to access information from its ENC. The component provider uses the standard deployment descriptor for the component to specify the required ENC entries. The entries are declarations of the information and resources the component requires at runtime.
- The container provides tools that allow a deployer of a component to map the ENC references made by the component developer to the deployment environment entity that satisfies the reference.
- The component deployer utilizes the container tools to ready a component for final deployment.
- The component container uses the deployment package information to build the complete component ENC at runtime

The complete specification regarding the use of JNDI in the J2EE platform can be found in Section 5 of the J2EE 1.3 specification. The J2EE specification is available at <http://java.sun.com/j2ee/download.html>.

An application component instance locates the ENC using the JNDI API. An application component instance creates a `javax.naming.InitialContext` object by using the no argument constructor and then looks up the naming environment under the name `java:comp/env`. The application component's environment entries are stored directly in the ENC, or in its subcontexts. Example 3.2 illustrates the prototypical lines of code a component uses to access its ENC.

Example 3.2. ENC access sample code

```
// Obtain the application component's ENC
Context iniCtx = new InitialContext();
Context compEnv = (Context) iniCtx.lookup("java:comp/env");
```

An application component environment is a local environment that is accessible only by the component when the application server container thread of control is interacting with the application component. This means that an EJB `Bean1` cannot access the ENC elements of EJB `Bean2`, and visa-versa. Similarly, Web application `Web1` cannot access the ENC elements of Web application `Web2` or `Bean1` or `Bean2` for that matter. Also, arbitrary client code, whether it is executing inside of the application server VM or externally cannot access a component's `java:comp` JNDI context. The purpose of the ENC is to provide an isolated, read-only namespace that the application component can rely on regardless of the type of environment in which the component is deployed. The ENC must be isolated from other components because each component defines its own ENC content. Components `A` and `B`, for example, may define the same name to refer to different objects. For example, EJB `Bean1` may define an environment entry `java:comp/env/red` to refer to the hexadecimal value for the RGB color for red, while Web application `Web1` may bind the same name to the deployment environment language locale representation of red.

There are three commonly used levels of naming scope in JBoss: names under `java:comp`, names under `java:`, and any other name. As discussed, the `java:comp` context and its subcontexts are only available to the application component associated with that particular context. Subcontexts and object bindings directly under `java:` are only visible within the JBoss server virtual machine and not to remote clients. Any other context or object binding is available to remote clients, provided the context or object supports serialization. You'll see how the isolation of these naming scopes is achieved in the Section 3.2.

An example of where the restricting a binding to the `java:` context is useful would be a `javax.sql.DataSource` connection factory that can only be used inside of the JBoss server where the associated database pool resides. On the other hand, an EJB home interface would be bound to a globally visible name that should be accessible by remote client.

3.1.2.1. ENC Usage Conventions

JNDI is used as the API for externalizing a great deal of information from an application component. The JNDI name that the application component uses to access the information is declared in the standard `ejb-jar.xml` deployment descriptor for EJB components, and the standard `web.xml` deployment descriptor for Web components. Several different types of information may be stored in and retrieved from JNDI including:

- Environment entries as declared by the `env-entry` elements
- EJB references as declared by `ejb-ref` and `ejb-local-ref` elements.

- Resource manager connection factory references as declared by the `resource-ref` elements
- Resource environment references as declared by the `resource-env-ref` elements

Each type of deployment descriptor element has a JNDI usage convention with regard to the name of the JNDI context under which the information is bound. Also, in addition to the standard deployment descriptor element, there is a JBoss server specific deployment descriptor element that maps the JNDI name as used by the application component to the deployment environment JNDI name.

3.1.2.1.1. The `ejb-jar.xml` ENC Elements

The EJB 2.0 deployment descriptor describes a collection of EJB components and their environment. Each of the three types of EJB components (session, entity, and message-driven) support the specification of an EJB local naming context. The `ejb-jar.xml` description is a logical view of the environment that the EJB needs to operate. Because the EJB component developer generally cannot know into what environment the EJB will be deployed, the developer describes the component environment in a deployment environment independent manner using logical names. It is the responsibility of a deployment administrator to link the EJB component logical names to the corresponding deployment environment resources.

Figure 3.1 gives a graphical view of the EJB deployment descriptor DTD without the non-ENC elements. Only the session element is shown fully expanded as the ENC elements for entity and message-driven are identical. The full `ejb-jar.xml` DTD is available from the Sun web site at http://java.sun.com/dtd/ejb-jar_2_0.dtd.

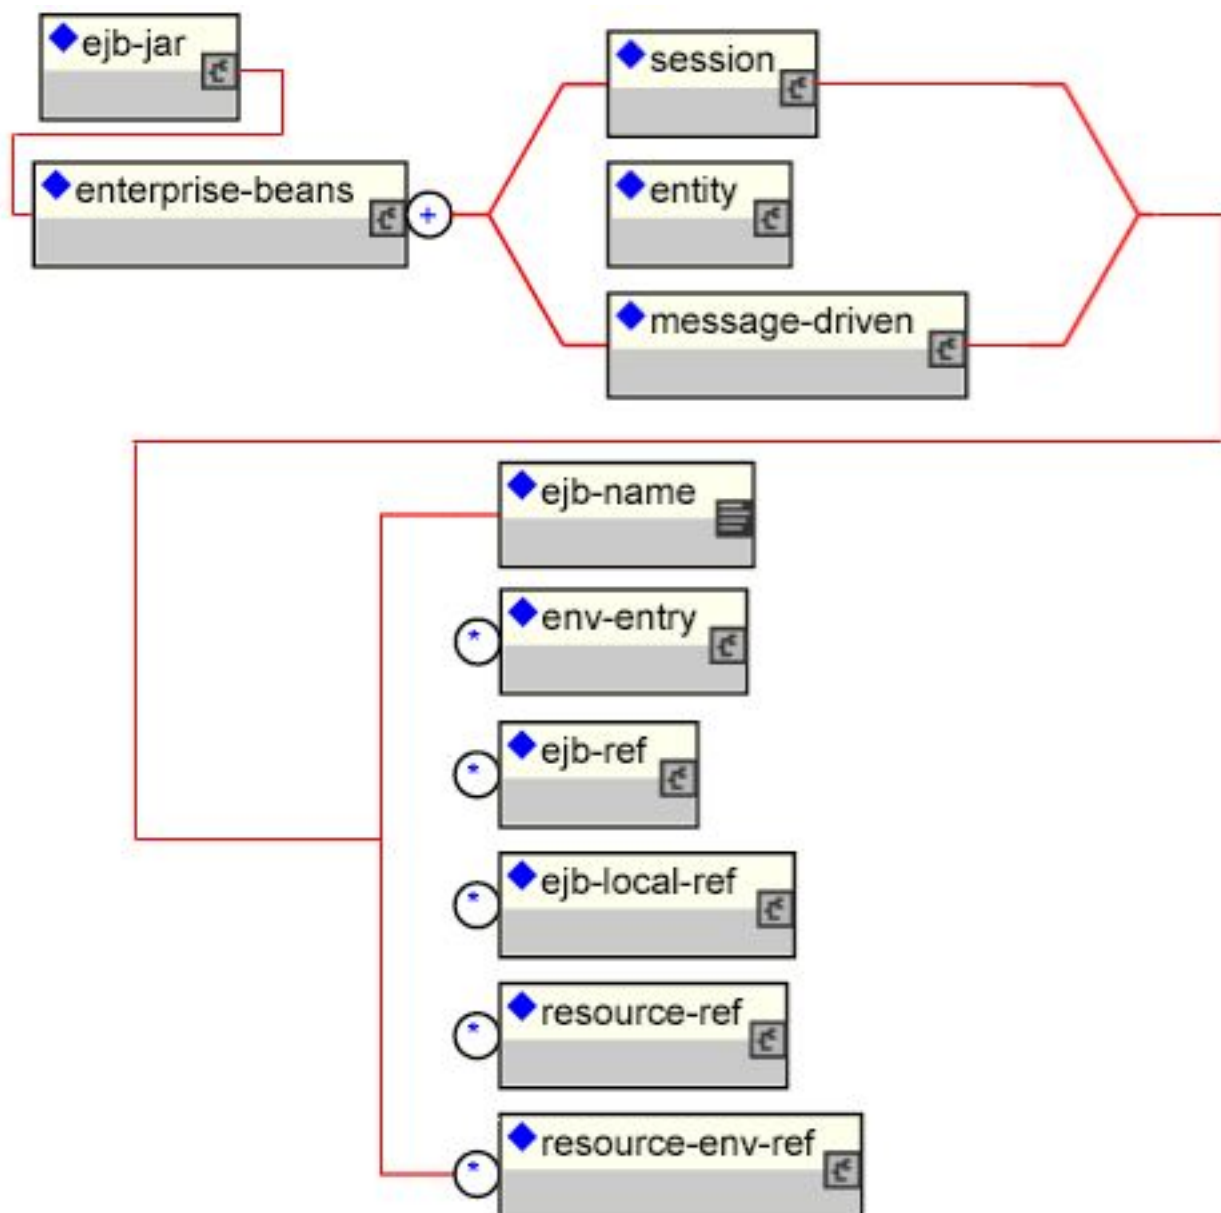


Figure 3.1. The ENC elements in the standard ejb-jar.xml 2.0 deployment descriptor.

3.1.2.1.2. The web.xml ENC Elements

The Servlet 2.3 deployment descriptor describes a collection of Web components and their environment. The ENC for a Web application is declared globally for all servlets and JSP pages in the Web application. Because the Web application developer generally cannot know into what environment the Web application will be deployed, the developer describes the component environment in a deployment environment independent manner using logical names. It is the responsibility of a deployment administrator to link the Web component logical names to the corresponding deployment environment resources.

Figure 3.2 gives a graphical view of the Web application deployment descriptor DTD without the non-ENC elements. The full web.xml DTD is available from the Sun Web site at http://java.sun.com/dtd/web-app_2_3.dtd.




Figure 3.2. The ENC elements in the standard servlet 2.3 web.xml deployment descriptor.

3.1.2.1.3. The jboss.xml ENC Elements

The JBoss EJB deployment descriptor provides the mapping from the EJB component ENC JNDI names to the actual deployed JNDI name. It is the responsibility of the application deployer to map the logical references made by the application component to the corresponding physical resource deployed in a given application server configuration. In JBoss, this is done for the `ejb-jar.xml` descriptor using the `jboss.xml` deployment descriptor. Figure 3.3 gives a graphical view of the JBoss EJB deployment descriptor DTD without the non-ENC elements. This is virtually identical to the corresponding elements of the `ejb-jar.xml` for the levels shown.

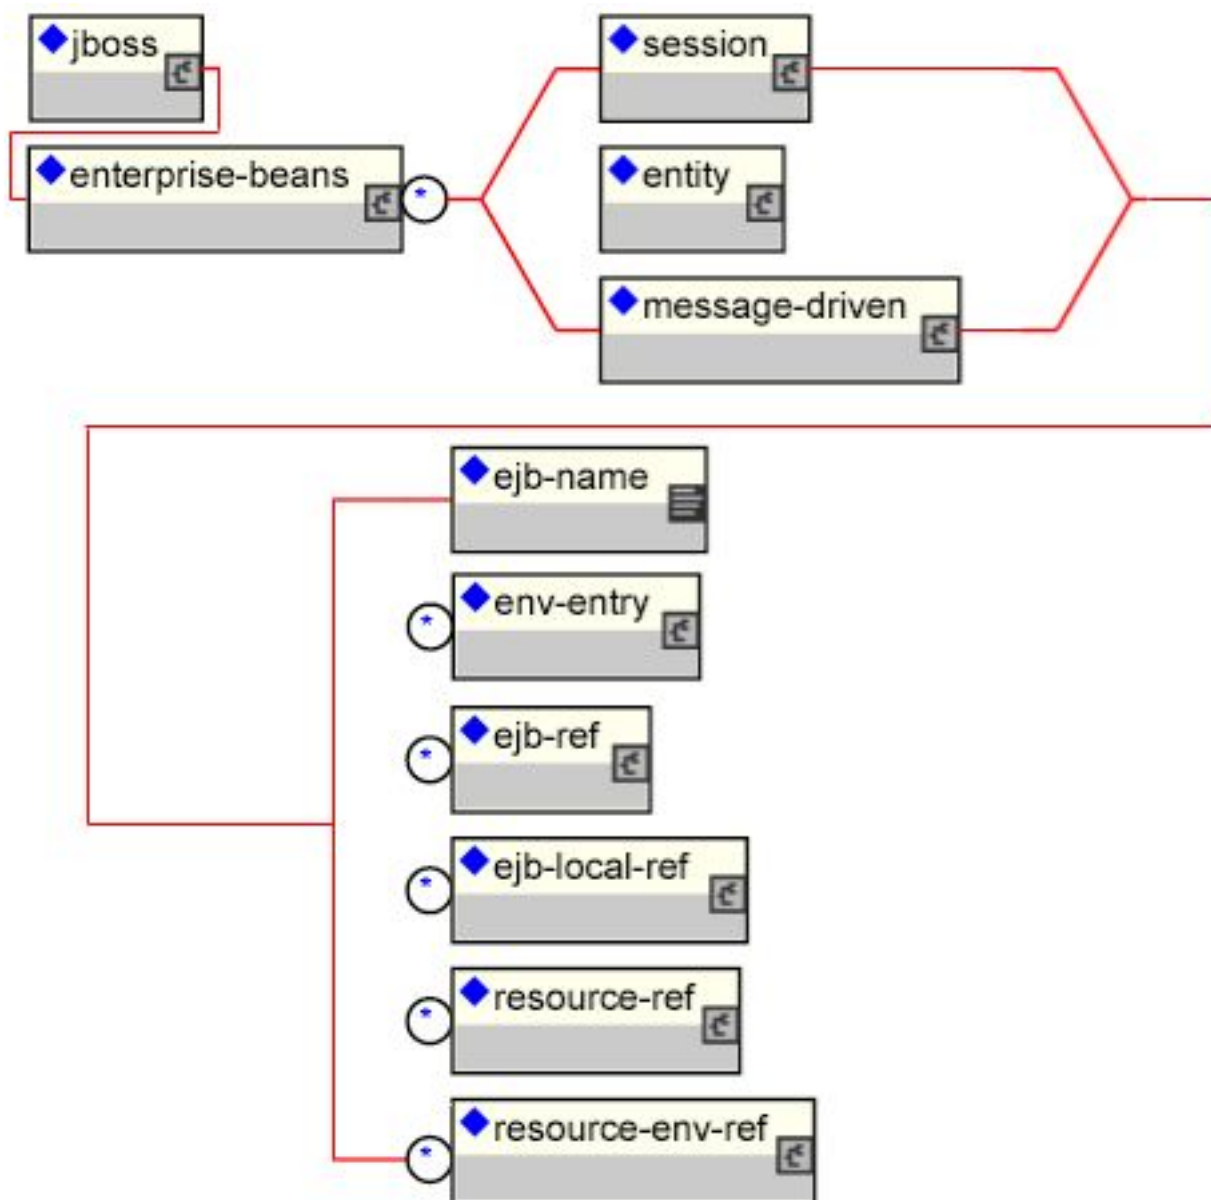


Figure 3.3. The ENC elements in the JBoss 3.2 jboss.xml deployment descriptor.

3.1.2.1.4. The jboss-web.xml ENC Elements

The JBoss Web deployment descriptor provides the mapping from the Web application ENC JNDI names to the actual deployed JNDI name. It is the responsibility of the application deployer to map the logical references made by the Web application to the corresponding physical resource deployed in a given application server configuration. In JBoss, this is done for the web.xml descriptor using the jboss-web.xml deployment descriptor. Figure 3.4 gives a graphical view of the JBoss Web deployment descriptor DTD without the non-ENC elements. The full jboss-web.xml DTD is available from the JBoss Web site at http://www.jboss.org/j2ee/dtd/jboss_web_3_2.dtd as well as the docs/dtd directory of the distribution.

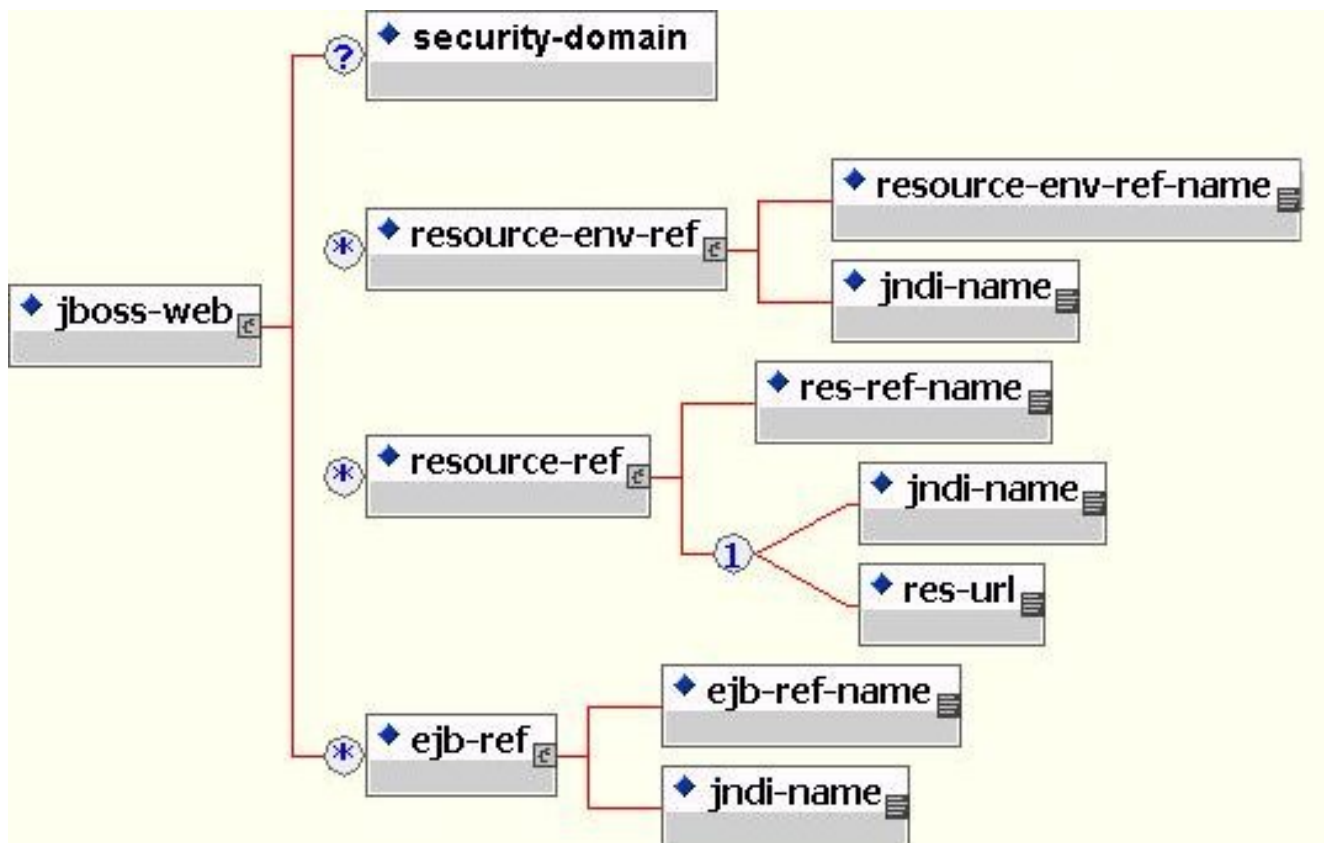


Figure 3.4. ENC elements in the JBoss 3.2 jboss-web.xml deployment descriptor.

3.1.2.1.5. Environment Entries

Environment entries are the simplest form of information stored in a component ENC, and are similar to operating system environment variables like those found on UNIX or Windows. Environment entries are a name-to-value binding that allows a component to externalize a value and refer to the value using a name.

An environment entry is declared using an `env-entry` element in the standard deployment descriptors. The `env-entry` element contains the following child elements:

- An optional `description` element that provides a description of the entry
- An `env-entry-name` element giving the name of the entry relative to `java:comp/env`
- An `env-entry-type` element giving the Java type of the entry value that must be one of:
 - `java.lang.Byte`
 - `java.lang.Boolean`
 - `java.lang.Character`
 - `java.lang.Double`
 - `java.lang.Float`
 - `java.lang.Integer`

- `java.lang.Long`
- `java.lang.Short`
- `java.lang.String`
- An `env-entry-value` element giving the value of entry as a string

An example of an `env-entry` fragment from an `ejb-jar.xml` deployment descriptor is given in Example 3.3. There is no JBoss specific deployment descriptor element because an `env-entry` is a complete name and value specification. Example 3.4 shows a sample code fragment for accessing the `maxExemptions` and `taxRate` `env-entry` values declared in the deployment descriptor.

Example 3.3. An example `ejb-jar.xml` `env-entry` fragment

```
<!-- ... -->
<session>
  <ejb-name>ASessionBean</ejb-name>
  <!-- ... -->
  <env-entry>
    <description>The maximum number of tax exemptions allowed </description>
    <env-entry-name>maxExemptions</env-entry-name>
    <env-entry-type>java.lang.Integer</env-entry-type>
    <env-entry-value>15</env-entry-value>
  </env-entry>
  <env-entry>
    <description>The tax rate </description>
    <env-entry-name>taxRate</env-entry-name>
    <env-entry-type>java.lang.Float</env-entry-type>
    <env-entry-value>0.23</env-entry-value>
  </env-entry>
</session>
<!-- ... -->
```

Example 3.4. ENC `env-entry` access code fragment

```
InitialContext iniCtx = new InitialContext();
Context envCtx = (Context) iniCtx.lookup("java:comp/env");
Integer maxExemptions = (Integer) envCtx.lookup("maxExemptions");
Float taxRate = (Float) envCtx.lookup("taxRate");
```

3.1.2.1.6. EJB References

It is common for EJBs and Web components to interact with other EJBs. Because the JNDI name under which an EJB home interface is bound is a deployment time decision, there needs to be a way for a component developer to declare a reference to an EJB that will be linked by the deployer. EJB references satisfy this requirement.

An EJB reference is a link in an application component naming environment that points to a deployed EJB home interface. The name used by the application component is a logical link that isolates the component from the actual name of the EJB home in the deployment environment. The J2EE specification recommends that all references to enterprise beans be organized in the `java:comp/env/ejb` context of the application component's environment.

An EJB reference is declared using an `ejb-ref` element in the deployment descriptor. Each `ejb-ref` element describes the interface requirements that the referencing application component has for the referenced enterprise bean. The `ejb-ref` element contains the following child elements:

- An optional `description` element that provides the purpose of the reference.
- An `ejb-ref-name` element that specifies the name of the reference relative to the `java:comp/env` context. To place the reference under the recommended `java:comp/env/ejb` context, use an `ejb/link-name` form for the `ejb-ref-name` value.
- An `ejb-ref-type` element that specifies the type of the EJB. This must be either `Entity` or `Session`.
- A `home` element that gives the fully qualified class name of the EJB home interface.
- A `remote` element that gives the fully qualified class name of the EJB remote interface.
- An optional `ejb-link` element that links the reference to another enterprise bean in the same EJB JAR or in the same J2EE application unit. The `ejb-link` value is the `ejb-name` of the referenced bean. If there are multiple enterprise beans with the same `ejb-name`, the value uses the path name specifying the location of the `ejb-jar` file that contains the referenced component. The path name is relative to the referencing `ejb-jar` file. The Application Assembler appends the `ejb-name` of the referenced bean to the path name separated by `#`. This allows multiple beans with the same name to be uniquely identified.

An EJB reference is scoped to the application component whose declaration contains the `ejb-ref` element. This means that the EJB reference is not accessible from other application components at runtime, and that other application components may define `ejb-ref` elements with the same `ejb-ref-name` without causing a name conflict. Example 3.5 provides an `ejb-jar.xml` fragment that illustrates the use of the `ejb-ref` element. A code sample that illustrates accessing the `ShoppingCartHome` reference declared in Example 3.5 is given in Example 3.6.

Example 3.5. An example `ejb-jar.xml` `ejb-ref` descriptor fragment

```
<!-- ... -->
<session>
  <ejb-name>ShoppingCartBean</ejb-name>
  <!-- ...-->
</session>

<session>
  <ejb-name>ProductBeanUser</ejb-name>
  <!-- ...-->
  <ejb-ref>
    <description>This is a reference to the store products entity </description>
    <ejb-ref-name>ejb/ProductHome</ejb-ref-name>
    <ejb-ref-type>Entity</ejb-ref-type>
    <home>org.jboss.store.ejb.ProductHome</home>
  </ejb-ref>
  <remote> org.jboss.store.ejb.Product</remote>
</session>

<session>
  <ejb-ref>
    <ejb-name>ShoppingCartUser</ejb-name>
    <!-- ...-->
    <ejb-ref-name>ejb/ShoppingCartHome</ejb-ref-name>
    <ejb-ref-type>Session</ejb-ref-type>
    <home>org.jboss.store.ejb.ShoppingCartHome</home>
    <remote> org.jboss.store.ejb.ShoppingCart</remote>
    <ejb-link>ShoppingCartBean</ejb-link>
```

```
</ejb-ref>
</session>

<entity>
  <description>The Product entity bean </description>
  <ejb-name>ProductBean</ejb-name>
  <!--...-->
</entity>

<!--...-->
```

Example 3.6. ENC ejb-ref access code fragment

```
InitialContext iniCtx = new InitialContext();
Context ejbCtx = (Context) iniCtx.lookup("java:comp/env/ejb");
ShoppingCartHome home = (ShoppingCartHome) ejbCtx.lookup("ShoppingCartHome");
```

3.1.2.1.7. EJB References with `jboss.xml` and `jboss-web.xml`

The JBoss specific `jboss.xml` EJB deployment descriptor affects EJB references in two ways. First, the `jndi-name` child element of the `session` and `entity` elements allows the user to specify the deployment JNDI name for the EJB home interface. In the absence of a `jboss.xml` specification of the `jndi-name` for an EJB, the home interface is bound under the `ejb-jar.xml` `ejb-name` value. For example, the session EJB with the `ejb-name` of `ShoppingCartBean` in Example 3.5 would have its home interface bound under the JNDI name `ShoppingCartBean` in the absence of a `jboss.xml` `jndi-name` specification.

The second use of the `jboss.xml` descriptor with respect to `ejb-refs` is the setting of the destination to which a component's ENC `ejb-ref` refers. The `ejb-link` element cannot be used to refer to EJBs in another enterprise application. If your `ejb-ref` needs to access an external EJB, you can specify the JNDI name of the deployed EJB home using the `jboss.xml` `ejb-ref/jndi-name` element.

The `jboss-web.xml` descriptor is used only to set the destination to which a Web application ENC `ejb-ref` refers. The content model for the JBoss `ejb-ref` is as follows:

- An `ejb-ref-name` element that corresponds to the `ejb-ref-name` element in the `ejb-jar.xml` or `web.xml` standard descriptor
- A `jndi-name` element that specifies the JNDI name of the EJB home interface in the deployment environment

Example 3.7 provides an example `jboss.xml` descriptor fragment that illustrates the following usage points:

- The `ProductBeanUser` `ejb-ref` link destination is set to the deployment name of `jboss/store/ProductHome`
- The deployment JNDI name of the `ProductBean` is set to `jboss/store/ProductHome`

Example 3.7. An example `jboss.xml` `ejb-ref` fragment

```
<!-- ... -->
<session>
  <ejb-name>ProductBeanUser</ejb-name>
  <ejb-ref>
```

```
<ejb-ref-name>ejb/ProductHome</ejb-ref-name>
<jndi-name>jboss/store/ProductHome</jndi-name>
</ejb-ref>
</session>

<entity>
  <ejb-name>ProductBean</ejb-name>
  <jndi-name>jboss/store/ProductHome</jndi-name>
  <!-- ... -->
</entity>
<!-- ... -->
```

3.1.2.1.8. EJB Local References

EJB 2.0 added local interfaces that do not use RMI call by value semantics. These interfaces use a call by reference semantic and therefore do not incur any RMI serialization overhead. An EJB local reference is a link in an application component naming environment that points to a deployed EJB local home interface. The name used by the application component is a logical link that isolates the component from the actual name of the EJB local home in the deployment environment. The J2EE specification recommends that all references to enterprise beans be organized in the `java:comp/env/ejb` context of the application component's environment.

An EJB local reference is declared using an `ejb-local-ref` element in the deployment descriptor. Each `ejb-local-ref` element describes the interface requirements that the referencing application component has for the referenced enterprise bean. The `ejb-local-ref` element contains the following child elements:

- An optional `description` element that provides the purpose of the reference.
- An `ejb-ref-name` element that specifies the name of the reference relative to the `java:comp/env` context. To place the reference under the recommended `java:comp/env/ejb` context, use an `ejb/link-name` form for the `ejb-ref-name` value.
- An `ejb-ref-type` element that specifies the type of the EJB. This must be either `Entity` or `Session`.
- A `local-home` element that gives the fully qualified class name of the EJB local home interface.
- A `local` element that gives the fully qualified class name of the EJB local interface.
- An `ejb-link` element that links the reference to another enterprise bean in the `ejb-jar` file or in the same J2EE application unit. The `ejb-link` value is the `ejb-name` of the referenced bean. If there are multiple enterprise beans with the same `ejb-name`, the value uses the path name specifying the location of the `ejb-jar` file that contains the referenced component. The path name is relative to the referencing `ejb-jar` file. The Application Assembler appends the `ejb-name` of the referenced bean to the path name separated by `#`. This allows multiple beans with the same name to be uniquely identified. An `ejb-link` element must be specified in JBoss to match the local reference to the corresponding EJB.

An EJB local reference is scoped to the application component whose declaration contains the `ejb-local-ref` element. This means that the EJB local reference is not accessible from other application components at runtime, and that other application components may define `ejb-local-ref` elements with the same `ejb-ref-name` without causing a name conflict. Example 3.8 provides an `ejb-jar.xml` fragment that illustrates the use of the `ejb-local-ref` element. A code sample that illustrates accessing the `ProbeLocalHome` reference declared in Example 3.8 is given in Example 3.9.

Example 3.8. An example `ejb-jar.xml` `ejb-local-ref` descriptor fragment

```

<!-- ... -->
<session>
  <ejb-name>Probe</ejb-name>
  <home>org.jboss.test.perf.interfaces.ProbeHome</home>
  <remote>org.jboss.test.perf.interfaces.Probe</remote>
  <local-home>org.jboss.test.perf.interfaces.ProbeLocalHome</local-home>
  <local>org.jboss.test.perf.interfaces.ProbeLocal</local>
  <ejb-class>org.jboss.test.perf.ejb.ProbeBean</ejb-class>
  <session-type>Stateless</session-type>
  <transaction-type>Bean</transaction-type>
</session>
<session>
  <ejb-name>PerfTestSession</ejb-name>
  <home>org.jboss.test.perf.interfaces.PerfTestSessionHome</home>
  <remote>org.jboss.test.perf.interfaces.PerfTestSession</remote>
  <ejb-class>org.jboss.test.perf.ejb.PerfTestSessionBean</ejb-class>
  <session-type>Stateless</session-type>
  <transaction-type>Container</transaction-type>
  <ejb-ref>
    <ejb-ref-name>ejb/ProbeHome</ejb-ref-name>
    <ejb-ref-type>Session</ejb-ref-type>
    <home>org.jboss.test.perf.interfaces.SessionHome</home>
    <remote>org.jboss.test.perf.interfaces.Session</remote>
    <ejb-link>Probe</ejb-link>
  </ejb-ref>
  <ejb-local-ref>
    <ejb-ref-name>ejb/ProbeLocalHome</ejb-ref-name>
    <ejb-ref-type>Session</ejb-ref-type>
    <local-home>org.jboss.test.perf.interfaces.ProbeLocalHome</local-home>
    <local>org.jboss.test.perf.interfaces.ProbeLocal</local>
    <ejb-link>Probe</ejb-link>
  </ejb-local-ref>
</session>
<!-- ... -->

```

Example 3.9. ENC ejb-local-ref access code fragment

```

InitialContext iniCtx = new InitialContext();
Context ejbCtx = (Context) iniCtx.lookup("java:comp/env/ejb");
ProbeLocalHome home = (ProbeLocalHome) ejbCtx.lookup("ProbeLocalHome");

```

3.1.2.1.9. Resource Manager Connection Factory References

Resource manager connection factory references allow application component code to refer to resource factories using logical names called resource manager connection factory references. Resource manager connection factory references are defined by the `resource-ref` elements in the standard deployment descriptors. The `Deployer` binds the resource manager connection factory references to the actual resource manager connection factories that exist in the target operational environment using the `jboss.xml` and `jboss-web.xml` descriptors.

Each `resource-ref` element describes a single resource manager connection factory reference. The `resource-ref` element consists of the following child elements:

- An optional `description` element that provides the purpose of the reference.
- A `res-ref-name` element that specifies the name of the reference relative to the `java:comp/env` context. The resource type based naming convention for which subcontext to place the `res-ref-name` into is discussed in the next paragraph.
- A `res-type` element that specifies the fully qualified class name of the resource manager connection fact-

ory.

- A `res-auth` element that indicates whether the application component code performs resource signon programmatically, or whether the container signs on to the resource based on the principal mapping information supplied by the Deployer. It must be one of `Application` or `Container`.
- An optional `res-sharing-scope` element. This currently is not supported by JBoss.

The J2EE specification recommends that all resource manager connection factory references be organized in the subcontexts of the application component's environment, using a different subcontext for each resource manager type. The recommended resource manager type to subcontext name is as follows:

- JDBC `DataSource` references should be declared in the `java:comp/env/jdbc` subcontext.
- JMS connection factories should be declared in the `java:comp/env/jms` subcontext.
- JavaMail connection factories should be declared in the `java:comp/env/mail` subcontext.
- URL connection factories should be declared in the `java:comp/env/url` subcontext.

Example 3.10 shows an example `web.xml` descriptor fragment that illustrates the `resource-ref` element usage. Example 3.11 provides a code fragment that an application component would use to access the `DefaultMail` resource declared by the `resource-ref`.

Example 3.10. A `web.xml` `resource-ref` descriptor fragment

```
<web>
  <!-- ... -->
  <servlet>
    <servlet-name>AServlet</servlet-name>
    <!-- ... -->
  </servlet>
  <!-- ... -->
  <!-- JDBC DataSources ( java:comp/env/jdbc ) -->
  <resource-ref>
    <description>The default DS</description>
    <res-ref-name>jdbc/DefaultDS</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
  </resource-ref>
  <!-- JavaMail Connection Factories ( java:comp/env/mail ) -->
  <resource-ref>
    <description>Default Mail</description>
    <res-ref-name>mail/DefaultMail</res-ref-name>
    <res-type>javax.mail.Session</res-type>
    <res-auth>Container</res-auth>
  </resource-ref>
  <!-- JMS Connection Factories ( java:comp/env/jms ) -->
  <resource-ref>
    <description>Default QueueFactory</description>
    <res-ref-name>jms/QueueFactory</res-ref-name>
    <res-type>javax.jms.QueueConnectionFactory</res-type>
    <res-auth>Container</res-auth>
  </resource-ref>
</web>
```

Example 3.11. ENC `resource-ref` access sample code fragment

```
Context initCtx = new InitialContext();
```

```
javax.mail.Session s = (javax.mail.Session)
initCtx.lookup("java:comp/env/mail/DefaultMail");
```

3.1.2.1.10. Resource Manager Connection Factory References with `jboss.xml` and `jboss-web.xml`

The purpose of the JBoss `jboss.xml` EJB deployment descriptor and `jboss-web.xml` Web application deployment descriptor is to provide the link from the logical name defined by the `res-ref-name` element to the JNDI name of the resource factory as deployed in JBoss. This is accomplished by providing a `resource-ref` element in the `jboss.xml` or `jboss-web.xml` descriptor. The JBoss `resource-ref` element consists of the following child elements:

- A `res-ref-name` element that must match the `res-ref-name` of a corresponding `resource-ref` element from the `ejb-jar.xml` or `web.xml` standard descriptors
- An optional `res-type` element that specifies the fully qualified class name of the resource manager connection factory
- A `jndi-name` element that specifies the JNDI name of the resource factory as deployed in JBoss
- A `res-url` element that specifies the URL string in the case of a `resource-ref` of type `java.net.URL`

Example 3.12 provides a sample `jboss-web.xml` descriptor fragment that shows sample mappings of the `resource-ref` elements given in Example 3.10.

Example 3.12. A sample `jboss-web.xml` resource-ref descriptor fragment

```
<jboss-web>
  <!-- ... -->
  <resource-ref>
    <res-ref-name>jdbc/DefaultDS</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <jndi-name>java:/DefaultDS</jndi-name>
  </resource-ref>
  <resource-ref>
    <res-ref-name>mail/DefaultMail</res-ref-name>
    <res-type>javax.mail.Session</res-type>
    <jndi-name>java:/Mail</jndi-name>
  </resource-ref>
  <resource-ref>
    <res-ref-name>jms/QueueFactory</res-ref-name>
    <res-type>javax.jms.QueueConnectionFactory</res-type>
    <jndi-name>QueueConnectionFactory</jndi-name>
  </resource-ref>
  <!-- ... -->
</jboss-web>
```

3.1.2.1.11. Resource Environment References

Resource environment references are elements that refer to administered objects that are associated with a resource (for example, JMS destinations) using logical names. Resource environment references are defined by the `resource-env-ref` elements in the standard deployment descriptors. The `Deployer` binds the resource environment references to the actual administered objects location in the target operational environment using the `jboss.xml` and `jboss-web.xml` descriptors.

Each `resource-env-ref` element describes the requirements that the referencing application component has for the referenced administered object. The `resource-env-ref` element consists of the following child elements:

- An optional `description` element that provides the purpose of the reference.
- A `resource-env-ref-name` element that specifies the name of the reference relative to the `java:comp/env` context. Convention places the name in a subcontext that corresponds to the associated resource factory type. For example, a JMS queue reference named `MyQueue` should have a `resource-env-ref-name` of `jms/MyQueue`.
- A `resource-env-ref-type` element that specifies the fully qualified class name of the referenced object. For example, in the case of a JMS queue, the value would be `javax.jms.Queue`.

Example 3.13 provides an example `resource-ref-env` element declaration by a session bean. Example 3.14 gives a code fragment that illustrates how to look up the `StockInfo` queue declared by the `resource-env-ref`.

Example 3.13. An example `ejb-jar.xml` `resource-env-ref` fragment

```
<session>
  <ejb-name>MyBean</ejb-name>

  <resource-env-ref>
    <description>This is a reference to a JMS queue used in the
      processing of Stock info
    </description>
    <resource-env-ref-name>jms/StockInfo</resource-env-ref-name>
    <resource-env-ref-type>javax.jms.Queue</resource-env-ref-type>
  </resource-env-ref>
  <!-- ... -->
</session>
```

Example 3.14. ENC `resource-env-ref` access code fragment

```
InitialContext iniCtx = new InitialContext();
javax.jms.Queue q = (javax.jms.Queue)
envCtx.lookup("java:comp/env/jms/StockInfo");
```

3.1.2.1.12. Resource Environment References and `jboss.xml`, `jboss-web.xml`

The purpose of the JBoss `jboss.xml` EJB deployment descriptor and `jboss-web.xml` Web application deployment descriptor is to provide the link from the logical name defined by the `resource-env-ref-name` element to the JNDI name of the administered object deployed in JBoss. This is accomplished by providing a `resource-env-ref` element in the `jboss.xml` or `jboss-web.xml` descriptor. The JBoss `resource-env-ref` element consists of the following child elements:

- A `resource-env-ref-name` element that must match the `resource-env-ref-name` of a corresponding `resource-env-ref` element from the `ejb-jar.xml` or `web.xml` standard descriptors
- A `jndi-name` element that specifies the JNDI name of the resource as deployed in JBoss

Example 3.15 provides a sample `jboss.xml` descriptor fragment that shows a sample mapping for the `StockInfo` `resource-env-ref`.

Example 3.15. A sample jboss.xml resource-env-ref descriptor fragment

```

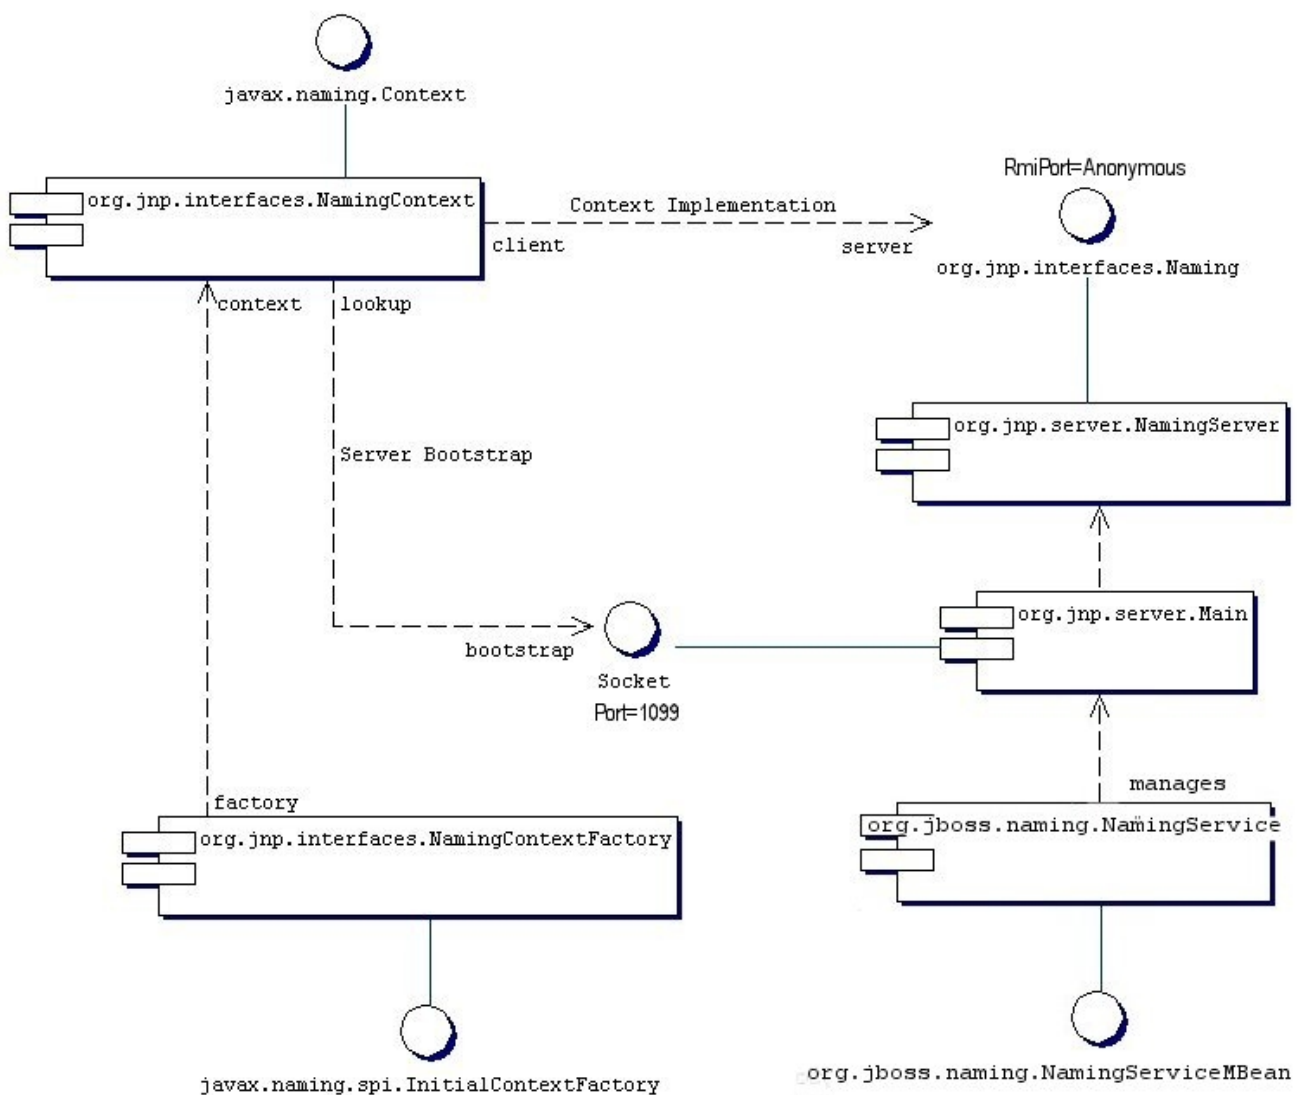
<session>
  <ejb-name>MyBean</ejb-name>

  <resource-env-ref>
    <resource-env-ref-name>jms/StockInfo</resource-env-ref-name>
    <jndi-name>queue/StockInfoQueue</jndi-name>
  </resource-env-ref>
  <!-- ... -->
</session>

```

3.2. The JBossNS Architecture

The JBossNS architecture is a Java socket/RMI based implementation of the `javax.naming.Context` interface. It is a client/server implementation that can be accessed remotely. The implementation is optimized so that access from within the same VM in which the JBossNS server is running does not involve sockets. Same VM access occurs through an object reference available as a global singleton. Figure 3.5 illustrates some of the key classes in the JBossNS implementation and their relationships.

**Figure 3.5. Key components in the JBossNS architecture.**

We will start with the `NamingService` MBean. The `NamingService` MBean provides the JNDI naming service. This is a key service used pervasively by the J2EE technology components. The configurable attributes for the `NamingService` are as follows.

- **Port:** The `jnp` protocol listening port for the `NamingService`. If not specified default is 1099, the same as the RMI registry default port.
- **RmiPort:** The RMI port on which the RMI Naming implementation will be exported. If not specified the default is 0 which means use any available port.
- **BindAddress:** The specific address the `NamingService` listens on. This can be used on a multi-homed host for a `java.net.ServerSocket` that will only accept connect requests on one of its addresses.
- **RmiBindAddress:** The specific address the RMI server portion of the `NamingService` listens on. This can be used on a multi-homed host for a `java.net.ServerSocket` that will only accept connect requests on one of its addresses. If this is not specified and the `BindAddress` is, the `RmiBindAddress` defaults to the `BindAddress` value.
- **Backlog:** The maximum queue length for incoming connection indications (a request to connect) is set to the `backlog` parameter. If a connection indication arrives when the queue is full, the connection is refused.
- **ClientSocketFactory:** An optional custom `java.rmi.server.RMIClientSocketFactory` implementation class name. If not specified the default `RMIClientSocketFactory` is used.
- **ServerSocketFactory:** An optional custom `java.rmi.server.RMIServerSocketFactory` implementation class name. If not specified the default `RMIServerSocketFactory` is used.
- **JNPServerSocketFactory:** An optional custom `javax.net.ServerSocketFactory` implementation class name. This is the factory for the `ServerSocket` used to bootstrap the download of the `JBossNS` Naming interface. If not specified the `javax.net.ServerSocketFactory.getDefault()` method value is used.

The `NamingService` also creates the `java:comp` context such that access to this context is isolated based on the context class loader of the thread that accesses the `java:comp` context. This provides the application component private ENC that is required by the J2EE specs. This segregation is accomplished by binding a `javax.naming.Reference` to a context that uses the `org.jboss.naming.ENCFactory` as its `javax.naming.ObjectFactory`. When a client performs a lookup of `java:comp`, or any subcontext, the `ENCFactory` checks the thread context `ClassLoader`, and performs a lookup into a map using the `ClassLoader` as the key.

If a context instance does not exist for the class loader instance, one is created and associated with that class loader in the `ENCFactory` map. Thus, correct isolation of an application component's ENC relies on each component receiving a unique `ClassLoader` that is associated with the component threads of execution.

The `NamingService` delegates its functionality to an `org.jnp.server.Main` MBean. The reason for the duplicate MBeans is because `JBossNS` started out as a stand-alone JNDI implementation, and can still be run as such. The `NamingService` MBean embeds the `Main` instance into the JBoss server so that usage of JNDI with the same VM as the JBoss server does not incur any socket overhead. The configurable attributes of the `NamingService` are really the configurable attributes of the `JBossNS` `Main` MBean. The setting of any attributes on the `NamingService` MBean simply set the corresponding attributes on the `Main` MBean the `NamingService` contains. When the `NamingService` is started, it starts the contained `Main` MBean to activate the JNDI naming service.

In addition, the `NamingService` exposes the `Naming` interface operations through a JMX detyped invoke operation. This allows the naming service to be accessed via JMX adaptors for arbitrary protocols. We will look at an

example of how HTTP can be used to access the naming service using the `invoke` operation later in this chapter.

The details of threads and the thread context class loader won't be explored here, but the JNDI tutorial provides a concise discussion that is applicable. See <http://java.sun.com/products/jndi/tutorial/beyond/misc/classloader.html> for the details.

When the `Main MBean` is started, it performs the following tasks:

- Instantiates an `org.jnp.naming.NamingService` instance and sets this as the local VM server instance. This is used by any `org.jnp.interfaces.NamingContext` instances that are created within the JBoss server VM to avoid RMI calls over TCP/IP.
- Exports the `NamingServer` instance's `org.jnp.naming.interfaces.Naming` RMI interface using the configured `RmiPort`, `ClientSocketFactory`, `ServerSocketFactory` attributes.
- Creates a socket that listens on the interface given by the `BindAddress` and `Port` attributes.
- Spawns a thread to accept connections on the socket.

3.2.1. The Naming InitialContext Factories

The JBoss JNDI provider currently supports three different `InitialContext` factory implementations. The most commonly used factory is the `org.jnp.interfaces.NamingContextFactory` implementation. Its properties include:

- **java.naming.factory.initial:** The name of the environment property for specifying the initial context factory to use. The value of the property should be the fully qualified class name of the factory class that will create an initial context. If it is not specified, a `javax.naming.NoInitialContextException` will be thrown when an `InitialContext` object is created.
- **java.naming.provider.url:** The name of the environment property for specifying the location of the JBoss JNDI service provider the client will use. The `NamingContextFactory` class uses this information to know which JBossNS server to connect to. The value of the property should be a URL string. For JBossNS the URL format is `jnp://host:port/[jndi_path]`. The `jnp:` portion of the URL is the protocol and refers to the socket/RMI based protocol used by JBoss. The `jndi_path` portion of the URL is an optional JNDI name relative to the root context, for example, `apps` or `apps/tmp`. Everything but the host component is optional. The following examples are equivalent because the default port value is 1099.
 - `jnp://www.jboss.org:1099/`
 - `www.jboss.org:1099`
 - `www.jboss.org`
- **java.naming.factory.url.pkgs:** The name of the environment property for specifying the list of package prefixes to use when loading in URL context factories. The value of the property should be a colon-separated list of package prefixes for the class name of the factory class that will create a URL context factory. For the JBoss JNDI provider this must be `org.jboss.naming:org.jnp.interfaces`. This property is essential for locating the `jnp:` and `java:` URL context factories of the JBoss JNDI provider.
- **jnp.socketFactory:** The fully qualified class name of the `javax.net.SocketFactory` implementation to use to create the bootstrap socket. The default value is `org.jnp.interfaces.TimedSocketFactory`. The `TimedSocketFactory` is a simple `SocketFactory` implementation that supports the specification of a connection and read timeout. These two properties are specified by:

- **jnp.timeout:** The connection timeout in milliseconds. The default value is 0 which means the connection will block until the VM TCP/IP layer times out.
- **jnp.sotimeout:** The connected socket read timeout in milliseconds. The default value is 0 which means reads will block. This is the value passed to the `Socket.setSoTimeout` on the newly connected socket.

When a client creates an `InitialContext` with these JBossNS properties available, the `org.jnp.interfaces.NamingContextFactory` object is used to create the `Context` instance that will be used in subsequent operations. The `NamingContextFactory` is the JBossNS implementation of the `javax.naming.spi.InitialContextFactory` interface. When the `NamingContextFactory` class is asked to create a `Context`, it creates an `org.jnp.interfaces.NamingContext` instance with the `InitialContext` environment and name of the context in the global JNDI namespace. It is the `NamingContext` instance that actually performs the task of connecting to the JBossNS server, and implements the `Context` interface. The `Context.PROVIDER_URL` information from the environment indicates from which server to obtain a `NamingServer` RMI reference.

The association of the `NamingContext` instance to a `NamingServer` instance is done in a lazy fashion on the first `Context` operation that is performed. When a `Context` operation is performed and the `NamingContext` has no `NamingServer` associated with it, it looks to see if its environment properties define a `Context.PROVIDER_URL`. A `Context.PROVIDER_URL` defines the host and port of the JBossNS server the `Context` is to use. If there is a provider URL, the `NamingContext` first checks to see if a `Naming` instance keyed by the host and port pair has already been created by checking a `NamingContext` class static map. It simply uses the existing `Naming` instance if one for the host port pair has already been obtained. If no `Naming` instance has been created for the given host and port, the `NamingContext` connects to the host and port using a `java.net.Socket`, and retrieves a `Naming` RMI stub from the server by reading a `java.rmi.MarshalledObject` from the socket and invoking its `get` method. The newly obtained `Naming` instance is cached in the `NamingContext` server map under the host and port pair. If no provider URL was specified in the JNDI environment associated with the context, the `NamingContext` simply uses the in VM `Naming` instance set by the Main MBean.

The `NamingContext` implementation of the `Context` interface delegates all operations to the `Naming` instance associated with the `NamingContext`. The `NamingServer` class that implements the `Naming` interface uses a `java.util.Hashtable` as the `Context` store. There is one unique `NamingServer` instance for each distinct JNDI Name for a given JBossNS server. There are zero or more transient `NamingContext` instances active at any given moment that refers to a `NamingServer` instance. The purpose of the `NamingContext` is to act as a `Context` to the `Naming` interface adaptor that manages translation of the JNDI names passed to the `NamingContext`. Because a JNDI name can be relative or a URL, it needs to be converted into an absolute name in the context of the JBossNS server to which it refers. This translation is a key function of the `NamingContext`.

3.2.1.1. Naming Discovery in Clustered Environments

When running in a clustered JBoss environment, you can choose not to specify a `Context.PROVIDER_URL` value and let the client query the network for available naming services. This only works with JBoss servers running with the `all` configuration, or an equivalent configuration that has `org.jboss.ha.framework.server.ClusterPartition` and `org.jboss.ha.jndi.HANamingService` services deployed. The discovery process consists of sending a multicast request packet to the discovery address/port and waiting for any node to respond. The response is a HA-RMI version of the `Naming` interface. The following `InitialContext` properties affect the discovery configuration:

- **jnp.partitionName:** The cluster partition name discovery should be restricted to. If you are running in an environment with multiple clusters, you may want to restrict the naming discovery to a particular cluster. There is no default value, meaning that any cluster response will be accepted.

- **jnp.discoveryGroup**: The multicast IP/address to which the discovery query is sent. The default is 230.0.0.4.
- **jnp.discoveryPort**: The port to which the discovery query is sent. The default is 1102.
- **jnp.discoveryTimeout**: The time in milliseconds to wait for a discovery query response. The default value is 5000 (5 seconds).
- **jnp.disableDiscovery**: A flag indicating if the discovery process should be avoided. Discovery occurs when either no `Context.PROVIDER_URL` is specified, or no valid naming service could be located among the URLs specified. If the `jnp.disableDiscovery` flag is true, then discovery will not be attempted.

3.2.1.2. The HTTP InitialContext Factory Implementation

The JNDI naming service can be accessed over HTTP. From a JNDI client's perspective this is a transparent change as they continue to use the JNDI `Context` interface. Operations through the `Context` interface are translated into HTTP posts to a servlet that passes the request to the `NamingService` using its `JMX` invoke operation. Advantages of using HTTP as the access protocol include better access through firewalls and proxies setup to allow HTTP, as well as the ability to secure access to the JNDI service using standard servlet role based security.

To access JNDI over HTTP you use the `org.jboss.naming.HttpNamingContextFactory` as the factory implementation. The complete set of support `InitialContext` environment properties for this factory are:

- **java.naming.factory.initial**: The name of the environment property for specifying the initial context factory, which must be `org.jboss.naming.HttpNamingContextFactory`.
- **java.naming.provider.url** (or `Context.PROVIDER_URL`): This must be set to the http URL of the JMX invoker servlet. It depends on the configuration of the `http-invoker.sar` and its contained WAR, but the default setup places the JMX invoker servlet under `/invoker/JMXInvokerServlet`. The full HTTP URL would be the public URL of the JBoss servlet container plus `/invoker/JMXInvokerServlet`. Examples include:
 - `http://www.jboss.org:8080/invoker/JMXInvokerServlet`
 - `http://www.jboss.org/invoker/JMXInvokerServlet`
 - `https://www.jboss.org/invoker/JMXInvokerServlet`

The first example accesses the servlet using the port 8080. The second uses the standard HTTP port 80, and the third uses an SSL encrypted connection to the standard HTTPS port 443.

- **java.naming.factory.url.pkgs**: For all JBoss JNDI provider this must be `org.jboss.naming:org.jnp.interfaces`. This property is essential for locating the `jnp:` and `java:` URL context factories of the JBoss JNDI provider.

The JNDI `Context` implementation returned by the `HttpNamingContextFactory` is a proxy that delegates invocations made on it to a bridge servlet which forwards the invocation to the `NamingService` through the JMX bus and marshalls the reply back over HTTP. The proxy needs to know what the URL of the bridge servlet is in order to operate. This value may have been bound on the server side if the JBoss web server has a well known public interface. If the JBoss web server is sitting behind one or more firewalls or proxies, the proxy cannot know what URL is required. In this case, the proxy will be associated with a system property value that must be set in the client VM. For more information on the operation of JNDI over HTTP see Section 3.2.2.

3.2.1.3. The Login InitialContext Factory Implementation

Historically JBoss has not supported providing login information via the `InitialContext` factory environment. The reason being that JAAS provides a much more flexible framework. For simplicity and migration from other application server environment that do make use of this mechanism, the `InitialContext` factory implementation that allows this. JAAS is still used under in the implementation, but there is no manifest use of the JAAS interfaces in the client application.

The factory class that provides this capability is the `org.jboss.security.jndi.LoginInitialContextFactory`. The complete set of support `InitialContext` environment properties for this factory are:

- **java.naming.factory.initial:** The name of the environment property for specifying the initial context factory, which must be `org.jboss.security.jndi.LoginInitialContextFactory`.
- **java.naming.provider.url:** This must be set to a `NamingContextFactory` provider URL. The `LoginInitialContext` is really just a wrapper around the `NamingContextFactory` that adds a JAAS login to the existing `NamingContextFactory` behavior.
- **java.naming.factory.url.pkgs:** For all JBoss JNDI provider this must be `org.jboss.naming:org.jnp.interfaces`. This property is essential for locating the `jnp:` and `java:` URL context factories of the JBoss JNDI provider.
- **java.naming.security.principal** (or `Context.SECURITY_PRINCIPAL`): The principal to authenticate. This may be either a `java.security.Principal` implementation or a string representing the name of a principal.
- **java.naming.security.credentials** (or `Context.SECURITY_CREDENTIALS`), The credentials that should be used to authenticate the principal, e.g., password, session key, etc.
- **java.naming.security.protocol:** (`Context.SECURITY_PROTOCOL`) This gives the name of the JAAS login module to use for the authentication of the principal and credentials.

3.2.2. Accessing JNDI over HTTP

In addition to the legacy RMI/JRMP with a socket bootstrap protocol, JBoss provides support for accessing its JNDI naming service over HTTP. This capability is provided by `http-invoker.sar`. The structure of the `http-invoker.sar` is:

```
http-invoker.sar
+- META-INF/jboss-service.xml
+- invoker.war
| +- WEB-INF/jboss-web.xml
| +- WEB-INF/classes/org/jboss/invokeation/http/servlet/InvokerServlet.class
| +- WEB-INF/classes/org/jboss/invokeation/http/servlet/NamingFactoryServlet.class
| +- WEB-INF/classes/org/jboss/invokeation/http/servlet/ReadOnlyAccessFilter.class
| +- WEB-INF/classes/roles.properties
| +- WEB-INF/classes/users.properties
| +- WEB-INF/web.xml
| +- META-INF/MANIFEST.MF
+- META-INF/MANIFEST.MF
```

The `jboss-service.xml` descriptor defines the `HttpInvoker` and `HttpInvokerHA` MBeans. These services handle the routing of methods invocations that are sent via HTTP to the appropriate target MBean on the JMX bus.

The `http-invoker.war` web application contains servlets that handle the details of the HTTP transport. The `NamingFactoryServlet` handles creation requests for the JBoss JNDI naming service `javax.naming.Context`

implementation. The `InvokerServlet` handles invocations made by RMI/HTTP clients. The `ReadOnlyAccessFilter` allows one to secure the JNDI naming service while making a single JNDI context available for read-only access by unauthenticated clients.

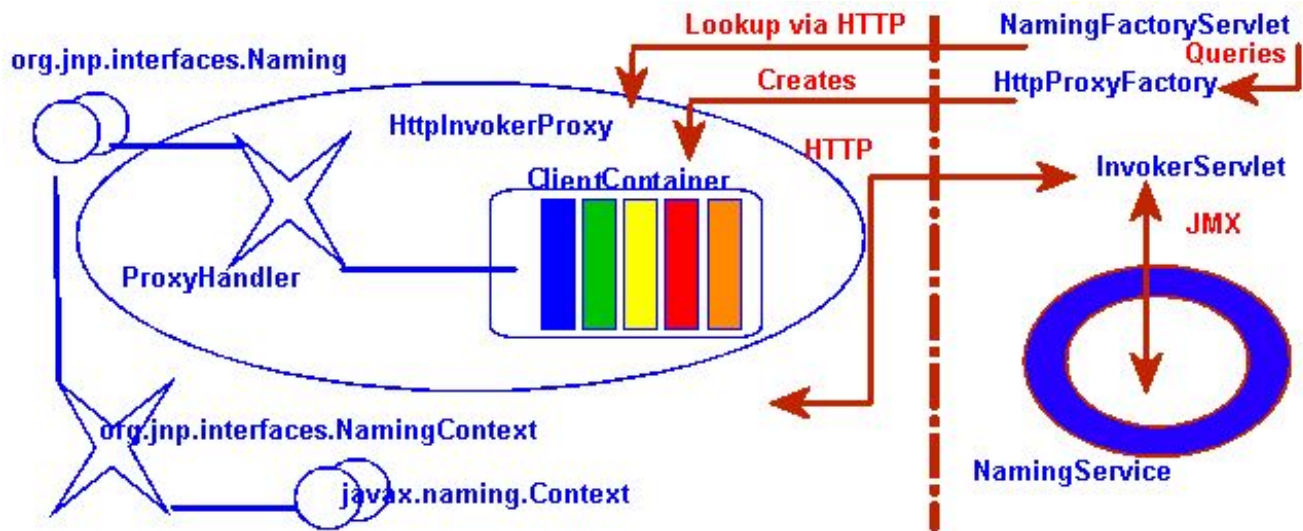


Figure 3.6. The HTTP invoker proxy/server structure for a JNDI Context

Before looking at the configurations let's look at the operation of the `http-invoker` services. Figure 3.6 shows a logical view of the structure of a JBoss JNDI proxy and its relationship to the JBoss server side components of the `http-invoker`. The proxy is obtained from the `NamingFactoryServlet` using an `InitialContext` with the `Context.INITIAL_CONTEXT_FACTORY` property set to `org.jboss.naming.HttpNamingContextFactory`, and the `Context.PROVIDER_URL` property set to the HTTP URL of the `NamingFactoryServlet`. The resulting proxy is embedded in an `org.jnp.interfaces.NamingContext` instance that provides the `Context` interface implementation.

The proxy is an instance of `org.jboss.invocation.http.interfaces.HttpInvokerProxy`, and implements the `org.jnp.interfaces.Naming` interface. Internally the `HttpInvokerProxy` contains an invoker that marshalls the `Naming` interface method invocations to the `InvokerServlet` via HTTP posts. The `InvokerServlet` translates these posts into JMX invocations to the `NamingService`, and returns the invocation response back to the proxy in the HTTP post response.

There are several configuration values that need to be set to tie all of these components together and Figure 3.7 illustrates the relationship between configuration files and the corresponding components.

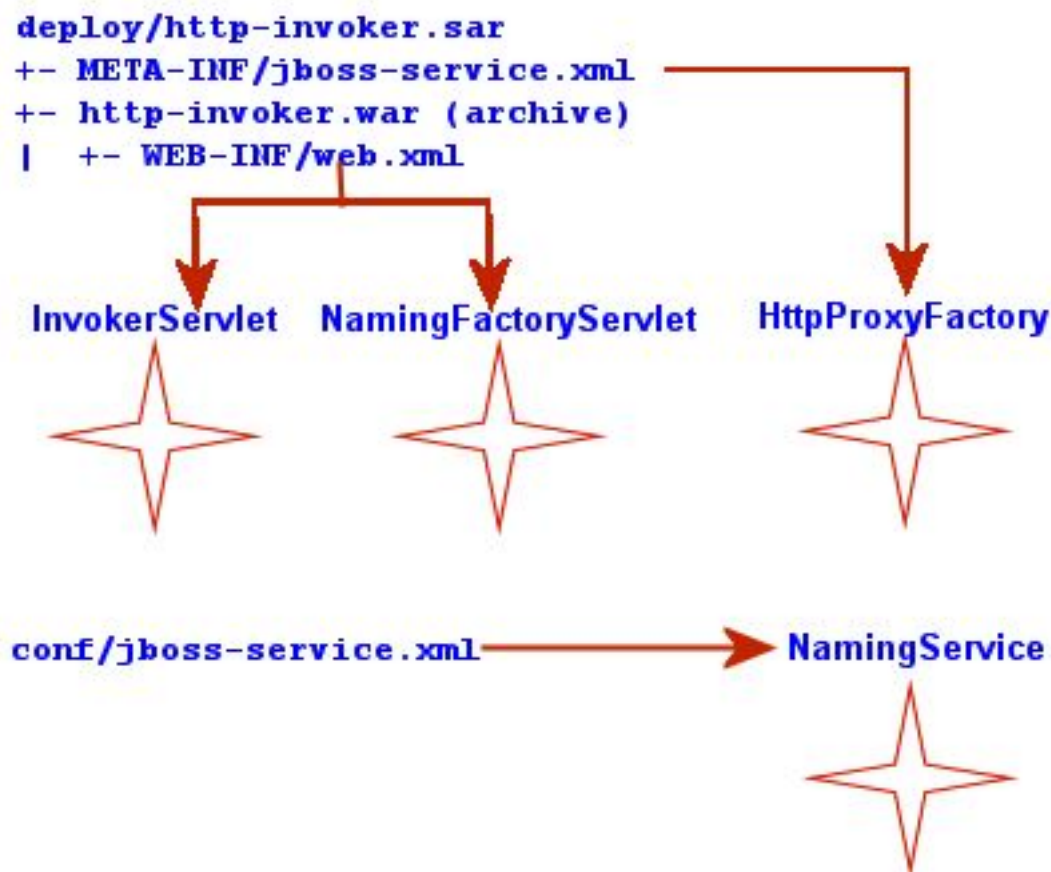


Figure 3.7. The relationship between configuration files and JNDI/HTTP component

The `http-invoker.sar/META-INF/jboss-service.xml` descriptor defines the `HttpProxyFactory` that creates the `HttpInvokerProxy` for the `NamingService`. The attributes that need to be configured for the `HttpProxyFactory` include:

- **InvokerName:** The JMX `ObjectName` of the `NamingService` defined in the `conf/jboss-service.xml` descriptor. The standard setting used in the JBoss distributions is `jboss:service=Naming`.
- **InvokerURL** or **InvokerURLPrefix** + **InvokerURLSuffix** + **UseHostName**. You can specify the full HTTP URL to the `InvokerServlet` using the `InvokerURL` attribute, or you can specify the hostname independent parts of the URL and have the `HttpProxyFactory` fill them in. An example `InvokerURL` value would be `http://jboss-host1.dot.com:8080/invoker/JMXInvokerServlet`. This can be broken down into:
 - **InvokerURLPrefix:** the URL prefix prior to the hostname. Typically this will be `http://` or `https://` if SSL is to be used.
 - **InvokerURLSuffix:** the URL suffix after the hostname. This will include the port number of the web server as well as the deployed path to the `InvokerServlet`. For the example `InvokerURL` value the `InvokerURLSuffix` would be `:8080/invoker/JMXInvokerServlet` without the quotes. The port number is determined by the web container service settings. The path to the `InvokerServlet` is specified in the `http-invoker.sar/invoker.war/WEB-INF/web.xml` descriptor.
 - **UseHostName:** a flag indicating if the hostname should be used in place of the host IP address when building the hostname portion of the full `InvokerURL`. If `true`, `InetAddress`

`dress.getLocalHost().getHostName` method will be used. Otherwise, the `InetAddress.getLocalHost().getHostAddress()` method is used.

- **ExportedInterface:** The `org.jnp.interfaces.Naming` interface the proxy will expose to clients. The actual client of this proxy is the JBoss JNDI implementation `NamingContext` class, which JNDI client obtain from `InitialContext` lookups when using the JBoss JNDI provider.
- **JndiName:** The name in JNDI under which the proxy is bound. This needs to be set to a blank/empty string to indicate the interface should not be bound into JNDI. We can't use the JNDI to bootstrap itself. This is the role of the `NamingFactoryServlet`.

The `http-invoker.sar/invoker.war/WEB-INF/web.xml` descriptor defines the mappings of the `NamingFactoryServlet` and `InvokerServlet` along with their initialization parameters. The configuration of the `NamingFactoryServlet` relevant to JNDI/HTTP is the `JNDIFactory` entry which defines:

- A `namingProxyMBean` initialization parameter that maps to the `HttpProxyFactory` MBean name. This is used by the `NamingFactoryServlet` to obtain the `Naming` proxy which it will return in response to HTTP posts. For the default `http-invoker.sar/META-INF/jboss-service.xml` settings the name `jboss:service=invoker,type=http,target=Naming`.
- A proxy initialization parameter that defines the name of the `namingProxyMBean` attribute to query for the `Naming` proxy value. This defaults to an attribute name of `Proxy`.
- The servlet mapping for the `JNDIFactory` configuration. The default setting for the unsecured mapping is `/JNDIFactory/*`. This is relative to the context root of the `http-invoker.sar/invoker.war`, which by default is the WAR name minus the `.war` suffix.

The configuration of the `InvokerServlet` relevant to JNDI/HTTP is the `JMXInvokerServlet` which defines:

- The servlet mapping of the `InvokerServlet`. The default setting for the unsecured mapping is `/JMXInvokerServlet/*`. This is relative to the context root of the `http-invoker.sar/invoker.war`, which by default is the WAR name minus the `.war` suffix.

3.2.3. Accessing JNDI over HTTPS

To be able to access JNDI over HTTP/SSL you need to enable an SSL connector on the web container. The details of this are covered in the Integrating Servlet Containers for Tomcat. We will demonstrate the use of HTTPS with a simple example client that uses an HTTPS URL as the JNDI provider URL. We will provide an SSL connector configuration for the example, so unless you are interested in the details of the SSL connector setup, the example is self contained.

We also provide a configuration of the `HttpProxyFactory` setup to use an HTTPS URL. shows the section of the `http-invoker.sar/jboss-service.xml` descriptor that the example installs to provide this configuration. All the has changed relative to the standard http configuration are the `InvokerURLPrefix` and `InvokerURLSuffix` attributes, which setup an HTTPS URL using the 8443 port.

```
<!-- Expose the Naming service interface via HTTPS -->
<mbean code="org.jboss.invocation.http.server.HttpProxyFactory"
  name="jboss:service=invoker,type=https,target=Naming">
  <!-- The Naming service we are proxying -->
  <attribute name="InvokerName">jboss:service=Naming</attribute>
  <!-- Compose the invoker URL from the cluster node address -->
  <attribute name="InvokerURLPrefix">https://</attribute>
```



```

<attribute name="InvokerURLSuffix">:8443/invoker/JMXInvokerServlet </attribute>
<attribute name="UseHostName">true</attribute>
<attribute name="ExportedInterface">org.jnp.interfaces.Naming </attribute>
<attribute name="JndiName"/>
<attribute name="ClientInterceptors">
    <interceptors>
        <interceptor>org.jboss.proxy.ClientMethodInterceptor </interceptor>
        <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
        <interceptor>org.jboss.naming.interceptors.ExceptionInterceptor </interceptor>
        <interceptor>org.jboss.invocation.InvokerInterceptor </interceptor>
    </interceptors>
</attribute>
</mbean>

```

At a minimum, a JNDI client using HTTPS requires setting up a https URL protocol handler. We will be using the Java Secure Socket Extension (JSSE) for HTTPS. The JSSE documentation does a good job of describing what is necessary to use https, and the following steps were needed to configure the example client shown in Example 3.16:

- A protocol handler for HTTPS URLs must be made available to Java. The JSSE release includes an HTTPS handler in the `com.sun.net.ssl.internal.www.protocol` package. To enable the use of https URLs you include this package in the standard URL protocol handler search property, `java.protocol.handler.pkgs`. We set the `java.protocol.handler.pkgs` property in the Ant script.
- The JSSE security provider must be installed in order for SSL to work. This can be done either by installing the JSSE jars as an extension package, or programatically. We use the programatic approach in the example since this is less intrusive. Line 18 of the `ExClient` code demonstrates how this is done.
- The JNDI provider URL must use HTTPS as the protocol. Lines 24-25 of the `ExClient` code specify an HTTP/SSL connection to the localhost on port 8443. The hostname and port are defined by the web container SSL connector.
- The validation of the HHTTPS URL hostname against the server certificate must be disabled. By default, the JSSE https protocol handler employs a strict validation of the hostname portion of the HTTPS URL against the common name of the server certificate. This is the same check done by web browsers when you connect to secured web site. We are using a self-signed server certificate that uses a common name of "Chapter8 SSL Example" rather than a particular hostname, and this is likely to be common in development environments or intranets. The JBoss `HttpInvokerProxy` will override the default hostname checking if a `org.jboss.security.ignoreHttpsHost` system property exists and has a value of true. We set the `org.jboss.security.ignoreHttpsHost` property to true in the Ant script.

Example 3.16. A JNDI client that uses HTTPS as the transport

```

package org.jboss.chap3.ex1;

import java.security.Security;
import java.util.Properties;
import javax.naming.Context;
import javax.naming.InitialContext;

/**
 * A simple JNDI client that uses HTTPS as the transport.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.8 $
 */
public class ExClient
{
    public static void main(String args[]) throws Exception

```

```

{
    // Install the Sun JSSE provider since we may not have JSSE installed
    Security.addProvider(new com.sun.net.ssl.internal.ssl.Provider());
    System.out.println("Added JSSE security provider");

    Properties env = new Properties();
    env.setProperty(Context.INITIAL_CONTEXT_FACTORY,
        "org.jboss.naming.HttpNamingContextFactory");
    env.setProperty(Context.PROVIDER_URL,
        "https://localhost:8443/invoker/JNDIFactory");
    Context ctx = new InitialContext(env);
    System.out.println("Created InitialContext, env="+env);
    Object data = ctx.lookup("jmx/rmi/RMIAdaptor");
    System.out.println("lookup(jmx/rmi/RMIAdaptor): "+data);
}
}

```

To test the client, first build the chapter 3 example to create the chap3 configuration fileset.

```

[nr@toki examples]$ ant -Dchap=chap3 config
Buildfile: build.xml

validate:
 [java] ImplementationTitle: JBoss [WonderLand]
 [java] ImplementationVendor: JBoss.org
 [java] ImplementationVersion: 3.2.6RC2 (build: CVSTag=Branch_3_2 date=200409270100)
 [java] SpecificationTitle: JBoss
 [java] SpecificationVendor: JBoss (http://www.jboss.org/)
 [java] SpecificationVersion: 3.2.6
 [java] JBoss version is: 3.2.6

fail_if_not_valid:

init:
 [echo] Using jboss.dist=/tmp/jboss-3.2.6

compile-src:

compile:

config:

config:
 [echo] Preparing chap3 configuration fileset
 [mkdir] Created dir: /tmp/jboss-3.2.6/server/chap3
 [copy] Copying 221 files to /tmp/jboss-3.2.6/server/chap3
 [copy] Copied 1 empty directory to /tmp/jboss-3.2.6/server/chap3
 [copy] Copying 1 file to /tmp/jboss-3.2.6/server/chap3/conf
 [copy] Copying 1 file to /tmp/jboss-3.2.6/server/chap3/conf
 [copy] Copying 1 file to /tmp/jboss-3.2.6/server/chap3/deploy/jbossweb-tomcat50.sar
 [copy] Copying 1 file to /tmp/jboss-3.2.6/server/chap3/deploy/http-invoker.sar/META
-INF
 [copy] Copying 1 file to /tmp/jboss-3.2.6/server/chap3/deploy/http-invoker.sar/invo
ker.war/WEB-INF
BUILD SUCCESSFUL
Total time: 6 seconds

```

Next, start the JBoss server using the chap3 configuration fileset:

```

[nr@toki bin]$ sh run.sh -c chap3
=====

JBoss Bootstrap Environment

JBOSS_HOME: /tmp/jboss-3.2.6
...

```

And finally, run the ExClient using:

```
[nr@toki examples]$ ant -Dchap=chap3 -Dex=1 run-example
Buildfile: build.xml

validate:
...

run-example:

run-example1:
    [java] JSSE already available
    [java] Created InitialContext, env={java.naming.provider.url=https://localhost:8443/
invoker/JNDIFactorySSL, java.naming.factory.initial=org.jboss.naming.HttpNamingContextFac
tory}
    [java] lookup(jmx/rmi/RMIAdaptor): org.jboss.invocation.jrmp.interfaces.JRMPInvokerP
roxy@873b9f
BUILD SUCCESSFUL
Total time: 6 seconds
```

3.2.4. Securing Access to JNDI over HTTP

One benefit to accessing JNDI over HTTP is that it is easy to secure access to the JNDI `InitialContext` factory as well as the naming operations using standard web declarative security. This is possible because the server side handling of the JNDI/HTTP transport is implemented with two servlets. These servlets are included in the `http-invoker.sar/invoker.war` directory found in the default and all configuration deploy directories as shown previously. To enable secured access to JNDI you need to edit the `invoker.war/WEB-INF/web.xml` descriptor and remove all unsecured servlet mappings. For example, the `web.xml` descriptor shown in Example 3.17 only allows access to the `invoker.war` servlets if the user has been authenticated and has a role of `HttpInvoker`.

Example 3.17. An example `web.xml` descriptor for secured access to the JNDI servlets

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE web-app PUBLIC
    "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
    "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>
    <!-- ### Servlets -->
    <servlet>
        <servlet-name>JMXInvokerServlet</servlet-name>
        <servlet-class>
            org.jboss.invocation.http.servlet.InvokerServlet
        </servlet-class>
        <load-on-startup>1</load-on-startup>
    </servlet>
    <servlet>
        <servlet-name>JNDIFactory</servlet-name>
        <servlet-class>
            org.jboss.invocation.http.servlet.NamingFactoryServlet
        </servlet-class>
        <init-param>
            <param-name>namingProxyMBean</param-name>
            <param-value>jboss:service=invoker,type=http,target=Naming</param-value>
        </init-param>
        <init-param>
            <param-name>proxyAttribute</param-name>
            <param-value>Proxy</param-value>
        </init-param>
        <load-on-startup>2</load-on-startup>
    </servlet>
    <!-- ### Servlet Mappings -->
    <servlet-mapping>
        <servlet-name>JNDIFactory</servlet-name>
```

```

    <url-pattern>/restricted/JNDIFactory/*</url-pattern>
  </servlet-mapping>
  <servlet-mapping>
    <servlet-name>JMXInvokerServlet</servlet-name>
    <url-pattern>/restricted/JMXInvokerServlet/*</url-pattern>
  </servlet-mapping>
  <security-constraint>
    <web-resource-collection>
      <web-resource-name>HttpInvokers</web-resource-name>
      <description>An example security config that only allows users with
        the role HttpInvoker to access the HTTP invoker servlets </description>
      <url-pattern>/restricted/*</url-pattern>
      <http-method>GET</http-method>
      <http-method>POST</http-method>
    </web-resource-collection>
    <auth-constraint>
      <role-name>HttpInvoker</role-name>
    </auth-constraint>
  </security-constraint>
  <login-config>
    <auth-method>BASIC</auth-method>
    <realm-name>JBoss HTTP Invoker</realm-name>
  </login-config>
  <security-role>
    <role-name>HttpInvoker</role-name>
  </security-role>
</web-app>

```

The `web.xml` descriptor only defines which sevlets are secured, and which roles are allowed to access the secured servlets. You must additionally define the security domain that will handle the authentication and authorization for the war. This is done through the `jboss-web.xml` descriptor, and an example that uses the `http-invoker` security domain is given below.

```

<jboss-web>
  <security-domain>java:/jaas/http-invoker</security-domain>
</jboss-web>

```

The `security-domain` element defines the name of the security domain that will be used for the JAAS login module configuration used for authentication and authorization. See Section 8.1.6 for additional details on the meaning and configuration of the security domain name.

3.2.5. Securing Access to JNDI with a Read-Only Unsecured Context

Another feature available for the JNDI/HTTP naming service is the ability to define a context that can be accessed by unauthenticated users in read-only mode. This can be important for services used by the authentication layer. For example, the `SRPLoginModule` needs to lookup the SRP server interface used to perform authentication. To enable this, some additional `web.xml` descriptor settings are needed. The following diagram shows the the additional `web.xml` descriptor elements needed for read-only access.

```

<web-app>
  <filter>
    <filter-name>ReadOnlyAccessFilter</filter-name>
    <filter-class>org.jboss.invocation.http.servlet.ReadOnlyAccessFilter</filter-class>
    <init-param>
      <param-name>readOnlyContext</param-name>
      <param-value>readonly-context</param-value>
    </init-param>
    <init-param>
      <param-name>invokerName</param-name>
      <param-value>jboss:service=Naming</param-value>
    </init-param>
  </filter>
  <filter-mapping>
    <filter-name>ReadOnlyAccessFilter</filter-name>

```

```

        <url-pattern>/readonly/*</url-pattern>
    </filter-mapping>
    <servlet>
        <servlet-name>ReadOnlyJNDIFactory</servlet-name>
        <servlet-class>
            org.jboss.invocation.http.servlet.NamingFactoryServlet
        </servlet-class>
        <init-param>
            <param-name>namingProxyMBean</param-name>
            <param-value>
                jboss:service=invoker,type=http,target=Naming,readonly=true
            </param-value>
        </init-param>
        <init-param>
            <param-name>proxyAttribute</param-name>
            <param-value>Proxy</param-value>
        </init-param>
        <load-on-startup>2</load-on-startup>
    </servlet>
    <!-- A mapping for the JMXInvokerServlet that only allows invocations
        of lookups under a read-only context. This is enforced by the
        ReadOnlyAccessFilter
    -->
    <servlet-mapping>
        <servlet-name>JMXInvokerServlet</servlet-name>
        <url-pattern>/readonly/JMXInvokerServlet/*</url-pattern>
    </servlet-mapping>
</web-app>

```

With these settings, one may perform `Context.lookup` operations on the `readonly-context` or its subcontexts, but no other operations on this context. Also, no operations of any kind may be performed on other contexts. Here is a code fragment for a lookup of the `readonly-context/data` binding:

```

Properties env = new Properties();
env.setProperty(Context.INITIAL_CONTEXT_FACTORY,
    "org.jboss.naming.HttpNamingContextFactory");
env.setProperty(Context.PROVIDER_URL,
    "http://localhost:8080/invoker/ReadOnlyJNDIFactory");

Context ctx2 = new InitialContext(env);
Object data = ctx2.lookup("readonly-context/data");

```

3.2.6. Additional Naming MBeans

In addition to the `NamingService` MBean that configures an embedded JBossNS server within JBoss, there are three additional MBean services related to naming that ship with JBoss. They are the `ExternalContext`, `NamingAlias`, and `JNDIView`.

3.2.6.1. org.jboss.naming.ExternalContext MBean

The `ExternalContext` MBean allows you to federate external JNDI contexts into the JBoss server JNDI namespace. The term `external` refers to any naming service external to the JBossNS naming service running inside of the JBoss server VM. You can incorporate LDAP servers, file systems, DNS servers, and so on, even if the JNDI provider root context is not serializable. The federation can be made available to remote clients if the naming service supports remote access.

To incorporate an external JNDI naming service, you have to add a configuration of the `ExternalContext` MBean service to the `jboss-service.xml` configuration file. The configurable attributes of the `ExternalContext` service are as follows:

- **JndiName:** The JNDI name under which the external context is to be bound.
- **RemoteAccess:** A boolean flag indicating if the external `InitialContext` should be bound using a `Serializable` form that allows a remote client to create the external `InitialContext`. When a remote client looks up the external context via the JBoss JNDI `InitialContext`, they effectively create an instance of the external `InitialContext` using the same `env` properties passed to the `ExternalContext` MBean. This will only work if the client can do a new `InitialContext(env)` remotely. This requires that the `Context.PROVIDER_URL` value of `env` is resolvable in the remote VM that is accessing the context. This should work for the LDAP example. For the file system example this most likely won't work unless the file system path refers to a common network path. If this property is not given it defaults to false.
- **CacheContext:** The `cacheContext` flag. When set to true, the external `Context` is only created when the MBean is started and then stored as an in memory object until the MBean is stopped. If `cacheContext` is set to false, the external `Context` is created on each lookup using the MBean properties and `InitialContext` class. When the uncached `Context` is looked up by a client, the client should invoke `close()` on the `Context` to prevent resource leaks.
- **InitialContext:** The fully qualified class name of the `InitialContext` implementation to use. Must be one of: `javax.naming.InitialContext`, `javax.naming.directory.InitialDirContext` or `javax.naming.ldap.InitialLdapContext`. In the case of the `InitialLdapContext` a null `Controls` array is used. The default is `javax.naming.InitialContext`.
- **Properties:** Set the `jndi.properties` information for the external `InitialContext`. This is either a URL, string or a classpath resource name. Examples are as follows:
 - `file:///config/myldap.properties`
 - `http://config.mycompany.com/myldap.properties`
 - `/conf/myldap.properties`
 - `myldap.properties`

The MBean definition below shows two configurations: one for an LDAP server, and the other for a local file system directory.

```
<!-- Bind a remote LDAP server -->
<mbean code="org.jboss.naming.ExternalContext"
  name="jboss.jndi:service=ExternalContext,jndiName=external/ldap/jboss">
  <attribute name="JndiName">external/ldap/jboss</attribute>
  <attribute name="Properties">jboss.ldap</attribute>
  <attribute name="InitialContext"> javax.naming.ldap.InitialLdapContext </attribute>
  <attribute name="RemoteAccess">true</attribute>
</mbean>

<!-- Bind the /usr/local file system directory -->
<mbean code="org.jboss.naming.ExternalContext"
  name="jboss.jndi:service=ExternalContext,jndiName=external/fs/usr/local">
  <attribute name="JndiName">external/fs/usr/local</attribute>
  <attribute name="Properties">local.props</attribute>
  <attribute name="InitialContext">javax.naming.IntialContext</attribute>
</mbean>
```

The first configuration describes binding an external LDAP context into the JBoss JNDI namespace under the name `external/ldap/jboss`. An example `jboss.ldap` properties file is as follows:

```
java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory
java.naming.provider.url=ldap://ldaphost.jboss.org:389/o=jboss.org
java.naming.security.principal=cn=Directory Manager
java.naming.security.authentication=simple
java.naming.security.credentials=secret
```

With this configuration, you can access the external LDAP context located at `ldap://ldaphost.jboss.org:389/o=jboss.org` from within the JBoss VM using the following code fragment:

```
InitialContext iniCtx = new InitialContext();
LdapContext ldapCtx = iniCtx.lookup("external/ldap/jboss");
```

Using the same code fragment outside of the JBoss server VM will work in this case because the `RemoteAccess` property was set to true. If it were set to false, it would not work because the remote client would receive a `Reference` object with an `ObjectFactory` that would not be able to recreate the external `InitialContext`.

The second configuration describes binding a local file system directory `/usr/local` into the JBoss JNDI namespace under the name `external/fs/usr/local`. An example `local.props` properties file is:

```
java.naming.factory.initial=com.sun.jndi.fscontext.RefFSContextFactory
java.naming.provider.url=file:///usr/local
```

With this configuration, you can access the external file system context located at `file:///usr/local` from within the JBoss VM using the following code fragment:

```
InitialContext iniCtx = new InitialContext();
Context ldapCtx = iniCtx.lookup("external/fs/usr/local");
```

3.2.6.2. The `org.jboss.naming.NamingAlias` MBean

The `NamingAlias` MBean is a simple utility service that allows you to create an alias in the form of a JNDI `javax.naming.LinkRef` from one JNDI name to another. This is similar to a symbolic link in the UNIX file system. To an alias you add a configuration of the `NamingAlias` MBean to the `jboss-service.xml` configuration file. The configurable attributes of the `NamingAlias` service are as follows:

- **FromName:** The location where the `LinkRef` is bound under JNDI.
- **ToName:** The to name of the alias. This is the target name to which the `LinkRef` refers. The name is a URL, or a name to be resolved relative to the `InitialContext`, or if the first character of the name is `.`, the name is relative to the context in which the link is bound.

The following example provides a mapping of the JNDI name `QueueConnectionFactory` to the name `ConnectionFactory`.

```
<mbean code="org.jboss.naming.NamingAlias"
      name="jboss.mq:service=NamingAlias,fromName=QueueConnectionFactory">
  <attribute name="ToName">ConnectionFactory</attribute>
  <attribute name="FromName">QueueConnectionFactory</attribute>
</mbean>
```

3.2.6.3. The `org.jboss.naming.JNDIView` MBean

The `JNDIView` MBean allows the user to view the JNDI namespace tree as it exists in the JBoss server using the JMX agent view interface. To view the JBoss JNDI namespace using the `JNDIView` MBean, you connect to the JMX Agent View using the http interface. The default settings put this at `http://localhost:8080/jmx-console/`. On this page you will see a section that lists the registered MBeans sorted by domain. It should look something like that shown in Figure 3.8.

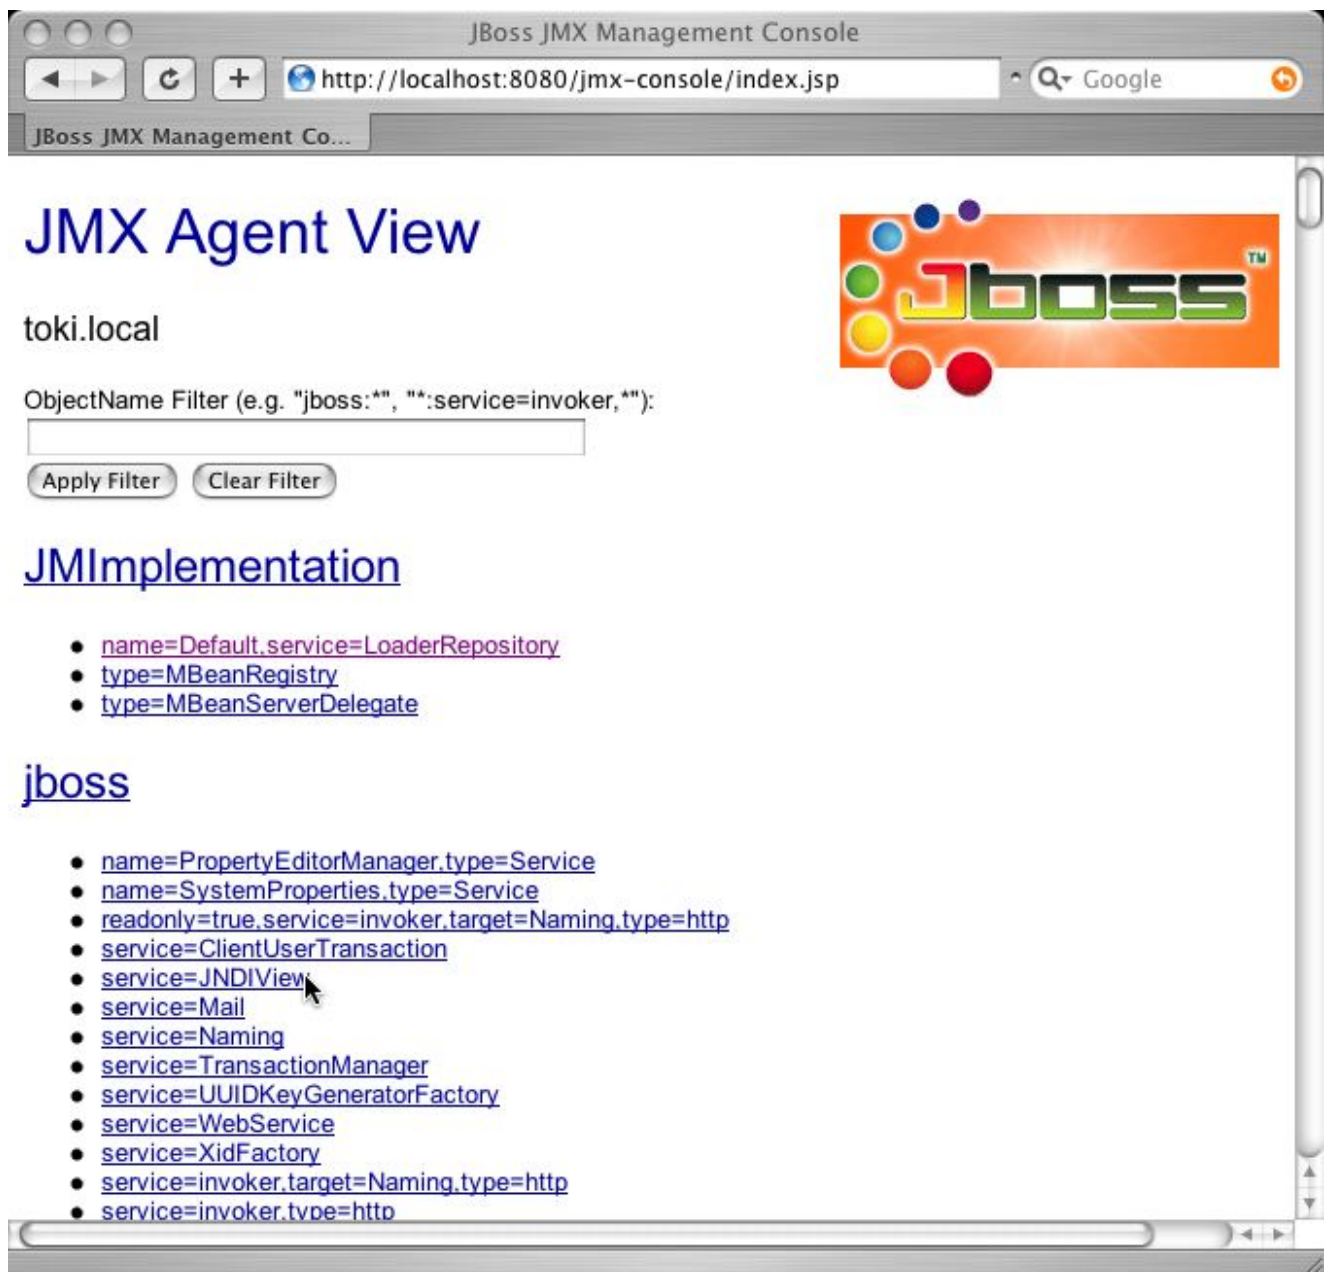


Figure 3.8. The JMX Console view of the configured JBoss MBeans

Selecting the JNDIView link takes you to the JNDIView MBean view, which will have a list of the JNDIView MBean operations. This view should look similar to that shown in Figure 3.9.

The screenshot shows the MBean Inspector window with the following content:

Attribute Name (Access) Type <i>Description</i>	Attribute Value
Name (R) java.lang.String <i>The class name of the MBean</i>	JNDIView
State (R) int <i>The status of the MBean</i>	3
StateString (R) java.lang.String <i>The status of the MBean in text form</i>	Started

Operation Name Return Type <i>Description</i>	Parameters
list java.lang.String <i>Output JNDI info as text</i>	verbose boolean <i>If true, list the class of each object in addition to its name</i> <input checked="" type="radio"/> True <input type="radio"/> False <input type="button" value="Invoke"/>
listXML java.lang.String <i>Output JNDI info in XML format</i>	<input type="button" value="Invoke"/>

Figure 3.9. The JMX Console view of the JNDIView MBean

The list operation dumps out the JBoss server JNDI namespace as an html page using a simple text view. As an example, invoking the list operation produces the view shown in Figure 3.10.

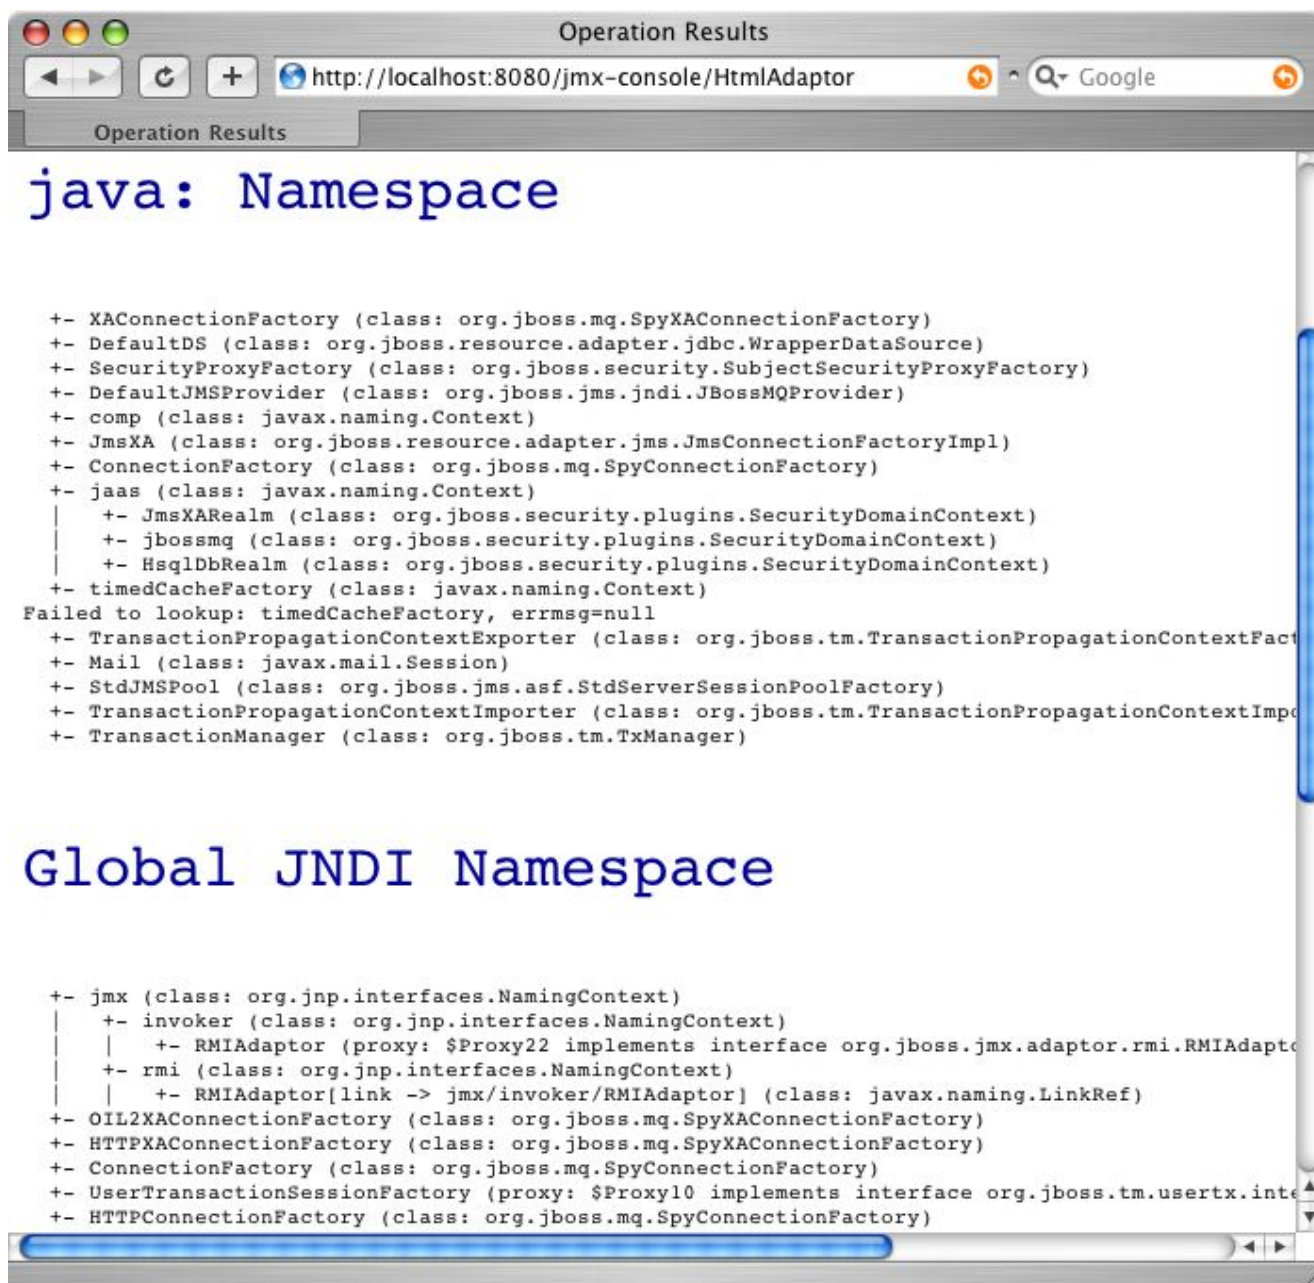


Figure 3.10. The JMX Console view of the JNDIView list operation output

Transactions on JBoss

The JTA Transaction Service

This chapter discusses transaction management in JBoss and the JBossTX architecture. The JBossTX architecture allows for any Java Transaction API (JTA) transaction manager implementation to be used. JBossTX includes a fast in-VM implementation of a JTA compatible transaction manager that is used as the default transaction manager. We will first provide an overview of the key transaction concepts and notions in the JTA to provide sufficient background for the JBossTX architecture discussion. We will then discuss the interfaces that make up the JBossTX architecture and conclude with a discussion of the MBeans available for integration of alternate transaction managers.

4.1. Transaction/JTA Overview

For the purpose of this discussion, we can define a transaction as a unit of work containing one or more operations involving one or more shared resources having ACID properties. ACID is an acronym for atomicity, consistency, isolation and durability, the four important properties of transactions. The meanings of these terms is:

- **Atomicity:** A transaction must be atomic. This means that either all the work done in the transaction must be performed, or none of it must be performed. Doing part of a transaction is not allowed.
- **Consistency:** When a transaction is completed, the system must be in a stable and consistent condition.
- **Isolation:** Different transactions must be isolated from each other. This means that the partial work done in one transaction is not visible to other transactions until the transaction is committed, and that each process in a multi-user system can be programmed as if it was the only process accessing the system.
- **Durability:** The changes made during a transaction are made persistent when it is committed. When a transaction is committed, its changes will not be lost, even if the server crashes afterwards.

To illustrate these concepts, consider a simple banking account application. The banking application has a database with a number of accounts. The sum of the amounts of all accounts must always be 0. An amount of money *M* is moved from account *A* to account *B* by subtracting *M* from account *A* and adding *M* to account *B*. This operation must be done in a transaction, and all four ACID properties are important.

The atomicity property means that both the withdrawal and deposit is performed as an indivisible unit. If, for some reason, both cannot be done nothing will be done.

The consistency property means that after the transaction, the sum of the amounts of all accounts must still be 0.

The isolation property is important when more than one bank clerk uses the system at the same time. A withdrawal or deposit could be implemented as a three-step process: First the amount of the account is read from the database; then something is subtracted from or added to the amount read from the database; and at last the new amount is written to the database. Without transaction isolation several bad things could happen. For ex-

ample, if two processes read the amount of account A at the same time, and each independently added or subtracted something before writing the new amount to the database, the first change would be incorrectly overwritten by the last.

The durability property is also important. If a money transfer transaction is committed, the bank must trust that some subsequent failure cannot undo the money transfer.

4.1.1. Pessimistic and optimistic locking

Transactional isolation is usually implemented by locking whatever is accessed in a transaction. There are two different approaches to transactional locking: Pessimistic locking and optimistic locking.

The disadvantage of pessimistic locking is that a resource is locked from the time it is first accessed in a transaction until the transaction is finished, making it inaccessible to other transactions during that time. If most transactions simply look at the resource and never change it, an exclusive lock may be overkill as it may cause lock contention, and optimistic locking may be a better approach. With pessimistic locking, locks are applied in a fail-safe way. In the banking application example, an account is locked as soon as it is accessed in a transaction. Attempts to use the account in other transactions while it is locked will either result in the other process being delayed until the account lock is released, or that the process transaction will be rolled back. The lock exists until the transaction has either been committed or rolled back.

With optimistic locking, a resource is not actually locked when it is first accessed by a transaction. Instead, the state of the resource at the time when it would have been locked with the pessimistic locking approach is saved. Other transactions are able to concurrently access the resource and the possibility of conflicting changes is possible. At commit time, when the resource is about to be updated in persistent storage, the state of the resource is read from storage again and compared to the state that was saved when the resource was first accessed in the transaction. If the two states differ, a conflicting update was made, and the transaction will be rolled back.

In the banking application example, the amount of an account is saved when the account is first accessed in a transaction. If the transaction changes the account amount, the amount is read from the store again just before the amount is about to be updated. If the amount has changed since the transaction began, the transaction will fail itself, otherwise the new amount is written to persistent storage.

4.1.2. The components of a distributed transaction

There are a number of participants in a distributed transaction. These include:

- **Transaction Manager:** This component is distributed across the transactional system. It manages and coordinates the work involved in the transaction. The transaction manager is exposed by the `javax.transaction.TransactionManager` interface in JTA.
- **Transaction Context:** A transaction context identifies a particular transaction. In JTA the corresponding interface is `javax.transaction.Transaction`.
- **Transactional Client:** A transactional client can invoke operations on one or more transactional objects in a single transaction. The transactional client that started the transaction is called the transaction originator. A transaction client is either an explicit or implicit user of JTA interfaces and has no interface representation in the JTA.
- **Transactional Object:** A transactional object is an object whose behavior is affected by operations performed on it within a transactional context. A transactional object can also be a transactional client. Most

Enterprise Java Beans are transactional objects.

- **Recoverable Resource:** A recoverable resource is a transactional object whose state is saved to stable storage if the transaction is committed, and whose state can be reset to what it was at the beginning of the transaction if the transaction is rolled back. At commit time, the transaction manager uses the two-phase XA protocol when communicating with the recoverable resource to ensure transactional integrity when more than one recoverable resource is involved in the transaction being committed. Transactional databases and message brokers like JBossMQ are examples of recoverable resources. A recoverable resource is represented using the `javax.transaction.xa.XAResource` interface in JTA.

4.1.3. The two-phase XA protocol

When a transaction is about to be committed, it is the responsibility of the transaction manager to ensure that either all of it is committed, or that all of it is rolled back. If only a single recoverable resource is involved in the transaction, the task of the transaction manager is simple: It just has to tell the resource to commit the changes to stable storage.

When more than one recoverable resource is involved in the transaction, management of the commit gets more complicated. Simply asking each of the recoverable resources to commit changes to stable storage is not enough to maintain the atomic property of the transaction. The reason for this is that if one recoverable resource has committed and another fails to commit, part of the transaction would be committed and the other part rolled back.

To get around this problem, the two-phase XA protocol is used. The XA protocol involves an extra prepare phase before the actual commit phase. Before asking any of the recoverable resources to commit the changes, the transaction manager asks all the recoverable resources to prepare to commit. When a recoverable resource indicates it is prepared to commit the transaction, it has ensured that it can commit the transaction. The resource is still able to rollback the transaction if necessary as well.

So the first phase consists of the transaction manager asking all the recoverable resources to prepare to commit. If any of the recoverable resources fails to prepare, the transaction will be rolled back. But if all recoverable resources indicate they were able to prepare to commit, the second phase of the XA protocol begins. This consists of the transaction manager asking all the recoverable resources to commit the transaction. Because all the recoverable resources have indicated they are prepared, this step cannot fail.

4.1.4. Heuristic exceptions

In a distributed environment communications failures can happen. If communication between the transaction manager and a recoverable resource is not possible for an extended period of time, the recoverable resource may decide to unilaterally commit or rollback changes done in the context of a transaction. Such a decision is called a heuristic decision. It is one of the worst errors that may happen in a transaction system, as it can lead to parts of the transaction being committed while other parts are rolled back, thus violating the atomicity property of transaction and possibly leading to data integrity corruption.

Because of the dangers of heuristic exceptions, a recoverable resource that makes a heuristic decision is required to maintain all information about the decision in stable storage until the transaction manager tells it to forget about the heuristic decision. The actual data about the heuristic decision that is saved in stable storage depends on the type of recoverable resource and is not standardized. The idea is that a system manager can look at the data, and possibly edit the resource to correct any data integrity problems.

There are several different kinds of heuristic exceptions defined by the JTA. The `javax.transaction.HeuristicCommitException` is thrown when a recoverable resource is asked to rollback

to report that a heuristic decision was made and that all relevant updates have been committed. On the opposite end is the `javax.transaction.HeuristicRollbackException`, which is thrown by a recoverable resource when it is asked to commit to indicate that a heuristic decision was made and that all relevant updates have been rolled back.

The `javax.transaction.HeuristicMixedException` is the worst heuristic exception. It is thrown to indicate that parts of the transaction were committed, while other parts were rolled back. The transaction manager throws this exception when some recoverable resources did a heuristic commit, while other recoverable resources did a heuristic rollback.

4.1.5. Transaction IDs and branches

In JTA, the identity of transactions is encapsulated in objects implementing the `javax.transaction.xa.Xid` interface. The transaction ID is an aggregate of three parts:

- The format identifier indicates the transaction family and tells how the other two parts should be interpreted.
- The global transaction id identified the global transaction within the transaction family.
- The branch qualifier denotes a particular branch of the global transaction.

Transaction branches are used to identify different parts of the same global transaction. Whenever the transaction manager involves a new recoverable resource in a transaction it creates a new transaction branch.

4.2. JBoss Transaction Internals

The JBoss application server is written to be independent of the actual transaction manager used. JBoss uses the JTA `javax.transaction.TransactionManager` interface as its view of the server transaction manager. Thus, JBoss may use any transaction manager which implements the JTA `TransactionManager` interface. Whenever a transaction manager is used it is obtained from the well-known JNDI location, `java:/TransactionManager`. This is the globally available access point for the server transaction manager.

If transaction contexts are to be propagated with RMI/JRMP calls, the transaction manager must also implement two simple interfaces for the import and export of transaction propagation contexts (TPCs). The interfaces are `org.jboss.tm.TransactionPropagationContextImporter`, and `org.jboss.tm.TransactionPropagationContextFactory`.

Being independent of the actual transaction manager used also means that JBoss does not specify the format of type of the transaction propagation contexts used. In JBoss, a TPC is of type `Object`, and the only requirement is that the TPC must implement the `java.io.Serializable` interface.

When using the RMI/JRMP protocol for remote calls, the TPC is carried as a field in the `org.jboss.ejb.plugins.jrmp.client.RemoteMethodInvocation` class that is used to forward remote method invocation requests.

4.2.1. Adapting a Transaction Manager to JBoss

A transaction manager has to implement the Java Transaction API to be easily integrated with JBoss. As almost everything in JBoss, the transaction manager is managed as an MBean. Like all JBoss services, it should implement `org.jboss.system.ServiceMBean` to ensure proper life-cycle management.

The primary requirement of the transaction manager service on startup is that it binds its implementation of the three required interfaces into JNDI. These interfaces and their JNDI locations are:

- The `javax.transaction.TransactionManager` interface is used by the application server to manage transactions on behalf of the transactional objects that use container managed transactions. It must be bound under the JNDI name `java:/TransactionManager`.
- The `transaction propagation context factory` interface `org.jboss.tm.TransactionPropagationContextFactory` is called by JBoss whenever a transaction propagation context is needed for for transporting a transaction with a remote method call. It must be bound under the JNDI name `java:/TransactionPropagationContextImporter`.
- The `transaction propagation context importer` interface `org.jboss.tm.TransactionPropagationContextImporter` is called by JBoss whenever a transaction propagation context from an incoming remote method invocation has to be converted to a transaction that can be used within the receiving JBoss server VM.

Establishing these JNDI bindings is all the transaction manager service needs to do to install its implementation as the JBoss server transaction manager.

4.2.2. The Default Transaction Manager

JBoss is by default configured to use the fast in-VM transaction manager. This transaction manager is very fast, but does have two limitations.

- It does not do transactional logging, and is thus incapable of automated recovery after a server crash.
- While it does support propagating transaction contexts with remote calls, it does not support propagating transaction contexts to other virtual machines, so all transactional work must be done in the same virtual machine as the JBoss server.

The corresponding default transaction manager MBean service is the `org.jboss.tm.TransactionManagerService` MBean. It has two configurable attributes:

- **TransactionTimeout:** The default transaction timeout in seconds. The default value is 300 seconds or 5 minutes.
- **XidFactory:** The JMX ObjectName of the MBean service that provides the `org.jboss.tm.XidFactoryMBean` implementation. The `XidFactoryMBean` interface is used to create `javax.transaction.xa.Xid` instances. This is a workaround for XA JDBC drivers that only work with their own Xid implementation. Examples of such drivers are the older Oracle XA drivers. If not specified a JBoss implementation of the Xid interface is used.

4.2.2.1. org.jboss.tm.XidFactory

The `XidFactory` MBean is a factory for `javax.transaction.xa.Xid` instances in the form of `org.jboss.tm.XidImpl`. The `XidFactory` allows for customization of the `XidImpl` that it constructs through the following attributes:

- **BaseGlobalId:** This is used for building globally unique transaction identifiers. This must be set individually if multiple JBoss instances are running on the same machine. The default value is the host name of the

JBoss server, followed by a slash.

- **GlobalIdNumber:** A long value used as initial transaction id. The default is 0.
- **Pad:** The pad value determines whether the `byte[]` returned by the `Xid` `getGlobalTransactionId` and `getBranchQualifier` methods should be equal to maximum 64 byte length or a variable value ≤ 64 . Some resource managers (Oracle for example) require ids that are max length in size.

4.2.3. UserTransaction Support

The JTA `javax.transaction.UserTransaction` interface allows applications to explicitly control transactions. For enterprise session beans that manage transaction themselves (BMT), a `UserTransaction` can be obtained by calling the `getUserTransaction` method on the bean context object, `javax.ejb.SessionContext`.

Note: For BMT beans, do not obtain the `UserTransaction` interface using a JNDI lookup. Doing this violates the EJB specification, and the returned `UserTransaction` object does not have the hooks the EJB container needs to make important checks.

To use the `UserTransaction` interface in other places, the `org.jboss.tm.usertx.server.ClientUserTransactionService` MBean must be configured and started. This MBean publishes a `UserTransaction` implementation under the JNDI name `UserTransaction`. This MBean is configured by default in the standard JBoss distributions and has no configurable attributes.

When the `UserTransaction` is obtained with a JNDI lookup from a stand-alone client (a client operating in a virtual machine than the server's, for example), a very simple `UserTransaction` suitable for thin clients is returned. This `UserTransaction` implementation only controls the transactions on the server the `UserTransaction` object was obtained from. Local transactional work done in the client is not done within the transactions started by this `UserTransaction` object.

When a `UserTransaction` object is obtained by looking up JNDI name `UserTransaction` in the same virtual machine as JBoss, a simple interface to the JTA `TransactionManager` is returned. This is suitable for web components running in web containers embedded in JBoss. When components are deployed in an embedded web server, the deployer will make a JNDI link from the standard `java:comp/UserTransaction` ENC name to the global `UserTransaction` binding so that the web components can lookup the `UserTransaction` instance under JNDI name as specified by the J2EE.

EJBs on JBoss

The EJB Container Configuration and Architecture

The JBoss EJB container architecture is a fourth generation design that emphasizes a modular plug-in approach. All key aspects of the EJB container may be replaced by custom versions of a plug-in and/or an interceptor by a developer. This approach allows for fine tuned customization of the EJB container behavior to optimally suite your needs. Most of the EJB container behavior is configurable through the EJB JAR `META-INF/jboss.xml` descriptor and the default server-wide equivalent `standardjboss.xml` descriptor. We will look at various configuration capabilities throughout this chapter as we explore the container architecture.

5.1. The EJB Client Side View

We will begin our tour of the EJB container by looking at the client view of an EJB through the home and remote proxies. It is the responsibility of the container provider to generate the `javax.ejb.EJBHome` and `javax.ejb.EJBObject` for an EJB implementation. A client never references an EJB bean instance directly, but rather references the `EJBHome` which implements the bean home interface, and the `EJBObject` which implements the bean remote interface. Figure 5.1 shows the composition of an EJB home proxy and its relation to the EJB deployment.

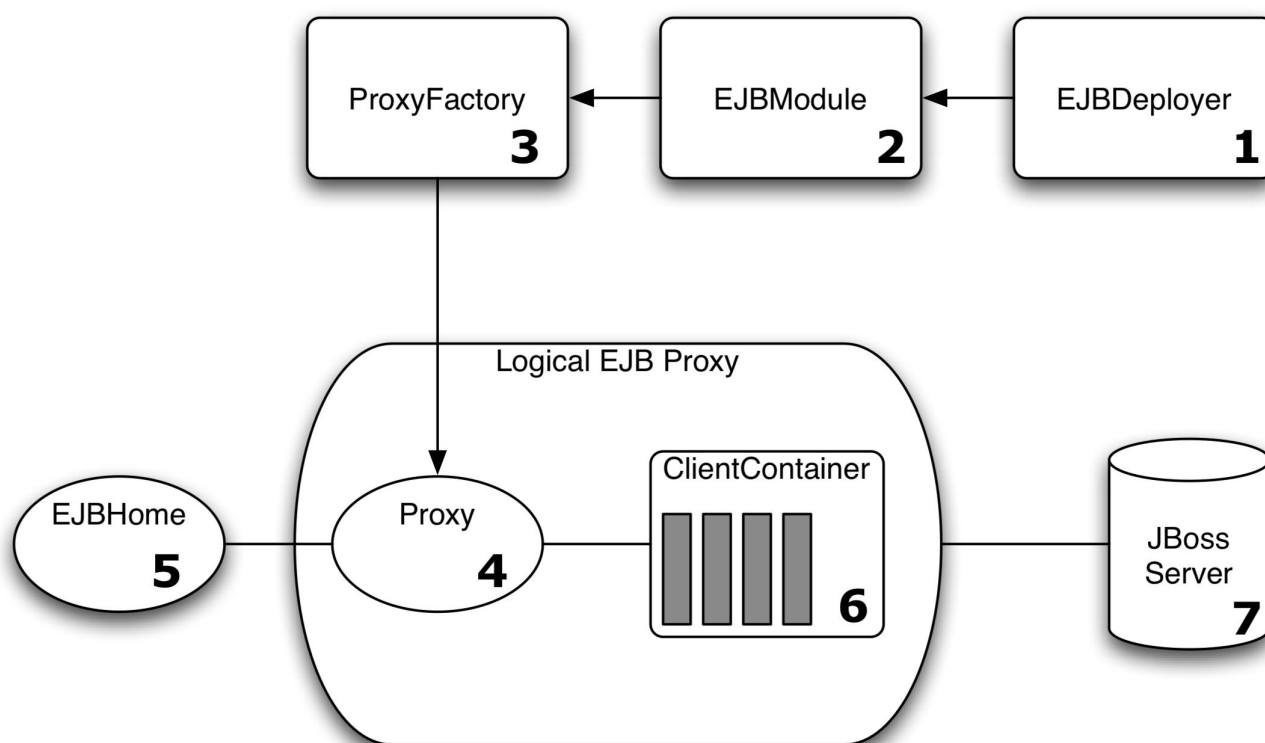


Figure 5.1. The composition of an EJBHome proxy in JBoss.

The numbered items in the figure are:

1. The `EJBDeployer` (`org.jboss.ejb.EJBDeployer`) is invoked to deploy an EJB JAR. An `EJBModule` (`org.jboss.ejb.EJBModule`) is created to encapsulate the deployment metadata.
2. The create phase of the `EJBModule` life cycle creates an `EJBProxyFactory` (`org.jboss.ejb.EJBProxyFactory`) that manages the creation of EJB home and remote interface proxies based on the `EJBModuleInvoker-proxy-bindings` metadata. There can be multiple proxy factories associated with an EJB and we will look at how this is defined shortly.
3. The `ProxyFactory` constructs the logical proxies and binds the homes into JNDI. A logical proxy is composed of a dynamic `Proxy` (`java.lang.reflect.Proxy`), the home interfaces of the EJB that the proxy exposes, the `ProxyHandler` (`java.lang.reflect.InvocationHandler`) implementation in the form of the `ClientContainer` (`org.jboss.proxy.ClientContainer`), and the client side interceptors.
4. The proxy created by the `EJBProxyFactory` is a JDK 1.3+ dynamic proxy. It is a serializable object that proxies the EJB home and remote interfaces as defined in the `EJBModule` metadata. The proxy translates requests made through the strongly typed EJB interfaces into a detyped invocation using the `ClientContainer` handler associated with the proxy. It is the dynamic proxy instance that is bound into JNDI as the EJB home interface that clients lookup. When a client does a lookup of an EJB home, the home proxy is transported into the client VM along with the `ClientContainer` and its interceptors. The use of dynamic proxies avoids the EJB specific compilation step required by many other EJB containers.
5. The EJB home interface is declared in the `ejb-jar.xml` descriptor and available from the `EJBModule` metadata. A key property of dynamic proxies is that they are seen to implement the interfaces they expose. This is true in the sense of Java's strong type system. A proxy can be cast to any of the home interfaces and reflection on the proxy provides the full details of the interfaces it proxies.
6. The proxy delegates calls made through any of its interfaces to the `ClientContainer` handler. The single method required of the handler is: `public Object invoke(Object proxy, Method m, Object[] args) throws Throwable`. The `EJBProxyFactory` creates a `ClientContainer` and assigns this as the `ProxyHandler`. The `ClientContainer`'s state consists of an `InvocationContext` (`org.jboss.invocation.InvocationContext`) and a chain of interceptors (`org.jboss.proxy.Interceptor`). The `InvocationContext` contains:
 - the JMX `ObjectName` of the EJB container MBean the `Proxy` is associated with
 - the `javax.ejb.EJBMetaData` for the EJB
 - the JNDI name of the EJB home interface
 - the transport specific invoker (`org.jboss.invocation.Invoker`)

The interceptor chain consists of the functional units that make up the EJB home or remote interface behavior. This is a configurable aspect of an EJB as we will see when we discuss the `jboss.xml` descriptor, and the interceptor makeup is contained in the `EJBModule` metadata. Interceptors (`org.jboss.proxy.Interceptor`) handle the different EJB types, security, transactions and transport. You can add your own interceptors as well.

7. The transport specific invoker associated with the proxy has an association to the server side detached invoker that handles the transport details of the EJB method invocation. The detached invoker is a JBoss server side component.

The configuration of the client side interceptors is done using the `jboss.xml` `client-interceptors` element. Figure 5.2 shows the subset of the `jboss.xml` DTD for the client interceptors. When the `ClientContainer` `invoke` method is called it creates an untyped `Invocation` (`org.jboss.invocation.Invocation`) to encapsulate

request. This is then passed through the interceptor chain. The last interceptor in the chain will be the transport handler that knows how to send the request to the server and obtain the reply, taking care of the transport specific details.

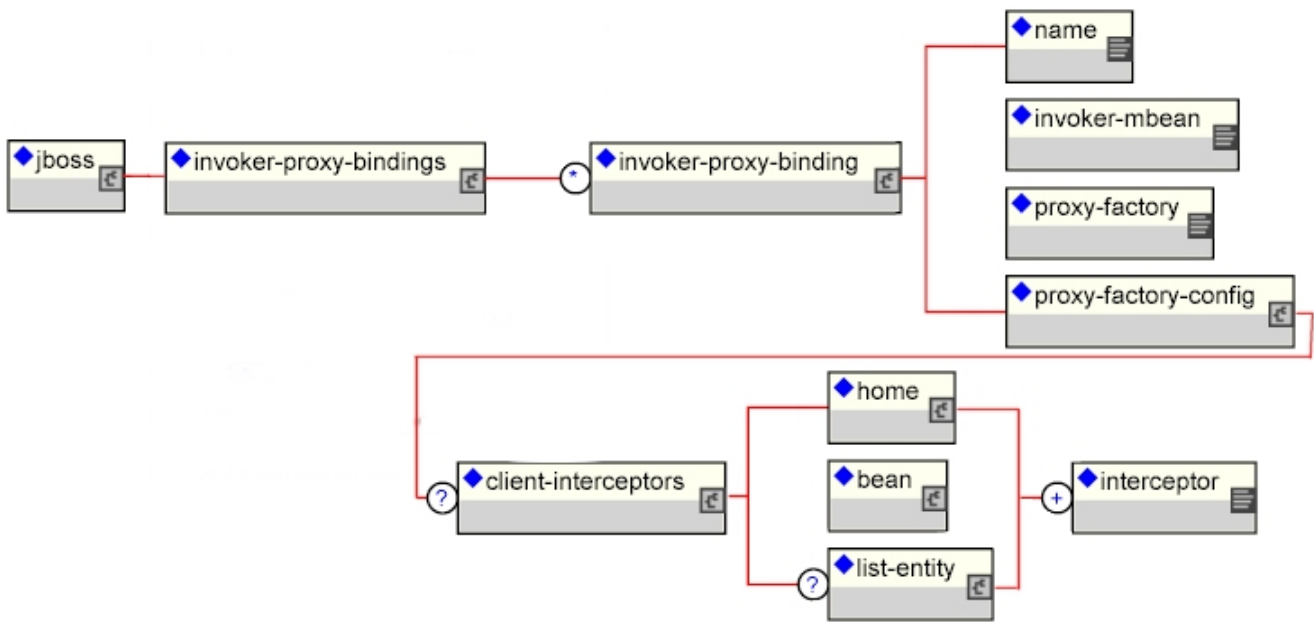


Figure 5.2. The jboss.xml descriptor client side interceptor configuration elements.

As an example of the client interceptor configuration usage, consider the default stateless session bean configuration found in the `server/default/standardjboss.xml` descriptor. Example 5.1 shows the `stateless-rmi-invoker` client interceptors configuration referenced by the Standard Stateless SessionBean.

Example 5.1. The client-interceptors from the Standard Stateless SessionBean configuration.

```

<invoker-proxy-bindings>
  <invoker-proxy-binding>
    <name>stateless-rmi-invoker</name>
    <invoker-mbean>jboss:service=invoker,type=jrmp</invoker-mbean>
    <proxy-factory>org.jboss.proxy.ejb.ProxyFactory</proxy-factory>
    <proxy-factory-config>
      <client-interceptors>
        <home>
          <interceptor>org.jboss.proxy.ejb.HomeInterceptor</interceptor>
          <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
          <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
          <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
        </home>
        <bean>
          <interceptor>
            org.jboss.proxy.ejb.StatelessSessionInterceptor
          </interceptor>
          <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
          <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
          <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
        </bean>
      </client-interceptors>
    </proxy-factory-config>
  </invoker-proxy-binding>
  <!-- ... -->
</invoker-proxy-bindings>
<container-configuration>
  <container-name>Standard Stateless SessionBean</container-name>
  <call-logging>>false</call-logging>
</container-configuration>

```

```
<invoker-proxy-binding-name>stateless-rmi-invoker</invoker-proxy-binding-name>
<!-- ... -->
</container-configuration>
</invoker-proxy-bindings>
</invoker-proxy-bindings>
```

This is the client interceptor configuration for stateless session beans that is used in the absence of an EJB JAR `META-INF/jboss.xml` configuration that overrides these settings. The functionality provided by each interceptor is:

- **org.jboss.proxy.ejb.HomeInterceptor:** this handles the `getHomeHandle`, `getEJBMetaData`, and `remove` methods of the `EJBHome` interface locally in the client VM. Any other methods are propagated to the next interceptor.
- **org.jboss.proxy.ejb.StatelessSessionInterceptor:** this handles the `toString`, `equals`, `hashCode`, `getHandle`, `getEJBHome` and `isIdentical` methods of the `EJBObject` interface locally in the client VM. Any other methods are propagated to the next interceptor.
- **org.jboss.proxy.SecurityInterceptor:** this associates the current security context with the method invocation for use by other interceptors or the server.
- **org.jboss.proxy.TransactionInterceptor:** this associates any active transaction with the invocation method invocation for use by other interceptors.
- **org.jboss.invocation.InvokerInterceptor:** this interceptor encapsulates the dispatch of the method invocation to the transport specific invoker. It knows if the client is executing in the same VM as the server and will optimally route the invocation to a by reference invoker in this situation. When the client is external to the server VM, this interceptor delegates the invocation to the transport invoker associated with the invocation context. In the case of the Example 5.1 configuration, this would be the invoker stub associated with the `jboss:service=invoker,type=jrmp`, the `JRMPInvoker` service.

5.1.1. Specifying the EJB Proxy Configuration

To specify the EJB invocation transport and the client proxy interceptor stack, you need to define an `invoker-proxy-binding` in either the EJB JAR `META-INF/jboss.xml` descriptor, or the server `standardjboss.xml` descriptor. There are several default `invoker-proxy-bindings` defined in the `standardjboss.xml` descriptor for the various default EJB container configurations and the standard RMI/JRMP and RMI/IIOP transport protocols. The current default proxy configurations are:

- **entity-rmi-invoker:** a RMI/JRMP configuration for entity beans
- **clustered-entity-rmi-invoker:** a RMI/JRMP configuration for clustered entity beans
- **stateless-rmi-invoker:** a RMI/JRMP configuration for stateless session beans
- **clustered-stateless-rmi-invoker:** a RMI/JRMP configuration for clustered stateless session beans
- **stateful-rmi-invoker:** a RMI/JRMP configuration for clustered stateful session beans
- **clustered-stateful-rmi-invoker:** a RMI/JRMP configuration for clustered stateful session beans
- **message-driven-bean:** a JMS invoker for message driven beans

- **iiop**: a RMI/IIOP for use with session and entity beans.

To introduce a new protocol binding, or customize the proxy factory, or the client side interceptor stack, requires defining a new `invoker-proxy-binding`. The full `invoker-proxy-binding` DTD fragment for the specification of the proxy configuration is given in Figure 5.3.

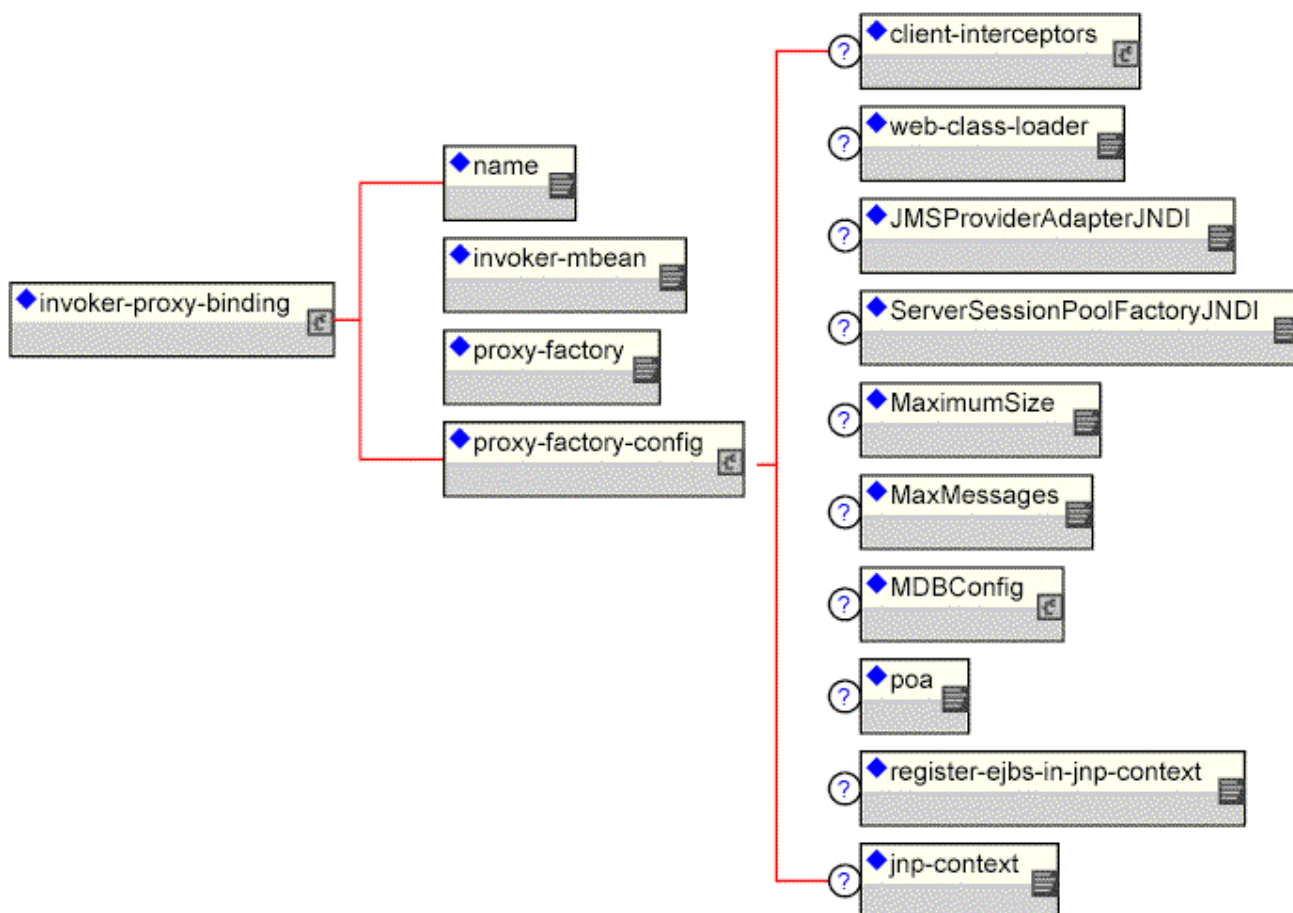


Figure 5.3. The `invoker-proxy-binding` schema

The `invoker-proxy-binding` child elements are:

- **name**: The `name` element gives a unique name for the `invoker-proxy-binding`. The name is used to reference the binding from the EJB container configuration when setting the default proxy binding as well as the EJB deployment level to specify addition proxy bindings. You will see how this is done when we look at the `jboss.xml` elements that control the server side EJB container configuration.
- **invoker-mbean**: The `invoker-mbean` element gives the JMX `ObjectName` string of the detached invoker MBean service the proxy invoker will be associated with.
- **proxy-factory**: The `proxy-factory` element specifies the fully qualified class name of the proxy factory, which must implement the `org.jboss.ejb.EJBProxyFactory` interface. The `EJBProxyFactory` handles the configuration of the proxy and the association of the protocol specific invoker and context. The current JBoss implementations of the `EJBProxyFactory` interface include:
 - **org.jboss.proxy.ejb.ProxyFactory**: The RMI/JRMP specific factory.
 - **org.jboss.proxy.ejb.ProxyFactoryHA**: The cluster RMI/JRMP specific factory.

- **org.jboss.ejb.plugins.jms.JMSContainerInvoker:** The JMS specific factory.
- **org.jboss.proxy.ejb.IORFactory:** The RMI/IIOP specific factory.
- **proxy-factory-config:** The proxy-factory-config element specifies additional information for the proxy-factory implementation. Unfortunately, its currently an unstructured collection of elements. Only a fraction of the elements apply to each type of proxy factory. The child elements break down into the three invocation protocols: RMI/RJMP, RMI/IIOP and JMS.

For the RMI/RJMP specific proxy factories, `org.jboss.proxy.ejb.ProxyFactory` and `org.jboss.proxy.ejb.ProxyFactoryHA` the following elements apply:

- **client-interceptors:** The client-interceptors define the home, remote and optionally the multi-valued proxy interceptor stacks.
- **web-class-loader:** The web class loader defines the instance of the `org.jboss.web.WebClassLoader` that should be associated with the proxy for dynamic class loading.

The following example gives a sample proxy-factory-config taken from the `standardjboss.xml` descriptor.

```
<invoker-proxy-bindings>
  <invoker-proxy-binding>
    <name>stateless-rmi-invoker</name>
    <invoker-mbean>jboss:service=invoker,type=jrmp</invoker-mbean>
    <proxy-factory>org.jboss.proxy.ejb.ProxyFactory</proxy-factory>
    <proxy-factory-config>
      <client-interceptors>
        <home>
          <interceptor>org.jboss.proxy.ejb.HomeInterceptor</interceptor>
          <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
          <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
          <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
        </home>
        <bean>
          <interceptor>
            org.jboss.proxy.ejb.StatelessSessionInterceptor
          </interceptor>
          <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
          <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
          <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
        </bean>
      </client-interceptors>
    </proxy-factory-config>
  </invoker-proxy-binding>
  <!-- ... -->
</invoker-proxy-bindings>
<container-configuration>
  <container-name>Standard Stateless SessionBean</container-name>
  <call-logging>false</call-logging>
  <invoker-proxy-binding-name>stateless-rmi-invoker</invoker-proxy-binding-name>
  <!-- ... -->
</container-configuration>
</invoker-proxy-bindings>
</invoker-proxy-bindings>
```

For the RMI/IIOP specific proxy factory, `org.jboss.proxy.ejb.IORFactory`, the following elements apply:

- **poa:** The portable object adapter usage, one of per-servent, shared
- **register-ejbs-in-jnp-context:** A flag indicating if the EJBs should register in JNDI.

- **jnp-context:** The JNDI context in which to register EJBs.
- **web-class-loader:** The web class loader defines the instance of the `org.jboss.web.WebClassLoader` that should be associated with the proxy for dynamic class loading.

Example 5.2 gives a sample `proxy-factory-config` fragment taken from the `standardjboss.xml` descriptor.

Example 5.2. A sample IOFactory proxy-factory-config

```
<proxy-factory-config>
  <web-class-loader>org.jboss.iiop.WebCL</web-class-loader>
  <poa>per-servant</poa>
  <register-ejbs-in-jnp-context>true</register-ejbs-in-jnp-context>
  <jnp-context>iiop</jnp-context>
</proxy-factory-config>
```

For the JMS specific proxy factory, `org.jboss.ejb.plugins.jms.JMSContainerInvoker`, there is an `MDB-Config`

- **MaximumSize:** This specifies the upper limit to the number of concurrent MDBs that will be allowed for the JMS destination associated with a given MDB deployment. This defaults to 15.
- **MaxMessages:** This specifies the `maxMessages` parameter value for the `createConnectionConsumer` method of `javax.jms.QueueConnection` and `javax.jms.TopicConnection` interfaces, as well as the `maxMessages` parameter value for the `createDurableConnectionConsumer` method of `javax.jms.TopicConnection`. It is the maximum number of messages that can be assigned to a server session at one time. This defaults to 1. This value should not be modified from the default unless your JMS provider indicates this is supported.
- **MDBConfig:** Configuration for the MDB JMS connection behavior. This include the reconnection interval and dead letter queue elements:
- **ReconnectIntervalSec:** The time to wait (in seconds) before trying to recover the connection to the JMS server.
- **DLQConfig:** Configuration for an MDB's dead letter queue, used when messages are redelivered too many times.
- **JMSProviderAdapterJNDI:** The JNDI name of the JMS provider adapter in the `java:/` namespace. This is mandatory for an MDB and must implement `org.jboss.jms.jndi.JMSProviderAdapter`.
- **ServerSessionPoolFactoryJNDI:** The JNDI name of the session pool in the `java:/` namespace of the JMS provider's session pool factory. This is mandatory for an MDB and must implement `org.jboss.jms.asf.ServerSessionPoolFactory`.

Example 5.3 gives a sample `proxy-factory-config` fragment taken from the `standardjboss.xml` descriptor.

Example 5.3. A sample JMSContainerInvoker proxy-factory-config

```
<proxy-factory-config>
  <JMSProviderAdapterJNDI>DefaultJMSProvider</JMSProviderAdapterJNDI>
  <ServerSessionPoolFactoryJNDI>StdJMSPool</ServerSessionPoolFactoryJNDI>
  <MaximumSize>15</MaximumSize>
  <MaxMessages>1</MaxMessages>
```

```

<MDBConfig>
  <ReconnectIntervalSec>10</ReconnectIntervalSec>
  <DLQConfig>
    <DestinationQueue>queue/DLQ</DestinationQueue>
    <MaxTimesRedelivered>10</MaxTimesRedelivered>
    <TimeToLive>0</TimeToLive>
  </DLQConfig>
</MDBConfig>
</proxy-factory-config>

```

5.2. The EJB Server Side View

Every EJB invocation must end up at a JBoss server hosted EJB container. In this section we will look at how invocations are transported to the JBoss server VM and find their way to the EJB container via the JMX bus.

5.2.1. Detached Invokers - The Transport Middlemen

We looked at the detached invoker architecture in the context of exposing RMI compatible interfaces of MBean services earlier. Here we will look at how detached invokers are used to expose the EJB container home and bean interfaces to clients. The generic view of the invoker architecture is presented in Figure 5.4.

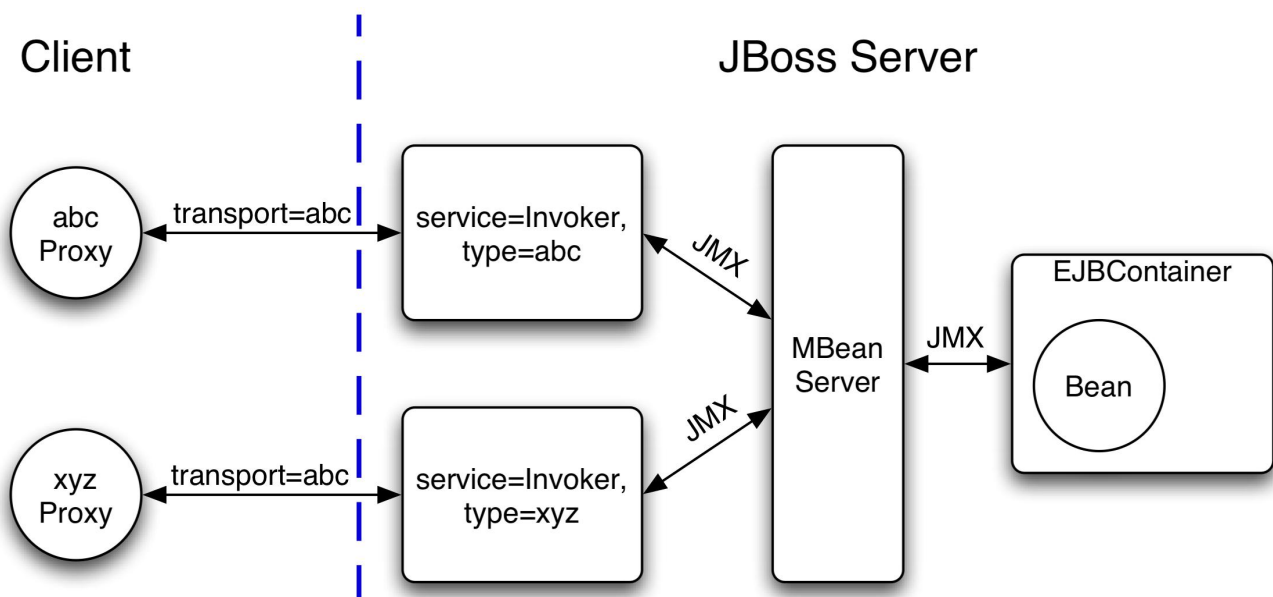


Figure 5.4. The transport invoker server side architecture

For each type of home proxy there is a binding to an invoker and its associated transport protocol. A container may have multiple invocation protocols active simultaneously. The `jboss.xml` DTD configuration fragments for the invoker configuration is given in Figure 5.5. The `invoker-proxy-binding-name` maps to an `invoker-proxy-binding/name` element. At the `container-configuration` level this specifies the default invoker that will be used for EJBs deployed to the container. At the bean level, the `invoker-bindings` specify one or more invokers to use with the EJB container MBean.

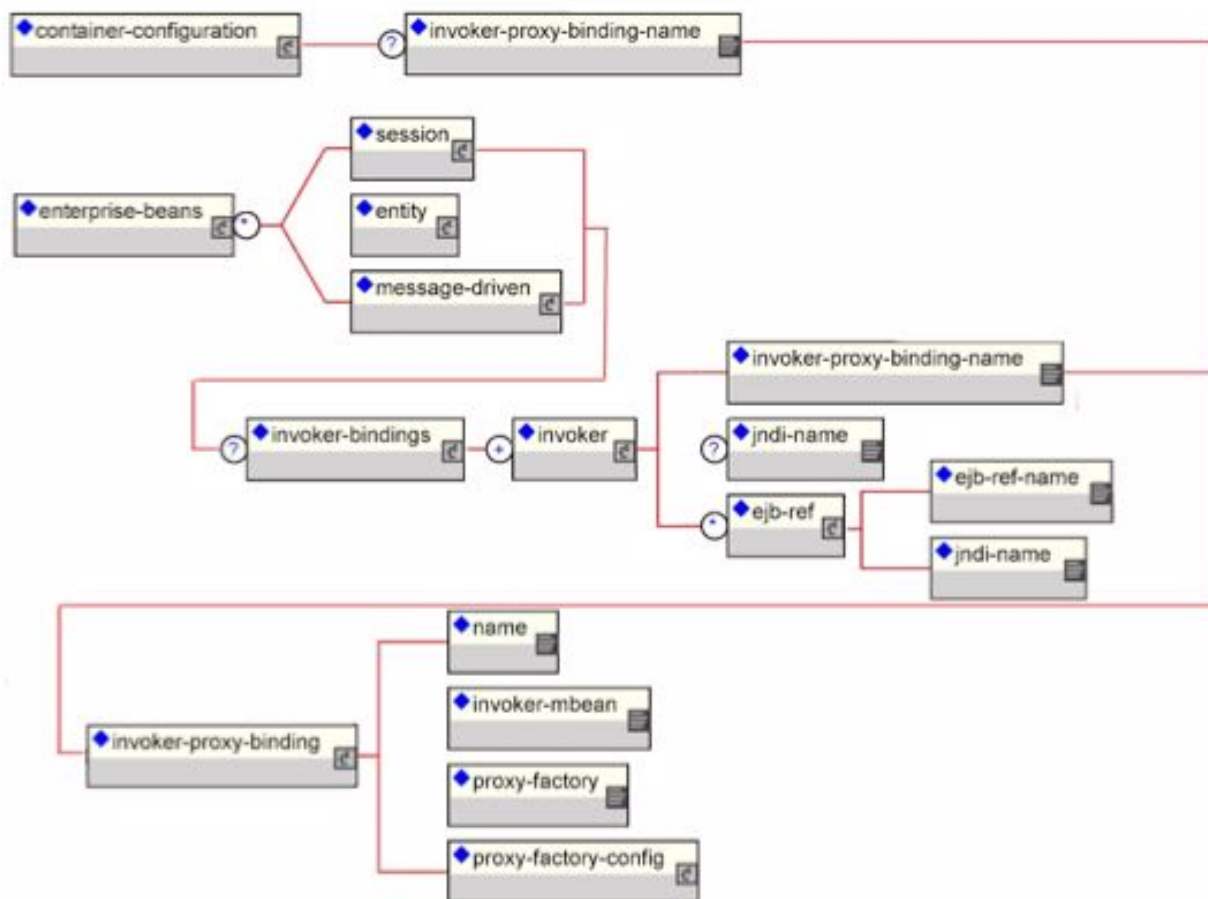


Figure 5.5. The jboss.xml descriptor invoker configuration elements.

When one specifies multiple invokers for a given EJB deployment, the home proxy must be given a unique JNDI binding location. This is specified by the `invoker/jndi-name` element value. Another issue when multiple invokers exist for an EJB is how to handle remote homes or interfaces obtained when the EJB calls other beans. Any such interfaces need to use the same invoker used to call the outer EJB in order for the resulting remote homes and interfaces to be compatible with the proxy the client has initiated the call through. The `invoker/ejb-ref` elements allow one to map from a protocol independent `ENC_ejb-ref` to the home proxy binding for `ejb-ref` target EJB home that matches the referencing invoker type.

An example of using a custom `JRMPInvoker` MBean that enables compressed sockets for session beans can be found in the `org.jboss.test.jrmp` package of the testsuite. The following example illustrates the custom `JRMPInvoker` configuration and its mapping to a stateless session bean.

Example 5.4. The custom `JRMPInvoker` `jboss-service.xml` descriptor

```

<server>
  <mbean code="org.jboss.invocation.jrmp.server.JRMPInvoker"
    name="jboss:service=invoker,type=jrmp,socketType=CompressionSocketFactory">
    <attribute name="RMIObjPort">4445</attribute>
    <attribute name="RMIClientSocketFactory">
      org.jboss.test.jrmp.ejb.CompressionClientSocketFactory
    </attribute>
    <attribute name="RMIServerSocketFactory">
      org.jboss.test.jrmp.ejb.CompressionServerSocketFactory
    </attribute>
  </mbean>

```

```
</server>
```

Example 5.5. The jboss.xml descriptor using the custom invoker

```
<?xml version="1.0"?>
<!DOCTYPE jboss PUBLIC
    "-//JBoss//DTD JBOSS 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jboss_3_2.dtd">
<!-- The jboss.xml descriptor for the jrmp-comp.jar ejb unit -->
<jboss>
    <enterprise-beans>
        <session>
            <ejb-name>StatelessSession</ejb-name>
            <configuration-name>Standard Stateless SessionBean</configuration-name>
            <invoker-bindings>
                <invoker>
                    <invoker-proxy-binding-name>
                        stateless-compression-invoker
                    </invoker-proxy-binding-name>
                    <jndi-name>jrmp-compressed/StatelessSession</jndi-name>
                </invoker>
            </invoker-bindings>
        </session>
    </enterprise-beans>
    <invoker-proxy-bindings>
        <invoker-proxy-binding>
            <name>stateless-compression-invoker</name>
            <invoker-mbean>
                jboss:service=invoker,type=jrmp,socketType=CompressionSocketFactory
            </invoker-mbean>
            <proxy-factory>org.jboss.proxy.ejb.ProxyFactory</proxy-factory>
            <proxy-factory-config>
                <client-interceptors>
                    <home>
                        <interceptor>org.jboss.proxy.ejb.HomeInterceptor</interceptor>
                        <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
                        <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
                        <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
                    </home>
                    <bean>
                        <interceptor>
                            org.jboss.proxy.ejb.StatelessSessionInterceptor
                        </interceptor>
                        <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
                        <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
                        <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
                    </bean>
                </client-interceptors>
            </proxy-factory-config>
        </invoker-proxy-binding>
    </invoker-proxy-bindings>
</jboss>
```

Here the default `JRMPInvoker` has been customized to bind to port 4445 and to use custom socket factories that enable compression at the transport level. The `StatelessSession` EJB `invoker-bindings` settings specify that the `stateless-compression-invoker` will be used with the home interface bound under the JNDI name `jrmp-compressed/StatelessSession`.

An example of using the `HttpInvoker` to configure a stateless session bean to use the RMI/HTTP protocol can be found in the `org.jboss.test.hello` testsuite package. Example 5.6 illustrates the custom settings.

Example 5.6. A sample jboss.xml descriptor for enabling RMI/HTTP for a stateless session bean.

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE jboss PUBLIC
    "-//JBoss//DTD JBOSS 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jboss_3_2.dtd">
<jboss>
  <enterprise-beans>
    <session>
      <ejb-name>HelloWorldViaHTTP</ejb-name>
      <jndi-name>helloworld/HelloHTTP</jndi-name>
      <invoker-bindings>
        <invoker>
          <invoker-proxy-binding-name>
            stateless-http-invoker
          </invoker-proxy-binding-name>
        </invoker>
      </invoker-bindings>
    </session>
  </enterprise-beans>
  <invoker-proxy-bindings>
    <!-- A custom invoker for RMI/HTTP -->
    <invoker-proxy-binding>
      <name>stateless-http-invoker</name>
      <invoker-mbean>jboss:service=invoker,type=http</invoker-mbean>
      <proxy-factory>org.jboss.proxy.ejb.ProxyFactory</proxy-factory>
      <proxy-factory-config>
        <client-interceptors>
          <home>
            <interceptor>org.jboss.proxy.ejb.HomeInterceptor</interceptor>
            <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
            <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
            <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
          </home>
          <bean>
            <interceptor>
              org.jboss.proxy.ejb.StatelessSessionInterceptor
            </interceptor>
            <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
            <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
            <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
          </bean>
        </client-interceptors>
      </proxy-factory-config>
    </invoker-proxy-binding>
  </invoker-proxy-bindings>
</jboss>

```

Here a custom invoker-proxy-binding named `stateless-http-invoker` is defined. It uses the `HttpInvoker` MBean as the detached invoker. The `jboss:service=invoker,type=http` name is the default name of the `HttpInvoker` MBean as found in the `http-inovker.sar/META-INF/jboss-service.xml` descriptor, and its service descriptor fragment is show here:

```

<!-- The HTTP invoker service configuration -->
<mbean code="org.jboss.invocation.http.server.HttpInvoker"
  name="jboss:service=invoker,type=http">
  <!-- Use a URL of the form http://<hostname>:8080/invoker/EJBInvokerServlet
    where <hostname> is InetAddress.getHostname value on which the server
    is running. -->
  <attribute name="InvokerURLPrefix">http://</attribute>
  <attribute name="InvokerURLSuffix">:8080/invoker/EJBInvokerServlet</attribute>
  <attribute name="UseHostName">true</attribute>
</mbean>

```

The client proxy posts the EJB invocation content to the `EJBInvokerServlet` URL specified in the `HttpInvoker` service configuration.

5.2.2. The HA JRMPInvoker - Clustered RMI/JRMP Transport

The `org.jboss.invocation.jrmp.server.JRMPInvokerHA` service is an extension of the `JRMPInvoker` that is a cluster aware invoker. The `JRMPInvokerHA` fully supports all of the attributes of the `JRMPInvoker`. This means that customized bindings of the port, interface and socket transport are available to clustered RMI/JRMP as well. For additional information on the clustering architecture and the implementation of the HA RMI proxies see the JBoss Clustering docs.

5.2.3. The HA HttpInvoker - Clustered RMI/HTTP Transport

The RMI/HTTP layer allows for software load balancing of the invocations in a clustered environment. An HA capable extension of the HTTP invoker has been added that borrows much of its functionality from the HA-RMI/JRMP clustering.

To enable HA-RMI/HTTP you need to configure the invokers for the EJB container. This is done through either a `jboss.xml` descriptor, or the `standardjboss.xml` descriptor. Example 5.7 shows is an example of a stateless session configuration taken from the `org.jboss.test.hello` testsuite package.

Example 5.7. A `jboss.xml` stateless session configuration for HA-RMI/HTTP

```
<jboss>
  <enterprise-beans>
    <session>
      <ejb-name>HelloWorldViaClusteredHTTP</ejb-name>
      <jndi-name>helloworld/HelloHA-HTTP</jndi-name>
      <invoker-bindings>
        <invoker>
          <invoker-proxy-binding-name>
            stateless-httpHA-invoker
          </invoker-proxy-binding-name>
        </invoker>
      </invoker-bindings>
      <clustered>true</clustered>
    </session>
  </enterprise-beans>
  <invoker-proxy-bindings>
    <invoker-proxy-binding>
      <name>stateless-httpHA-invoker</name>
      <invoker-mbean>jboss:service=invoker,type=httpHA</invoker-mbean>
      <proxy-factory>org.jboss.proxy.ejb.ProxyFactoryHA</proxy-factory>
      <proxy-factory-config>
        <client-interceptors>
          <home>
            <interceptor>org.jboss.proxy.ejb.HomeInterceptor</interceptor>
            <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
            <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
            <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
          </home>
          <bean>
            <interceptor>
              org.jboss.proxy.ejb.StatelessSessionInterceptor
            </interceptor>
            <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
            <interceptor>org.jboss.proxy.TransactionInterceptor</interceptor>
            <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
          </bean>
        </client-interceptors>
      </proxy-factory-config>
    </invoker-proxy-binding>
  </invoker-proxy-bindings>
</jboss>
```

```
        </proxy-factory-config>
    </invoker-proxy-binding>
</invoker-proxy-bindings>
</jboss>
```

The `stateless-httpHA-invoker invoker-proxy-binding` references the `jboss:service=invoker,type=httpHA` invoker service. This service is configured in the `http-invoker.sar/META-INF/jboss-service.xml` descriptor, and its default configuration from the SARdescriptor is:

```
<mbean code="org.jboss.invocation.http.server.HttpInvokerHA"
  name="jboss:service=invoker,type=httpHA">
  <!-- Use a URL of the form
    http://<hostname>:8080/invoker/EJBInvokerHAServlet
    where <hostname> is InetAddress.getHostname value on which the server
    is running.
  -->
  <attribute name="InvokerURLPrefix">http://</attribute>
  <attribute name="InvokerURLSuffix">:8080/invoker/EJBInvokerHAServlet</attribute>
  <attribute name="UseHostName">true</attribute>
</mbean>
```

The URL used by the invoker proxy is the `EJBInvokerHAServlet` mapping as deployed on the cluster node. The `HttpInvokerHA` instances across the cluster form a collection of candidate http URLs that are made available to the client side proxy for failover and/or load balancing.

5.3. The EJB Container

An EJB container is the component that manages a particular class of EJB. In JBoss there is one instance of the `org.jboss.ejb.Container` created for each unique configuration of an EJB that is deployed. The actual object that is instantiated is a subclass of `Container` and the creation of the container instance is managed by the `EJBDeployer MBean`.

5.3.1. EJBDeployer MBean

The `org.jboss.ejb.EJBDeployer MBean` is responsible for the creation of EJB containers. Given an EJB JAR that is ready for deployment, the `EJBDeployer` will create and initialize the necessary EJB containers, one for each type of EJB. The configurable attributes of the `EJBDeployer` are:

- **VerifyDeployments:** a boolean flag indicating if the EJB verifier should be run. This validates that the EJBs in a deployment unit conform to the EJB 2.0 specification. Setting this to true is useful for ensuring your deployments are valid.
- **VerifierVerbose:** A boolean that controls the verbosity of any verification failures/warnings that result from the verification process.
- **StrictVerifier:** A boolean that enables/disables strict verification. When strict verification is enable an EJB will deploy only if verifier reports no errors.
- **ValidateDTDs:** a boolean flag that indicates if the `ejb-jar.xml` and `jboss.xml` descriptors should be validated against their declared DTDs. Setting this to true is useful for ensuring your deployment descriptors are valid.
- **MetricsEnabled:** a boolean flag that controls whether container interceptors marked with an `metricsEn-`

`abled=true` attribute should be included in the configuration. This allows one to define a container interceptor configuration that includes metrics type interceptors that can be toggled on and off.

- **WebServiceName:** The JMX ObjectName string of the web service MBean that provides support for the dynamic class loading of EJB classes.
- **TransactionManagerServiceName:** The JMX ObjectName string of the JTA transaction manager service. This must have an attribute named `TransactionManager` that returns that `javax.transaction.TransactionManager` instance.

The deployer contains two central methods: `deploy` and `undeploy`. The `deploy` method takes a URL, which either points to an EJB JAR, or to a directory whose structure is the same as a valid EJB JAR (which is convenient for development purposes). Once a deployment has been made, it can be undeployed by calling `undeploy` on the same URL. A call to `deploy` with an already deployed URL will cause an undeploy, followed by deployment of the URL, such as a re-deploy. JBoss has support for full re-deployment of both implementation and interface classes, and will reload any changed classes. This will allow you to develop and update EJBs without ever stopping a running server.

During the deployment of the EJB JAR the `EJBDeployer` and its associated classes perform three main functions, verify the EJBs, create a container for each unique EJB, initialize the container with the deployment configuration information. We will talk about each function in the following sections.

5.3.1.1. Verifying EJB deployments

When the `VerifyDeployments` attribute of the that the `EJBDeployer` is true, the deployer performs a verification of EJBs in the deployment. The verification checks that an EJB meets EJB specification compliance. This entails validating that the EJB deployment unit contains the required home and remote, local home and local interfaces, and that the objects appearing in these interfaces are of the proper types, and that the required methods are present in the implementation class. This is a useful behavior that is enabled by default since there are a number of steps that an EJB developer and deployer must perform correctly to construct a proper EJB JAR, and it is easy to make a mistake. The verification stage attempts to catch any errors and fail the deployment with an error that indicates what needs to be corrected.

Probably the most problematic aspect of writing EJBs is the fact that there is a disconnection between the bean implementation and its remote and home interfaces, as well as its deployment descriptor configuration. It is easy to have these separate elements get out of synch. One tool that helps eliminate this problem is XDoclet, an extension of the standard JavaDoc Doclet engine. It works off of custom JavaDoc tags in the EJB bean implementation class and creates the remote and home interfaces as well as the deployment descriptors. See the XDoclet home page, <http://sourceforge.net/projects/xdoclet> for additional details.

5.3.1.2. Deploying EJBs Into Containers

The most important role performed by the `EJBDeployer` is the creation of an EJB container and the deployment of the EJB into the container. The deployment phase consists of iterating over EJBs in an EJB JAR, and extracting the bean classes and their metadata as described by the `ejb-jar.xml` and `jboss.xml` deployment descriptors. For each EJB in the EJB jar, the following steps are performed:

- Create subclass of `org.jboss.ejb.Container` depending on the type of the EJB: stateless, stateful, BMP entity, CMP entity, or message driven. The container is assigned a unique `ClassLoader` from which it can load local resources. The uniqueness of the `ClassLoader` is also used to isolate the standard `java:comp` JNDI namespace from other J2EE components.
- Set all container configurable attributes from a merge of the `jboss.xml` and `standardjboss.xml`

descriptors.

- Create and add the container interceptors as configured for the container.
- Associate the container with an application object. This application object represents a J2EE enterprise application and may contain multiple EJBs and web contexts.

If all EJBs are successfully deployed, the application is started which in turn starts all containers and makes the EJBs available to clients. If any EJB fails to deploy, a deployment exception is thrown and the deployment module is failed.

5.3.1.3. Container configuration information

JBoss externalizes most if not all of the setup of the EJB containers using an XML file that conforms to the `jboss_3_2.dtd`. The section of the `jboss_3_2` DTD that relates to container configuration information is shown in Figure 5.6.

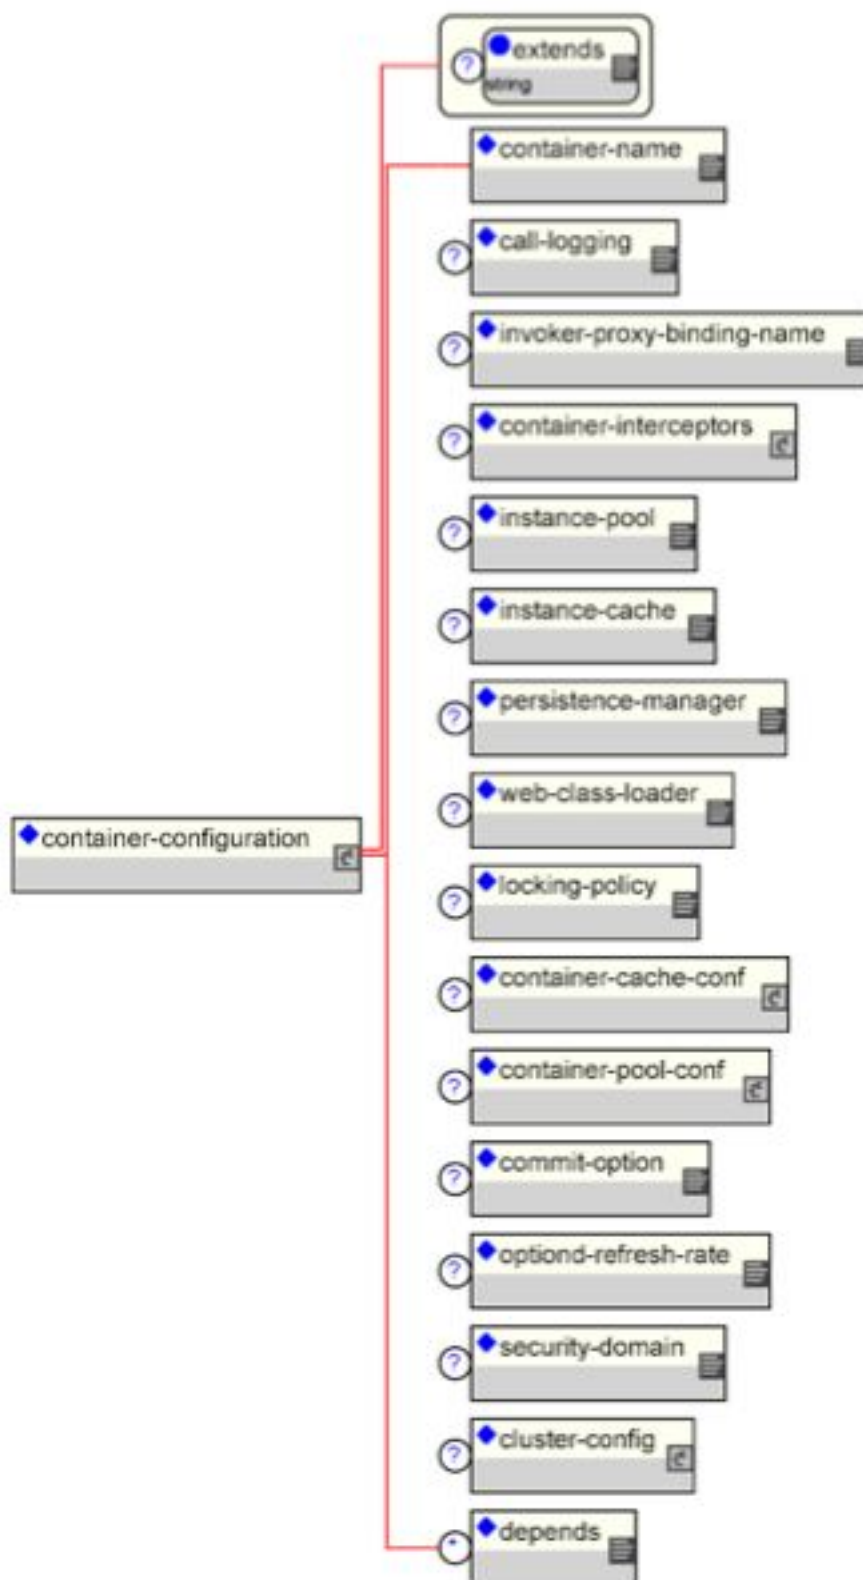


Figure 5.6. The jboss_3_2 DTD elements related to container configuration.

The `container-configuration` element and its subelements specify container configuration settings for a type of container as given by the `container-name` element. Each configuration specifies information such as the default invoker type, the container interceptor makeup, instance caches/pools and their sizes, persistence manager, security, and so on. Because this is a large amount of information that requires a detailed understanding of the JBoss container architecture, JBoss ships with a standard configuration for the four types of EJBs. This configuration file is called `standardjboss.xml` and it is located in the `conf` directory of any configuration file set that

uses EJBs. Example 5.8 gives a sample of a configuration from the standardjboss.xml.

Example 5.8. An example of a complex container-configuration element from the server/default/conf/standardjboss.xml file.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE jboss PUBLIC
    "-//JBoss//DTD JBOSS 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jboss_3_2.dtd">
<jboss>
  <!-- ... -->
  <container-configurations>
    <container-configuration>
      <container-name>Standard CMP 2.x EntityBean</container-name>
      <call-logging>false</call-logging>
      <invoker-proxy-binding-name>entity-rmi-invoker</invoker-proxy-binding-name>
      <sync-on-commit-only>false</sync-on-commit-only>
      <container-interceptors>
        <interceptor>
          org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor
        </interceptor>
        <interceptor>org.jboss.ejb.plugins.LogInterceptor</interceptor>
        <interceptor>org.jboss.ejb.plugins.SecurityInterceptor</interceptor>
        <interceptor>org.jboss.ejb.plugins.TxInterceptorCMT</interceptor>
        <interceptor metricsEnabled="true">
          org.jboss.ejb.plugins.MetricsInterceptor
        </interceptor>
        <interceptor>org.jboss.ejb.plugins.EntityCreationInterceptor</interceptor>
        <interceptor>org.jboss.ejb.plugins.EntityLockInterceptor</interceptor>
        <interceptor>org.jboss.ejb.plugins.EntityInstanceInterceptor</interceptor>
        <interceptor>
          org.jboss.ejb.plugins.EntityReentranceInterceptor
        </interceptor>
        <interceptor>
          org.jboss.resource.connectionmanager.CachedConnectionInterceptor
        </interceptor>
        <interceptor>
          org.jboss.ejb.plugins.EntitySynchronizationInterceptor
        </interceptor>
        <interceptor>
          org.jboss.ejb.plugins.cmp.jdbc.JDBCRelationInterceptor
        </interceptor>
      </container-interceptors>
      <instance-pool>
        org.jboss.ejb.plugins.EntityInstancePool
      </instance-pool>
      <instance-cache>
        org.jboss.ejb.plugins.InvalidableEntityInstanceCache
      </instance-cache>
      <persistence-manager>
        org.jboss.ejb.plugins.cmp.jdbc.JDBCStoreManager
      </persistence-manager>
      <locking-policy>
        org.jboss.ejb.plugins.lock.QueuedPessimisticEJBLock
      </locking-policy>
      <container-cache-conf>
        <cache-policy>
          org.jboss.ejb.plugins.LRUEnterpriseContextCachePolicy
        </cache-policy>
        <cache-policy-conf>
          <min-capacity>50</min-capacity>
          <max-capacity>1000000</max-capacity>
          <overager-period>300</overager-period>
          <max-bean-age>600</max-bean-age>
          <resizer-period>400</resizer-period>
          <max-cache-miss-period>60</max-cache-miss-period>
          <min-cache-miss-period>1</min-cache-miss-period>
          <cache-load-factor>0.75</cache-load-factor>
        </cache-policy-conf>
      </container-cache-conf>
    </container-configuration>
  </container-configurations>
</jboss>
```

```

    </container-cache-conf>
    <container-pool-conf>
        <MaximumSize>100</MaximumSize>
    </container-pool-conf>
    <commit-option>B</commit-option>
</container-configuration>
<!-- ... -->
</container-configurations>
</jboss>

```

These two examples demonstrate how extensive the container configuration options are. The container configuration information can be specified at two levels. The first is in the `standardjboss.xml` file contained in the configuration file set directory. The second is at the EJB JAR level. By placing a `jboss.xml` file in the EJB JAR `META-INF` directory, you can specify either overrides for container configurations in the `standardjboss.xml` file, or entirely new named container configurations. This provides great flexibility in the configuration of containers. As you have seen, all container configuration attributes have been externalized and as such are easily modifiable. Knowledgeable developers can even implement specialized container components, such as instance pools or caches, and easily integrate them with the standard container configurations to optimize behavior for a particular application or environment.

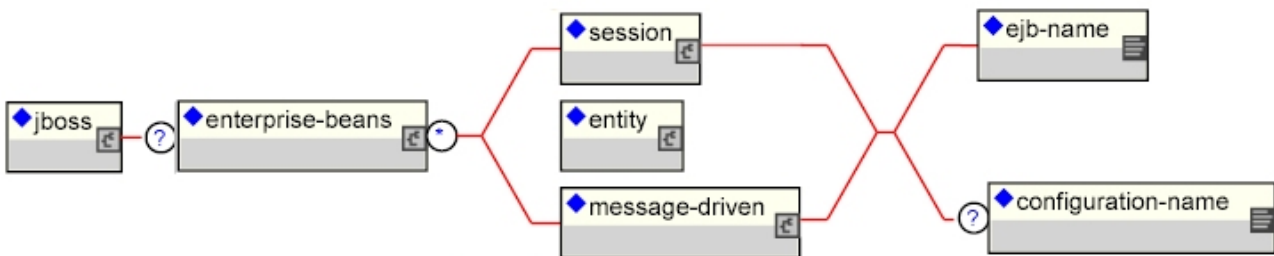


Figure 5.7. The `jboss.xml` descriptor EJB to container configuration mapping elements

How an EJB deployment chooses its container configuration is based on the explicit or implicit `jboss/enterprise-beans/<type>/configuration-name` element. Figure 5.7 shows the `jboss.xml` DTD fragment that shows how an EJB can declare which container configuration it should use.

The `configuration-name` element is a link to a `container-configurations/container-configuration` element in Figure 5.6. It specifies which container configuration to use for the referring EJB. The link is from a `configuration-name` element to a `container-name` element. You are able to specify container configurations per class of EJB by including a `container-configuration` element in the EJB definition. Typically one does not define completely new container configurations, although this is supported. The typical usage of a `jboss.xml` level `container-configuration` is to override one or more aspects of a `container-configuration` coming from the `standardjboss.xml` descriptor. This is done by specifying `container-configuration` that references the name of an existing `standardjboss.xml` `container-configuration/container-name` as the value for the `container-configuration/extends` attribute. Example 5.9 shows an example of defining a new Secured Stateless SessionBean configuration that is an extension of the standard stateless session configuration whose name is `Standard Stateless SessionBean`.

Example 5.9. An example of overriding the `standardjboss.xml` container stateless session beans configuration to enable secured access.

```

<?xml version="1.0"?>
<jboss>
    <enterprise-beans>
        <session>

```

```

        <ejb-name>EchoBean</ejb-name>
        <configuration-name>Secured Stateless SessionBean</configuration-name>
        <!-- ... -->
    </session>
</enterprise-beans>
<container-configurations>
    <container-configuration extends="Standard Stateless SessionBean">
        <container-name>Secured Stateless SessionBean</container-name>
        <!-- Override the container security domain -->
        <security-domain>java:/jaas/my-security-domain</security-domain>
    </container-configuration>
</container-configurations>
</jboss>

```

If an EJB does not provide a container configuration specification in the deployment unit EJB JAR, the container factory chooses a container configuration from the `standardjboss.xml` descriptor based on the type of the EJB. So, in reality there is an implicit `configuration-name` element for every type of EJB, and the mappings from the EJB type to default container configuration name are as follows:

- container-managed persistence entity version 2.0 = Standard CMP 2.x EntityBean
- container-managed persistence entity version 1.1 = Standard CMP EntityBean
- bean-managed persistence entity = Standard BMP EntityBean
- stateless session = Standard Stateless SessionBean
- stateful session = Standard Stateful SessionBean
- message driven = Standard Message Driven Bean

It is not necessary to indicate which container configuration an EJB is using if you want to use the default based on the bean type. It probably provides for a more self-contained descriptor to include the `configuration-name` element, but this is a matter of style.

Now that you know how to specify which container configuration an EJB is using, and that you can define a deployment unit level override, the question is what are all of those `container-configuration` child elements? This question will be addressed element by element in the following sections. A number of the elements specify interface class implementations whose configuration is affected by other elements, so before starting in on the configuration elements you need to understand the `org.jboss.metadata.XmlLoadable` interface.

The `XmlLoadable` interface is a simple interface that consists of a single method. The interface definition is:

```

import org.w3c.dom.Element;
public interface XmlLoadable
{
    public void importXml(Element element) throws Exception;
}

```

Classes implement this interface to allow their configuration to be specified via an XML document fragment. The root element of the document fragment is what would be passed to the `importXml` method. You will see a few examples of this as the container configuration elements are described in the following sections.

5.3.1.3.1. The container-name Element

The `container-name` element specifies a unique name for a given configuration. EJBs link to a particular container configuration by setting their `configuration-name` element to the value of the `container-name` for the

container configuration.

5.3.1.3.2. The call-logging Element

The `call-logging` element expects a boolean (true or false) as its value to indicate whether or not the `LogInterceptor` should log method calls to a container. This is somewhat obsolete with the change to `log4j`, which provides a fine-grained logging API.

5.3.1.3.3. The invoker-proxy-binding-name Element

The `invoker-proxy-binding-name` element specifies the name of the default invoker to use. In the absence of a bean level `invoker-bindings` specification, the `invoker-proxy-binding` whose name matches the `invoker-proxy-binding-name` element value will be used to create home and remote proxies.

5.3.1.3.4. The container-interceptors Element

The `container-interceptors` element specifies one or more interceptor elements that are to be configured as the method interceptor chain for the container. The value of the interceptor element is a fully qualified class name of an `org.jboss.ejb.Interceptor` interface implementation. The container interceptors form a linked-list structure through which EJB method invocations pass. The first interceptor in the chain is invoked when the `MBeanServer` passes a method invocation to the container. The last interceptor invokes the business method on the bean. We will discuss the `Interceptor` interface latter in this chapter when we talk about the container plugin framework. Generally, care must be taken when changing an existing standard EJB interceptor configuration as the EJB contract regarding security, transactions, persistence, and thread safety derive from the interceptors.

5.3.1.3.5. The instance-pool and container-pool-conf Elements

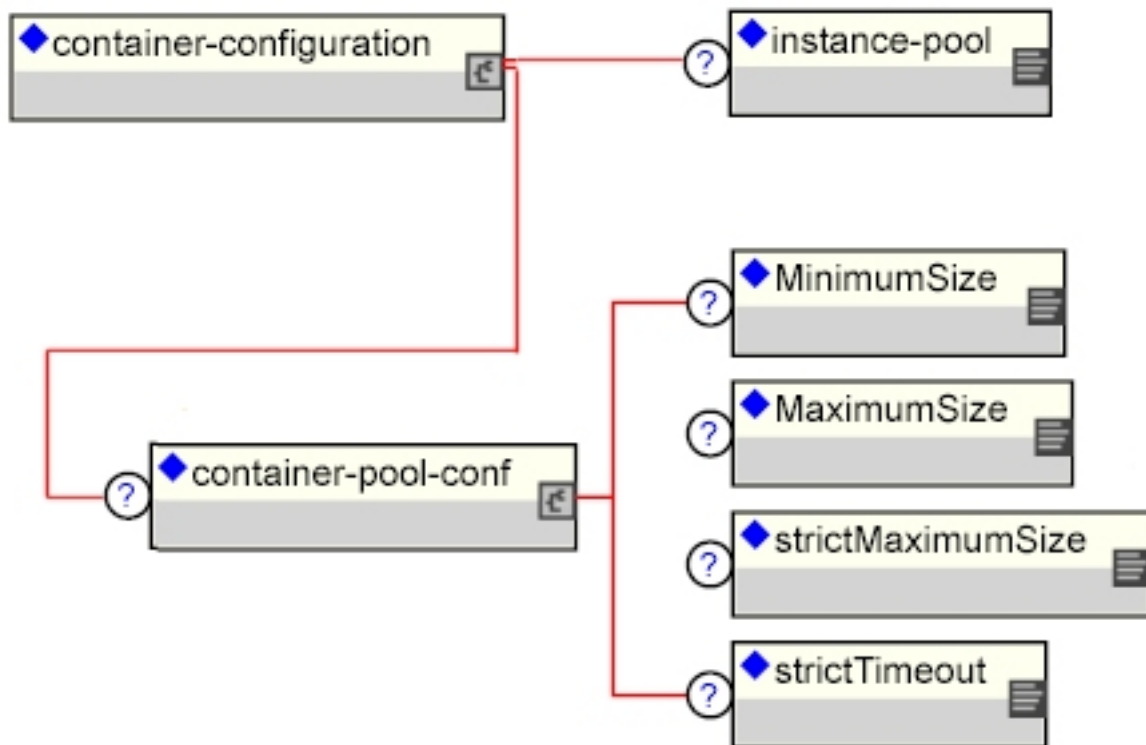


Figure 5.8. The instance-pool and container-pool-conf elements

The `instance-pool` element specifies the fully qualified class name of an `org.jboss.ejb.InstancePool` interface implementation to use as the container `InstancePool`. We will discuss the `InstancePool` interface in detail later in this chapter when we talk about the container plugin framework.

The `container-pool-conf` is passed to the `InstancePool` implementation class given by the `instance-pool` element if it implements `XmlLoadable` interface. All current JBoss `InstancePool` implementations derive from the `org.jboss.ejb.plugins.AbstractInstancePool` class and it provides support for the `MinimumSize`, `MaximumSize`, `strictMaximumSize` and `strictTimeout` `container-pool-conf` child elements. The `MinimumSize` element gives the minimum number of instances to keep in the pool, although JBoss does not currently seed an `InstancePool` to the `MinimumSize` value.

The `MaximumSize` specifies the maximum number of pool instances that are allowed. The default use of `MaximumSize` may not be what you expect. The pool `MaximumSize` is the maximum number of EJB instances that are kept available, but additional instances can be created if the number of concurrent requests exceeds the `MaximumSize` value. If you want to limit the maximum concurrency of an EJB to the pool `MaximumSize`, you need to set the `strictMaximumSize` element to `true`. When `strictMaximumSize` is `true`, only `MaximumSize` EJB instances may be active. When there are `MaximumSize` active instances, any subsequent requests will be blocked until an instance is freed back to the pool. The default value for `strictMaximumSize` is `false`. How long a request blocks waiting for an instance pool object is controlled by the `strictTimeout` element. The `strictTimeout` defines the time in milliseconds to wait for an instance to be returned to the pool when there are `MaximumSize` active instances. A value less than or equal to 0 will mean not to wait at all. When a request times out waiting for an instance a `java.rmi.ServerException` is generated and the call aborted. This is parsed as a `Long` so the maximum possible wait time is 9,223,372,036,854,775,807 or about 292,471,208 years, and this is the default value.

5.3.1.3.6. The instance-cache and container-cache-conf Elements

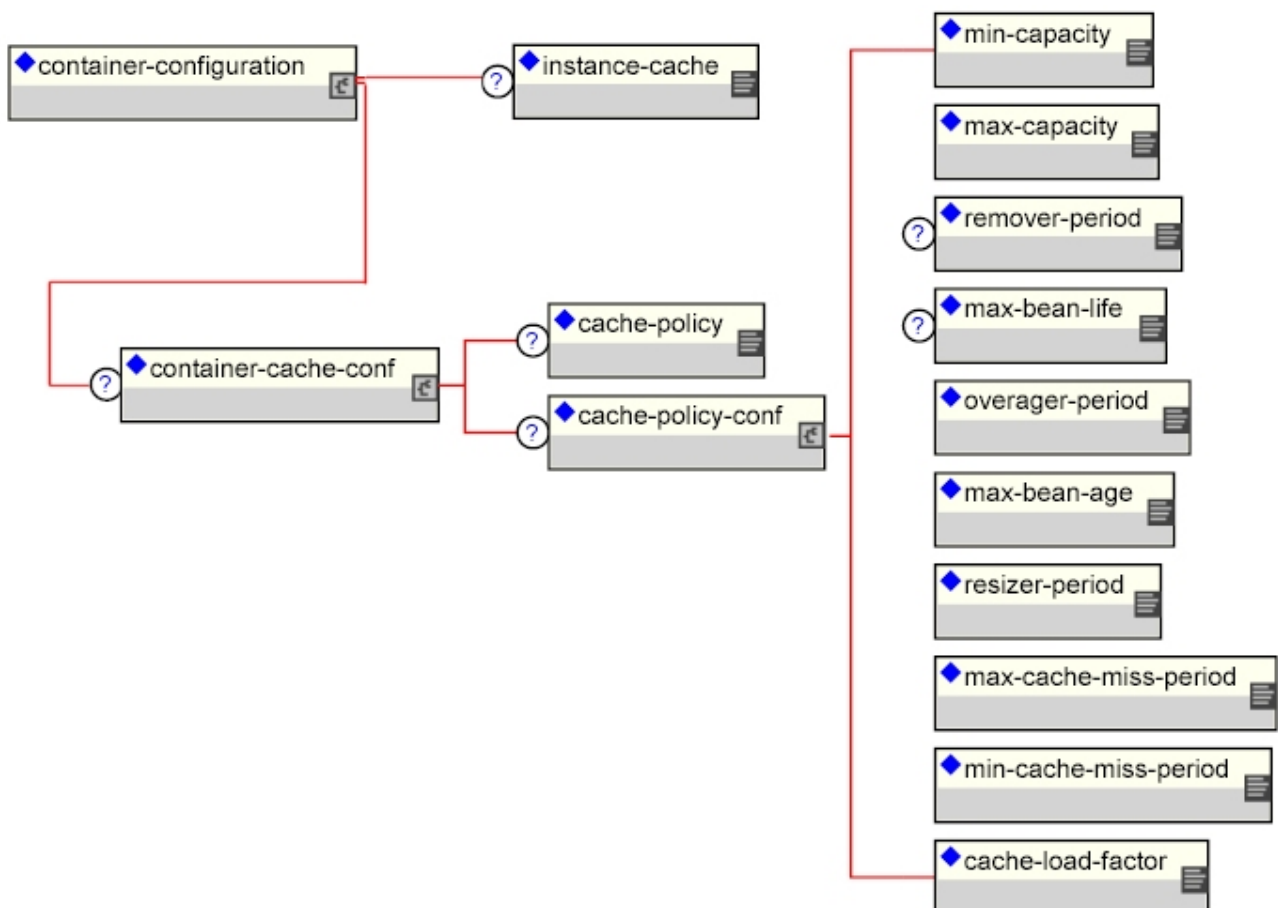


Figure 5.9. The instance-cache and container-cache-conf and related elements

The `instance-cache` element specifies the fully qualified class name of the `org.jboss.ejb.InstanceCache` interface implementation. This element is only meaningful for entity and stateful session beans as these are the only EJB types that have an associated identity. We will discuss the `InstanceCache` interface in detail latter in this chapter when we talk about the container plugin framework.

The `container-cache-conf` element is passed to the `InstanceCache` implementation if it supports the `XmlLoadable` interface. All current JBoss `InstanceCache` implementations derive from the `org.jboss.ejb.plugins.AbstractInstanceCache` class and it provides support for the `XmlLoadable` interface and uses the `cache-policy` child element as the fully qualified class name of an `org.jboss.util.CachePolicy` implementation that is used as the instance cache store. The `cache-policy-conf` child element is passed to the `CachePolicy` implementation if it supports the `XmlLoadable` interface. If it does not, the `cache-policy-conf` will silently be ignored.

There are two JBoss implementations of `CachePolicy` used by the `standardjboss.xml` configuration that support the current array of `cache-policy-conf` child elements. The classes are `org.jboss.ejb.plugins.LRUEnterpriseContextCachePolicy` and `org.jboss.ejb.plugins.LRUStatefulContextCachePolicy`. The `LRUEnterpriseContextCachePolicy` is used by entity bean containers while the `LRUStatefulContextCachePolicy` is used by stateful session bean containers. Both cache policies support the following `cache-policy-conf` child elements:

- **min-capacity:** specifies the minimum capacity of this cache
- **max-capacity:** specifies the maximum capacity of the cache, which cannot be less than `min-capacity`.
- **overager-period:** specifies the period in seconds between runs of the overager task. The purpose of the overager task is to see if the cache contains beans with an age greater than the `max-bean-age` element value. Any beans meeting this criterion will be passivated.
- **max-bean-age:** specifies the maximum period of inactivity in seconds a bean can have before it will be passivated by the overager process.
- **resizer-period:** specifies the period in seconds between runs of the resizer task. The purpose of the resizer task is to contract or expand the cache capacity based on the remaining three element values in the following way. When the resizer task executes it checks the current period between cache misses, and if the period is less than the `min-cache-miss-period` value the cache is expanded up to the `max-capacity` value using the `cache-load-factor`. If instead the period between cache misses is greater than the `max-cache-miss-period` value the cache is contracted using the `cache-load-factor`.
- **max-cache-miss-period:** specifies the time period in seconds in which a cache miss should signal that the cache capacity be contracted. It is equivalent to the minimum miss rate that will be tolerated before the cache is contracted.
- **min-cache-miss-period:** specifies the time period in seconds in which a cache miss should signal that the cache capacity be expanded. It is equivalent to the maximum miss rate that will be tolerated before the cache is expanded.
- **cache-load-factor:** specifies the factor by which the cache capacity is contracted and expanded. The factor should be less than 1. When the cache is contracted the capacity is reduced so that the current ratio of beans to cache capacity is equal to the `cache-load-factor` value. When the cache is expanded the new capacity is determined as `current-capacity * 1/cache-load-factor`. The actual expansion factor may be as high as

2 based on an internal algorithm based on the number of cache misses. The higher the cache miss rate the closer the true expansion factor will be to 2.

The `LRUStatefulContextCachePolicy` also supports the remaining child elements:

- **remover-period:** specifies the period in seconds between runs of the remover task. The remover task removes passivated beans that have not been accessed in more than `max-bean-life` seconds. This task prevents stateful session beans that were not removed by users from filling up the passivation store.
- **max-bean-life:** specifies the maximum period of inactivity in seconds that a bean can exist before being removed from the passivation store.

An alternative cache policy implementation is the `org.jboss.ejb.plugins.NoPassivationCachePolicy` class, which simply never passivates instances. It uses an in-memory `HashMap` implementation that never discards instances unless they are explicitly removed. This class does not support any of the `cache-policy-conf` configuration elements.

5.3.1.3.7. The persistence-manager Element

The `persistence-manager` element value specifies the fully qualified class name of the persistence manager implementation. The type of the implementation depends on the type of EJB. For stateful session beans it must be an implementation of the `org.jboss.ejb.StatefulSessionPersistenceManager` interface. For BMP entity beans it must be an implementation of the `org.jboss.ejb.EntityPersistenceManager` interface, while for CMP entity beans it must be an implementation of the `org.jboss.ejb.EntityPersistenceStore` interface.

5.3.1.3.8. The web-class-loader Element

The `web-class-loader` element specifies a subclass of `org.jboss.web.WebClassLoader` that is used in conjunction with the `WebService` MBean to allow dynamic loading of resources and classes from deployed ears, EJB JARs and wars. A `WebClassLoader` is associated with a `Container` and must have an `org.jboss.mx.loading.UnifiedClassLoader` as its parent. It overrides the `getURLs()` method to return a different set of URLs for remote loading than what is used for local loading.

`WebClassLoader` has two methods meant to be overridden by subclasses: `getKey()` and `getBytes()`. The latter is a no-op in this implementation and should be overridden by subclasses with bytecode generation ability, such as the classloader used by the `iiop` module.

`WebClassLoader` subclasses must have a constructor with the same signature as the `WebClassLoader(ObjectName containerName, UnifiedClassLoader parent)` constructor.

5.3.1.3.9. The locking-policy Element

The `locking-policy` element gives the fully qualified class name of the EJB lock implementation to use. This class must implement the `org.jboss.ejb.BeanLock` interface. The current JBoss versions include:

- **org.jboss.ejb.plugins.lock.QueuedPessimisticEJBLock:** an implementation that holds threads awaiting the transactional lock to be freed in a fair FIFO queue. Non-transactional threads are also put into this wait queue as well. This class pops the next waiting transaction from the queue and notifies only those threads waiting associated with that transaction. The `QueuedPessimisticEJBLock` is the current default used by the standard configurations.
- **org.jboss.ejb.plugins.lock.SimpleReadWriteEJBLock:** This lock allows multiple read locks concurrently. Once a writer has requested the lock, future read-lock requests whose transactions do not already

have the read lock will block until all writers are done; then all the waiting readers will concurrently go (depending on the `reentrant` setting / `methodLock`). A reader who promotes gets first crack at the write lock, ahead of other waiting writers. If there is already a reader that is promoting, we throw an inconsistent read exception. Of course, writers have to wait for all read-locks to release before taking the write lock.

- **org.jboss.ejb.plugins.lock.NoLock**: an anti-locking policy used with the instance per transaction container configurations.

We will talk in more detail about the locking policy usage in Section 5.4.

5.3.1.3.10. The `commit-option` and `optiond-refresh-rate` Element

The `commit-option` value specifies the EJB entity bean persistent storage commit option. It must be one of A, B, C or D. The meaning of the option values is:

- **A**: the container caches the beans state between transactions. This option assumes that the container is the only user accessing the persistent store. This assumption allows the container to synchronize the in-memory state from the persistent storage only when absolutely necessary. This occurs before the first business method executes on a found bean or after the bean is passivated and reactivated to serve another business method. This behavior is independent of whether the business method executes inside a transaction context.
- **B**: the container caches the bean state between transactions. However, unlike option A the container does not assume exclusive access to the persistent store. Therefore, the container will synchronize the in-memory state at the beginning of each transaction. Thus, business methods executing in a transaction context don't see much benefit from the container caching the bean, whereas business methods executing outside a transaction context (transaction attributes `Never`, `NotSupported` or `Supports`) access the cached (and potentially invalid) state of the bean.
- **C**: the container does not cache bean instances. The in-memory state must be synchronized on every transaction start. For business methods executing outside a transaction the synchronization is still performed, but the `ejbLoad` executes in the same transaction context as that of the caller.
- **D**: is a JBoss specific feature which is not described in the EJB specification. It is a lazy read scheme where bean state is cached between transactions as with option A, but the state is periodically resynchronized with that of the persistent store. The default time between reloads is 30 seconds, but may configured using the `optiond-refresh-rate` element.

5.3.1.3.11. The `security-domain` Element

Inside the EJB `org.jboss.ejb.Container` class, the `security-domain` element specifies the JNDI name of the object that implements the `org.jboss.security.AuthenticationManager` and `org.jboss.security.RealmMapping` interfaces. Usually one specifies the `security-domain` globally under the `jboss` root element so that all EJBs in a given deployment are secured. The details of the security manager interfaces and configuring the security layer are discussed in Chapter 8.

5.3.1.3.12. `cluster-config`

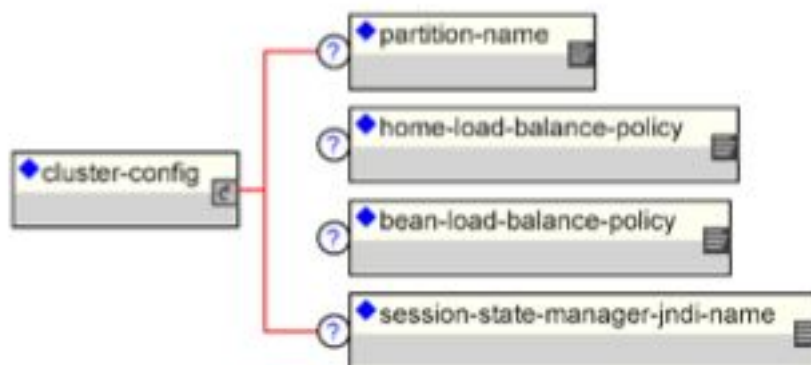


Figure 5.10. The cluster-config and related elements

The `cluster-config` element allows to specify cluster specific settings for all EJBs that use the container configuration. Specification of the cluster configuration may be done at the container configuration level or at the individual EJB deployment level.

The `partition-name` element indicates where to find the `org.jboss.ha.framework.interfaces.HAPartition` interface to be used by the container to exchange clustering information. This is not the full JNDI name under which `HAPartition` is bound. Rather, it should correspond to the `PartitionName` attribute of the `ClusterPartitionMBean` service that is managing the desired cluster. The actual JNDI name of the `HAPartition` binding will be formed by appending `/HASessionState/` to the `partition-name` value. The default value is `DefaultPartition`.

The `home-load-balance-policy` element indicates the Java class name to be used to load balance calls made on the home proxy. The class must implement the `org.jboss.ha.framework.interface.LoadBalancePolicy` interface. The default policy is `org.jboss.ha.framework.interfaces.RoundRobin`.

The `bean-load-balance-policy` element indicates the java class name to be used to load balance calls in the bean proxy. The class must implement the `org.jboss.ha.framework.interface.LoadBalancePolicy` interface. For entity beans and stateful session beans, the default is `org.jboss.ha.framework.interfaces.FirstAvailavble`. For stateless session beans, `org.jboss.ha.framework.interfaces.RoundRobin`.

The `session-state-manager-jndi-name` element indicates the name of the `org.jboss.ha.framework.interfaces.HASessionState` to be used by the container as a backend for state session management in the cluster. Unlike the `partition-name` element, this is a JNDI name under which the `HASessionState` implementation is bound. The default location used is `/HASessionState/Default`.

5.3.1.3.13. depends

The `depends` element gives a JMX `ObjectName` of a service on which the container or EJB depends. Specification of explicit dependencies on other services avoids having to rely on the deployment order being after the required services are started.

5.3.2. Container Plug-in Framework

The JBoss EJB container uses a framework pattern that allows one to change implementations of various aspects of the container behavior. The container itself does not perform any significant work other than connecting the various behavioral components together. Implementations of the behavioral components are referred to as plugins, because you can plug in a new implementation by changing a container configuration. Examples of

plug-in behavior you may want to change include persistence management, object pooling, object caching, container invokers and interceptors. There are four subclasses of the `org.jboss.ejb.Container` class, each one implementing a particular bean type:

- **`org.jboss.ejb.EntityContainer`** handles `javax.ejb.EntityBean` types
- **`org.jboss.ejb.StatelessSessionContainer`** handles Stateless `javax.ejb.SessionBean` types
- **`org.jboss.ejb.StatefulSessionContainer`** handles Stateful `javax.ejb.SessionBean` types
- **`org.jboss.ejb.MessageDrivenContainer`** handles `javax.ejb.MessageDrivenBean` types

The EJB containers delegate much of their behavior to components known as container plug-ins. The interfaces that make up the container plugin points include the following:

- `org.jboss.ejb.ContainerPlugin`
- `org.jboss.ejb.ContainerInvoker`
- `org.jboss.ejb.Interceptor`
- `org.jboss.ejb.InstancePool`
- `org.jboss.ejb.InstanceCache`
- `org.jboss.ejb.EntityPersistenceManager`
- `org.jboss.ejb.EntityPersistenceStore`
- `org.jboss.ejb.StatefulSessionPersistenceManager`

The container's main responsibility is to manage its plug-ins. This means ensuring that the plug-ins have all the information they need to implement their functionality.

5.3.2.1. `org.jboss.ejb.ContainerPlugin`

The `ContainerPlugin` interface is the parent interface of all container plug-in interfaces. It provides a callback that allows a container to provide each of its plug-ins a pointer to the container the plug-in is working on behalf of. The `ContainerPlugin` interface is given below.

Example 5.10. The `org.jboss.ejb.ContainerPlugin` interface

```
public interface ContainerPlugin
    extends org.jboss.system.Service
{
    /**
     * This callback is set by the container so that the plugin
     * may access its container
     *
     * @param con the container which owns the plugin
     */
    public void setContainer(Container con);
}
```

5.3.2.2. `org.jboss.ejb.Interceptor`

The `Interceptor` interface enables one to build a chain of method interceptors through which each EJB method invocation must pass. The `Interceptor` interface is given below.

Example 5.11. The `org.jboss.ejb.Interceptor` interface

```
import org.jboss.invocation.Invocation;

public interface Interceptor
    extends ContainerPlugin
{
    public void setNext(Interceptor interceptor);
    public Interceptor getNext();
    public Object invokeHome(Invocation mi) throws Exception;
    public Object invoke(Invocation mi) throws Exception;
}
```

All interceptors defined in the container configuration are created and added to the container interceptor chain by the `EJBDeployer`. The last interceptor is not added by the deployer but rather by the container itself because this is the interceptor that interacts with the EJB bean implementation.

The order of the interceptor in the chain is important. The idea behind ordering is that interceptors that are not tied to a particular `EnterpriseContext` instance are positioned before interceptors that interact with caches and pools.

Implementers of the `Interceptor` interface form a linked-list like structure through which the `Invocation` object is passed. The first interceptor in the chain is invoked when an invoker passes a `Invocation` to the container via the JMX bus. The last interceptor invokes the business method on the bean. There are usually on the order of five interceptors in a chain depending on the bean type and container configuration. `Interceptor` semantic complexity ranges from simple to complex. An example of a simple interceptor would be `LoggingInterceptor`, while a complex example is `EntitySynchronizationInterceptor`.

One of the main advantages of an interceptor pattern is flexibility in the arrangement of interceptors. Another advantage is the clear functional distinction between different interceptors. For example, logic for transaction and security is cleanly separated between the `TXInterceptor` and `SecurityInterceptor` respectively.

If any of the interceptors fail, the call is terminated at that point. This is a fail-quickly type of semantic. For example, if a secured EJB is accessed without proper permissions, the call will fail as the `SecurityInterceptor` before any transactions are started or instances caches are updated.

5.3.2.3. org.jboss.ejb.InstancePool

An `InstancePool` is used to manage the EJB instances that are not associated with any identity. The pools actually manage subclasses of the `org.jboss.ejb.EnterpriseContext` objects that aggregate unassociated bean instances and related data.

Example 5.12. The org.jboss.ejb.InstancePool interface

```
public interface InstancePool
    extends ContainerPlugin
{
    /**
     * Get an instance without identity. Can be used
     * by finders and create-methods, or stateless beans
     *
     * @return Context /w instance
     * @exception RemoteException
     */
    public EnterpriseContext get() throws Exception;

    /** Return an anonymous instance after invocation.
     *
     * @param ctx
     */
}
```

```

public void free(EnterpriseContext ctx);

/**
 * Discard an anonymous instance after invocation.
 * This is called if the instance should not be reused,
 * perhaps due to some exception being thrown from it.
 *
 * @param ctx
 */
public void discard(EnterpriseContext ctx);

/**
 * Return the size of the pool.
 *
 * @return the size of the pool.
 */
public int getCurrentSize();

/**
 * Get the maximum size of the pool.
 *
 * @return the size of the pool.
 */
public int getMaxSize();
}

```

Depending on the configuration, a container may choose to have a certain size of the pool contain recycled instances, or it may choose to instantiate and initialize an instance on demand.

The pool is used by the `InstanceCache` implementation to acquire free instances for activation, and it is used by interceptors to acquire instances to be used for Home interface methods (create and finder calls).

5.3.2.4. `org.jboss.ejb.InstanceCache`

The container `InstanceCache` implementation handles all EJB-instances that are in an active state, meaning bean instances that have an identity attached to them. Only entity and stateful session beans are cached, as these are the only bean types that have state between method invocations. The cache key of an entity bean is the bean primary key. The cache key for a stateful session bean is the session id.

Example 5.13. The `org.jboss.ejb.InstanceCache` interface

```

public interface InstanceCache
    extends ContainerPlugin
{
    /**
     * Gets a bean instance from this cache given the identity.
     * This method may involve activation if the instance is not
     * in the cache.
     * Implementation should have O(1) complexity.
     * This method is never called for stateless session beans.
     *
     * @param id the primary key of the bean
     * @return the EnterpriseContext related to the given id
     * @exception RemoteException in case of illegal calls
     * (concurrent / reentrant), NoSuchObjectException if
     * the bean cannot be found.
     * @see #release
     */
    public EnterpriseContext get(Object id)
        throws RemoteException, NoSuchObjectException;

    /**
     * Inserts an active bean instance after creation or activation.
     */
}

```

```

    * Implementation should guarantee proper locking and O(1) complexity.
    *
    * @param ctx the EnterpriseContext to insert in the cache
    * @see #remove
    */
    public void insert(EnterpriseContext ctx);

    /**
     * Releases the given bean instance from this cache.
     * This method may passivate the bean to get it out of the cache.
     * Implementation should return almost immediately leaving the
     * passivation to be executed by another thread.
     *
     * @param ctx the EnterpriseContext to release
     * @see #get
     */
    public void release(EnterpriseContext ctx);

    /**
     * Removes a bean instance from this cache given the identity.
     * Implementation should have O(1) complexity and guarantee
     * proper locking.
     *
     * @param id the primary key of the bean
     * @see #insert
     */
    public void remove(Object id);

    /**
     * Checks whether an instance corresponding to a particular
     * id is active
     *
     * @param id the primary key of the bean
     * @see #insert
     */
    public boolean isActive(Object id);
}

```

In addition to managing the list of active instances, the `InstanceCache` is also responsible for activating and passivating instances. If an instance with a given identity is requested, and it is not currently active, the `InstanceCache` must use the `InstancePool` to acquire a free instance, followed by the persistence manager to activate the instance. Similarly, if the `InstanceCache` decides to passivate an active instance, it must call the persistence manager to passivate it and release the instance to the `InstancePool`.

5.3.2.5. org.jboss.ejb.EntityPersistenceManager

The `EntityPersistenceManager` is responsible for the persistence of `EntityBeans`. This includes the following:

- Creating an EJB instance in a storage
- Loading the state of a given primary key into an EJB instance
- Storing the state of a given EJB instance
- Removing an EJB instance from storage
- Activating the state of an EJB instance
- Passivating the state of an EJB instance

Example 5.14. The `org.jboss.ejb.EntityPersistenceManager` interface

```

public interface EntityPersistenceManager
    extends ContainerPlugin
{
    /**
     * Returns a new instance of the bean class or a subclass of the

```

```

    * bean class.
    *
    * @return the new instance
    */
Object createBeanClassInstance() throws Exception;

/**
 * This method is called whenever an entity is to be created. The
 * persistence manager is responsible for calling the ejbCreate method
 * on the instance and to handle the results properly wrt the persistent
 * store.
 *
 * @param m the create method in the home interface that was
 * called
 * @param args any create parameters
 * @param instance the instance being used for this create call
 */
void createEntity(Method m,
                  Object[] args,
                  EntityEnterpriseContext instance)
    throws Exception;

/**
 * This method is called whenever an entity is to be created. The
 * persistence manager is responsible for calling the ejbPostCreate method
 * on the instance and to handle the results properly wrt the persistent
 * store.
 *
 * @param m the create method in the home interface that was
 * called
 * @param args any create parameters
 * @param instance the instance being used for this create call
 */
void postCreateEntity(Method m,
                     Object[] args,
                     EntityEnterpriseContext instance)
    throws Exception;

/**
 * This method is called when single entities are to be found. The
 * persistence manager must find out whether the wanted instance is
 * available in the persistence store, and if so it shall use the
 * ContainerInvoker plugin to create an EJBObject to the instance, which
 * is to be returned as result.
 *
 * @param finderMethod the find method in the home interface that was
 * called
 * @param args any finder parameters
 * @param instance the instance to use for the finder call
 * @return an EJBObject representing the found entity
 */
Object findEntity(Method finderMethod,
                  Object[] args,
                  EntityEnterpriseContext instance)
    throws Exception;

/**
 * This method is called when collections of entities are to be
 * found. The persistence manager must find out whether the wanted
 * instances are available in the persistence store, and if so it
 * shall use the ContainerInvoker plugin to create EJBObjects to
 * the instances, which are to be returned as result.
 *
 * @param finderMethod the find method in the home interface that was
 * called
 * @param args any finder parameters
 * @param instance the instance to use for the finder call
 * @return an EJBObject collection representing the found
 * entities
 */
Collection findEntities(Method finderMethod,

```

```

        Object[] args,
        EntityEnterpriseContext instance)
        throws Exception;

/**
 * This method is called when an entity shall be activated. The
 * persistence manager must call the ejbActivate method on the
 * instance.
 *
 * @param instance the instance to use for the activation
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void activateEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is called whenever an entity shall be load from the
 * underlying storage. The persistence manager must load the state
 * from the underlying storage and then call ejbLoad on the
 * supplied instance.
 *
 * @param instance the instance to synchronize
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void loadEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is used to determine if an entity should be stored.
 *
 * @param instance the instance to check
 * @return true, if the entity has been modified
 * @throws Exception thrown if some system exception occurs
 */
boolean isModified(EntityEnterpriseContext instance) throws Exception;

/**
 * This method is called whenever an entity shall be stored to the
 * underlying storage. The persistence manager must call ejbStore
 * on the supplied instance and then store the state to the
 * underlying storage.
 *
 * @param instance the instance to synchronize
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void storeEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is called when an entity shall be passivate. The
 * persistence manager must call the ejbPassivate method on the
 * instance.
 *
 * @param instance the instance to passivate
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void passivateEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is called when an entity shall be removed from the
 * underlying storage. The persistence manager must call ejbRemove
 * on the instance and then remove its state from the underlying
 * storage.
 *
 * @param instance the instance to remove
 *

```

```

    * @throws RemoteException thrown if some system exception occurs
    * @throws RemoveException thrown if the instance could not be removed
    */
    void removeEntity(EntityEnterpriseContext instance)
        throws RemoteException, RemoveException;
}

```

As per the EJB 2.0 specification, JBoss supports two entity bean persistence semantics: Container Managed Persistence (CMP) and Bean Managed Persistence (BMP). The CMP implementation uses an implementation of the `org.jboss.ejb.EntityPersistenceStore` interface. By default this is the `org.jboss.ejb.plugins.cmp.jdbc.JDBCStoreManager` which is the entry point for the CMP2 persistence engine. The `EntityPersistenceStore` interface is shown below.

Example 5.15. The `org.jboss.ejb.EntityPersistenceStore` interface

```

public interface EntityPersistenceStore
    extends ContainerPlugin
{
    /**
     * Returns a new instance of the bean class or a subclass of the
     * bean class.
     *
     * @return the new instance
     *
     * @throws Exception
     */
    Object createBeanClassInstance() throws Exception;

    /**
     * Initializes the instance context.
     *
     * <p>This method is called before createEntity, and should
     * reset the value of all cmpFields to 0 or null.
     *
     * @param ctx
     *
     * @throws RemoteException
     */
    void initEntity(EntityEnterpriseContext ctx);

    /**
     * This method is called whenever an entity is to be created. The
     * persistence manager is responsible for handling the results
     * properly wrt the persistent store.
     *
     * @param m the create method in the home interface that was
     * called
     * @param args any create parameters
     * @param instance the instance being used for this create call
     * @return The primary key computed by CMP PM or null for BMP
     *
     * @throws Exception
     */
    Object createEntity(Method m,
                        Object[] args,
                        EntityEnterpriseContext instance)
        throws Exception;

    /**
     * This method is called when single entities are to be found. The
     * persistence manager must find out whether the wanted instance
     * is available in the persistence store, if so it returns the
     * primary key of the object.
     *
     * @param finderMethod the find method in the home interface that was

```



```

* called
* @param args any finder parameters
* @param instance the instance to use for the finder call
* @return a primary key representing the found entity
*
* @throws RemoteException thrown if some system exception occurs
* @throws FinderException thrown if some heuristic problem occurs
*/
Object findEntity(Method finderMethod,
                  Object[] args,
                  EntityEnterpriseContext instance)
    throws Exception;

/**
 * This method is called when collections of entities are to be
 * found. The persistence manager must find out whether the wanted
 * instances are available in the persistence store, and if so it
 * must return a collection of primaryKeys.
 *
 * @param finderMethod the find method in the home interface that was
 * called
 * @param args any finder parameters
 * @param instance the instance to use for the finder call
 * @return an primary key collection representing the found
 * entities
 *
 * @throws RemoteException thrown if some system exception occurs
 * @throws FinderException thrown if some heuristic problem occurs
 */
Collection findEntities(Method finderMethod,
                       Object[] args,
                       EntityEnterpriseContext instance)
    throws Exception;

/**
 * This method is called when an entity shall be activated.
 *
 * <p>With the PersistenceManager factorization most EJB
 * calls should not exists However this calls permits us to
 * introduce optimizations in the persistence store. Particularly
 * the context has a "PersistenceContext" that a PersistenceStore
 * can use (JAWS does for smart updates) and this is as good a
 * callback as any other to set it up.
 * @param instance the instance to use for the activation
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void activateEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is called whenever an entity shall be load from the
 * underlying storage. The persistence manager must load the state
 * from the underlying storage and then call ejbLoad on the
 * supplied instance.
 *
 * @param instance the instance to synchronize
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void loadEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is used to determine if an entity should be stored.
 *
 * @param instance the instance to check
 * @return true, if the entity has been modified
 * @throws Exception thrown if some system exception occurs
 */
boolean isModified(EntityEnterpriseContext instance) throws Exception;

```

```

/**
 * This method is called whenever an entity shall be stored to the
 * underlying storage. The persistence manager must call ejbStore
 * on the supplied instance and then store the state to the
 * underlying storage.
 *
 * @param instance the instance to synchronize
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void storeEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is called when an entity shall be passivate. The
 * persistence manager must call the ejbPassivate method on the
 * instance.
 *
 * <p>See the activate discussion for the reason for
 * exposing EJB callback * calls to the store.
 *
 * @param instance the instance to passivate
 *
 * @throws RemoteException thrown if some system exception occurs
 */
void passivateEntity(EntityEnterpriseContext instance)
    throws RemoteException;

/**
 * This method is called when an entity shall be removed from the
 * underlying storage. The persistence manager must call ejbRemove
 * on the instance and then remove its state from the underlying
 * storage.
 *
 * @param instance the instance to remove
 *
 * @throws RemoteException thrown if some system exception occurs
 * @throws RemoveException thrown if the instance could not be removed
 */
void removeEntity(EntityEnterpriseContext instance)
    throws RemoteException, RemoveException;
}

```

The default BMP implementation of the `EntityPersistenceManager` interface is `org.jboss.ejb.plugins.BMPPersistenceManager`. The BMP persistence manager is fairly simple since all persistence logic is in the entity bean itself. The only duty of the persistence manager is to perform container callbacks.

5.3.2.6. `org.jboss.ejb.StatefulSessionPersistenceManager`

The `StatefulSessionPersistenceManager` is responsible for the persistence of stateful `SessionBeans`. This includes the following:

- Creating stateful sessions in a storage
- Activating stateful sessions from a storage
- Passivating stateful sessions to a storage
- Removing stateful sessions from a storage

The `StatefulSessionPersistenceManager` interface is shown below.

Example 5.16. The `org.jboss.ejb.StatefulSessionPersistenceManager` interface

```
public interface StatefulSessionPersistenceManager
    extends ContainerPlugin
{
    public void createSession(Method m, Object[] args,
                             StatefulSessionEnterpriseContext ctx)
        throws Exception;

    public void activateSession(StatefulSessionEnterpriseContext ctx)
        throws RemoteException;

    public void passivateSession(StatefulSessionEnterpriseContext ctx)
        throws RemoteException;

    public void removeSession(StatefulSessionEnterpriseContext ctx)
        throws RemoteException, RemoveException;

    public void removePassivated(Object key);
}
```

The default implementation of the `StatefulSessionPersistenceManager` interface is `org.jboss.ejb.plugins.StatefulSessionFilePersistenceManager`. As its name implies, `StatefulSessionFilePersistenceManager` utilizes the file system to persist stateful session beans. More specifically, the persistence manager serializes beans in a flat file whose name is composed of the bean name and session id with a `.ser` extension. The persistence manager restores a bean's state during activation and respectively stores its state during passivation from the bean's `.ser` file.

5.4. Entity Bean Locking and Deadlock Detection

This section provides information on what entity bean locking is and how entity beans are accessed and locked within JBoss. It also describes the problems you may encounter as you use entity beans within your system and how to combat these issues. Deadlocking is formally defined and examined. And, finally, we walk you through how to fine tune your system in terms of entity bean locking.

5.4.1. Why JBoss Needs Locking

Locking is about protecting the integrity of your data. Sometimes you need to be sure that only one user can update critical data at one time. Sometimes, access to sensitive objects in your system need to be serialized so that data is not corrupted by concurrent reads and writes. Databases traditionally provide this sort of functionality with transactional scopes and table and row locking facilities.

Entity beans are a great way to provide an object-oriented interface to relational data. Beyond that, they can improve performance by taking the load off of the database through caching and delaying updates until absolutely needed so that the database efficiency can be maximized. But, with caching, data integrity is a problem, so some form of application server level locking is needed for entity beans to provide the transaction isolation properties that you are used to with traditional databases.

5.4.2. Entity Bean Lifecycle

With the default configuration of JBoss there is only one active instance of a given entity bean in memory at one time. This applies for every cache configuration and every type of `commit-option`. The lifecycle for this instance is different for every `commit-option` though.

- For commit option `A`, this instance is cached and used between transactions.

- For commit option *B*, this instance is cached and used between transactions, but is marked as dirty at the end of a transaction. This means that at the start of a new transaction `ejbLoad` must be called.
- For commit option *C*, this instance is marked as dirty, released from the cache, and marked for passivation at the end of a transaction.
- For commit option *D*, a background refresh thread periodically calls `ejbLoad` on stale beans within the cache. Otherwise, this option works in the same way as *A*.

When a bean is marked for passivation, the bean is placed in a passivation queue. Each entity bean container has a passivation thread that periodically passivates beans that have been placed in the passivation queue. A bean is pulled out of the passivation queue and reused if the application requests access to a bean of the same primary key.

On an exception or transaction rollback, the entity bean instance is thrown out of cache entirely. It is not put into the passivation queue and is not reused by an instance pool. Except for the passivation queue, there is no entity bean instance pooling.

5.4.3. Default Locking Behavior

Entity bean locking is totally decoupled from the entity bean instance. The logic for locking is totally isolated and managed in a separate lock object. Because there is only one allowed instance of a given entity bean active at one time, JBoss employs two types of locks to ensure data integrity and to conform to the EJB spec.

- **Method Lock:** The method lock ensures that only one thread of execution at a time can invoke on a given Entity Bean. This is required by the EJB spec.
- **Transaction Lock:** A transaction lock ensures that only one transaction at a time has access to a give Entity Bean. This ensures the ACID properties of transactions at the application server level. Since, by default, there is only one active instance of any given Entity Bean at one time, JBoss must protect this instance from dirty reads and dirty writes. So, the default entity bean locking behavior will lock an entity bean within a transaction until it completes. This means that if any method at all is invoked on an entity bean within a transaction, no other transaction can have access to this bean until the holding transaction commits or is rolled back.

5.4.4. Pluggable Interceptors and Locking Policy

We saw that the basic entity bean lifecycle and behavior is defined by the container configuration defined in `standardjboss.xml` descriptor. The following container configuration shows the `container-interceptors` definition for the *Standard CMP 2.x EntityBean* configuration.

Example 5.17. The Standard CMP 2.x EntityBean interceptor definition

```
<container-configuration>
  <container-name>Standard CMP 2.x EntityBean</container-name>
  <!-- ... -->
  <container-interceptors>
    <interceptor>org.jboss.ejb.plugins.LogInterceptor</interceptor>
    <interceptor>org.jboss.ejb.plugins.SecurityInterceptor</interceptor>
    <interceptor>org.jboss.ejb.plugins.TxInterceptorCMT</interceptor>
    <interceptor>org.jboss.ejb.plugins.MetricsInterceptor</interceptor>
    <interceptor>org.jboss.ejb.plugins.EntityCreationInterceptor</interceptor>
    <interceptor>org.jboss.ejb.plugins.EntityLockInterceptor</interceptor>
    <interceptor>org.jboss.ejb.plugins.EntityInstanceInterceptor</interceptor>
  </container-interceptors>
</container-configuration>
```

```

<interceptor>
  org.jboss.resource.connectionmanager.CachedConnectionInterceptor
</interceptor>
<interceptor>org.jboss.ejb.plugins.EntitySynchronizationInterceptor</interceptor>
<interceptor>org.jboss.ejb.plugins.cmp.jdbc.JDBCRelationInterceptor</interceptor>
</container-interceptors>
</container-configuration>

```

The interceptors shown above define most of the behavior of the entity bean. Below is an explanation of the interceptors that are relevant to this section.

- **EntityLockInterceptor:** This interceptor's role is to schedule any locks that must be acquired before the invocation is allowed to proceed. This interceptor is very lightweight and delegates all locking behavior to a pluggable locking policy.
- **EntityInstanceInterceptor:** The job of this interceptor is to find the entity bean within the cache or create a new one. This interceptor also ensures that there is only one active instance of a bean in memory at one time.
- **EntitySynchronizationInterceptor:** The role of this interceptor is to synchronize the state of the cache with the underlying storage. It does this with the `ejbLoad` and `ejbStore` semantics of the EJB specification. In the presence of a transaction this is triggered by transaction demarcation. It registers a callback with the underlying transaction monitor through the JTA interfaces. If there is no transaction the policy is to store state upon returning from invocation. The synchronization policies *A*, *B* and *C* of the specification are taken care of here as well as the JBoss specific commit-option *D*.

5.4.5. Deadlock

Finding deadlock problems and resolving them is the topic of this section. We will describe what deadlocking MBeans, how you can detect it within your application, and how you can resolve deadlocks. Deadlock can occur when two or more threads have locks on shared resources. Figure 5.11 illustrates a simple deadlock scenario. Here, Thread 1 has the lock for Bean A, and Thread 2 has the lock for Bean B. At a later time, Thread 1 tries to lock Bean B and blocks because Thread 2 has it. Likewise, as Thread 2 tries to lock A it also blocks because Thread 1 has the lock. At this point both threads are deadlocked waiting for access to the resource already locked by the other thread.




Figure 5.11. Deadlock definition example

The default locking policy of JBoss is to lock an Entity bean when an invocation occurs in the context of a transaction until the transaction completes. Because of this, it is very easy to encounter deadlock if you have long running transactions that access many entity beans, or if you are not careful about ordering the access to them. Various techniques and advanced configurations can be used to avoid deadlocking problems. They are discussed later in this section.

5.4.5.1. Dedlock Detection

Fortunately, JBoss is able to perform deadlock detection. JBoss holds a global internal graph of waiting transactions and what transactions they are blocking on. Whenever a thread determines that it cannot acquire an entity bean lock, it figures out what transaction currently holds the lock on the bean and add itself to the blocked transaction graph. An example of what the graph may look like is given in Table 5.1.

Table 5.1. An example blocked transaction table

Blocking TX	Tx that holds needed lock
Tx1	Tx2
Tx3	Tx4
Tx4	Tx1

Before the thread actually blocks it tries to detect whether there is deadlock problem. It does this by traversing the block transaction graph. As it traverses the graph, it keeps track of what transactions are blocked. If it sees a blocked node more than once in the graph, then it knows there is deadlock and will throw an `ApplicationDeadlockException`. This exception will cause a transaction rollback which will cause all locks that transaction holds to be released. The algorithm for the deadlock dection is found in the `BeanLockSupport` `deadlockDetection` method. The following code shows this method.

Example 5.18. The `org.jboss.ejb.plugins.lock.BeanLockSupport` `deadlockDetection` method

```
// This following is for deadlock detection
protected static HashMap waiting = new HashMap();

public void deadlockDetection(Transaction miTx) throws Exception
{
    HashSet set = new HashSet();
    set.add(miTx);

    Object checkTx = this.tx;
    synchronized(waiting) {
        while (checkTx != null) {
            Object waitingFor = waiting.get(checkTx);
            if (waitingFor != null) {
                if (set.contains(waitingFor)) {
                    log.error("Application deadlock detected: " + miTx +
                        " has deadlock conditions");
                    throw new ApplicationDeadlockException("application deadlock detected");
                }
                set.add(waitingFor);
            }
            checkTx = waitingFor;
        }
    }
}
```

5.4.5.2. Catching ApplicationDeadlockException

Since JBoss can detect application deadlock, you should write your application so that it can retry a transaction if the invocation fails because of the `ApplicationDeadlockException`. Unfortunately, this exception can be deeply embedded within a `RemoteException`, so you have to search for it in your catch block. For example:

```
try {
    // ...
} catch (RemoteException ex) {
    Throwable cause = null;
    RemoteException rex = ex;
    while (rex.detail != null) {
        cause = rex.detail;
        if (cause instanceof ApplicationDeadlockException) {
            // ... We have deadlock, force a retry of the transaction.
            break;
        }
        if (cause instanceof RemoteException) {
            rex = (RemoteException)cause;
        }
    }
}
```

5.4.5.3. Viewing Lock Information

The `EntityLockMonitor` MBean service allows one to view basic locking statistics as well as printing out the state of the transaction locking table. To enable this monitor uncomment its configuration in the `conf/jboss-service.xml`:

```
<mbean code="org.jboss.monitor.EntityLockMonitor"
       name="jboss.monitor:name=EntityLockMonitor"/>
```

The `EntityLockMonitor` has no configurable attributes. It does have the following read-only attributes:

- **MedianWaitTime:** The median value of all times threads had to wait to acquire a lock.
- **AverageContenters:** The ratio of the total number of contentions to the sum of all threads that had to wait for a lock.
- **TotalContentions:** The total number of threads that had to wait to acquire the transaction lock. This happens when a thread attempts to acquire a lock that is associated with another transaction
- **MaxContenters:** The maximum number of threads that were waiting to acquire the transaction lock.

It also has the following operations:

- **clearMonitor:** This operation resets the lock monitor state by zeroing all counters.
- **printLockMonitor:** This operation prints out a table of all EJB locks that lists the `ejbName` of the bean, the total time spent waiting for the lock, the count of times the lock was waited on and the number of transactions that timed out waiting for the lock.

5.4.6. Advanced Configurations and Optimizations

The default locking behavior of entity beans can cause deadlock. Since access to an entity bean locks the bean into the transaction, this also can present a huge performance/throughput problem for your application. This

section walks through various techniques and configurations that you can use to optimize performance and reduce the possibility of deadlock.

5.4.6.1. Short-lived Transactions

Make your transactions as short-lived and fine-grained as possible. The shorter the transaction you have, the less likelihood you will have concurrent access collisions and your application throughput will go up.

5.4.6.2. Ordered Access

Ordering the access to your entity beans can help lessen the likelihood of deadlock. This means making sure that the entity beans in your system are always accessed in the same exact order. In most cases, user applications are just too complicated to use this approach and more advanced configurations are needed.

5.4.6.3. Read-Only Beans

Entity beans can be marked as read-only. When a bean is marked as read-only, it never takes part in a transaction. This means that it is never transactionally locked. Using commit-option *D* with this option is sometimes very useful when your read-only bean's data is sometimes updated by an external source.

To mark a bean as read-only, use the `read-only` flag in the `jboss.xml` deployment descriptor.

Example 5.19. Marking an entity bean read-only using `jboss.xml`

```
<jboss>
  <enterprise-beans>
    <entity>
      <ejb-name>MyEntityBean</ejb-name>
      <jndi-name>MyEntityHomeRemote</jndi-name>
      <read-only>True</read-only>
    </entity>
  </enterprise-beans>
</jboss>
```

5.4.6.4. Explicitly Defining Read-Only Methods

After reading and understanding the default locking behavior of entity beans, you're probably wondering, "Why lock the bean if its not modifying the data?" JBoss allows you to define what methods on your entity bean are read only so that it will not lock the bean within the transaction if only these types of methods are called. You can define these read only methods within a `jboss.xml` deployment descriptor. Wildcards are allowed for method names. The following is an example of declaring all getter methods and the `anotherReadOnlyMethod` as read-only.

Example 5.20. Defining entity bean methods as read only

```
<jboss>
  <enterprise-beans>
    <entity>
      <ejb-name>nextgen.EnterpriseEntity</ejb-name>
      <jndi-name>nextgen.EnterpriseEntity</jndi-name>
      <method-attributes>
        <method>
          <method-name>get*</method-name>
          <read-only>true</read-only>
        </method>
      </method-attributes>
    </entity>
  </enterprise-beans>
</jboss>
```



```

        </method>
        <method>
            <method-name>anotherReadOnlyMethod</method-name>
            <read-only>true</read-only>
        </method>
    </method-attributes>
</entity>
</enterprise-beans>
</jboss>

```

5.4.6.5. Instance Per Transaction Policy

The Instance Per Transaction policy is an advanced configuration that can totally wipe away deadlock and throughput problems caused by JBoss's default locking policy. The default Entity Bean locking policy is to only allow one active instance of a bean. The Instance Per Transaction policy breaks this requirement by allocating a new instance of a bean per transaction and dropping this instance at the end of the transaction. Because each transaction has its own copy of the bean, there is no need for transaction based locking.

This option does sound great but does have some drawbacks right now. First, the transactional isolation behavior of this option is equivalent to `READ_COMMITTED`. This can create repeatable reads when they are not desired. In other words, a transaction could have a copy of a stale bean. Second, this configuration option currently requires commit-option *B* or *C* which can be a performance drain since an `ejbLoad` must happen at the beginning of the transaction. But, if your application currently requires commit-option *B* or *C* anyways, then this is the way to go. The JBoss developers are currently exploring ways to allow commit-option *A* as well (which would allow the use of caching for this option).

JBoss has container configurations named `Instance Per Transaction CMP 2.x EntityBean` and `Instance Per Transaction BMP EntityBean` defined in the `standardjboss.xml` that implement this locking policy. To use this configuration, you just have to reference the name of the container configuration to use with your bean in the `jboss.xml` deployment descriptor as show below.

Example 5.21. An example of using the Instance Per Transaction policy.

```

<jboss>
  <enterprise-beans>
    <entity>
      <ejb-name>MyCMP2Bean</ejb-name>
      <jndi-name>MyCMP2</jndi-name>
      <configuration-name>
        Instance Per Transaction CMP 2.x EntityBean
      </configuration-name>
    </entity>
    <entity>
      <ejb-name>MyBMPBean</ejb-name>
      <jndi-name>MyBMP</jndi-name>
      <configuration-name>
        Instance Per Transaction BMP EntityBean
      </configuration-name>
    </entity>
  </enterprise-beans>
</jboss>

```

5.4.7. Running Within a Cluster

Currently there is no distributed locking capability for entity beans within the cluster. This functionality has been delegated to the database and must be supported by the application developer. For clustered entity beans,

it is suggested to use commit-option *B* or *C* in combination with a row locking mechanism. For CMP, there is a row-locking configuration option. This option will use use a SQL `select for update` when the bean is loaded from the database. With commit-option *B* or *C*, this implements a transactional lock that can be used across the cluster. For BMP, you must explicitly implement the select for update invocation within the BMP's `ejbLoad` method.

5.4.8. Troubleshooting

This section will describe some common locking problems and their solution.

5.4.8.1. Locking Behavior Not Working

There are many emails on the the JBoss User email list which sometimes state that the locking is not working and they are having concurrent access to their beans, and thus dirty reads. Here are some common reasons for this:

- If you have custom container-configurations, make sure you have updated these configurations.
- Make absolutely sure that you have implemented `equals` and `hashCode` correctly from custom/complex primary key classes.
- Make absolutely sure that your custom/complex primary key classes serialize correctly. One common mistake is assuming that member variable initializations will be executed when a primary key is unmarshalled.

5.4.8.2. `IllegalStateException`

An `IllegalStateException` with the message "removing bean lock and it has tx set!" usually means that you have not implemented `equals` and/or `hashCode` correctly for your custom/complex primary key class, or that your primary key class is not implemented correctly for serialization.

5.4.8.3. Hangs and Transaction Timeouts

One long outstanding bug of JBoss is that on a transaction timeout, that transaction is only marked for a roll-back and not actually rolled back. This responsibility is delegated to the invocation thread. This can cause major problems if the invocation thread hangs indefinitely since things like entity bean locks will never be released. The solution to this problem is not a good one. You really just need to avoid doing stuff within a transaction that could hang indefinitely. One common mistake is making connections across the internet or running a web-crawler within a transaction.

Messaging on JBoss

JMS Configuration and Architecture

The JMS API stands for Java Message Service Application Programming Interface, and it is used by applications to send asynchronous *business-quality* messages to other applications. In the JMS world, messages are not sent directly to other applications. Instead, messages are sent to destinations, also known as queues or topics. Applications sending messages do not need to worry if the receiving applications are up and running, and conversely, receiving applications do not need to worry about the sending application's status. Both senders, and receivers only interact with the destinations.

The JMS API is the standardized interface to a JMS provider, sometimes called a Message Oriented Middleware (MOM) system. JBoss comes with a JMS 1.0.2b compliant JMS provider called JBoss Messaging or JBossMQ. When you use the JMS API with JBoss, you are using the JBoss Messaging engine transparently. JBoss Messaging fully implements the JMS specification; therefore, the best JBoss Messaging user guide is the JMS specification. For more information about the JMS API please visit the JMS Tutorial or JMS Downloads & Specifications.

This chapter focuses on the JBoss specific aspects of using JMS and message driven beans as well as the JBoss Messaging configuration and MBeans.

6.1. JMS Examples

In this section we discuss the basics needed to use the JBoss JMS implementation. JMS leaves the details of accessing JMS connection factories and destinations as provider specific details. What you need to know to use the JBoss Messaging layer is:

- The location of the `javax.jms.QueueConnectionFactory` and `javax.jms.TopicConnectionFactory`. In JBoss both connection factory implementations are located under the JNDI name `ConnectionFactory`.
- How to lookup JMS destinations (`javax.jms.Queue` and `javax.jms.Topic`). Destinations are configured via MBeans as we will see when we discuss the messaging MBeans. JBoss comes with a few queues and topics preconfigured. You can find them under the `jboss.mq.destination` domain in the JMX Console..
- The JBoss Messaging JARs. These include `concurrent.jar`, `jbossmq-client.jar`, `jboss-common-client.jar`, `jboss-system-client.jar`, `jnp-client.jar`, `log4j.jar` and `jnet.jar` (`jnet.jar` is only needed for JDK 1.3)

In the following subsections we will look at examples of the various JMS messaging models and message driven beans. The chapter example source is located under the `src/main/org/jboss/chap6` directory of the book examples.

6.1.1. A Point-To-Point Example

Let's start out with a point-to-point (P2P) example. In the P2P model, a sender delivers messages to a queue and a single receiver pulls the message off of the queue. The receiver does not need to be listening to the queue at the time the message is sent. Example 6.1 shows a complete P2P example that sends a `javax.jms.TextMessage` to a the queue `queue/testQueue` and asynchronously receives the message from the same queue.

Example 6.1. A P2P JMS client example

```
package org.jboss.chap6.ex1;

import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageListener;
import javax.jms.Queue;
import javax.jms.QueueConnection;
import javax.jms.QueueConnectionFactory;
import javax.jms.QueueReceiver;
import javax.jms.QueueSender;
import javax.jms.QueueSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

import edu.oswego.cs.dl.util.concurrent.CountDown;

/**
 * A complete JMS client example program that sends a TextMessage to
 * a Queue and asynchronously receives the message from the same
 * Queue.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.10 $
 */
public class SendRecvClient
{
    static CountDown done = new CountDown(1);
    QueueConnection conn;
    QueueSession session;
    Queue que;

    public static class ExListener
        implements MessageListener
    {
        public void onMessage(Message msg)
        {
            {
                done.release();
                TextMessage tm = (TextMessage) msg;
                try {
                    System.out.println("onMessage, recv text=" + tm.getText());
                } catch(Throwable t) {
                    t.printStackTrace();
                }
            }
        }
    }

    public void setupPTP()
        throws JMSEException,
              NamingException
    {
        InitialContext iniCtx = new InitialContext();
        Object tmp = iniCtx.lookup("ConnectionFactory");
        QueueConnectionFactory qcf = (QueueConnectionFactory) tmp;
        conn = qcf.createQueueConnection();
        que = (Queue) iniCtx.lookup("queue/testQueue");
        session = conn.createQueueSession(false,
                                           QueueSession.AUTO_ACKNOWLEDGE);

        conn.start();
    }
}
```

```

public void sendRecvAsync(String text)
    throws JMSEException,
           NamingException
{
    System.out.println("Begin sendRecvAsync");
    // Setup the PTP connection, session
    setupPTP();

    // Set the async listener
    QueueReceiver recv = session.createReceiver(que);
    recv.setMessageListener(new ExListener());

    // Send a text msg
    QueueSender send = session.createSender(que);
    TextMessage tm = session.createTextMessage(text);
    send.send(tm);
    System.out.println("sendRecvAsync, sent text=" + tm.getText());
    send.close();
    System.out.println("End sendRecvAsync");
}

public void stop()
    throws JMSEException
{
    conn.stop();
    session.close();
    conn.close();
}

public static void main(String args[])
    throws Exception
{
    SendRecvClient client = new SendRecvClient();
    client.sendRecvAsync("A text msg");
    client.done.acquire();
    client.stop();
    System.exit(0);
}
}

```

The client may be run using the following command line:

```

[nr@toki examples]$ ant -Dchap=chap6 -Dex=lp2p run-example
...
run-examplelp2p:
[java] [INFO,SendRecvClient] Begin SendRecvClient, now=1098416473521
[java] [INFO,SendRecvClient] Begin sendRecvAsync
[java] [INFO,SendRecvClient] onMessage, recv text=A text msg
[java] [INFO,SendRecvClient] sendRecvAsync, sent text=A text msg
[java] [INFO,SendRecvClient] End sendRecvAsync
[java] [INFO,SendRecvClient] End SendRecvClient

```

6.1.2. A Pub-Sub Example

The JMS publish/subscribe (Pub-Sub) message model is a one-to-many model. A publisher sends a message to a topic and all active subscribers of the topic receive the message. Subscribers that are not actively listening to the topic will miss the published message. shows a complete JMS client that sends a `javax.jms.TextMessage` to a topic and asynchronously receives the message from the same topic.

Example 6.2. A Pub-Sub JMS client example

```
package org.jboss.chap6.ex1;
```

```

import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageListener;
import javax.jms.Topic;
import javax.jms.TopicConnection;
import javax.jms.TopicConnectionFactory;
import javax.jms.TopicPublisher;
import javax.jms.TopicSubscriber;
import javax.jms.TopicSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

import EDU.oswego.cs.dl.util.concurrent.CountDown;

/**
 * A complete JMS client example program that sends a TextMessage to
 * a Topic and asynchronously receives the message from the same
 * Topic.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.10 $
 */
public class TopicSendRecvClient
{
    static CountDown done = new CountDown(1);
    TopicConnection conn = null;
    TopicSession session = null;
    Topic topic = null;

    public static class ExListener implements MessageListener
    {
        public void onMessage(Message msg)
        {
            done.release();
            TextMessage tm = (TextMessage) msg;
            try {
                System.out.println("onMessage, recv text=" + tm.getText());
            } catch (Throwable t) {
                t.printStackTrace();
            }
        }
    }

    public void setupPubSub()
        throws JMSEException, NamingException
    {
        InitialContext iniCtx = new InitialContext();
        Object tmp = iniCtx.lookup("ConnectionFactory");
        TopicConnectionFactory tcf = (TopicConnectionFactory) tmp;
        conn = tcf.createTopicConnection();
        topic = (Topic) iniCtx.lookup("topic/testTopic");
        session = conn.createTopicSession(false,
                                           TopicSession.AUTO_ACKNOWLEDGE);

        conn.start();
    }

    public void sendRecvAsync(String text)
        throws JMSEException, NamingException
    {
        System.out.println("Begin sendRecvAsync");
        // Setup the PubSub connection, session
        setupPubSub();
        // Set the async listener

        TopicSubscriber recv = session.createSubscriber(topic);
        recv.setMessageListener(new ExListener());
        // Send a text msg
        TopicPublisher send = session.createPublisher(topic);

```

```

        TextMessage tm = session.createTextMessage(text);
        send.publish(tm);
        System.out.println("sendRecvAsync, sent text=" + tm.getText());
        send.close();
        System.out.println("End sendRecvAsync");
    }

    public void stop() throws JMSEException
    {
        conn.stop();
        session.close();
        conn.close();
    }

    public static void main(String args[]) throws Exception
    {
        System.out.println("Begin TopicSendRecvClient, now=" +
            System.currentTimeMillis());
        TopicSendRecvClient client = new TopicSendRecvClient();
        client.sendRecvAsync("A text msg, now="+System.currentTimeMillis());
        client.done.acquire();
        client.stop();
        System.out.println("End TopicSendRecvClient");
        System.exit(0);
    }
}

```

The client may be run using the following command line:

```

[nr@toki examples]$ ant -Dchap=chap6 -Dex=1ps run-example
...
run-example1ps:
    [java] Begin TopicSendRecvClient, now=1098416563162
    [java] Begin sendRecvAsync
    [java] onMessage, recv text=A text msg, now=1098416563171
    [java] sendRecvAsync, sent text=A text msg, now=1098416563171
    [java] End sendRecvAsync
    [java] End TopicSendRecvClient

```

Now let's break the publisher and subscribers into separate programs to demonstrate that subscribers only receive messages while they are listening to a topic. Example 6.3 shows a variation of the previous pub-sub client that only publishes messages to the `topic/testTopic` topic. The subscriber only client is shown in Example 6.3.

Example 6.3. A JMS publisher client

```

package org.jboss.chap6.ex1;

import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageListener;
import javax.jms.Topic;
import javax.jms.TopicConnection;
import javax.jms.TopicConnectionFactory;
import javax.jms.TopicPublisher;
import javax.jms.TopicSubscriber;
import javax.jms.TopicSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

/**
 * A JMS client example program that sends a TextMessage to a Topic

```

```

*
*  @author Scott.Stark@jboss.org
*  @version $Revision: 1.10 $
*/
public class TopicSendClient
{
    TopicConnection conn = null;
    TopicSession session = null;
    Topic topic = null;

    public void setupPubSub()
        throws JMSEException, NamingException
    {
        InitialContext iniCtx = new InitialContext();
        Object tmp = iniCtx.lookup("ConnectionFactory");
        TopicConnectionFactory tcf = (TopicConnectionFactory) tmp;
        conn = tcf.createTopicConnection();
        topic = (Topic) iniCtx.lookup("topic/testTopic");
        session = conn.createTopicSession(false,
                                           TopicSession.AUTO_ACKNOWLEDGE);

        conn.start();
    }

    public void sendAsync(String text)
        throws JMSEException, NamingException
    {
        System.out.println("Begin sendAsync");
        // Setup the pub/sub connection, session
        setupPubSub();
        // Send a text msg
        TopicPublisher send = session.createPublisher(topic);
        TextMessage tm = session.createTextMessage(text);
        send.publish(tm);
        System.out.println("sendAsync, sent text=" + tm.getText());
        send.close();
        System.out.println("End sendAsync");
    }

    public void stop()
        throws JMSEException
    {
        conn.stop();
        session.close();
        conn.close();
    }

    public static void main(String args[])
        throws Exception
    {
        System.out.println("Begin TopicSendClient, now=" +
                           System.currentTimeMillis());
        TopicSendClient client = new TopicSendClient();
        client.sendAsync("A text msg, now="+System.currentTimeMillis());
        client.stop();
        System.out.println("End TopicSendClient");
        System.exit(0);
    }
}

```

Example 6.4. A JMS subscriber client

```

package org.jboss.chap6.ex1;

import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageListener;

```



```

import javax.jms.Topic;
import javax.jms.TopicConnection;
import javax.jms.TopicConnectionFactory;
import javax.jms.TopicPublisher;
import javax.jms.TopicSubscriber;
import javax.jms.TopicSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

/**
 * A JMS client example program that synchronously receives a message a Topic
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.10 $
 */
public class TopicRecvClient
{
    TopicConnection conn = null;
    TopicSession session = null;
    Topic topic = null;

    public void setupPubSub()
        throws JMSEException, NamingException
    {
        InitialContext iniCtx = new InitialContext();
        Object tmp = iniCtx.lookup("ConnectionFactory");
        TopicConnectionFactory tcf = (TopicConnectionFactory) tmp;
        conn = tcf.createTopicConnection();
        topic = (Topic) iniCtx.lookup("topic/testTopic");
        session = conn.createTopicSession(false,
                                           TopicSession.AUTO_ACKNOWLEDGE);

        conn.start();
    }

    public void recvSync()
        throws JMSEException, NamingException
    {
        System.out.println("Begin recvSync");
        // Setup the pub/sub connection, session
        setupPubSub();

        // Wait upto 5 seconds for the message
        TopicSubscriber recv = session.createSubscriber(topic);
        Message msg = recv.receive(5000);
        if (msg == null) {
            System.out.println("Timed out waiting for msg");
        } else {
            System.out.println("TopicSubscriber.recv, msgt="+msg);
        }
    }

    public void stop()
        throws JMSEException
    {
        conn.stop();
        session.close();
        conn.close();
    }

    public static void main(String args[])
        throws Exception
    {
        System.out.println("Begin TopicRecvClient, now=" +
                           System.currentTimeMillis());
        TopicRecvClient client = new TopicRecvClient();
        client.recvSync();
        client.stop();
        System.out.println("End TopicRecvClient");
        System.exit(0);
    }
}

```

```
}
```

Run the `TopicSendClient` followed by the `TopicRecvClient` as follows:

```
[nr@toki examples]$ ant -Dchap=chap6 -Dex=lps2 run-example
...
run-examplelps2:
    [java] Begin TopicSendClient, now=1098416676618
    [java] Begin sendAsync
    [java] sendAsync, sent text=A text msg, now=1098416676621
    [java] End sendAsync
    [java] End TopicSendClient
    [java] Begin TopicRecvClient, now=1098416683857
    [java] Begin recvSync
    [java] Timed out waiting for msg
    [java] End TopicRecvClient
```

The output shows that the topic subscriber client (`TopicRecvClient`) fails to receive the message sent by the publisher due to a timeout.

6.1.3. A Pub-Sub With Durable Topic Example

JMS supports a messaging model that is a cross between the P2P and pub-sub models. When a pub-sub client wants to receive all messages posted to the topic it subscribes to even when it is not actively listening to the topic, the client may achieve this behavior using a durable topic. Let's look at a variation of the preceding subscriber client that uses a durable topic to ensure that it receives all messages, include those published when the client is not listening to the topic. Example 6.5 shows the durable topic client with the key differences between the Example 6.4 client highlighted in bold.

Example 6.5. A durable topic JMS client example

```
package org.jboss.chap6.ex1;

import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageListener;
import javax.jms.Topic;
import javax.jms.TopicConnection;
import javax.jms.TopicConnectionFactory;
import javax.jms.TopicPublisher;
import javax.jms.TopicSubscriber;
import javax.jms.TopicSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

/**
 * A JMS client example program that synchronously receives a message a Topic
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.10 $
 */
public class DurableTopicRecvClient
{
    TopicConnection conn = null;
    TopicSession session = null;
    Topic topic = null;

    public void setupPubSub()
        throws JMSEException, NamingException
```

```

{
    InitialContext iniCtx = new InitialContext();
    Object tmp = iniCtx.lookup("ConnectionFactory");

    TopicConnectionFactory tcf = (TopicConnectionFactory) tmp;
    conn = tcf.createTopicConnection("john", "needle");
    topic = (Topic) iniCtx.lookup("topic/testTopic");

    session = conn.createTopicSession(false,
                                      TopicSession.AUTO_ACKNOWLEDGE);
    conn.start();
}

public void recvSync()
    throws JMSEException, NamingException
{
    System.out.println("Begin recvSync");
    // Setup the pub/sub connection, session
    setupPubSub();
    // Wait upto 5 seconds for the message
    TopicSubscriber recv = session.createDurableSubscriber(topic, "chap6-ex1dtps");
    Message msg = recv.receive(5000);
    if (msg == null) {
        System.out.println("Timed out waiting for msg");
    } else {
        System.out.println("DurableTopicRecvClient.recv, msgt=" + msg);
    }
}

public void stop()
    throws JMSEException
{
    conn.stop();
    session.close();
    conn.close();
}

public static void main(String args[])
    throws Exception
{
    System.out.println("Begin DurableTopicRecvClient, now=" +
                      System.currentTimeMillis());
    DurableTopicRecvClient client = new DurableTopicRecvClient();
    client.recvSync();
    client.stop();
    System.out.println("End DurableTopicRecvClient");
    System.exit(0);
}
}

```

Now run the previous topic publisher with the durable topic subscriber as follows:

```

[nr@toki examples]$ ant -Dchap=chap6 -Dex=1psdt run-example
run-example1psdt:
[java] Begin DurableTopicSetup
[java] End DurableTopicSetup
[java] Begin TopicSendClient, now=1098420531772
[java] Begin sendAsync
[java] sendAsync, sent text=A text msg, now=1098420531775
[java] End sendAsync
[java] End TopicSendClient
[java] Begin DurableTopicRecvClient, now=1098420538269
[java] Begin recvSync
[java] DurableTopicRecvClient.recv, msgt=SpyTextMessage {
[java] Header {
[java]     jmsDestination : TOPIC.testTopic.DurableSubscription[clientId=DurableSubscriberExamp1
[java]     jmsDeliveryMode : 2
[java]     jmsExpiration   : 0

```

```
[java]      jmsPriority      : 4
[java]      jmsMessageID    : ID:29-10984205372121
[java]      jmsTimeStamp     : 1098420537212
[java]      jmsCorrelationID: null
[java]      jmsReplyTo       : null
[java]      jmsType          : null
[java]      jmsRedelivered   : false
[java]      jmsProperties    : {}
[java]      jmsPropReadWrite : false
[java]      msgReadOnly      : true
[java]      producerClientId: ID:29
[java] }
[java] Body {
[java]   text      :A text msg, now=1098420531775
[java] }
[java] }
[java] End DurableTopicRecvClient
```

Items of note for the durable topic example include:

- The `TopicConnectionFactory` creation in the durable topic client used a username and password, and the `TopicSubscriber` creation was done using the `createDurableSubscriber(Topic, String)` method. This is a requirement of durable topic subscribers. The messaging server needs to know what client is requesting the durable topic and what the name of the durable topic subscription is. We will discuss the details of durable topic setup in the configuration section.
- An `org.jboss.chap6.DurableTopicSetup` client was run prior to the `TopicSendClient`. The reason for this is a durable topic subscriber must have registered a subscription at some point in the past in order for the messaging server to save messages. JBoss supports dynamic durable topic subscribers and the `DurableTopicSetup` client simply creates a durable subscription receiver and then exits. This leaves an active durable topic subscriber on the `topic/testTopic` and the messaging server knows that any messages posted to this topic must be saved for latter delivery.
- The `TopicSendClient` does not change for the durable topic. The notion of a durable topic is a subscriber only notion.
- The `DurableTopicRecvClient` sees the message published to the `topic/testTopic` even though it was not listening to the topic at the time the message was published.

6.1.4. A Point-To-Point With MDB Example

The EJB 2.0 specification added the notion of message driven beans (MDB). A MDB is a business component that may be invoked asynchronously. As of the EJB 2.0 specification, JMS was the only mechanism by which MDBs could be accessed. Example 6.6 shows an MDB that transforms the `TextMessages` it receives and sends the transformed messages to the queue found in the incoming message `JMSReplyTo` header.

Example 6.6. A `TextMessage` processing MDB

```
package org.jboss.chap6.ex2;

import javax.ejb.MessageDrivenBean;
import javax.ejb.MessageDrivenContext;
import javax.ejb.EJBException;
import javax.jms.JMSException;
import javax.jms.Message;
import javax.jms.MessageListener;
import javax.jms.Queue;
import javax.jms.QueueConnection;
```

```

import javax.jms.QueueConnectionFactory;
import javax.jms.QueueSender;
import javax.jms.QueueSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

/**
 * An MDB that transforms the TextMessages it receives and send the
 * transformed messages to the Queue found in the incoming message
 * JMSReplyTo header.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.10 $
 */
public class TextMDB
    implements MessageDrivenBean, MessageListener
{
    private MessageDrivenContext ctx = null;
    private QueueConnection conn;
    private QueueSession session;

    public TextMDB()
    {
        System.out.println("TextMDB.ctor, this="+hashCode());
    }

    public void setMessageDrivenContext(MessageDrivenContext ctx)
    {
        this.ctx = ctx;
        System.out.println("TextMDB.setMessageDrivenContext, this=" +
            hashCode());
    }

    public void ejbCreate()
    {
        System.out.println("TextMDB.ejbCreate, this="+hashCode());
        try {
            setupPTP();
        } catch (Exception e) {
            throw new EJBException("Failed to init TextMDB", e);
        }
    }

    public void ejbRemove()
    {
        System.out.println("TextMDB.ejbRemove, this="+hashCode());
        ctx = null;
        try {
            if (session != null) {
                session.close();
            }
            if (conn != null) {
                conn.close();
            }
        } catch (JMSEException e) {
            e.printStackTrace();
        }
    }

    public void onMessage(Message msg)
    {
        System.out.println("TextMDB.onMessage, this="+hashCode());
        try {
            TextMessage tm = (TextMessage) msg;
            String text = tm.getText() + "processed by: " + hashCode();
            Queue dest = (Queue) msg.getJMSReplyTo();
            sendReply(text, dest);
        } catch (Throwable t) {
            t.printStackTrace();
        }
    }
}

```

```

    }

    private void setupPTP()
        throws JMSEException, NamingException
    {
        InitialContext iniCtx = new InitialContext();
        Object tmp = iniCtx.lookup("java:comp/env/jms/QCF");
        QueueConnectionFactory qcf = (QueueConnectionFactory) tmp;
        conn = qcf.createQueueConnection();
        session = conn.createQueueSession(false,
                                           QueueSession.AUTO_ACKNOWLEDGE);

        conn.start();
    }

    private void sendReply(String text, Queue dest)
        throws JMSEException
    {
        System.out.println("TextMDB.sendReply, this=" +
                           hashCode() + ", dest="+dest);
        QueueSender sender = session.createSender(dest);
        TextMessage tm = session.createTextMessage(text);
        sender.send(tm);
        sender.close();
    }
}

```

The MDB ejb-jar.xml and jboss.xml deployment descriptors are shown in Example 6.7.

Example 6.7. The MDB ejb-jar.xml descriptor

```

<?xml version="1.0"?>
<!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>
    <enterprise-beans>
        <message-driven>
            <ejb-name>TextMDB</ejb-name>
            <ejb-class>org.jboss.chap6.ex2.TextMDB</ejb-class>
            <transaction-type>Container</transaction-type>
            <acknowledge-mode>AUTO_ACKNOWLEDGE</acknowledge-mode>
            <message-driven-destination>
                <destination-type>javax.jms.Queue</destination-type>
            </message-driven-destination>
            <res-ref-name>jms/QCF</res-ref-name>
            <resource-ref>
                <res-type>javax.jms.QueueConnectionFactory</res-type>
                <res-auth>Container</res-auth>
            </resource-ref>
        </message-driven>
    </enterprise-beans>
</ejb-jar>

```

Example 6.8. The MDB jboss.xml descriptor

```

<?xml version="1.0"?>
<jboss>
    <enterprise-beans>
        <message-driven>
            <ejb-name>TextMDB</ejb-name>
            <destination-jndi-name>queue/B</destination-jndi-name>
            <resource-ref>
                <res-ref-name>jms/QCF</res-ref-name>
            </resource-ref>
        </message-driven>
    </enterprise-beans>

```

```

        <jndi-name>ConnectionFactory</jndi-name>
    </resource-ref>
</message-driven>
</enterprise-beans>
</jboss>

```

Example 6.9 shows a variation of the P2P client that sends several messages to the `queue/B` destination and asynchronously receives the messages as modified by `TextMDB` from queue A.

Example 6.9. A JMS client that interacts with the TextMDB

```

package org.jboss.chap6.ex2;

import javax.jms.JMSEException;
import javax.jms.Message;
import javax.jms.MessageListener;
import javax.jms.Queue;
import javax.jms.QueueConnection;
import javax.jms.QueueConnectionFactory;
import javax.jms.QueueReceiver;
import javax.jms.QueueSender;
import javax.jms.QueueSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.NamingException;

import edu.oswego.cs.dl.util.concurrent.CountDown;

/**
 * A complete JMS client example program that sends N TextMessages to
 * a Queue B and asynchronously receives the messages as modified by
 * TextMDB from Queue A.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.10 $
 */
public class SendRecvClient
{
    static final int N = 10;
    static CountDown done = new CountDown(N);

    QueueConnection conn;
    QueueSession session;
    Queue queA;
    Queue queB;

    public static class ExListener
        implements MessageListener
    {
        public void onMessage(Message msg)
        {
            {
                done.release();
                TextMessage tm = (TextMessage) msg;
                try {
                    System.out.println("onMessage, recv text="+tm.getText());
                } catch(Throwable t) {
                    t.printStackTrace();
                }
            }
        }
    }

    public void setupPTP()
        throws JMSEException, NamingException
    {
        InitialContext iniCtx = new InitialContext();
        Object tmp = iniCtx.lookup("ConnectionFactory");
    }
}

```

```

    QueueConnectionFactory qcf = (QueueConnectionFactory) tmp;
    conn = qcf.createQueueConnection();
    queA = (Queue) iniCtx.lookup("queue/A");
    queB = (Queue) iniCtx.lookup("queue/B");
    session = conn.createQueueSession(false,
                                      QueueSession.AUTO_ACKNOWLEDGE);
    conn.start();
}

public void sendRecvAsync(String textBase)
    throws JMSException, NamingException, InterruptedException
{
    System.out.println("Begin sendRecvAsync");

    // Setup the PTP connection, session
    setupPTP();

    // Set the async listener for queA
    QueueReceiver recv = session.createReceiver(queA);
    recv.setMessageListener(new ExListener());

    // Send a few text msgs to queB
    QueueSender send = session.createSender(queB);

    for(int m = 0; m < 10; m++) {
        TextMessage tm = session.createTextMessage(textBase+"#" +m);
        tm.setJMSReplyTo(queA);
        send.send(tm);
        System.out.println("sendRecvAsync, sent text="+tm.getText());
    }
    System.out.println("End sendRecvAsync");
}

public void stop()
    throws JMSException
{
    conn.stop();
    session.close();
    conn.close();
}

public static void main(String args[])
    throws Exception
{
    System.out.println("Begin SendRecvClient,now=" +
                      System.currentTimeMillis());
    SendRecvClient client = new SendRecvClient();
    client.sendRecvAsync("A text msg");
    client.done.acquire();
    client.stop();
    System.exit(0);
    System.out.println("End SendRecvClient");
}
}

```

Run the client as follows:

```

[nr@toki examples]$ ant -Dchap=chap6 -Dex=2 run-example
...
run-example2:
[copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
[echo] Waiting 5 seconds for deploy...
[java] Begin SendRecvClient, now=1098419197580
[java] Begin sendRecvAsync
[java] onMessage, recv text=A text msg#0processed by: 13929978
[java] sendRecvAsync, sent text=A text msg#0
[java] sendRecvAsync, sent text=A text msg#1
[java] onMessage, recv text=A text msg#2processed by: 5495387

```



```
[java] sendRecvAsync, sent text=A text msg#2
[java] sendRecvAsync, sent text=A text msg#3
[java] onMessage, recv text=A text msg#1processed by: 13929978
[java] sendRecvAsync, sent text=A text msg#4
[java] sendRecvAsync, sent text=A text msg#5
[java] onMessage, recv text=A text msg#5processed by: 5495387
[java] sendRecvAsync, sent text=A text msg#6
[java] sendRecvAsync, sent text=A text msg#7
[java] onMessage, recv text=A text msg#4processed by: 13929978
[java] sendRecvAsync, sent text=A text msg#8
[java] sendRecvAsync, sent text=A text msg#9
[java] End sendRecvAsync
[java] onMessage, recv text=A text msg#3processed by: 15690844
[java] onMessage, recv text=A text msg#8processed by: 15690844
[java] onMessage, recv text=A text msg#7processed by: 13929978
[java] onMessage, recv text=A text msg#6processed by: 5495387
[java] onMessage, recv text=A text msg#9processed by: 14089812
```

The corresponding JBoss server console output is:

```
23:26:36,720 INFO [EjbModule] Deploying TextMDB
23:26:37,073 INFO [EJBDeployer] Deployed: file:/private/tmp/jboss-3.2.6/server/default/deploy/chap6-ex2.jar
23:26:43,216 INFO [TextMDB] TextMDB.ctor, this=13929978
23:26:43,224 INFO [TextMDB] TextMDB.setMessageDrivenContext, this=13929978
23:26:43,299 INFO [TextMDB] TextMDB.ejbCreate, this=13929978
23:26:43,401 INFO [TextMDB] TextMDB.onMessage, this=13929978
23:26:43,408 INFO [TextMDB] TextMDB.sendReply, this=13929978, dest=QUEUE.A
23:26:43,494 INFO [TextMDB] TextMDB.onMessage, this=13929978
23:26:43,518 INFO [TextMDB] TextMDB.sendReply, this=13929978, dest=QUEUE.A
23:26:43,571 INFO [TextMDB] TextMDB.ctor, this=5495387
23:26:43,573 INFO [TextMDB] TextMDB.setMessageDrivenContext, this=5495387
23:26:43,574 INFO [TextMDB] TextMDB.ejbCreate, this=5495387
23:26:43,596 INFO [TextMDB] TextMDB.onMessage, this=5495387
23:26:43,597 INFO [TextMDB] TextMDB.sendReply, this=5495387, dest=QUEUE.A
23:26:43,802 INFO [TextMDB] TextMDB.onMessage, this=13929978
23:26:43,803 INFO [TextMDB] TextMDB.sendReply, this=13929978, dest=QUEUE.A
23:26:43,825 INFO [TextMDB] TextMDB.onMessage, this=5495387
23:26:43,825 INFO [TextMDB] TextMDB.sendReply, this=5495387, dest=QUEUE.A
23:26:43,880 INFO [TextMDB] TextMDB.ctor, this=15690844
23:26:43,884 INFO [TextMDB] TextMDB.setMessageDrivenContext, this=15690844
23:26:43,887 INFO [TextMDB] TextMDB.ejbCreate, this=15690844
23:26:43,944 INFO [TextMDB] TextMDB.onMessage, this=15690844
23:26:43,945 INFO [TextMDB] TextMDB.sendReply, this=15690844, dest=QUEUE.A
23:26:44,022 INFO [TextMDB] TextMDB.onMessage, this=13929978
23:26:44,022 INFO [TextMDB] TextMDB.sendReply, this=13929978, dest=QUEUE.A
23:26:44,041 INFO [TextMDB] TextMDB.onMessage, this=15690844
23:26:44,041 INFO [TextMDB] TextMDB.sendReply, this=15690844, dest=QUEUE.A
23:26:44,065 INFO [TextMDB] TextMDB.ctor, this=14089812
23:26:44,069 INFO [TextMDB] TextMDB.setMessageDrivenContext, this=14089812
23:26:44,069 INFO [TextMDB] TextMDB.ejbCreate, this=14089812
23:26:44,200 INFO [TextMDB] TextMDB.onMessage, this=5495387
23:26:44,201 INFO [TextMDB] TextMDB.sendReply, this=5495387, dest=QUEUE.A
23:26:44,249 INFO [TextMDB] TextMDB.onMessage, this=14089812
23:26:44,250 INFO [TextMDB] TextMDB.sendReply, this=14089812, dest=QUEUE.A
```

Items of note in this example include:

- The JMS client has no explicit knowledge that it is dealing with an MDB. The client simply uses the standard JMS APIs to send messages to a queue and receive messages from another queue.
- The MDB declares whether it will listen to a queue or topic in the `ejb-jar.xml` descriptor. The name of the queue or topic must be specified using a `jboss.xml` descriptor. In this example the MDB also sends messages to a JMS queue. MDBs may act as queue senders or topic publishers within their `onMessage` callback.

- The messages received by the client include a "processed by: NNN" suffix, where NNN is the `hashCode` value of the MDB instance that processed the message. This shows that many MDBs may actively process messages posted to a destination. Concurrent processing is one of the benefits of MDBs.

6.2. JBoss Messaging Overview

JBossMQ is composed of several services working together to provide JMS API level services to client applications. The services that make up the JBossMQ JMS implementation are introduced in this section.

6.2.1. Invocation Layer

The Invocation Layer (IL) services are responsible for handling the communication protocols that clients use to send and receive messages. JBossMQ can support running different types of Invocation Layers concurrently. All Invocation Layers support bidirectional communication which allows clients to send and receive messages concurrently. ILs only handle the transport details of messaging. They delegate messages to the JMS server JMX gateway service known as the invoker. This is similar to how the detached invokers expose the EJB container via different transports.

Each IL service binds a JMS connection factory to a specific location in the JNDI tree. Clients choose the protocol they wish to use by the JNDI location used to obtain the JMS connection factory. JBossMQ currently has six different invocation layers, and they are introduced in the following sections.

6.2.1.1. RMI IL (deprecated)

The first IL that was developed was based on Java's Remote Method Invocation (RMI). This is a robust IL since it is based on standard RMI technology, but it has a high overhead compared to other ILs and will likely be dropped in future releases.

NOTE: This IL will try to establish a TCP/IP socket from the server to the client. Therefore, clients that sit behind firewalls or have security restrictions prohibiting the use of `SeverSockets` should not use this IL.

6.2.1.2. OIL IL (deprecated)

The next IL that was developed was the Optimized IL (OIL). The OIL uses a custom TCP/IP protocol and serialization protocol that has very low overhead. This was the recommended socket based protocol until the addition of the UIL2 protocol.

NOTE: This IL will try to establish a TCP/IP socket from the server to the client. Therefore, clients that sit behind firewalls or have security restrictions prohibiting the use of `SeverSockets` should not use this IL.

6.2.1.3. UIL IL (deprecated)

The Unified Invocation Layer (UIL) was developed to allow clients that cannot have a connection created from the server back to the client due to firewall or other restrictions. It is almost identical to the OIL protocol except that a multiplexing layer is used to provide the bidirectional communication. The multiplexing layer creates two virtual sockets over one physical socket. This IL is slower than the OIL due to the higher overhead incurred by the multiplexing layer. This invocation layer is now deprecated in favor of UIL2.

6.2.1.4. UIL2 IL

The Unified version 2 Invocation Layer (UIL2) is a variation of the UIL protocol that also uses a single socket

between the client and server. However, unlike all other socket based invocation layers like RMI, UIL and OIL which use a blocking round-trip message at the socket level, the UIL2 protocol uses true asynchronous send and receive messaging at the transport level. This provides for improved throughput and utilization and as such, it is the preferred socket invocation layer.

6.2.1.5. JVM IL

The Java Virtual Machine (JVM) Invocation Layer was developed to cut out the TCP/IP overhead when the JMS client is running in the same JVM as the server. This IL uses direct method calls for the server to service the client requests. This increases efficiency since no sockets are created and there is no need for the associated worker threads. This is the IL that should be used by Message Driven Beans (MDB) or any other component that runs in the same virtual machine as the server such as servlets, MBeans, or EJBs.

6.2.1.6. HTTP IL

The HTTP Invocation Layer (HTTPIL) allows for accessing the JBossMQ service over the HTTP or HTTPS protocols. This IL relies on the servlet deployed in the `deploy/jms/jbossmq-httpil.sar` to handle the http traffic. This IL is useful for access to JMS through a firewall when the only port allowed requires HTTP.

6.2.2. Security Manager

The JBossMQ SecurityManager is the service that enforces an access control list to guard access to your destinations. This subsystem works closely with the StateManager service.

6.2.3. Destination Manager

The DestinationManager can be thought as being the central service in JBossMQ. It keeps track of all the destinations that have been created on the server. It also keeps track of the other key services such as the `MessageCache`, `StateManager`, and `PersistenceManager`.

6.2.4. Message Cache

Messages created in the server are passed to the `MessageCache` for memory management. JVM memory usage goes up as messages are added to a destination that does not have any receivers. These messages are held in the main memory until the receiver picks them up. If the `MessageCache` notices that the JVM memory usage starts passing the defined limits, the `MessageCache` starts moving those messages from memory to persistent storage on disk. The `MessageCache` uses a least recently used (LRU) algorithm to determine which messages should go to disk.

6.2.5. State Manager

The `StateManager` (SM) is in charge of keeping track of who is allowed to log into the server and what their durable subscriptions are.

6.2.6. Persistence Manager

The `PersistenceManager` (PM) is used by a destination to store messages marked as being persistent. JBossMQ has several different implementations of the persistent manager, but only one can be enabled per server instance. You should enable the persistence manager that best matches your requirements.

6.2.6.1. File PM

The File PM is a robust persistence manager that comes with JBossMQ. It creates separate directories for each of the destination created on the server, and stores each persistent message as a separate file in the appropriate directory. It has poor performance characteristics since it is frequently opening and closing files.

6.2.6.2. Rolling Logged PM

The Rolling Logged PM is also a file based persistence manager that has better performance than the File PM because it stores multiple messages in one file, reducing the overhead of opening/closing multiple files. This is a very fast PM but it is less transactionally reliable than the File PM due to its use of the `FileOutputStream.flush()` method call. On some operating systems/JVMs the `FileOutputStream.flush()` method does not guarantee that the data has been written to disk by the time the call returns.

6.2.6.3. JDBC2 PM

The JDBC2 PM is the second version of the original JDBC PM in JBossMQ 2.4.x. It has been substantially simplified and improved. This PM allows you to store persistent messages to a relational database using JDBC. The performance of this PM is directly related to the performance that can be obtained from the database. This PM has a very low memory overhead compared to the other persistence managers. Furthermore it is also highly integrated with the `MessageCache` to provide efficient persistence on a system that has a very active `MessageCache`.

6.2.7. Destinations

A destination is the object on the JBossMQ server that clients use to send and receive messages. There are two types of destination objects, `Queues` and `Topics`. References to the destinations created by JBossMQ are stored in JNDI.

6.2.7.1. Queues

Clients that are in the Point-to-Point paradigm typically use `Queues`. They expect that message sent to a `Queue` will be receive by only one other client once and only once. If multiple clients are receiving messages from a single queue, the messages will be load balanced across the receivers. `Queue` objects, by default, will be stored under the JNDI `queue/` sub context.

6.2.7.2. Topics

`Topics` are used in the publish-subscribe paradigm. When a client publishes a message to a topic, he expects that a copy of the message will be delivered to each client that has subscribed to the topic. Topic messages are delivered in the same manner a television show is delivered. Unless you have the TV on and are watching the show, you will miss it. Similarly, if the client is not up, running and receiving messages from the topics, it will miss messages published to the topic. To get around this problem of missing messages, clients can start a durable subscription. This is like having a VCR record a show you cannot watch at its scheduled time so that you can see what you missed when you turn your TV back on.

6.3. JBoss Messaging Configuration and MBeans

This section defines the MBean services that correspond to the components introduced in the previous section along with their MBean attributes. The configuration and service files that make up the JBossMQ system in-

clude:

- **conf/jbossmq-state.xml**: the configuration file read by the `org.jboss.mq.sm.file.DynamicStateManager` MBean. This is the default security store for the JMS valid username/passwords used to authenticate connections, as well as the active durable topic subscriptions.
- **deploy/jms/jbossmq-destinations-service.xml**: This service describes defines default JMS queue and topic destination configurations used by the testsuite unit tests. You can add/remove destinations to this file, or deploy another `*-service.xml` descriptor with the destination configurations.
- **deploy/jms/jbossmq-service.xml**: This service descriptor configures the core JBossMQ MBeans like the `Invoker`, `SecurityManager`, `DynamicStateManager`, and core interceptor stack. It also defines the MDB default dead letter queue `DLQ`.
- **deploy/jms/jms-ra.rar**: This is a JCA resource adaptor for JMS providers.
- **deploy/jms/jms-ds.xml**: This is a JCA connection factory and JMS provider MDB integration services configuration which sets JBossMQ as the JMS provider.
- **deploy/jms/hsqldb-jdbc2-service.xml**: This service descriptor configures the `DestinationManager`, `MessageCache`, and `jdbc2 PersistenceManager` for `hsqldb`.
- **deploy/jms/jvm-il-service.xml**: This service descriptor configures the `JVMServerILService` which provides the JVM IL transport.
- **deploy/jms/oil-service.xml**: This service descriptor configures the `OILServerILService` which provides the OIL transport. The queue and topic connection factory for this IL is bound under the JNDI name `ConnectionFactory`.
- **deploy/jms/oil2-service.xml**: This is an experimental version OIL transport that should not be used. Remove this descriptor as it will be dropped in the next release.
- **deploy/jms/rmi-il-service.xml**: This service descriptor configures the `RMIServerILService` which provides the RMI IL. The queue and topic connection factory for this IL is bound under the name `RMIConnectionFactory`.
- **deploy/jms/uil2-service.xml**: This service descriptor configures the `UILServerILService` which provides the UIL2 transport. The queue and topic connection factory for this IL is bound under the name `UIL2ConnectionFactory` as well as `UILConnectionFactory` to replace the deprecated version 1 UIL service.

We will discuss the associated MBeans in the following subsections.

6.3.1. org.jboss.mq.il.jvm.JVMServerILService

The `org.jboss.mq.il.jvm.JVMServerILService` MBean is used to configure the JVM IL. The configurable attributes are as follows:

- **Invoker**: This attribute specifies JMX ObjectName of the JMS entry point service that is used to pass incoming requests to the JMS server. This attribute should be setup via a `<depends optional-attribute-name="Invoker">` tag. This is not something you would typically change from the `jboss.mq:service=Invoker` setting unless you change the entry point service.

- **ConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `ConnectionFactory` setup to use this IL.
- **XAConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `XAConnectionFactory` setup to use this IL.
- **PingPeriod:** How often, in milliseconds, the client should send a ping message to the server to validate that the connection is still valid. If this is set to zero, then no ping message will be sent. Since it is impossible for JVM IL connection to go bad, it is recommended that you keep this set to 0.

6.3.2. org.jboss.mq.il.rmi.RMIServerILService (deprecated)

The `org.jboss.mq.il.rmi.RMIServerILService` is used to configure the RMI IL. The configurable attributes are as follows:

- **Invoker:** This attribute specifies JMX `ObjectName` of the JMS entry point service that is used to pass incoming requests to the JMS server. This attribute should be setup via a `<depends optional-attribute-name="Invoker">` tag. This is not something you would typically change from the `jboss.mq:service=Invoker` setting unless you change the entry point service.
- **ConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `ConnectionFactory` setup to use this IL.
- **XAConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `XAConnectionFactory` setup to use this IL.
- **PingPeriod:** How often, in milliseconds, the client should send a ping message to the server to validate that the connection is still valid. If this is set to zero, then no ping message will be sent.

6.3.3. org.jboss.mq.il.oil.OILServerILService (deprecated)

The `org.jboss.mq.il.oil.OILServerILService` is used to configure the OIL IL. The configurable attributes are as follows:

- **Invoker:** This attribute specifies JMX `ObjectName` of the JMS entry point service that is used to pass incoming requests to the JMS server. This attribute should be setup via a `<depends optional-attribute-name="Invoker">` tag. This is not something you would typically change from the `jboss.mq:service=Invoker` setting unless you change the entry point service.
- **ConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `ConnectionFactory` setup to use this IL.
- **XAConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `XAConnectionFactory` setup to use this IL.
- **PingPeriod:** How often, in milliseconds, the client should send a ping message to the server to validate that the connection is still valid. If this is set to zero, then no ping message will be sent.
- **ReadTimeout:** The period in milliseconds is passed onto as the `SoTimeout` value of the `UIL2` socket. This allows detection of dead sockets that are not responsive and are not capable of receiving ping messages. Note that this setting should be longer in duration than the `PingPeriod` setting.

- **ServerBindPort:** The protocol listening port for this IL. If not specified default is 0, which means that a random port will be chosen.
- **BindAddress:** The specific address this IL listens on. This can be used on a multi-homed host for a `java.net.ServerSocket` that will only accept connection requests on one of its addresses.
- **EnableTcpNoDelay:** If set to true, then the `TcpNoDelay` option is enabled. This improves request response times since TCP/IP packets are sent as soon as the request is flushed. Otherwise request packets may be buffered by the operating system to create larger IP packets.
- **ServerSocketFactory:** The `javax.net.ServerSocketFactory` implementation class name to use to create the service `java.net.ServerSocket`. If not specified the default factory will be obtained from `javax.net.ServerSocketFactory.getDefault()`.
- **ClientSocketFactory:** The `javax.net.SocketFactory` implementation class name to use on the client. If not specified the default factory will be obtained from `javax.net.SocketFactory.getDefault()`.
- **SecurityDomain:** Specify the security domain name to use with JBoss SSL aware socket factories. This is the JNDI name of the security manager implementation as described for the `security-domain` element of the `jboss.xml` and `jboss-web.xml` descriptors in *Enabling Declarative Security in JBoss Revisited*.

6.3.4. org.jboss.mq.il.uil.UILServerILService (deprecated)

The `org.jboss.mq.il.uil.UILServerILService` is used to configure the UIL IL. Note that this service has been removed from the default distribution in JBoss 3.2.2, but an example configuration file can be found in the `docs/examples/jca` directory.

The configurable attributes of the `UILServerILService` are as follows:

- **Invoker:** This attribute specifies JMX `ObjectName` of the JMS entry point service that is used to pass incoming requests to the JMS server. This attribute should be setup via a `<depends optional-attribute-name="Invoker">` tag. This is not something you would typically change from the `jboss.mq:service=Invoker` setting unless you change the entry point service.
- **ConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `ConnectionFactory` setup to use this IL.
- **XAConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `XAConnectionFactory` setup to use this IL.
- **PingPeriod:** How often, in milliseconds, the client should send a ping message to the server to validate that the connection is still valid. If this is set to zero, then no ping message will be sent.
- **ServerBindPort:** The protocol listening port for this IL. If not specified default is 0, which means that a random port will be chosen.
- **BindAddress:** The specific address this IL listens on. This can be used on a multi-homed host for a `java.net.ServerSocket` that will only accept connection requests on one of its addresses.
- **EnableTcpNoDelay:** If set to true, then the `TcpNoDelay` option is enabled. This improves request response times since TCP/IP packets are sent as soon as the request is flushed. Otherwise request packets may be buffered by the operating system to create larger IP packets.
- **ServerSocketFactory:** The `javax.net.ServerSocketFactory` implementation class name to use to cre-

ate the service `java.net.ServerSocket`. If not specified the default factory will be obtained from `javax.net.ServerSocketFactory.getDefault()`.

- **ClientSocketFactory:** The `javax.net.SocketFactory` implementation class name to use on the client. If not specified the default factory will be obtained from `javax.net.SocketFactory.getDefault()`.
- **SecurityDomain:** Specify the security domain name to use with JBoss SSL aware socket factories. This is the JNDI name of the security manager implementation as described for the `security-domain` element of the `jboss.xml` and `jboss-web.xml` descriptors in Section 8.3.1.

6.3.5. org.jboss.mq.il.util2.UILServerILService

The `org.jboss.mq.il.util2.UILServerILService` is used to configure the UIL2 IL. The configurable attributes are as follows:

- **Invoker:** This attribute specifies JMX `ObjectName` of the JMS entry point service that is used to pass incoming requests to the JMS server. This attribute should be setup via a `<depends optional-attribute-name="Invoker">` tag. This is not something you would typically change from the `jboss.mq:service=Invoker` setting unless you change the entry point service.
- **ConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `ConnectionFactory` setup to use this IL.
- **XAConnectionFactoryJNDIRef:** The JNDI location that this IL will bind a `XAConnectionFactory` setup to use this IL.
- **PingPeriod:** How often, in milliseconds, the client should send a ping message to the server to validate that the connection is still valid. If this is set to zero, then no ping message will be sent.
- **ReadTimeout:** The period in milliseconds is passed onto as the `SoTimeout` value of the UIL2 socket. This allows detection of dead sockets that are not responsive and are not capable of receiving ping messages. Note that this setting should be longer in duration than the `PingPeriod` setting.
- **BufferSize:** The size in bytes used as the buffer over the basic socket streams. This corresponds to the `java.io.BufferedOutputStream` buffer size.
- **ChunkSize:** The size in bytes between stream listener notifications. The UIL2 layer uses the `org.jboss.util.stream.NotifyingBufferedOutputStream` and `NotifyingBufferedInputStream` implementations that support the notion of a heartbeat that is triggered based on data read/written to the stream. Whenever `ChunkSize` bytes are read/written to a stream. This allows serves as a ping or keepalive notification when large reads or writes require a duration greater than the `PingPeriod`.
- **ServerBindPort:** The protocol listening port for this IL. If not specified default is 0, which means that a random port will be chosen.
- **BindAddress:** The specific address this IL listens on. This can be used on a multi-homed host for a `java.net.ServerSocket` that will only accept connection requests on one of its addresses.
- **EnableTcpNoDelay:** If set to true, then the `TcpNoDelay` option is enabled. This improves request response times since TCP/IP packets are sent as soon as the request is flushed. Otherwise request packets may be buffered by the operating system to create larger IP packets.
- **ServerSocketFactory:** The the `javax.net.ServerSocketFactory` implementation class name to use to cre-

ate the service `java.net.ServerSocket`. If not specified the default factory will be obtained from `javax.net.ServerSocketFactory.getDefault()`.

- **ClientAddress:** The address passed to the client as the address that should be used to connect to the server.
- **ClientSocketFactory:** The `javax.net.SocketFactory` implementation class name to use on the client. If not specified the default factory will be obtained from `javax.net.SocketFactory.getDefault()`.
- **SecurityDomain:** Specify the security domain name to use with JBoss SSL aware socket factories. This is the JNDI name of the security manager implementation as described for the `security-domain` element of the `jboss.xml` and `jboss-web.xml` descriptors in Section 8.3.1.

6.3.5.1. Configuring ILs for SSL

The UIL2 and OIL services support the use of SSL through custom socket factories that integrate JSSE using the security domain associated with the IL service. An example UIL2 service descriptor fragment that illustrates the use of the custom JBoss SSL socket factories is shown in Example 6.10.

Example 6.10. An example UIL2 config fragment for using SSL

```
<mbean code="org.jboss.mq.il.uid2.UILServerILService"
  name="jboss.mq:service=InvocationLayer,type=HTTPSUIL2">
  <depends optional-attribute-name="Invoker">jboss.mq:service=Invoker</depends>
  <attribute name="ConnectionFactoryJNDIRef">SSLConnectionFactory</attribute>
  <attribute name="XAConnectionFactoryJNDIRef">SSLXAConnectionFactory</attribute>

  <!-- ... -->

  <!-- SSL Socket Factories -->
  <attribute name="ClientSocketFactory">
    org.jboss.security.ssl.ClientSocketFactory
  </attribute>
  <attribute name="ServerSocketFactory">
    org.jboss.security.ssl.DomainServerSocketFactory
  </attribute>
  <!-- Security domain - see below -->
  <attribute name="SecurityDomain">java:/jaas/SSL</attribute>
</mbean>

<!-- Configures the keystore on the "SSL" security domain
  This mbean is better placed in conf/jboss-service.xml where it
  can be used by other services, but it will work from anywhere.
  Use keytool from the sdk to create the keystore. -->

<mbean code="org.jboss.security.plugins.JaasSecurityDomain"
  name="jboss.security:service=JaasSecurityDomain,domain=SSL">
  <!-- This must correlate with the java:/jaas/SSL above -->
  <constructor>
    <arg type="java.lang.String" value="SSL"/>
  </constructor>
  <!-- The location of the keystore resource: loads from the
    classpath and the server conf dir is a good default -->
  <attribute name="KeyStoreURL">resource:uil2.keystore</attribute>
  <attribute name="KeyStorePass">changeme</attribute>
</mbean>
```

6.3.5.2. JMS client properties for the UIL2 transport

There are several system properties that a JMS client using the UIL2 transport can set to control the client connection back to the server

- **org.jboss.mq.il.util2.useServerHost:** This system property allows a client to connect to the server `InetAddress.getHostName` rather than the `InetAddress.getHostAddress` value. This will only make a difference if there is a different address to name resolution between the server and client environments.
- **org.jboss.mq.il.util2.localAddr:** This system property allows a client to define the local interface to which its sockets should be bound.
- **org.jboss.mq.il.util2.localPort:** This system property allows a client to define the local port to which its sockets should be bound.
- **org.jboss.mq.il.util2.serverAddr:** This system property allows a client to override the address to which it attempts to connect to. This is useful for networks where NAT is occurring between the client and JMS server.
- **org.jboss.mq.il.util2.serverPort:** This system property allows a client to override the port to which it attempts to connect. This is useful for networks where port forwarding is occurring between the client and JMS server.
- **org.jboss.mq.il.util2.retryCount:** This system property controls the number of attempts to retry connecting to the JMS server. Retries are only made for `java.net.ConnectException` failures. A value ≤ 0 means no retry attempts will be made.
- **org.jboss.mq.il.util2.retryDelay:** This system property controls the delay in milliseconds between retries due to `ConnectException` failures.

6.3.6. org.jboss.mq.il.http.HTTPServerILService

The `org.jboss.mq.il.http.HTTPServerILService` is used to manage the HTTP/S IL. This IL allows for the use of the JMS service over HTTP or HTTPS connections. It relies on the servlet deployed in the `deploy/jms/jbossmq-httpil.sar` to handle the HTTP traffic. The configurable attributes are as follows:

- **TimeOut:** The default timeout in seconds that the client HTTP requests will wait for messages. This can be overridden on the client by setting the system property `org.jboss.mq.il.http.timeout` to the number of seconds.
- **RestInterval:** The number of seconds the client will sleep after each request. The default is 0, but you can set this value in conjunction with the `TimeOut` value to implement a pure timed based polling mechanism. For example, you could simply do a short lived request by setting the `TimeOut` value to 0 and then setting the `RestInterval` to 60. This would cause the client to send a single non-blocking request to the server, return any messages if available, then sleep for 60 seconds, before issuing another request. Like the `TimeOut` value, this can be explicitly overridden on a given client by specifying the `org.jboss.mq.il.http.restinterval` with the number of seconds you wish to wait between requests.
- **URL:** Set the servlet URL. This value takes precedence over any individual values set (i.e. the `URLPrefix`, `URLSuffix`, `URLPort`, etc.) It may be an actual URL or a property name which will be used on the client side to resolve the proper URL by calling `System.getProperty(propertyName)`. If not specified the URL will be formed from `URLPrefix + URLHostName + ":" + URLPort + "/" + URLSuffix`.
- **URLPrefix:** The prefix portion of the servlet URL.
- **URLHostName:** The hostname portion of the servlet URL.
- **URLPort:** The port portion of the URL.

- **URLSuffix:** The trailing path portion of the URL.
- **UseHostName:** A flag that if set to true the default setting for the `URLHostName` attribute will be taken from `InetAddress.getLocalHost().getHostName()`. If false the default setting for the `URLHostName` attribute will be taken from `InetAddress.getLocalHost().getHostAddress()`.

6.3.7. org.jboss.mq.server.jmx.Invoker

The `org.jboss.mq.server.jmx.Invoker` is used to pass IL requests down to the destination manager service through an interceptor stack. The configurable attributes are as follows:

- **NextInterceptor:** The JMX `ObjectName` of the next request interceptor. This attribute is used by all the interceptors to create the interceptor stack. The last interceptor in the chain should be the `DestinationManager`. This attribute should be setup via a `<depends optional-attribute-name="NextInterceptor">` tag.

6.3.8. org.jboss.mq.server.jmx.InterceptorLoader

The `org.jboss.mq.server.jmx.InterceptorLoader` is used to load a generic interceptor and make it part of the interceptor stack. This MBean is typically used to load custom interceptors like `org.jboss.mq.server.TracingInterceptor`, which is can be used to efficiently log all client requests via trace level log messages. The configurable attributes are as follows:

- **NextInterceptor:** The JMX `ObjectName` of the next request interceptor. This attribute is used by all the interceptors to create the interceptor stack. The last interceptor in the chain should be the `DestinationManager`. This attribute should be setup via a `<depends optional-attribute-name="NextInterceptor">` XML tag.
- **InterceptorClass:** The class name of the interceptor that will be loaded and made part of the interceptor stack. This class specified here must extend the `org.jboss.mq.server.JMSServerInterceptor` class.

6.3.9. org.jboss.mq.sm.file.DynamicStateManager

The `org.jboss.mq.sm.file.DynamicStateManager` MBean is used as the default state manager assigned to the `DestinationManager` service. It manages an XML user security store that provides the authentication, authorization and durable subscriber information. The configurable attributes are as follows:

- **StateFile:** The file used to store state information such as created durable subscriptions. This is an XML file that the server reads and writes data to, and the content model is shown in Figure 6.1. You should never edit the XML file while the server is running. The default is the `conf/jbossmq-state.xml` file.
- **User/Name:** the username that corresponds to the `Connection.createConnection(username, password)` method.
- **User/Password:** the password that corresponds to the `Connection.createConnection(username, password)` method.
- **User/Id:** the clientID that will be associated with the connection for the username. This limits the client to a single active connection.
- **DurableSubscriptions/DurableSubscription/ClientID:** the unique client connection id associated with the

durable subscription.

- **DurableSubscriptions/DurableSubscription/Name:** the name of the durable subscription. This is the value passed in as the name parameter to the `TopicSession.createDurableSubscriber(Topic, name)` method.
- **DurableSubscriptions/DurableSubscription/TopicName:** the name of the Topic currently associated with the durable subscription.
- **HasSecurityManager:** A boolean flag indicating whether the JAAS `SecurityManager` service has been configured as part of the core JMS server interceptor stack. If false, this service performs that connection authentication. The default is true.

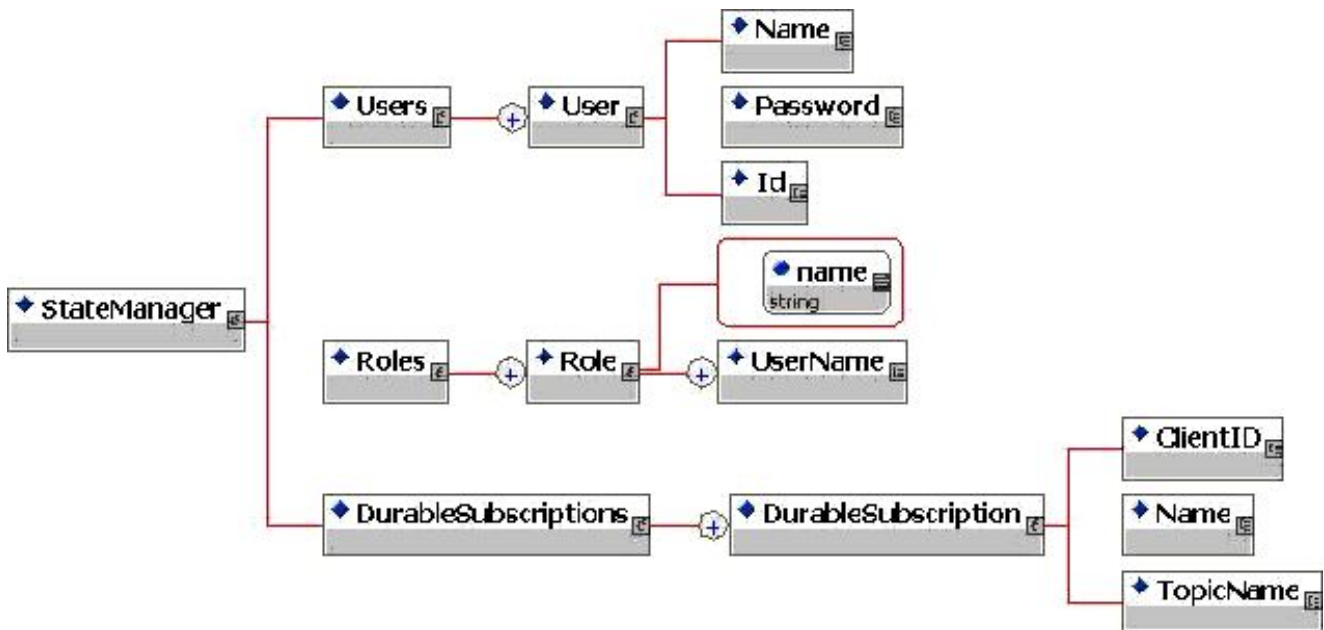


Figure 6.1. The jbossmq-state.xml content model

6.3.10. org.jboss.mq.security.SecurityManager

If the `org.jboss.mq.security.SecurityManager` is part of the interceptor stack, then it will enforce the access control lists assigned to the destinations. The `SecurityManager` uses JAAS, and as such requires that an application policy be setup for in the JBoss `login-config.xml` file. The default configuration is shown below.

```
<application-policy name="jbossmq">
  <authentication>
    <login-module code="org.jboss.mq.sm.file.DynamicLoginModule"
      flag="required">
      <module-option name="unauthenticatedIdentity">guest</module-option>
      <module-option name="sm.objectname">
        jboss.mq:service=StateManager
      </module-option>
    </login-module>
  </authentication>
</application-policy>
```

This integrates the `DynamicStateManager` `jbossmq-state.xml` security store into the JAAS based framework through the `org.jboss.mq.sm.file.DynamicLoginModule`. The configuration also maps any unauthenticated JBossMQ client to the `guest` role.

The configurable attributes of the SecurityManager are as follows:

- **NextInterceptor:** The JMX `ObjectName` of the next request interceptor. This attribute is used by all the interceptors to create the interceptor stack. The last interceptor in the chain should be the `DestinationManager`. This attribute should be setup via a `<depends optional-attribute-name="NextInterceptor">` tag.
- **DefaultSecurityConfig:** This element specifies the default security configuration settings for destinations. This applies to temporary queues and topics as well as queues and topics that do not specifically specify a security configuration. The content model of this element is shown in Figure 6.2.



Figure 6.2. The destination security config content model

- **role:** Each role that is allowed access to a destination is represented by a role element.
- **role@name:** The name attribute defines the name of the role.
- **role@create:** The create attribute is a true/false enum that indicates whether the role has the ability to create durable subscriptions on the topic.
- **role@read:** The read attribute is a true/false enum that indicates whether the role can receive messages from the destination.
- **role@write:** The write attribute is a true/false enum that indicates whether the role can send messages to the destination.
- **SecurityDomain:** Specify the security domain name to use for authentication and role based authorization. This is the JNDI name of the security manager implementation as described for the `security-domain` element of the `jboss.xml` and `jboss-web.xml` descriptors in Section 8.3.1. Note however, that this attribute value cannot have the standard `java:/jaas` prefix and that this prefix is currently an assumed and hard-coded value.

You may be uncomfortable having to maintain your authentication and authorization information in an XML file. You can use any standard security store such as a database or LDAP server by simply updating the JAAS `login-config.xml` to provide the same username to password and user to role mappings as the `DynamicStateManager`. For example, to use a JDBC database, the following sample database tables and `login-config.xml` entry would work.

Table 6.1. An example JMSPasswords username to password table

username	password
jduke	theduke

Table 6.2. An example JMSRoles username to acl roles table

username	role
jduke	create
jduke	read
jduke	write

Example 6.11. An alternate login-config.xml configuration for JBossMQ

```
<application-policy name="jbossmq">
  <authentication>
    <login-module
      code="org.jboss.security.auth.spi.DatabaseServerLoginModule" flag="required">
      <module-option name="unauthenticatedIdentity">guest </module-option>
      <module-option name="dsJndiName">java:/DefaultDS </module-option>
      <module-option name="principalsQuery"> select password from
        JMSPasswords where username = ? </module-option>
      <module-option name="rolesQuery"> select role, "Roles" from JMSRoles
        where username= ? </module-option>
    </login-module>
  </authentication>
</application-policy>
```

For a complete description of the `DatabaseServerLoginModule` see Section 8.4.6.4.

6.3.11. org.jboss.mq.server.jmx.DestinationManager

The `org.jboss.mq.server.jmx.DestinationManager` must be the last interceptor in the interceptor stack. The configurable attributes are as follows:

- **PersistenceManager:** The JMX `ObjectName` of the persistence manager service the server should use. This attribute should be setup via a `<depends optional-attribute-name="PersistenceManager">` XML tag.
- **StateManager:** The JMX `ObjectName` of the state manager service the server should use. This attribute should be setup via a `<depends optional-attribute-name="StateManager">` XML tag.
- **MessageCache:** The JMX `ObjectName` of the message cache service the server should use. This attribute should be setup via a `<depends optional-attribute-name="StateManager">` XML tag.

Additional read-only attributes and operations that support monitoring include:

- **ClientCount:** The number of clients connected to the server.
- **Clients:** A `java.util.Map<org.jboss.mq.ConnectionToken, org.jboss.mq.server.ClientConsumer>` instances for the clients connected to the server.
- **MessageCounter:** An array of `org.jboss.mq.server.MessageCounter` instances that provide statistics for a JMS destination.
- **String listMessageCounter():** This operation generates an HTML table that contains:
 - **Type:** Either Queue or Topic indicating the destination type.

- **Name:** The name of the destination.
- **Subscription:** The subscription ID for a topic.
- **Durable:** A boolean indicating if the topic subscription is durable.
- **Count:** The number of message delivered to the destination.
- **CountDelta:** The change in message count since the previous access of count.
- **Depth:** The number of messages in the destination.
- **DepthDelta:** The change in the number of messages in the destination since the previous access of depth.
- **Last Add:** The date/time string in `DateFormat.SHORT`/`DateFormat.MEDIUM` format of the last time a message was added to the destination.
- **void resetMessageCounter():** This zeros all destination counts and last added times.

6.3.12. org.jboss.mq.server.MessageCache

The server determines when to move messages to secondary storage by using the `org.jboss.mq.server.MessageCache` MBean. The configurable attributes are as follows:

- **CacheStore:** The JMX `ObjectName` of the service that will act as the cache store. The cache store is used by the `MessageCache` to move messages to persistent storage. The value you set here typically depends on the type of persistence manager you are using. This attribute should be setup via a `<depends optional-attribute-name="CacheStore">` XML tag.
- **HighMemoryMark:** The amount of JVM heap memory in megabytes that must be reached before the `MessageCache` starts to move messages to secondary storage.
- **MaxMemoryMark:** The maximum amount of JVM heap memory in megabytes that the `MessageCache` considers to be the max memory mark. As memory usage approaches the max memory mark, the `MessageCache` will move messages to persistent storage so that the number of messages kept in memory approaches zero.

Additional read-only cache attribute that provide statistics include:

- **CacheHits:** The number of times a message was requested and it was found to be in memory.
- **CacheMisses:** The number of times a message was requested and it was not found in memory so a read from persistent storage was required to retrieve the message.
- **HardRefCacheSize:** The number of messages the cache forcing to stay inn memory by using a hard reference.
- **SoftRefCacheSize:** The number of messages the cache has persisted but is still lingering around in memory as soft references due to the garbage collector not being eager to free up space.
- **TotalCacheSize:** The total number of messages that are being managed by the cache.

6.3.13. org.jboss.mq.pm.file.CacheStore

The `org.jboss.mq.pm.file.CacheStore` MBean should be used as the cache store for the `MessageCache` service when you are using the File or Rolling Logged PM. The configurable attributes are as follows:

- **DataDirectory:** The directory used to store messages for the `MessageCache`.

6.3.14. org.jboss.mq.pm.file.PersistenceManager

The `org.jboss.mq.pm.file.PersistenceManager` should be used as the Persistence Manager assigned to the `DestinationManager` if you wish to use the File PM. The configurable attributes are as follows:

- **MessageCache:** The JMX `ObjectName` of the `MessageCache` that has been assigned to the `DestinationManager`. This attribute should be setup via a `<depends optional-attribute-name="MessageCache">` XML tag.
- **DataDirectory:** The directory used to store persistent messages.

6.3.15. org.jboss.mq.pm.rollinglogged.PersistenceManager

The `org.jboss.mq.pm.rollinglogged.PersistenceManager` should be used as the `PersistenceManager` assigned to the `DestinationManager` if you wish to use the Rolling Logged PM. The configurable attributes are as follows:

- **DataDirectory:** The directory used to store persistent messages.
- **RollOverSize:** The maximum number of messages until log rolls over.

6.3.16. org.jboss.mq.pm.jdbc2.PersistenceManager

The `org.jboss.mq.pm.jdbc.PersistenceManager` should be used as the persistence manager assigned to the `DestinationManager` if you wish to store messages in a database. This PM has been tested against the HyperSQL, MS SQL, Oracle, MySQL and Postgres databases. The configurable attributes are as follows:

- **MessageCache:** The JMX `ObjectName` of the `MessageCache` that has been assigned to the `DestinationManager`. This attribute should be setup via a `<depends optional-attribute-name="MessageCache">` XML tag.
- **ConnectionManager:** The JMX `ObjectName` of the JCA data source that will be used to obtain JDBC connections. This attribute should be setup via a `<depends optional-attribute-name="DataSource">` XML tag. You may also need to add another `<depends>` XML tag to wait for the data source `datasource` connection manager service to be started before this PM is started.
- **ConnectionRetryAttempts:** An integer count used to allow the PM to retry attempts at getting a connection to the JDBC store. There is a 1500 millisecond delay between each connection failed connection attempt and the next attempt. This must be greater than or equal to 1 and defaults to 5.
- **SqlProperties:** A property list is used to define the SQL Queries and other JDBC2 Persistence Manager options. You will need to adjust these properties if you which to run against another database other than HypersonicSQL. Example 6.12 shows default setting for this attribute for the Hypersonic DB, while Ex-

ample 6.13 shows an alternate setting for Oracle. Additional examples can be found in the `docs/examples/jms` directory of the distribution.

Example 6.12. Default JDBC2 PersistenceManager SqlProperties

```
<attribute name="SqlProperties">
  BLOB_TYPE=OBJECT_BLOB
  INSERT_TX = INSERT INTO JMS_TRANSACTIONS (TXID) values(?)
  INSERT_MESSAGE = INSERT INTO JMS_MESSAGES (MESSAGEID, DESTINATION,
    MESSAGEBLOB, TXID, TXOP) VALUES(?,?,?,?,?)
  SELECT_ALL_UNCOMMITTED_TXS = SELECT TXID FROM JMS_TRANSACTIONS
  SELECT_MAX_TX = SELECT MAX(TXID) FROM JMS_MESSAGES
  SELECT_MESSAGES_IN_DEST = SELECT MESSAGEID, MESSAGEBLOB FROM JMS_MESSAGES \
    WHERE DESTINATION=?
  SELECT_MESSAGE = SELECT MESSAGEID, MESSAGEBLOB FROM JMS_MESSAGES WHERE \
    MESSAGEID=? AND DESTINATION=?
  MARK_MESSAGE = UPDATE JMS_MESSAGES SET TXID=?, TXOP=? WHERE MESSAGEID=? AND \
    DESTINATION=?
  UPDATE_MESSAGE = UPDATE JMS_MESSAGES SET MESSAGEBLOB=? WHERE MESSAGEID=? AND \
    DESTINATION=?
  UPDATE_MARKED_MESSAGES = UPDATE JMS_MESSAGES SET TXID=?, TXOP=? WHERE TXOP=?
  UPDATE_MARKED_MESSAGES_WITH_TX = UPDATE JMS_MESSAGES SET TXID=?, TXOP=? WHERE
    TXOP=? AND TXID=?
  DELETE_MARKED_MESSAGES_WITH_TX = DELETE FROM JMS_MESSAGES WHERE TXID IN \
    (SELECT TXID FROM JMS_TRANSACTIONS) AND TXOP=?
  DELETE_TX = DELETE FROM JMS_TRANSACTIONS WHERE TXID = ?
  DELETE_MARKED_MESSAGES = DELETE FROM JMS_MESSAGES WHERE TXID=? AND TXOP=?
  DELETE_MESSAGE = DELETE FROM JMS_MESSAGES WHERE MESSAGEID=? AND DESTINATION=?
  CREATE_MESSAGE_TABLE = CREATE TABLE JMS_MESSAGES ( MESSAGEID INTEGER NOT NULL, \
    DESTINATION VARCHAR(255) NOT NULL, TXID INTEGER, TXOP CHAR(1), \
    MESSAGEBLOB OBJECT, PRIMARY KEY (MESSAGEID, DESTINATION) )
  CREATE_TX_TABLE = CREATE TABLE JMS_TRANSACTIONS ( TXID INTEGER )
</attribute>
```

Example 6.13. A sample JDBC2 PersistenceManager SqlProperties for Oracle

```
<attribute name="SqlProperties">
  BLOB_TYPE=BINARYSTREAM_BLOB
  INSERT_TX = INSERT INTO JMS_TRANSACTIONS (TXID) values(?)
  INSERT_MESSAGE = INSERT INTO JMS_MESSAGES (MESSAGEID, DESTINATION, \
    MESSAGEBLOB, TXID, TXOP) VALUES(?,?,?,?,?)
  SELECT_ALL_UNCOMMITTED_TXS = SELECT TXID FROM JMS_TRANSACTIONS
  SELECT_MAX_TX = SELECT MAX(TXID) FROM JMS_MESSAGES
  SELECT_MESSAGES_IN_DEST = SELECT MESSAGEID, MESSAGEBLOB FROM JMS_MESSAGES \
    WHERE DESTINATION=?
  SELECT_MESSAGE = SELECT MESSAGEID, MESSAGEBLOB FROM JMS_MESSAGES WHERE \
    MESSAGEID=? AND DESTINATION=?
  MARK_MESSAGE = UPDATE JMS_MESSAGES SET TXID=?, TXOP=? WHERE MESSAGEID=? \
    AND DESTINATION=?
  UPDATE_MESSAGE = UPDATE JMS_MESSAGES SET MESSAGEBLOB=? WHERE MESSAGEID=? \
    AND DESTINATION=?
  UPDATE_MARKED_MESSAGES = UPDATE JMS_MESSAGES SET TXID=?, TXOP=? WHERE TXOP=?
  UPDATE_MARKED_MESSAGES_WITH_TX = UPDATE JMS_MESSAGES SET TXID=?, TXOP=? WHERE \
    TXOP=? AND TXID=?
  DELETE_MARKED_MESSAGES_WITH_TX = DELETE FROM JMS_MESSAGES WHERE TXID IN \
    (SELECT TXID FROM JMS_TRANSACTIONS) AND TXOP=?
  DELETE_TX = DELETE FROM JMS_TRANSACTIONS WHERE TXID = ?
  DELETE_MARKED_MESSAGES = DELETE FROM JMS_MESSAGES WHERE TXID=? AND TXOP=?
  DELETE_MESSAGE = DELETE FROM JMS_MESSAGES WHERE MESSAGEID=? AND DESTINATION=?
  CREATE_MESSAGE_TABLE = CREATE TABLE JMS_MESSAGES ( MESSAGEID INTEGER NOT NULL, \
    DESTINATION VARCHAR(255) NOT NULL, TXID INTEGER, TXOP CHAR(1), \
    MESSAGEBLOB BLOB, PRIMARY KEY (MESSAGEID, DESTINATION) )
  CREATE_TX_TABLE = CREATE TABLE JMS_TRANSACTIONS ( TXID INTEGER )
</attribute>
```

```
CREATE_TX_TABLE = CREATE TABLE JMS_TRANSACTIONS ( TXID INTEGER )
</attribute>
```

6.3.17. Destination MBeans

This section describes the destination MBeans used in the `jbossmq-destinations-service.xml` and `jbossmq-service.xml` descriptors.

6.3.17.1. org.jboss.mq.server.jmx.Queue

The `org.jboss.mq.server.jmx.Queue` is used to define a Queue Destination on the JBossMQ server. The `name` attribute of the JMX object name of this MBean is used to determine the destination name. For example, if the JMX MBean begins with:

```
<mbean code="org.jboss.mq.server.jmx.Queue"
      name="jboss.mq.destination:service=Queue,name=testQueue">
```

Then, the JMX object name is `jboss.mq.destination:service=Queue,name=testQueue` and the name of the queue is "testQueue". The configurable attributes are as follows:

- **DestinationManager:** The JMX `ObjectName` of the destination manager service for the server. This attribute should be set via a `<depends optional-attribute-name="DestinationManager">` XML tag.
- **SecurityManager:** The JMX `ObjectName` of the security manager service that is being used to validate client requests. This attribute should be set via a `<depends optional-attribute-name="SecurityManager">` XML tag.
- **SecurityConf:** This element specifies a XML fragment which describes the access control list to be used by the `SecurityManager` to authorize client operations against the destination. The content model is the same as for the `SecurityManagerSecurityConf` attribute.
- **JNDIName:** The location in JNDI to which the queue object will be bound. If this is not set it will default to `queue/queue-name`.
- **MaxDepth:** The `MaxDepth` is an upper limit to the backlog of messages that can exist for a destination, and if exceeded, attempts to add new messages will result in a `org.jboss.mq.DestinationFullException`. The `MaxDepth` can still be exceeded in a number of situations, e.g. when a message is knacked back into the queue. Also transactions performing read committed processing, look at the current size of queue, ignoring any messages that may be added as a result of the current transaction or other transactions. This is because we don't want the transaction to fail during the commit phase when the message is physically added to the queue.
- **MessageCounterHistoryDayLimit:** Sets the destination message counter history day limit with a value `< 0` indicating unlimited history, a `0` value disabling history, and a value `> 0` giving the history days count.

Additional read-only attributes that provide statistics information include:

- **MessageCounter:** An array of `org.jboss.mq.server.MessageCounter` instances that provide statistics for this destination.
- **QueueDepth:** The current backlog of waiting messages.

- **ReceiversCount:** The number of receivers currently associated with the queue.
- **ScheduledMessageCount:** The number of messages waiting in the queue for their scheduled delivery time to arrive.
- **String listMessageCounter():** This operation generates an HTML table that contains:
 - **Type:** Either `Queue` or `Topic` indicating the destination type.
 - **Name:** The name of the destination.
 - **Subscription:** The subscription ID for a topic.
 - **Durable:** A boolean indicating if the topic subscription is durable.
 - **Count:** The number of message delivered to the destination.
 - **CountDelta:** The change in message count since the previous access of count.
 - **Depth:** The number of messages in the destination.
 - **DepthDelta:** The change in the number of messages in the destination since the previous access of depth.
 - **Last Add:** The date/time string in `DateFormat.SHORT/DateFormat.MEDIUM` format of the last time a message was added to the destination.
- **void resetMessageCounter():** This zeros all destination counts and last added times.
- **String listMessageCounterHistory():** This operation display an HTML table showing the hourly message counts per hour for each day of history.
- **void resetMessageCounterHistory():** This operation resets the day history message counts.

6.3.17.2. org.jboss.mq.server.jmx.Topic

The `org.jboss.mq.server.jmx.Topic` is used to define a topic destination on the JBossMQ server. The name attribute of the JMX object name of this MBean is used to determine the destination name. For example, if the JMX MBean begins with:

```
<mbean code="org.jboss.mq.server.jmx.Topic"
      name="jboss.mq.destination:service=Topic,name=testTopic">
```

Then, the JMX object name is `jboss.mq.destination:service=Topic,name=testTopic` and the name of the topic is `testTopic`. The configurable attributes are as follows:

- **DestinationManager:** The JMX object name of the destination manager configured for the server. This attribute should be setup via a `<depends optional-attribute-name="DestinationManager">` XML tag.
- **SecurityManager:** The JMX object name of the security manager that is being used to validate client requests. This attribute should be setup via a `<depends optional-attribute-name="SecurityManager">` XML tag.
- **SecurityConf:** This element specifies a XML fragment which describes the access control list to be used by the `SecurityManager` to authorize client operations against the destination. The content model is the same

as that for the `SecurityManagerSecurityConf` attribute.

- **JNDIName:** The location in JNDI to which the queue object will be bound. If this is not set it will default to `topic/topic-name`.
- **MaxDepth:** The `MaxDepth` is an upper limit to the backlog of messages that can exist for a destination, and if exceeded, attempts to add new messages will result in a `org.jboss.mq.DestinationFullException`. The `MaxDepth` can still be exceeded in a number of situations, e.g. when a message is knacked back into the queue. Also transactions performing read committed processing, look at the current size of queue, ignoring any messages that may be added as a result of the current transaction or other transactions. This is because we don't want the transaction to fail during the commit phase when the message is physically added to the topic.
- **MessageCounterHistoryDayLimit:** Sets the destination message counter history day limit with a value < 0 indicating unlimited history, a 0 value disabling history, and a value > 0 giving the history days count.

Additional read-only attributes that provide statistics information include:

- **AllMessageCount:** The message count across all queue types associated with the topic.
- **AllSubscriptionsCount:** The count of durable and non-durable subscriptions.
- **DurableMessageCount:** The count of messages in durable subscription queues.
- **DurableSubscriptionsCount:** The count of durable subscribers.
- **MessageCounter:** An array of `org.jboss.mq.server.MessageCounter` instances that provide statistics for this destination.
- **NonDurableMessageCount:** The count on messages in non-durable subscription queues.
- **NonDurableSubscriptionsCount:** The count of non-durable subscribers.
- **String listMessageCounter():** This operation generates an HTML table that contains
 - **Type:** Either Queue or Topic indicating the destination type.
 - **Name:** The name of the destination.
 - **Subscription:** The subscription ID for a topic.
 - **Durable:** A boolean indicating if the topic subscription is durable.
 - **Count:** The number of message delivered to the destination.
 - **CountDelta:** The change in message count since the previous access of count.
 - **Depth:** The number of messages in the destination.
 - **DepthDelta:** The change in the number of messages in the destination since the previous access of depth.
 - **Last Add:** The date/time string in `DateFormat.SHORT`/`DateFormat.MEDIUM` format of the last time a message was added to the destination.

- **void resetMessageCounter():** This zeros all destination counts and last added times.
- **String listMessageCounterHistory():** This operation display an HTML table showing the hourly message counts per hour for each day of history.
- **void resetMessageCounterHistory():** This operation resets the day history message counts.

6.3.18. Administration Via JMX

JBossMQ statistics and several management functions are accessible via JMX. JMX can be accessed interactively via a Web Application or programmatically via the JMX API. It is recommended that you use the `http://localhost:8080/jmx-console` web application to get familiar with all the JBossMQ JMX MBeans running inside the server and how to invoke methods on those MBeans via the JMX Console. This section will outline the most common runtime management tasks that administrators must perform.

6.3.18.1. Creating Queues At Runtime

Applications that require the dynamic creation of queues at runtime can use the Destination Manager's MBean `createQueue` method: `void createQueue(String name, String jndiLocation)`

This method creates a queue with the given `name` and binds it in JNDI at the `jndiLocation`. Queues created via this method exist until the server is restarted. To destroy a previously created Queue, you would issue a `void destroyQueue(String name)`

6.3.18.2. Creating Topics At Runtime

Applications that require the dynamic creation of topics at runtime can use the Destination Manager's MBean `createTopic` method: `void createTopic(String name, String jndiLocation)`

This method creates a topic with the given `name` and binds it in JNDI at the `jndiLocation`. Topics created via this method exist until the server is restarted. To destroy a previously created Topic, you would issue a: `void destroyTopic(String name)`

6.3.18.3. Managing a JBossMQ User IDs at Runtime

The `org.jboss.mq.sm.file.DynamicStateManager`'s MBean can be used to add and remove user ids and roles at runtime. To add a user id, you would use: `void addUser(String name, String password, String clientID)`

This method creates a user id with the given `name` and `password` and configures him to have the given `clientID`. To remove a previously created user id, you would call the following method: `void removeUser(String name)`

To manage the roles that the user ids belong to, you would use the following set of methods to create roles, remove roles, add users to roles, and remove users from roles:

- `void addRole(String name)`
- `void removeRole(String name)`
- `void addUserToRole(String roleName, String user)`
- `void removeUserFromRole(String roleName, String user)`

6.4. Specifying the MDB JMS Provider

Up to this point we have looked at the standard JMS client/server architecture. The JMS specification defines an advanced set of interfaces that allow for concurrent processing of a destination's messages, and collectively this functionality is referred to as application server facilities (ASF). Two of the interfaces that support concurrent message processing, `javax.jms.ServerSessionPool` and `javax.jms.ServerSession`, must be provided by the application server in which the processing will occur. Thus, the set of components that make up the JBossMQ ASF involves both JBossMQ components as well as JBoss server components. The JBoss server MDB container utilizes the JMS service's ASF to concurrently process messages sent to MDBs.

The responsibilities of the ASF domains are well defined by the JMS specification and so we won't go into a discussion of how the ASF components are implemented. Rather, we want to discuss how ASF components used by the JBoss MDB layer are integrated using MBeans that allow either the application server interfaces, or the JMS provider interfaces to be replaced with alternate implementations.

Let's start with the `org.jboss.jms.jndi.JMSProviderLoader` MBean. This MBean is responsible for loading an instance of the `org.jboss.jms.jndi.JMSProviderAdaptor` interface into the JBoss server and binding it into JNDI. The `JMSProviderAdaptor` interface is an abstraction that defines how to get the root JNDI context for the JMS provider, and an interface for getting and setting the JNDI names for the `Context.PROVIDER_URL` for the root `InitialContext`, and the `QueueConnectionFactory` and `TopicConnectionFactory` locations in the root context. This is all that is really necessary to bootstrap use of a JMS provider. By abstracting this information into an interface, alternate JMS ASF provider implementations can be used with the JBoss MDB container. The `org.jboss.jms.jndi.JBossMQProvider` is the default implementation of `JMSProviderAdaptor` interface, and provides the adaptor for the JBossMQ JMS provider. To replace the JBossMQ provider with an alternate JMS ASF implementation, simply create an implementation of the `JMSProviderAdaptor` interface and configure the `JMSProviderLoader` with the class name of the implementation. We'll see an example of this in the configuration section.

In addition to being able to replace the JMS provider used for MDBs, you can also replace the `javax.jms.ServerSessionPool` interface implementation. This is possible by configuring the class name of the `org.jboss.jms.asf.ServerSessionPoolFactory` implementation using the `org.jboss.jms.asf.ServerSessionPoolLoader` MBean `PoolFactoryClass` attribute. The default `ServerSessionPoolFactory` factory implementation is the JBoss `org.jboss.jms.asf.StdServerSessionPoolFactory` class.

6.4.1. org.jboss.jms.jndi.JMSProviderLoader MBean

The `JMSProviderLoader` MBean service creates a JMS provider adaptor and binds it into JNDI. A JMS provider adaptor is a class that implements the `org.jboss.jms.jndi.JMSProviderAdapter` interface. It is used by the message driven bean container to access a JMS service provider in a provider independent manner. The configurable attributes of the `JMSProviderLoader` service are:

- **ProviderName:** A unique name for the JMS provider. This will be used to bind the `JMSProviderAdapter` instance into JNDI under `java: /<ProviderName>` unless overridden by the `AdapterJNDIName` attribute.
- **ProviderAdapterClass:** The fully qualified class name of the `org.jboss.jms.jndi.JMSProviderAdapter` interface to create an instance of. To use an alternate JMS provider like SonicMQ, one would create an implementation of the `JMSProviderAdaptor` interface that allows the administration of the `InitialContext` provider URL, and the locations of the `QueueConnectionFactory` and `TopicConnectionFactory` in JNDI.
- **AdapterJNDIName:** Specify the exact name into JNDI under which the `JMSProviderAdapter` instance will

be bound.

- **ProviderURL:** The JNDI `Context.PROVIDER_URL` value to use when creating the JMS provider root `InitialContext`.
- **QueueFactoryRef:** The JNDI name under which the provider `javax.jms.QueueConnectionFactory` will be bound.
- **TopicFactoryRef:** The JNDI name under which the `javax.jms.TopicConnectionFactory` will be bound.

Example 6.14. A JMSProviderLoader for accessing a remote JBossMQ server

```
<mbean code="org.jboss.jms.jndi.JMSProviderLoader"
  name="jboss.mq:service=JMSProviderLoader,name=RemoteJBossMQProvider">
  <attribute name="ProviderName">RemoteJMSProvider</attribute>
  <attribute name="ProviderUrl">jnp://remotehost:1099</attribute>
  <attribute name="ProviderAdapterClass">
    org.jboss.jms.jndi.JBossMQProvider
  </attribute>
  <attribute name="QueueFactoryRef">XAConnectionFactory</attribute>
  <attribute name="TopicFactoryRef">XAConnectionFactory</attribute>
</mbean>
```

The `RemoteJMSProvider` can be referenced on the `mdb invoker config` as shown in the `jboss.xml` fragment given in Example 6.15.

Example 6.15. A `jboss.xml` fragment for specifying the MDB JMS provider adaptor

```
<proxy-factory-config>
  <JMSProviderAdapterJNDI>RemoteJMSProvider</JMSProviderAdapterJNDI>
  <ServerSessionPoolFactoryJNDI>StdJMSPool</ServerSessionPoolFactoryJNDI>
  <MaximumSize>15</MaximumSize>
  <MaxMessages>1</MaxMessages>
  <MDBConfig>
    <ReconnectIntervalSec>10</ReconnectIntervalSec>
    <DLQConfig>
      <DestinationQueue>queue/DLQ</DestinationQueue>
      <MaxTimesRedelivered>10</MaxTimesRedelivered>
      <TimeToLive>0</TimeToLive>
    </DLQConfig>
  </MDBConfig>
</proxy-factory-config>
```

Incidentally, because one can specify multiple `invoker-proxy-binding` elements, this allows an MDB to listen to the same queue/topic on multiple servers by configuring multiple bindings with different `JMSProviderAdapterJNDI` settings.

Alternatively, one can integrate the JMS provider using JCA configuration like that shown in Example 6.16.

Example 6.16. A `jms-ds.xml` descriptor for integrating a JMS provider adaptor via JCA

```
<tx-connection-factory>
  <jndi-name>RemoteJmsXA</jndi-name>
  <xa-transaction/>
  <adapter-display-name>JMS Adapter</adapter-display-name>
  <config-property name="JMSProviderAdapterJNDI"
```

```
        type="java.lang.String">RemoteJMSProvider</config-property>
    <config-property name="SessionDefaultType"
        type="java.lang.String">javax.jms.Topic</config-property>

    <security-domain-and-application>JmsXARealm</security-domain-and-application>
</tx-connection-factory>
```

6.4.2. org.jboss.jms.asf.ServerSessionPoolLoader MBean

The `ServerSessionPoolLoader` MBean service manages a factory for `javax.jms.ServerSessionPool` objects used by the message driven bean container. The configurable attributes of the `ServerSessionPoolLoader` service are:

- **PoolName:** A unique name for the session pool. This will be used to bind the `ServerSessionPoolFactory` instance into JNDI under `java:/PoolName`.
- **PoolFactoryClass:** The fully qualified class name of the `org.jboss.jms.asf.ServerSessionPoolFactory` interface to create an instance of.
- **XidFactory:** The JMX ObjectName of the service to use for generating `javax.transaction.xa.Xid` values for local transactions when two phase commit is not required. The `XidFactory` MBean must provide an `Instance` operation which returns a `org.jboss.tm.XidFactoryMBean` instance.

6.4.3. Integrating non-JBoss JMS Providers

We have mentioned that one can replace the JBossMQ JMS implementation with a foreign implementation. Here we summarize the various approaches one can take to do the replacement:

- Replace the `JMSProviderLoader` `JBossMQProvider` class with one that instantiates the correct JNDI context for communicating with the foreign JMS providers managed objects.
- Use the `ExternalContext` MBean to federate the foreign JMS providers managed objects into the JBoss JNDI tree.
- Use MBeans to instantiate the foreign JMS objects into the JBoss JNDI tree. An example of this approach can be found for Websphere MQ at http://sourceforge.net/tracker/index.php?func=detail&aid=753022&group_id=22866&atid=376687

Connectors on JBoss

The JCA Configuration and Architecture

This chapter discusses the JBoss server implementation of the J2EE Connector Architecture (JCA). JCA is a resource manager integration API whose goal is to standardize access to non-relational resources in the same way the JDBC API standardized access to relational data. The purpose of this chapter is to introduce the utility of the JCA APIs and then describe the architecture of JCA in JBoss

7.1. JCA Overview

J2EE 1.3 contains a connector architecture (JCA) specification that allows for the integration of transacted and secure resource adaptors into a J2EE application server environment. The full JCA specification is available from the JCA home page, <http://java.sun.com/j2ee/connector/>. The JCA specification describes the notion of such resource managers as Enterprise Information Systems (EIS). Examples of EIS systems include enterprise resource planning packages, mainframe transaction processing, non-Java legacy applications, etc.

The reason for focusing on EIS is primarily because the notions of transactions, security, and scalability are requirements in enterprise software systems. However, the JCA is applicable to any resource that needs to integrate into JBoss in a secure, scalable and transacted manner. In this introduction we will focus on resource adapters as a generic notion rather than something specific to the EIS environment.

The connector architecture defines a standard SPI (Service Provider Interface) for integrating the transaction, security and connection management facilities of an application server with those of a resource manager. The SPI defines the system level contract between the resource adaptor and the application server.

The connector architecture also defines a Common Client Interface (CCI) for accessing resources. The CCI is targeted at EIS development tools and other sophisticated users of integrated resources. The CCI provides a way to minimize the EIS specific code required by such tools. Typically J2EE developers will access a resource using such a tool, or a resource specific interface rather than using CCI directly. The reason is that the CCI is not a type specific API. To be used effectively it must be used in conjunction with metadata that describes how to map from the generic CCI API to the resource manager specific data types used internally by the resource manager.

The purpose of the connector architecture is to enable a resource vendor to provide a standard adaptor for its product. A resource adaptor is a system-level software driver that is used by a Java application to connect to resource. The resource adaptor plugs into an application server and provides connectivity between the resource manager, the application server, and the enterprise application. A resource vendor need only implement a JCA compliant adaptor once to allow use of the resource manager in any JCA capable application server.

An application server vendor extends its architecture once to support the connector architecture and is then assured of seamless connectivity to multiple resource managers. Likewise, a resource manager vendor provides one standard resource adaptor and it has the capability to plug in to any application server that supports the connector architecture.

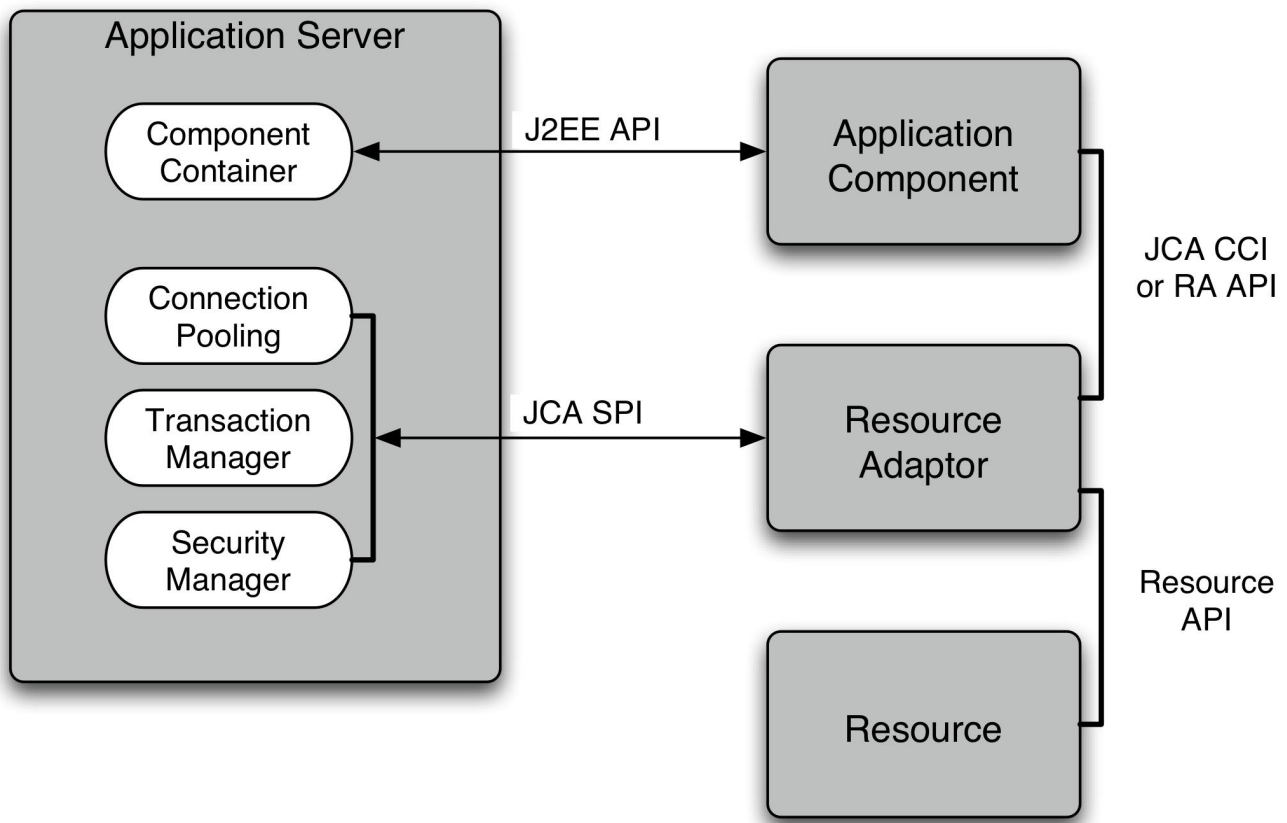


Figure 7.1. The relationship between a J2EE application server and a JCA resource adaptor

Figure 7.1 illustrates that the application server is extended to provide support for the JCA SPI to allow a resource adaptor to integrate with the server connection pooling, transaction management and security management facilities. This integration API defines a three part system contract.

- **Connection management:** a contract that allows the application server to pool resource connections. The purpose of the pool management is to allow for scalability. Resource connections are typically expensive objects to create and pooling them allows for more effective reuse and management.
- **Transaction Management:** a contract that allows the application server transaction manager to manage transactions that engage resource managers.
- **Security Management:** a contract that enables secured access to resource managers.

The resource adaptor implements the resource manager side of the system contract. This entails using the application server connection pooling, providing transaction resource information and using the security integration information. The resource adaptor also exposes the resource manager to the application server components. This can be done using the CCI and/or a resource adaptor specific API.

The application component integrates into the application server using a standard J2EE container to component contract. For an EJB component this contract is defined by the EJB specification. The application component interacts with the resource adaptor in the same way as it would with any other standard resource factory, for example, a `javax.sql.DataSource` JDBC resource factory. The only difference with a JCA resource adaptor is that the client has the option of using the resource adaptor independent CCI API if the resource adaptor supports this.

Figure 7.2 (from the JCA 1.0 specification) illustrates the relationship between the JCA architecture parti-

cipants in terms of how they relate to the JCA SPI, CCI and JTA packages.

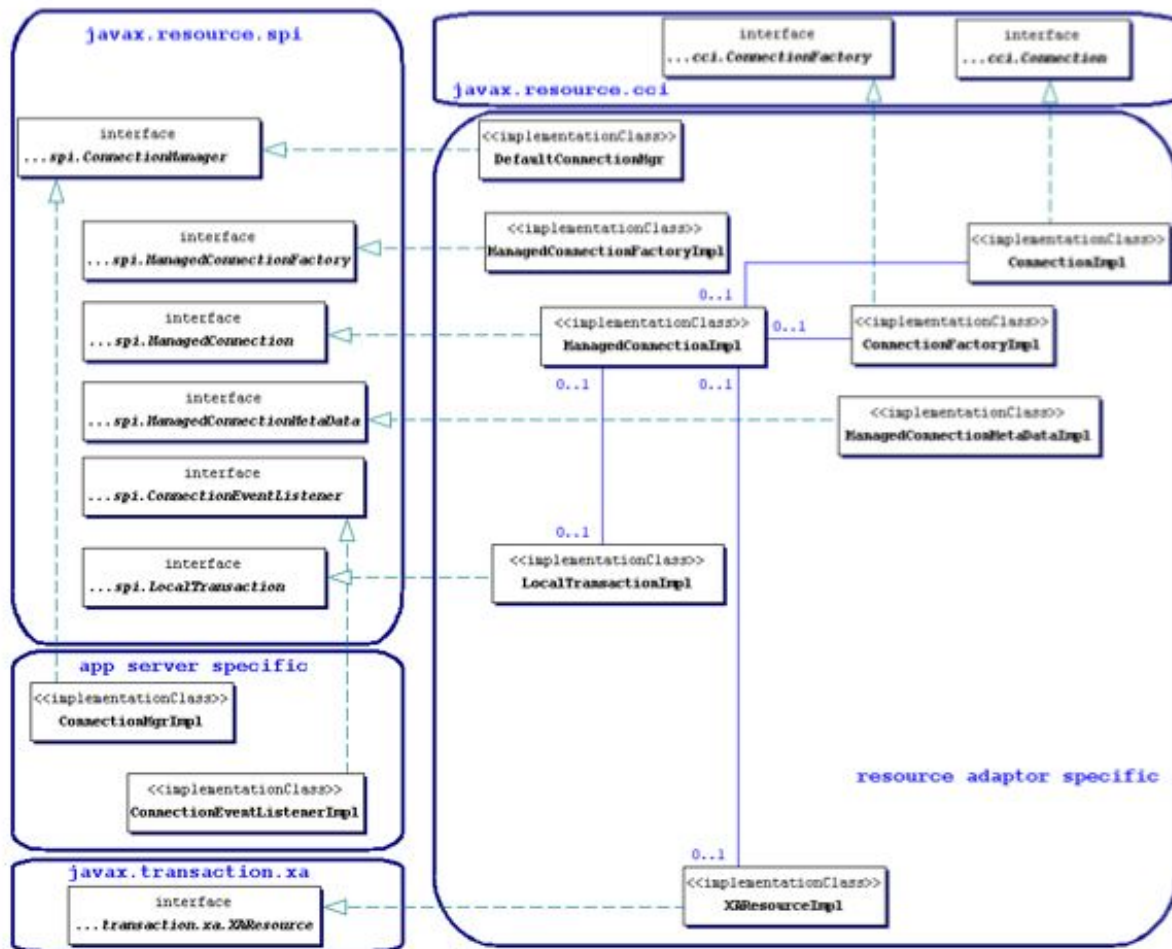


Figure 7.2. The JCA 1.0 specification class diagram for the connection management architecture.

The JBossCX architecture provides the implementation of the application server specific classes. Figure 7.2 shows that this comes down to the implementation of the `javax.resource.spi.ConnectionManager` and `javax.resource.spi.ConnectionEventListener` interfaces. The key aspects of this implementation are discussed in the following section on the JBossCX architecture.

7.2. An Overview of the JBossCX Architecture

The JBossCX framework provides the application server architecture extension required for the use of JCA resource adaptors. This is primarily a connection pooling and management extension along with a number of MBeans for loading resource adaptors into the JBoss server. Figure 7.3 expands the generic view given by Figure 7.2 to illustrate how the JBoss JCA layer implements the application server specific extension along with an example file system resource adaptor that we will look at latter in this chapter.

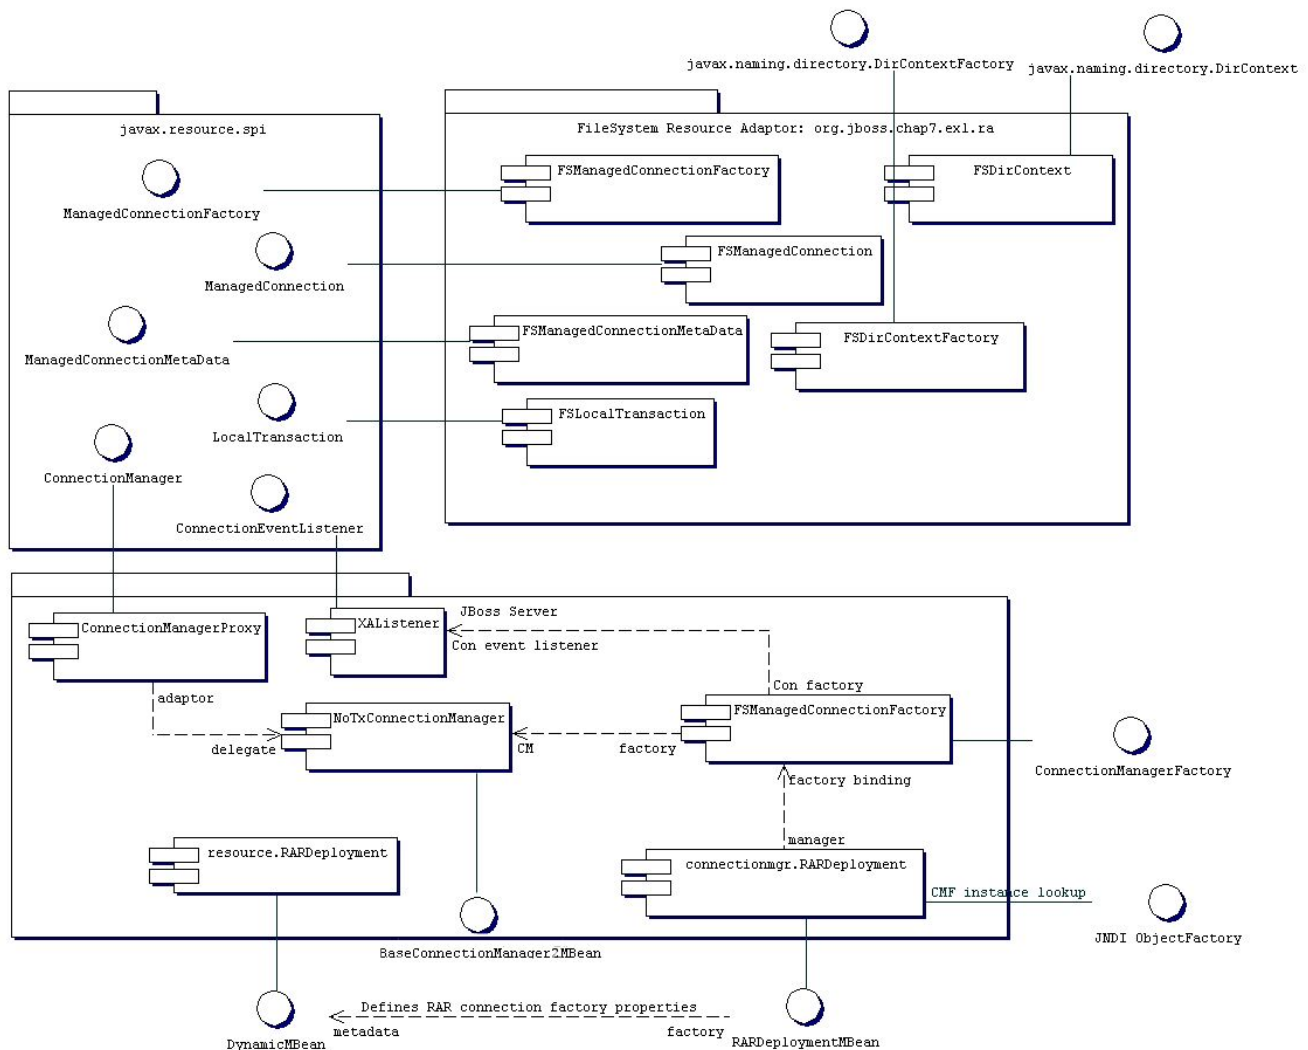


Figure 7.3. The JBoss JCA implementation components

There are three coupled MBeans that make up a RAR deployment. These are the `org.jboss.resource.RARDeployment`, `org.jboss.resource.connectionmanager.RARDeployment`, and `org.jboss.resource.connectionmanager.BaseConnectionManager2`. The `org.jboss.resource.RARDeployment` is simply an encapsulation of the metadata of a RAR `META-INF/ra.xml` descriptor. It exposes this information as a `DynamicMBean` simply to make it available to the `org.jboss.resource.connectionmanager.RARDeployment` MBean.

The `RARDeployer` service handles the deployment of archives files containing resource adaptors (RARs). It creates the `org.jboss.resource.RARDeployment` MBeans when a RAR file is deployed. Deploying the RAR file is the first step in making the resource adaptor available to application components. For each deployed RAR, one or more connection factories must be configured and bound into JNDI. This task performed using a JBoss service descriptor that sets up a `org.jboss.resource.connectionmanager.BaseConnectionManager2` MBean implementation with a `org.jboss.resource.connectionmgr.RARDeployment` dependent.

7.2.1. BaseConnectionManager2 MBean

The `org.jboss.resource.connectionmanager.BaseConnectionManager2` MBean is a base class for the various types of connection managers required by the JCA spec. Subclasses include (`org.jboss.resource.connectionmanager`) `NoTxConnectionManager`, `LocalTxConnectionManager` and `XATx-`

ConnectionManager. These correspond to resource adaptors that support no transactions, local transaction and XA transaction respectively. You choose which subclass to use based on the type of transaction semantics you want, provided the JCA resource adaptor supports the corresponding transaction capability.

The common attributes supported by the BaseConnectionManager2 MBean are:

- **ManagedConnectionFactoryName:** This specifies the `ObjectName` of the MBean that creates `javax.resource.spi.ManagedConnectionFactory` instances. Normally this is configured as an embedded mbean in a `depends` element rather than a separate MBean reference using the `RARDeployment` MBean. The MBean must provide an operation with the following signature:
`javax.resource.spi.ManagedConnectionFactory startManagedConnectionFactory(javax.resource.spi.ConnectionManager)`
- **ManagedConnectionPool:** This specifies the `ObjectName` of the MBean representing the pool for this connection manager. The MBean must have an `ManagedConnectionPool` attribute that is an implementation of the `org.jboss.resource.connectionmanager.ManagedConnectionPool` interface. Normally it will be an embedded MBean in a `depends` tag rather than an `ObjectName` reference to an existing MBean. The default MBean for use is the `org.jboss.resource.connectionmanager.JBossManagedConnectionPool`. Its configurable attributes are discussed below.
- **CachedConnectionManager:** This specifies the `ObjectName` of the `(org.jboss.resource.connectionmanager) CachedConnectionManager` MBean implementation used by the connection manager. Normally this will be specified using a `depends` tag with the `ObjectName` of the unique `CachedConnectionManager` for the server. The name `jboss.jca:service=CachedConnectionManager` is the standard setting to use.
- **SecurityDomainJndiName:** This specifies the JNDI name of the security domain to use for authentication and authorization of resource connections. This is typically of the form `java:/jaas/<domain>` where the `<domain>` value is the name of an entry in the `conf/login-config.xml` JAAS login module configuration file. This defines which JAAS login modules execute to perform authentication. Chapter 8 has more information on the security settings.
- **JaasSecurityManagerService:** This is the `ObjectName` of the security manager service. This should be set to the security manager MBean name as defined in the `conf/jboss-service.xml` descriptor, and currently this is `jboss.security:service=JaasSecurityManager`. This attribute will likely be removed in the future.

7.2.2. RARDeployment MBean

The `org.jboss.resource.connectionmanager.RARDeployment` MBean manages configuration and instantiation `ManagedConnectionFactory` instance. It does this using the resource adaptor metadata settings from the `RAR META-INF/ra.xml` descriptor along with the `RARDeployment` attributes. The configurable attributes are:

- **OldRarDeployment:** This is the `ObjectName` of the `org.jboss.resource.RarDeployment` MBean that contains the resource adaptor metadata. The form of this name is `jboss.jca:service=RARDeployment,name=<ra-display-name>` where the `<ra-display-name>` is the `ra.xml` descriptor `display-name` attribute value. This is created by the `RARDeployer` when it deploys a RAR file. This attribute will likely be removed in the future.
- **ManagedConenctionFactoryProperties:** This is a collection of (name, type, value) tripples that define attributes of the `ManagedConnectionFactory` instance. Therefore, the names of the attributes depend on the resource adaptor `ManagedConnectionFactory` instance. The following example shows the structure of the content of this attribute.

```
<properties>
  <config-property>
    <config-property-name>Attr0Name</config-property-name>
    <config-property-type>Attr0Type</config-property-type>
    <config-property-value>Attr0Value</config-property-value>
  </config-property>
  <config-property>
    <config-property-name>Attr1Name</config-property-name>
    <config-property-type>Attr2Type</config-property-type>
    <config-property-value>Attr2Value</config-property-value>
  </config-property>
  ...
</properties>
```

AttrXName is the Xth attribute name, AttrXType is the fully qualified Java type of the attribute, and AttrXValue is the string representation of the value. The conversion from string to AttrXType is done using the `java.beans.PropertyEditor` class for the AttrXType.

- **JndiName:** This is the JNDI name under which the will be made available. Clients of the resource adaptor use this name to obtain either the `javax.resource.cci.ConnectionFactory` or resource adaptor specific connection factory. The full JNDI name will be `java:/<JndiName>` meaning that the `JndiName` attribute value will be prefixed with `java:/`. This prevents use of the connection factory outside of the JBoss server VM. In the future this restriction may be configurable.

7.2.3. JBossManagedConnectionPool MBean

The `org.jboss.resource.connectionmanager.JBossManagedConnectionPool` MBean is a connection pooling MBean. It is typically used as the embedded MBean value of the `BaseConnectionManger2` `ManagedConnectionPool` attribute. When you setup a connection manager MBean you typically embed the pool configuration in the connection manager descriptor. The configurable attributes of the `JBossManagedConnectionPool` are:

- **MinSize:** This attribute indicates the minimum number of connections this pool should hold. These are not created until a `Subject` is known from a request for a connection. `MinSize` connections will be created for each sub-pool.
- **MaxSize:** This attribute indicates the maximum number of connections for a pool. No more than `MaxSize` connections will be created in each sub-pool.
- **BlockingTimeoutMillis:** This attribute indicates the maximum time to blockwhile waiting for a connection before throwing an exception. Note that this blocks only while waiting for a permit for a connection, and will never throw an exception if creating a new connection takes an inordinately long time.
- **IdleTiemoutMinutes:** This attribute indicates the maximum time a connection may be idle before being closed. The actual maximum time depends also on the idle remover thread scan time, which is 1/2 the smallest idle timeout of any pool.
- **NoTxSeperatePools:** Setting this to true doubles the available pools. One pool is for connections used outside a transaction the other inside a transaction. The actual pools are lazily constructed on first use. This is only relevent when setting the pool parameters associated with the `LocalTxConnectionManager` and `XATxConnectionManager`. It's use case is for Oracle (and possibly other vendors) XA implementations that don't like using an XA connection with and without a JTA transaction.
- **Criteria:** This attribute indicates if the JAAS `javax.security.auth.Subject` from security domain associated with the connection, or app supplied parameters (such as from `getConnection(user, pw)`) are used to

distinguish connections in the pool. The allowed values are:

- **ByContainer:** use `Subject`
- **ByApplication:** use app supplied params only
- **ByContainerAndApplication:** use both
- **ByNothing:** all connections are equivalent, usually if adapter supports reauthentication

7.2.4. CachedConnectionManager MBean

The `org.jboss.resource.connectionmanager.CachedConnectionManager` MBean manages associations between meta-aware objects (those accessed through interceptor chains) and connection handles, as well as between user transactions and connection handles. Normally there should only be one such MBean, and this is configured in the core `jboss-service.xml` descriptor. It is used by `org.jboss.resource.connectionmanager.CachedConnectionInterceptor`, JTA `javax.transaction.UserTransaction` implementation, and all `BaseConnectionManager2` instances. The configurable attributes of the `CachedConnectionManager` MBean are:

- **SpecCompliant:** Enable this boolean attribute for spec compliant non-shareable connections reconnect processing. This allows a connection to be opened in one call and used in another. Note that specifying this behavior disables connection close processing.
- **Debug:** Enable this boolean property for connection close processing. At the completion of an EJB method invocation, unclosed connections are registered with a transaction synchronization. If the transaction ends without the connection being closed, an error is reported and JBoss closes the connection. This is a development feature that should be turned off in production for optimal performance.
- **TransactionManagerServiceName:** This attribute specifies the JMX `ObjectName` of the JTA transaction manager service. Connection close processing is now synchronized with the transaction manager and this attribute specifies the transaction manager to use.

7.2.5. A Sample Skeleton JCA Resource Adaptor

To conclude our discussion of the JBoss JCA framework we will create and deploy a single non-transacted resource adaptor that simply provides a skeleton implementation that stubs out the required interfaces and logs all method calls. We will not discuss the details of the requirements of a resource adaptor provider as these are discussed in detail in the JCA specification. The purpose of the adaptor is to demonstrate the steps required to create and deploy a RAR in JBoss, and to see how JBoss interacts with the adaptor.

The adaptor we will create could be used as the starting point for a non-transacted file system adaptor. The source to the example adaptor can be found in the `src/main/org/jboss/chap7/ex1` directory of the book examples. A class diagram that shows the mapping from the required `javax.resource.spi` interfaces to the resource adaptor implementation is given in Figure 7.4.

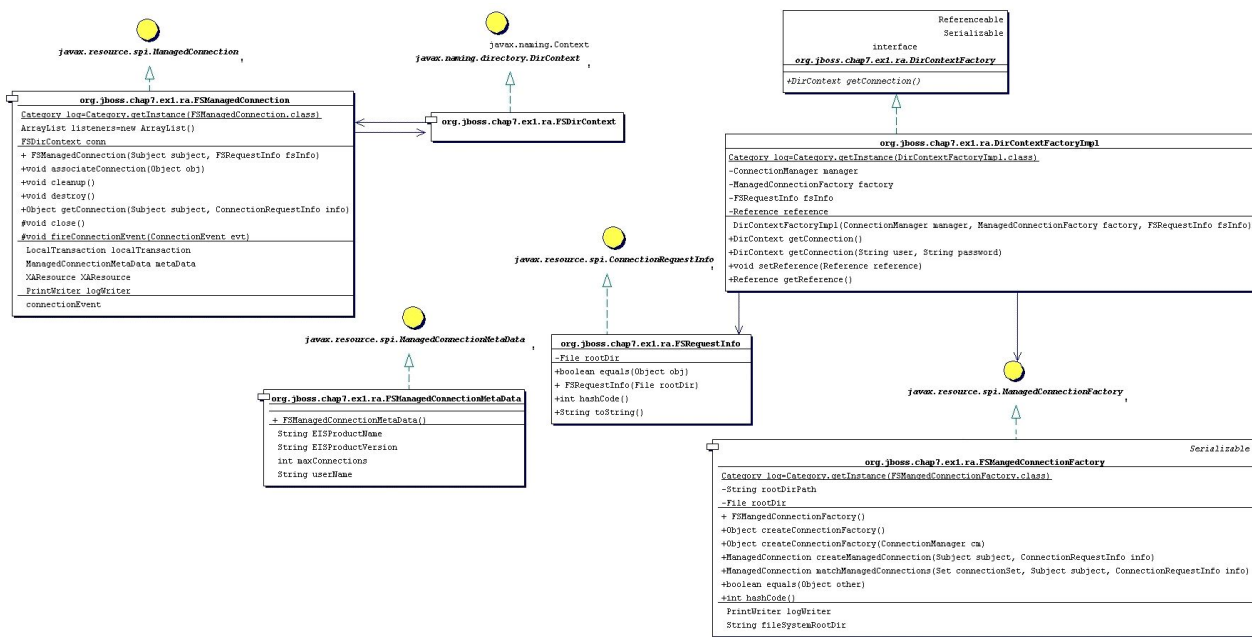


Figure 7.4. The file system RAR class diagram

We will build the adaptor, deploy it to the JBoss server and then run an example client against an EJB that uses the resource adaptor to demonstrate the basic steps in a complete context. We'll then take a look at the JBoss server log to see how the JBoss JCA framework interacts with the resource adaptor to help you better understand the components in the JCA system level contract.

To build the example and deploy the RAR to the JBoss server `deploy/lib` directory, execute the following Ant command in the book examples directory.

```
[nr@tokil]$ ant -Dchap=chap7 build-chap
Buildfile: build.xml

...
prepare:
  [mkdir] Created dir: /Users/orb/proj/jboss/education/books/admin-devel/examples/output/chap7

chap7-ex1-rar:
  [jar] Building jar: /Users/orb/proj/jboss/education/books/admin-devel/examples/output/chap7/ra.
  [jar] Building jar: /Users/orb/proj/jboss/education/books/admin-devel/examples/output/chap7/chap

prepare:

chap7-ex1-jar:
  [jar] Building jar: /Users/orb/proj/jboss/education/books/admin-devel/examples/output/chap7/chap

BUILD SUCCESSFUL
```

The deployed files include a `chap7-ex1.sar` and a `notxfs-service.xml` service descriptor. The example resource adaptor deployment descriptor is shown in Example 7.1 while the connection manager MBeans service descriptor is shown in Example 7.2.

Example 7.1. The nontransactional file system resource adaptor deployment descriptor.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE connector PUBLIC
  "-//Sun Microsystems, Inc.//DTD Connector 1.0//EN"
  "http://java.sun.com/dtd/connector_1_0.dtd">
```



```

<connector>
  <display-name>File System Adapter</display-name>
  <vendor-name>JBoss Group</vendor-name>
  <spec-version>1.0</spec-version>
  <version>1.0</version>
  <eis-type>FileSystem</eis-type>
  <license>
    <description>LGPL</description>
    <license-required>false</license-required>
  </license>
  <resourceadapter>
    <managedconnectionfactory-class>
      org.jboss.chap7.ex1.ra.FSMangedConnectionFactory
    </managedconnectionfactory-class>
    <connectionfactory-interface>
      org.jboss.chap7.ex1.ra.DirContextFactory
    </connectionfactory-interface>

    <connectionfactory-impl-class>
      org.jboss.chap7.ex1.ra.DirContextFactoryImpl
    </connectionfactory-impl-class>
    <connection-interface>javax.naming.directory.DirContext</connection-interface>
    <connection-impl-class>org.jboss.chap7.ex1.ra.FSDirContext</connection-impl-class>
    <transaction-support>NoTransaction</transaction-support>
    <config-property>
      <config-property-name>FileSystemRootDir</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value>/tmp/db/fs_store</config-property-value>
    </config-property>
    <config-property>
      <config-property-name>UserName</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value/>
    </config-property>
    <config-property>
      <config-property-name>Password</config-property-name>
      <config-property-type>java.lang.String</config-property-type>
      <config-property-value/>
    </config-property>
    <authentication-mechanism>
      <authentication-mechanism-type>BasicPassword</authentication-mechanism-type>
      <credential-interface>
        javax.resource.security.PasswordCredential
      </credential-interface>
    </authentication-mechanism>
    <reauthentication-support>true</reauthentication-support>
    <security-permission>
      <description>
        Read/Write access is required to the contents of the FileSystemRootDir
      </description>
      <security-permission-spec>permission java.io.FilePermission
        "/tmp/db/fs_store/*", "read,write"</security-permission-spec>
    </security-permission>
  </resourceadapter>
</connector>

```

Example 7.2. The notxfds.xml resource adaptor MBeans service descriptor.

```

<!-- The non-transaction FileSystem resource adaptor service configuration -->
<connection-factories>
  <no-tx-connection-factory>
    <jndi-name>NoTransFS</jndi-name>
    <adapter-display-name>File System Adapter</adapter-display-name>
    <config-property name="FileSystemRootDir"
      type="java.lang.String">/tmp/db/fs_store</config-property>
  </no-tx-connection-factory>
</connection-factories>

```

The key items in the resource adaptor deployment descriptor are highlighted in bold. These define the classes of the resource adaptor, and the elements are:

- **display-name**: Recall from our discussion of the connection manager factory MBeans that the association between the factory and the resource adaptor classes was done via a `RARDeploymentDynamicMBean` located by name. The name was based on the `display-name` value found in the `ra.xml` descriptor. Here the name is `File System Adaptor`. We will use it in the connection manager service descriptor.
- **managedconnectionfactory-class**: The implementation of the `ManagedConnectionFactory` interface, `org.jboss.chap7.ex1.ra.FSMangedConnectionFactory`
- **connectionfactory-interface**: The interface that clients will obtain when they lookup the connection factory instance from JNDI, here a proprietary resource adaptor value, `org.jboss.chap7.ex1.ra.DirContextFactory`
- **connectionfactory-impl-class**: The class that provides the implementation of the `connectionfactory-interface`, `org.jboss.chap7.ex1.ra.DirContextFactoryImpl`
- **connection-interface**: The interface for the connections returned by the resource adaptor connection factory, here the JNDI `javax.naming.directory.DirContext` interface.
- **connection-impl-class**: The class that provides the `connection-interface` implementation, `org.jboss.chap7.ex1.ra.FSDirContext`
- **transaction-support**: The level of transaction support, here defined as `NoTransaction`, meaning the file system resource adaptor does not do transactional work.

You can see the JCA 1.0 spec, or the book *J2EE Connector Architecture and Enterprise Application Integration* by Sharma, Stearns and Ng for the full details of the `ra.xml` descriptor elements.

The RAR classes and deployment descriptor only define a resource adaptor. To use the resource adaptor it must be integrated into the JBoss application server. As we have discussed this is done with a connection factory MBeans. A simplified descriptor format is available for configuring the JCA services of the application server, and this is described below in Section 7.3.2. The `notxfs-ds.xml` descriptor shown in Example 7.2, and the following notes apply.

- The `jndi-name` element is used to specify where the connection factory will be bound into JNDI. For this deployment that binding will be `java:/NoTransFS`.
- The `adapter-display-name` element specifies the same value as the `ra.xml` `display-name` element. This is how the JCA layer knows how to associate which RAR which this connection factory configuration.
- `ManagedConnectionFactoryProperties` may be specified using `config-property` elements to provide non-default settings to the resource adaptor connection factory. Here the `FileSystemRootDir` of type `java.lang.String` attribute is being set to `/tmp/db/fs_store`.

To deploy the RAR and connection manager configuration to the JBoss server, run the following:

```
[nr@toki examples]$ ant -Dchap=chap7 config
Buildfile: build.xml
...
config:
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
```

```
BUILD SUCCESSFUL
```

The server console will display some logging output indicating that the resource adaptor has been deployed.

Now we want to test access of the resource adaptor by a J2EE component. To do this we have created a trivial stateless session bean that has a single method called `echo`. Inside of the `echo` method the EJB accesses the resource adaptor connection factory, creates a connection, and then immediately closes the connection. The `echo` method code is shown in Example 7.3.

Example 7.3. The stateless session bean `echo` method code which shows the access of the resource adaptor connection factory.

```
public String echo(String arg)
{
    log.debug("echo, arg="+arg);
    try {
        InitialContext iniCtx = new InitialContext();
        Context enc = (Context) iniCtx.lookup("java:comp/env");
        Object ref = enc.lookup("ra/DirContextFactory");
        log.debug("echo, ra/DirContextFactory="+ref);

        DirContextFactory dcf = (DirContextFactory) ref;
        log.debug("echo, found dcf="+dcf);
        DirContext dc = dcf.getConnection();
        log.debug("echo, lookup dc="+dc);

        dc.close();
    } catch(NamingException e) {
        log.error("Failed during JNDI access", e);
    }
    return arg;
}
```

The EJB is not using the CCI interface to access the resource adaptor. Rather, it is using the resource adaptor specific API based on the proprietary `DirContextFactory` interface that returns a JNDI `DirContext` object as the connection object. The example EJB is simply exercising the system contract layer by looking up the resource adaptor connection factory, creating a connection to the resource and closing the connection. The EJB does not actually do anything with the connection, as this would only exercise the resource adaptor implementation since this is a non-transactional resource.

Run the test client which calls the `EchoBean.echo` method by running Ant as follows from the examples directory:

```
[nr@toki examples]$ ant -Dchap=chap7 -Dex=1 run-example
Buildfile: build.xml
...
run-example1:
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
  [echo] Waiting for deploy...
  [java] Created Echo
  [java] Echo.echo('Hello') = Hello

BUILD SUCCESSFUL
```

Now let's look at the output that has been logged by the resource adaptor to understand the interaction between the adaptor and the JBoss JCA layer. The output is in the `server/default/log/server.log` file of the JBoss server distribution. We'll summarize the events seen in the log using a sequence diagram.

Those are the steps involved with making the resource adaptor connection factory available to application server components. The remaining log messages are the result of the example client invoking the `EchoBean.echo` method and this method's interaction with the resource adaptor connection factory. Figure 7.5 is a sequence diagram that summarizes the events that occur when the `EchoBean` accesses the resource adaptor connection factory from JNDI and creates a connection.

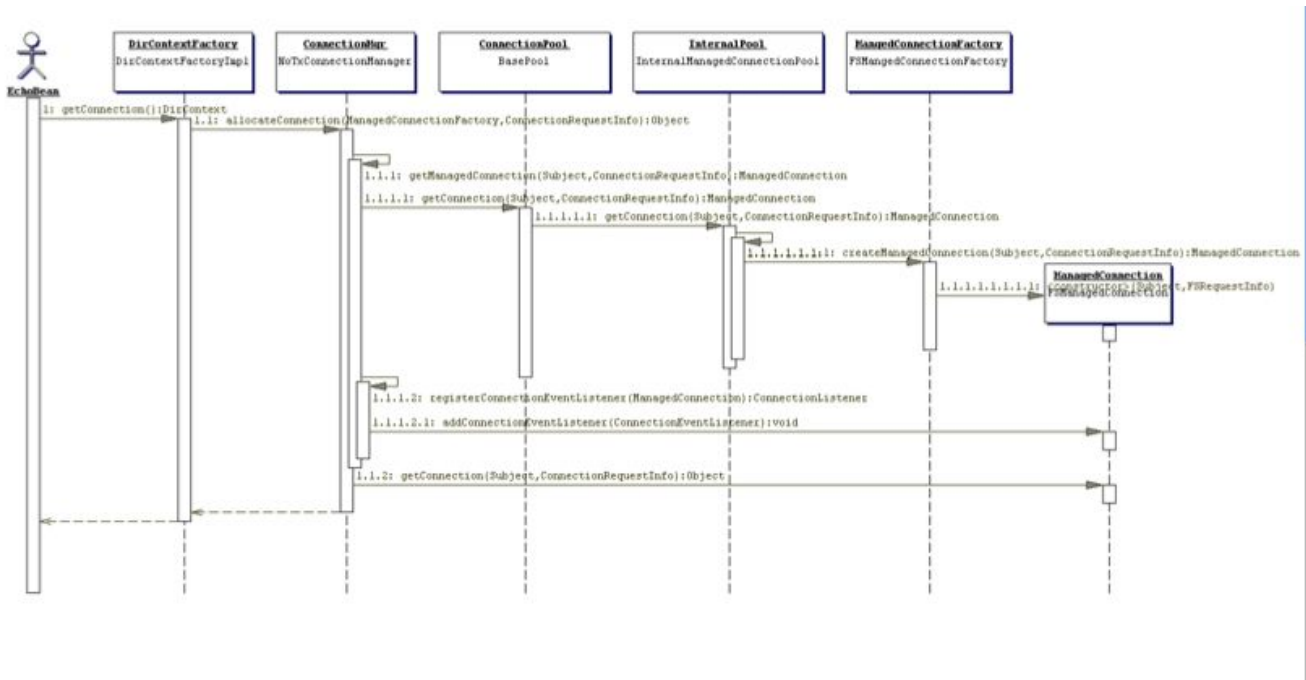


Figure 7.5. A sequence diagram illustrating the key interactions between the JBossCX framework and the example resource adaptor that result when the `EchoBean` accesses the resource adaptor connection factory.

The starting point is the client's invocation of the `EchoBean.echo` method. For the sake of conciseness of the diagram, the client is shown directly invoking the `EchoBean.echo` method when in reality the JBoss EJB container handles the invocation. There are three distinct interactions between the `EchoBean` and the resource adaptor; the lookup of the connection factory, the creation of a connection, and the close of the connection.

The lookup of the resource adaptor connection factory is illustrated by the 1.1 sequences of events. The events are:

- 1, the echo method invokes the `getConnection` method on the resource adaptor connection factory obtained from the JNDI lookup on the `java:comp/env/ra/DirContextFactory` name which is a link to the `java:/NoTransFS` location.
- 1.1, the `DirContextFactoryImpl` class asks its associated `ConnectionManager` to allocate a connection. It passes in the `ManagedConnectionFactory` and `FSRequestInfo` that were associated with the `DirContextFactoryImpl` during its construction.
- 1.1.1, the `ConnectionManager` invokes its `getManagedConnection` method with the current `Subject` and `FSRequestInfo`.
- 1.1.1.1, the `ConnectionManager` asks its object pool for a connection object. The `JBossManagedConnectionPool$BasePool` is get the key for the connection and then asks the matching `InternalPool` for a connection.

- 1.1.1.1.1, Since no connections have been created the pool must create a new connection. This is done by requesting a new managed connection from the `ManagedConnectionFactory`. The Subject associated with the pool as well as the `FSRequestInfo` data are passed as arguments to the `createManagedConnection` method invocation.
- 1.1.1.1.1.1, the `FSManagedConnectionFactory` creates a new `FSManagedConnection` instance and passes in the Subject and `FSRequestInfo` data.
- 1.1.1.2, a `javax.resource.spi.ConnectionListener` instance is created. The type of listener created is based on the type of `ConnectionManager`. In this case it is an `org.jboss.resource.connectionmgr.BaseConnectionManager2$NoTransactionListener` instance.
- 1.1.1.2.1, the listener registers as a `javax.resource.spi.ConnectionEventListener` with the `ManagedConnection` instance created in 1.2.1.1.
- 1.1.2, the `ManagedConnection` is asked for the underlying resource manager connection. The Subject and `FSRequestInfo` data are passed as arguments to the `getConnection` method invocation.
- The resulting connection object is cast to a `javax.naming.directory.DirContext` instance since this is the public interface defined by the resource adaptor.
- After the `EchoBean` has obtained the `DirContext` for the resource adaptor, it simply closes the connection to indicate its interaction with the resource manager is complete.

This concludes the resource adaptor example. Our investigation into the interaction between the JBossCX layer and a trivial resource adaptor should give you sufficient understanding of the steps required to configure any resource adaptor. The example adaptor can also serve as a starting point for the creation of your own custom resource adaptors if you need to integrate non-JDBC resources into the JBoss server environment.

7.3. Configuring JCA Adaptors

Configuration of the JCA resource adaptors may be done by configuring the JBoss JCA services along with the JCA resource adaptor as shown in the previous section. JBoss 3.2 provides an alternate simplified schema that avoids having to specify so much redundant configuration information.

7.3.1. Configuring JDBC DataSources

The syntax for configuring JCA JDBC connection factories has been simplified in 3.2. Rather than configuring the connection manager factory related MBeans discussed in the previous section via a mbean services deployment descriptor, an abbreviated datasource centric descriptor is used. This is transformed into the standard `jboss-service.xml` MBean services deployment descriptor using a XSL transform applied by the `org.jboss.deployment.XSLSubDeployer` included in the `jboss-jca.sar` deployment. The simplified configuration descriptor is deployed the same as other deployable components. The descriptor must be named using a `*-ds.xml` pattern in order to be recognized by the `XSLSubDeployer`.

The schema for the top-level datasource elements of the `*-ds.xml` configuration deployment file is shown in Figure 7.6.




Figure 7.6. The simplified JCA DataSource configuration descriptor top-level schema elements

Multiple datasource configurations may be specified in a configuration deployment file. The child elements of the `datasources` root are:

- **mbean:** Any number `mbean` elements may be specified to define MBean services that should be included in the `jboss-service.xml` descriptor that results from the transformation. This may be used to configure services used by the datasources.
- **no-tx-datasource:** this element is used to specify the `(org.jboss.resource.connectionmanager) NoTxConnectionManager` service configuration. `NoTxConnectionManager` is a JCA connection manager with no transaction support. The `no-tx-datasource` child element schema is given in Figure 7.7.
- **local-tx-datasource:** this element is used to specify the `(org.jboss.resource.connectionmanager) LocalTxConnectionManager` service configuration. `LocalTxConnectionManager` implements a `ConnectionEventListener` that implements `XAResource` to manage transactions through the transaction manager. To ensure that all work in a local transaction occurs over the same `ManagedConnection`, it includes a `xid` to `ManagedConnection` map. When a `Connection` is requested or a transaction started with a connection handle in use, it checks to see if a `ManagedConnection` already exists enrolled in the global transaction and uses it if found. Otherwise, a free `ManagedConnection` has its `LocalTransaction` started and is used. The `local-tx-datasource` child element schema is given in Figure 7.8
- **xa-datasource:** this element is used to specify the `(org.jboss.resource.connectionmanager) XATxConnectionManager` service configuration. `XATxConnectionManager` implements a `ConnectionEventListener` that obtains the `XAResource` to manage transactions through the transaction manager from the adaptor `ManagedConnection`. To ensure that all work in a local transaction occurs over the same `ManagedConnection`, it includes a `xid` to `ManagedConnection` map. When a `Connection` is requested or a transaction started with a connection handle in use, it checks to see if a `ManagedConnection` already exists enrolled in the global transaction and uses it if found. Otherwise, a free `ManagedConnection` has its `LocalTransaction` started and is used. The `xa-datasource` child element schema is given in Figure 7.9.

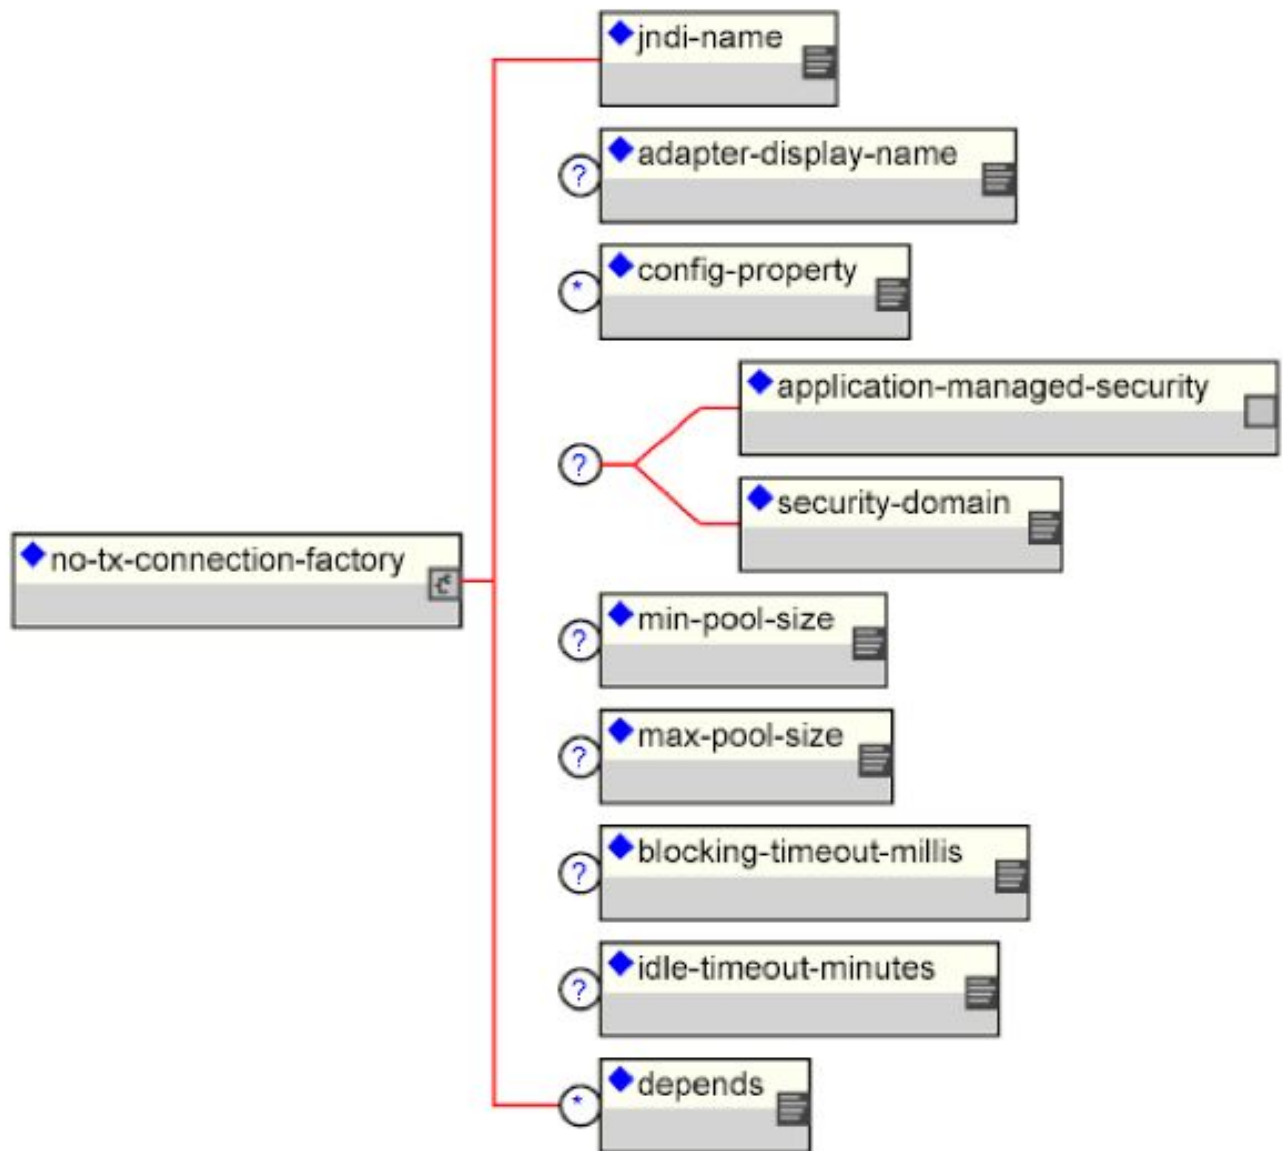


Figure 7.7. The non-transactional DataSource configuration schema

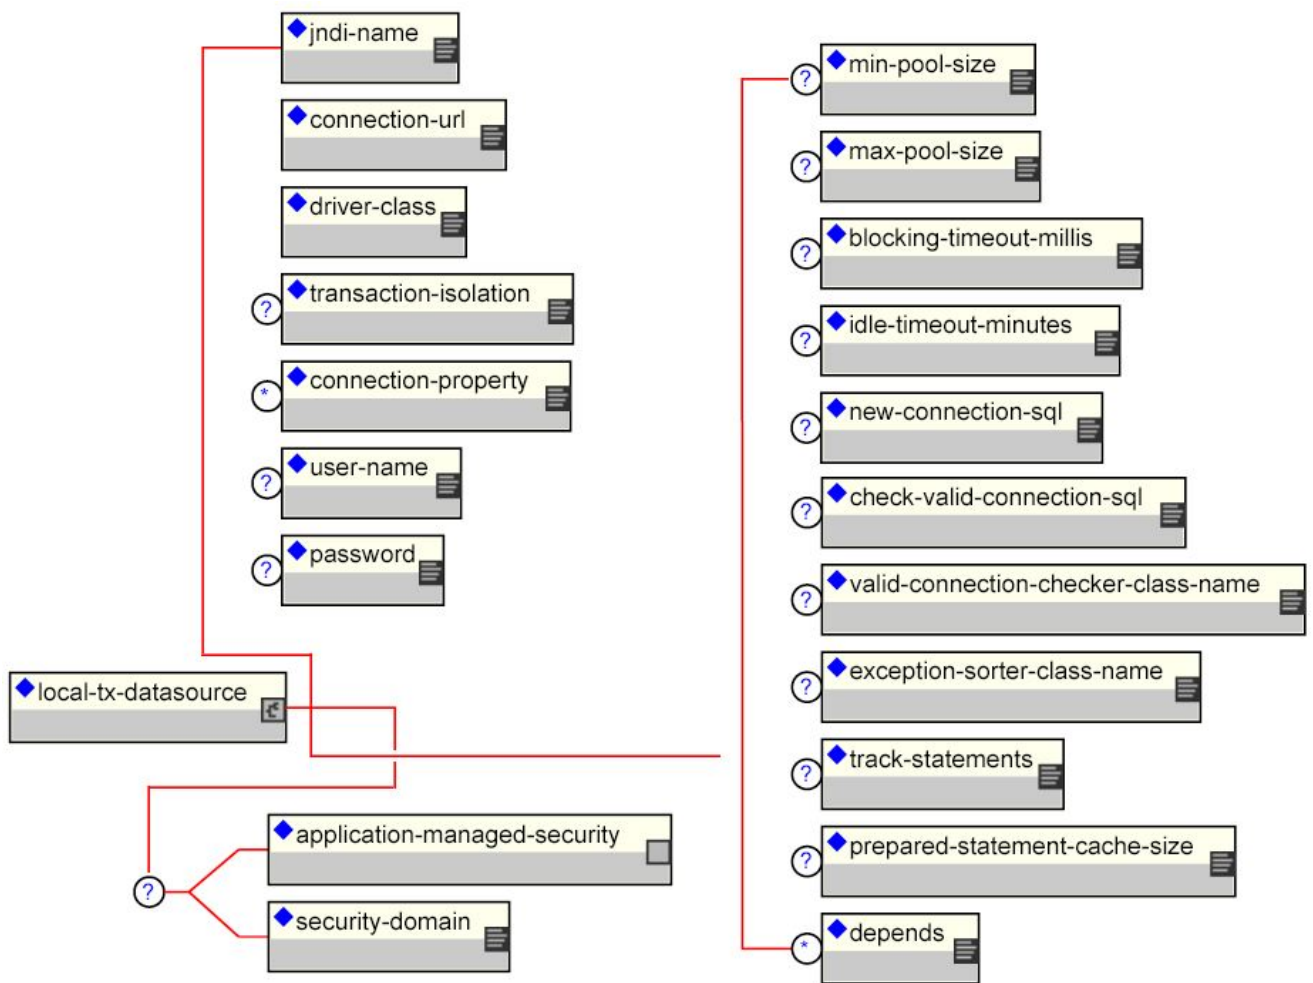


Figure 7.8. The non-XA DataSource configuration schema

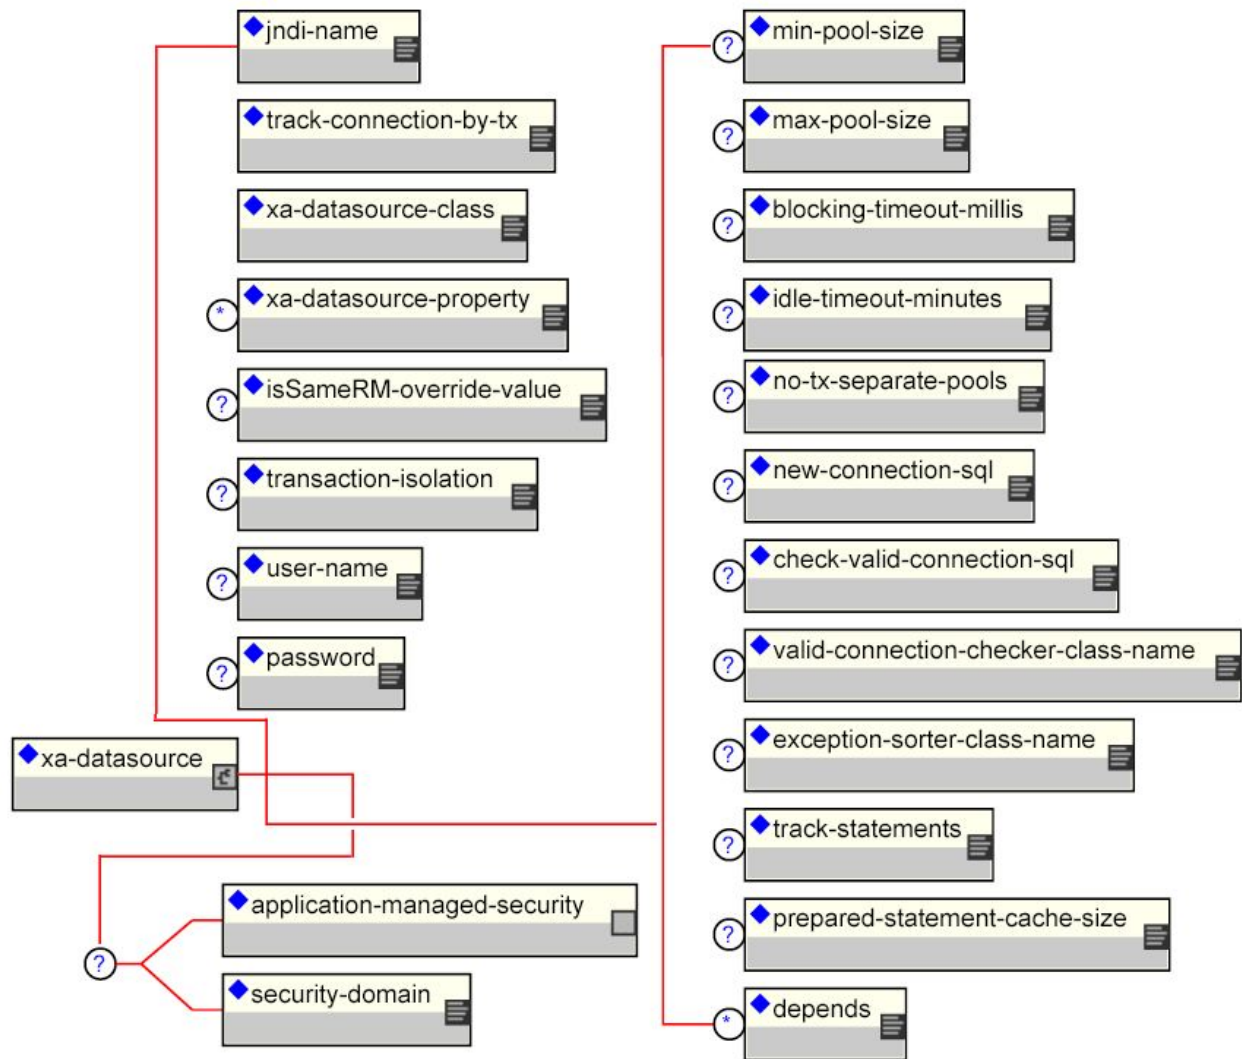


Figure 7.9. The XADatasource configuration schema

Elements that are common to all datasources include:

- **jndi-name:** The JNDI name under which the `DataSource` wrapper will be bound. Note that this name is relative to the `java:/` prefix. The full JNDI name of the `DataSource` will be `java:/ + jndi-name`. `DataSource` wrappers are bound under the `java:/` prefix since they are not usable outside of the server VM.
- **application-managed-security:** Specifying this element indicates that application code supplied parameters, such as from `getConnection(user, pw)`, are used to distinguish connections in the pool.
- **security-domain:** Specifying this element indicates that either application code supplied parameters, or JAAS Subject based information is to distinguish connections in the pool. The content of the `security-domain` is the name of the JAAS security manager that will handle authentication. This name correlates to the JAAS `login-config.xml` descriptor `application-policy/name` attribute.
- **min-pool-size:** This element specifies the minimum number of connections a pool should hold. These pool instances are not created until an initial request for a connection is made. This default to 0.
- **max-pool-size:** This element specifies the maximum number of connections for a pool. No more than the `max-pool-size` number of connections will be created in a pool. This defaults to 20.

- **blocking-timeout-millis:** This element specifies the maximum time in milliseconds to block while waiting for a connection before throwing an exception. Note that this blocks only while waiting for a permit for a connection, and will never throw an exception if creating a new connection takes an inordinately long time. The default is 5000.
- **idle-timeout-minutes:** This element specifies the maximum time in minutes a connection may be idle before being closed. The actual maximum time depends also on the `IdleRemover` scan time, which is 1/2 the smallest idle-timeout-minutes of any pool.
- **depends:** The depends element specifies the JMX `ObjectName` string of a service that the connection manager services depend on. The connection manager service will not be started until the dependent services have been started.

Additional common child elements for both no-tx-datasource and local-tx-datasource include:

- **connection-url:** The JDBC driver connection URL string, for example, `jdbc:hsqldb:hsql://localhost:1701`.
- **driver-class:** The fully qualified name of the JDBC driver class, for example, `org.hsqldb.jdbcDriver`.
- **connection-property:** The `connection-property` element allows you to pass in arbitrary connection properties to the `java.sql.Driver.connect(url, props)` method. Each `connection-property` specifies a string name/value pair with the property name coming from the name attribute and the value coming from the element content.
- **user-name:** This element specifies the default username used when creating a new connection. The actual username may be overridden by the application code `getConnection` parameters or the connection creation context JAAS Subject.
- **password:** This element specifies the default password used when creating a new connection. The actual password may be overridden by the application code `getConnection` parameters or the connection creation context JAAS Subject.

Elements in common to the local-tx-datasource and xa-datasource are:

- **transaction-isolation:** This element specifies the `java.sql.Connection` transaction isolation level to use. The constants defined in the `Connection` interface are the possible element content values and include:
 - `TRANSACTION_READ_UNCOMMITTED`
 - `TRANSACTION_READ_COMMITTED`
 - `TRANSACTION_REPEATABLE_READ`
 - `TRANSACTION_SERIALIZABLE`
 - `TRANSACTION_NONE`
- **new-connection-sql:** A SQL statement that should be executed when a new connection is created. This can be used to configure a connection with database specific settings not configurable via connection properties.
- **check-valid-connection-sql:** A SQL statement that should be run on a connection before it is returned from the pool to test its validity. This allows for the detection of stale pool connections. An example statement could be `"select count(*) from x"`.
- **exception-sorter-class-name:** This specifies a class that implements the `org.jboss.resource.adapter.jdbc.ExceptionSorter` interface to filter `SQLExceptions` as to whether or not a connection error event should be generated. Current implementations include:

- `org.jboss.resource.adapter.jdbc.vendor.OracleExceptionSorter`
- `org.jboss.resource.adapter.jdbc.vendor.SybaseExceptionSorter`
- ***valid-connection-checker-class-name***: This specifies a class that implements the `org.jboss.resource.adapter.jdbc.ValidConnectionChecker` interface to provide a `SQLException is-ValidConnection(Connection c)` method that is called with a connection that is to be returned from the pool to test its validity. This overrides the `check-valid-connection-sql` when present. Current implementations include:
 - `org.jboss.resource.adapter.jdbc.vendor.OracleValidConnectionChecker`
- **track-statements**: This boolean element specifies whether to check for unclosed statements when a connection is returned to the pool. If true, a warning message is issued for each unclosed statement. If the log4j category `org.jboss.resource.adapter.jdbc.WrappedConnection` has trace level enabled, a stack trace of the connection close call is logged as well. This is a debug feature that can be turned off in production.
- **prepared-statement-cache-size**: This element specifies the number of prepared statements per connection in an LRU cache which is keyed by the SQL query. Setting this to zero disables the cache.

The unique xa-datasource child elements are:

- **track-connection-by-tx**: Specifying a true value for this element makes the connection manager keep an xid to connection map and only put the connection back in the pool when the transaction completes and all the connection handles are closed or disassociated (by the method calls returning). As a side effect, we never suspend and resume the xid on the connection's `XAResource`. This is the same connection tracking behavior used for local transactions.

The XA spec implies that any connection may be enrolled in any transaction using any xid for that transaction at any time from any thread (suspending other transactions if necessary). The original JCA implementation assumed this and aggressively delisted connections and put them back in the pool as soon as control left the EJB they were used in or handles were closed. Since some other transaction could be using the connection the next time work needed to be done on the original transaction, there is no way to get the original connection back. It turns out that most `XADataSource` driver vendors do not support this, and require that all work done under a particular xid go through the same connection.

- **xa-datasource-class**: The fully qualified name of the `javax.sql.XADataSource` implementation class, for example, `com.informix.jdbcx.IfxXADataSource`.
- **xa-datasource-property**: The `xa-datasource-property` element allows for specification of the properties to assign to the `XADataSource` implementation class. Each property is identified by the name attribute and the property value is given by the `xa-datasource-property` element content. The property is mapped onto the `XADataSource` implementation by looking for a JavaBeans style getter method for the property name. If found, the value of the property is set using the JavaBeans setter with the element text translated to the true property type using the `java.beans.PropertyEditor` for the type.
- **isSameRM-override-value**: A boolean flag that allows one to override the behavior of the `javax.transaction.xa.XAResource.isSameRM(XAResource xaRes)` method behavior on the XA managed connection. If specified, this value is used unconditionally as the `isSameRM(xaRes)` return value regardless of the `xaRes` parameter.
- **no-tx-seperate-pools**: The presence of this element indicates that two connection pools are required to isolate connections used with JTA transaction from thoses used without a JTA transaction. The pools are lazily constructed on first use. It's use case is for Oracle (and possibly other vendors) XA implementations that

don't like using an XA connection with and without a JTA transaction.

7.3.2. Configuring Generic JCA Adaptors

The XSLSubDeployer also supports the deployment of arbitrary non-JDBC JCA resource adaptors using an alternate abbreviated syntax. The schema for the top-level connection factory elements of the `*-ds.xml` configuration deployment file is shown in Figure 7.10.




Figure 7.10. The simplified JCA adaptor connection factory configuration descriptor top-level schema elements

Multiple connection factory configurations may be specified in a configuration deployment file. The child elements of the `connection-factories` root are:

- **mbean:** Any number mbean elements may be specified to define MBean services that should be included in the `jboss-service.xml` descriptor that results from the transformation. This may be used to configure services used by the adaptor.
- **no-tx-connection-factory:** this element is used to specify the `(org.jboss.resource.connectionmanager) NoTxConnectionManager` service configuration. `NoTxConnectionManager` is a JCA connection manager with no transaction support. The `no-tx-connection-factory` child element schema is given in Figure 7.11.
- **tx-connection-factory:** this element is used to specify the `(org.jboss.resource.connectionmanager) TxConnectionManager` service configuration. The `tx-connection-factory` child element schema is given in Figure 7.12.

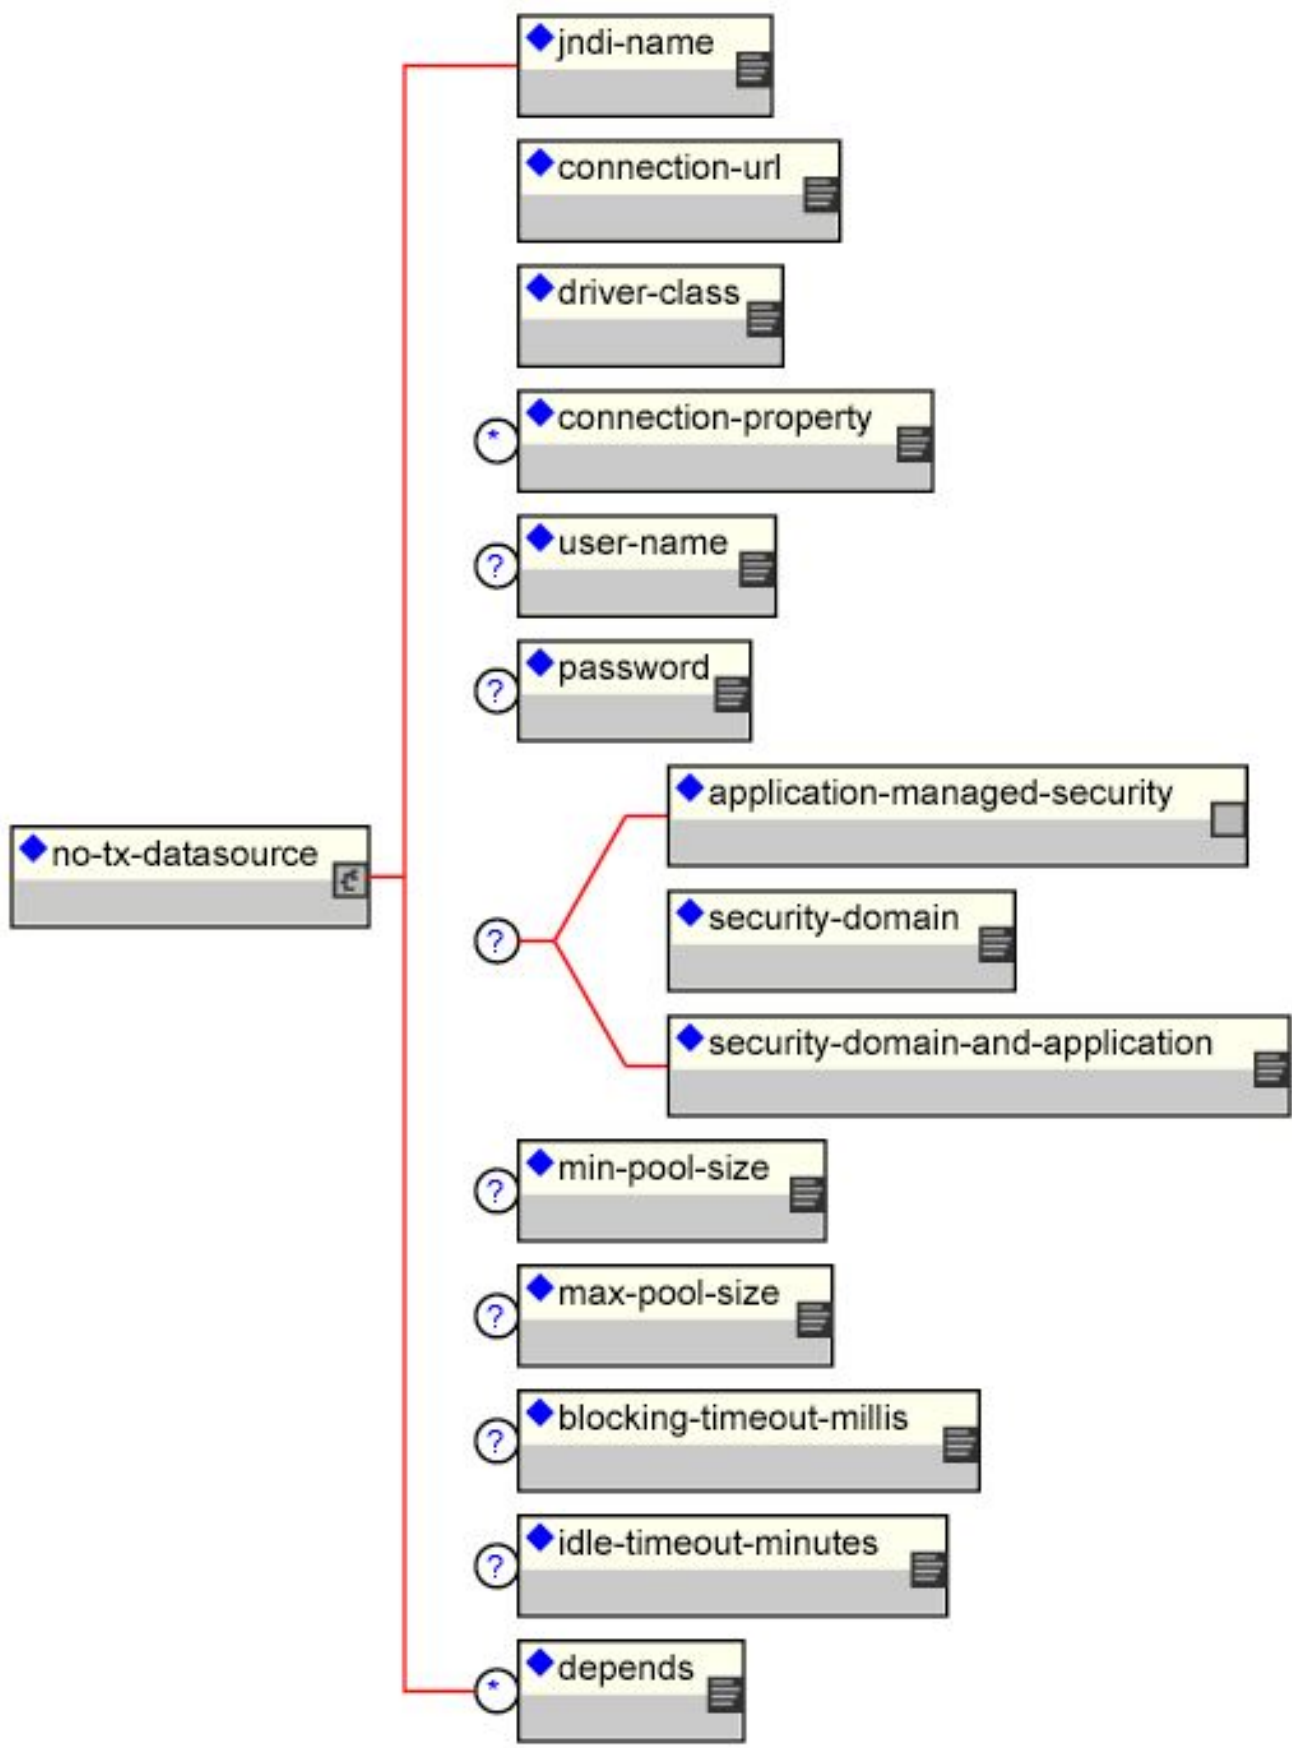


Figure 7.11. The no-tx-connection-factory element schema

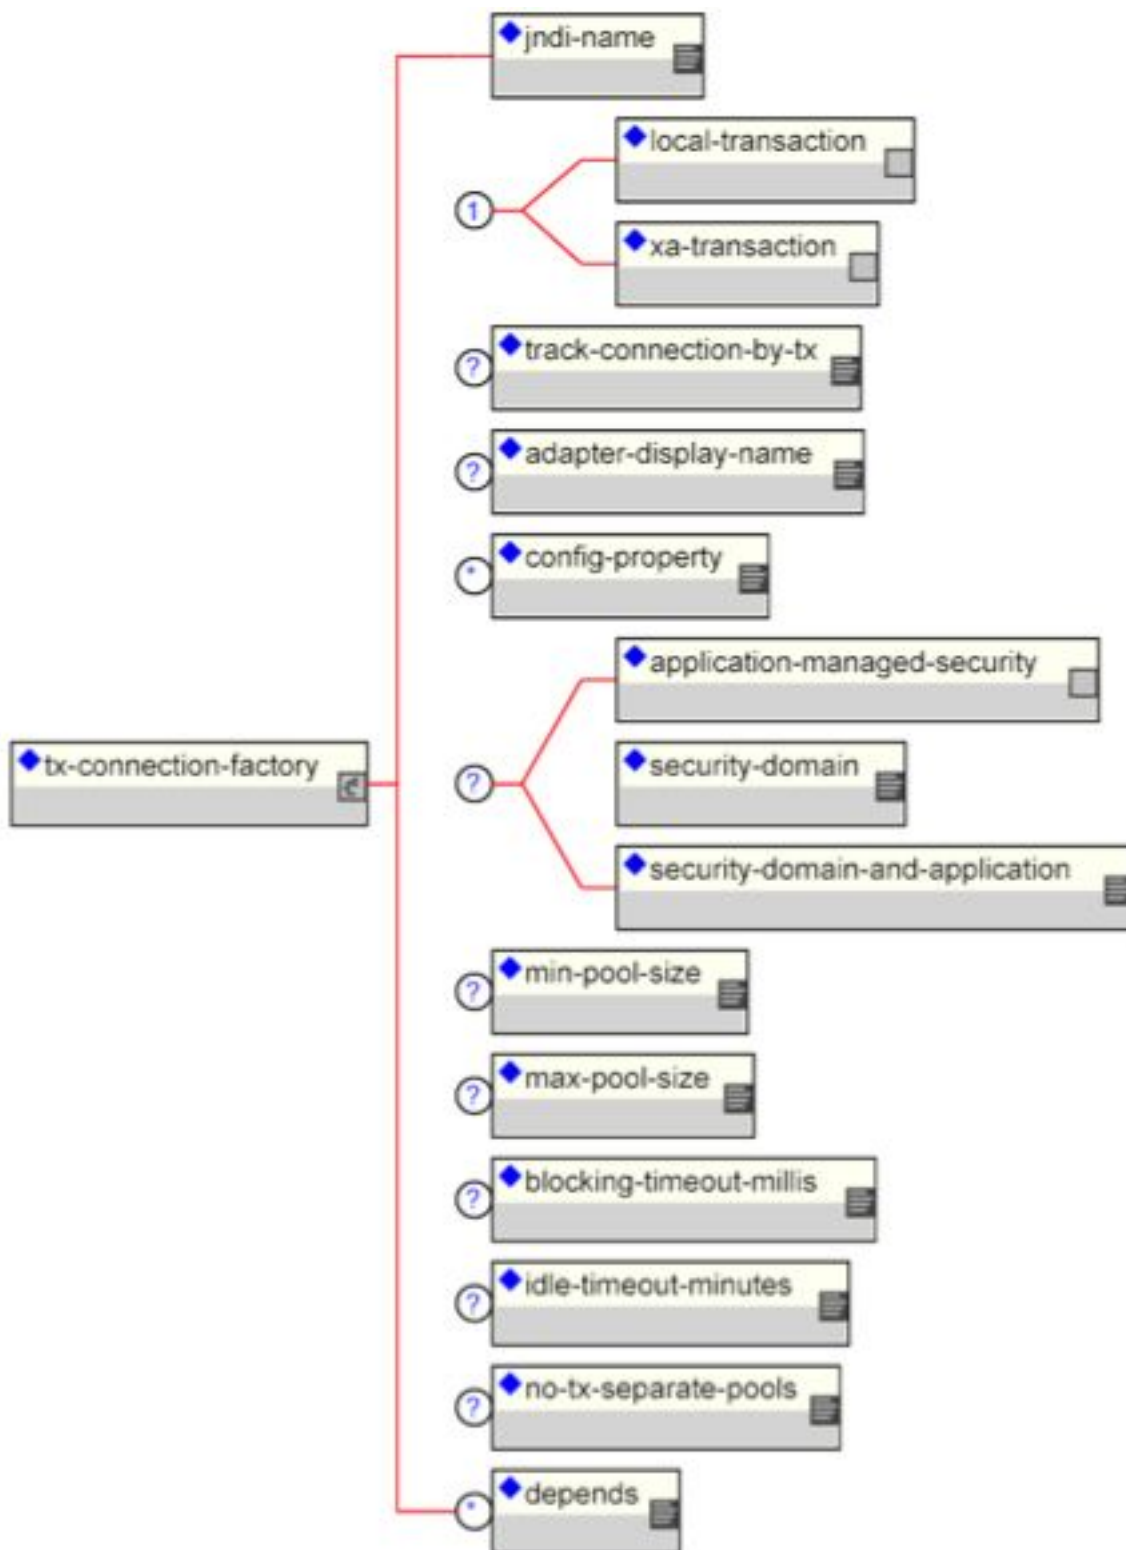


Figure 7.12. The tx-connection-factory element schema

The majority of the elements are the same as those of the datasources configuration. The element unique to the connection factory configuration include:

- **adaptor-display-name:** A human readable display name to assign to the connection manager MBean.

- **config-property**: Any number of properties to supply to the `ManagedConnectionFactory` (MCF) MBean service configuration. Each `config-property` element specifies the value of a MCF property. The `config-property` element has two required attributes:
 - **name**: The name of the property
 - **type**: The fully qualified type of the property
- The content of the `config-property` element provides the string representation of the property value. This will be converted to the true property type using the associated type `PropertyEditor`.
- **local-transaction** | **xa-transaction**: These element specify whether the `tx-connection-factory` supports local transaction or XA transactions.

7.3.3. Sample Configurations

Example configurations of many third-party JDBC drivers is included in the `JBOSS_DIST/docs/examples/jca` directory. Current example configurations include:

- `asapxcess-jb3.2-ds.xml`
- `cicsr9s-service.xml`
- `db2-ds.xml`
- `db2-xa-ds.xml`
- `facets-ds.xml`
- `fast-objects-jboss32-ds.xml`
- `firebird-ds.xml`
- `firstsql-ds.xml`
- `firstsql-xa-ds.xml`
- `generic-ds.xml`
- `hsqldb-ds.xml`
- `informix-ds.xml`
- `informix-xa-ds.xml`
- `jdatastore-ds.xml`
- `jms-ds.xml`
- `jsql-ds.xml`
- `lido-versant-service.xml`
- `mimer-ds.xml`
- `mimer-xa-ds.xml`
- `msaccess-ds.xml`
- `mssql-ds.xml`
- `mssql-xa-ds.xml`
- `mysql-ds.xml`
- `oracle-ds.xml`
- `oracle-xa-ds.xml`
- `postgres-ds.xml`
- `sapdb-ds.xml`
- `sapr3-ds.xml`
- `solid-ds.xml`
- `sybase-ds.xml`

Security on JBoss

J2EE Security Configuration and Architecture

Security is a fundamental part of any enterprise application. You need to be able to restrict who is allowed to access your applications and control what operations application users may perform. The J2EE specifications define a simple role-based security model for EJBs and web components. The JBoss component framework that handles security is the JBossSX extension framework. The JBossSX security extension provides support for both the role-based declarative J2EE security model as well as integration of custom security via a security proxy layer. The default implementation of the declarative security model is based on Java Authentication and Authorization Service (JAAS) login modules and subjects. The security proxy layer allows custom security that cannot be described using the declarative model to be added to an EJB in a way that is independent of the EJB business object. Before getting into the JBoss security implementation details, we will review EJB and Servlet specification security models as well as JAAS to establish the foundation for these details.

8.1. J2EE Declarative Security Overview

The security model advocated by the J2EE specification is a declarative model. It is declarative in that you describe the security roles and permissions using a standard XML descriptor rather than embedding security into your business component. This isolates security from business-level code because security tends to be a more a function of where the component is deployed, rather than an inherent aspect of the component's business logic. For example, consider an ATM component that is to be used to access a bank account. The security requirements, roles and permissions will vary independent of how one accesses the bank account based on what bank is managing the account, where the ATM machine is deployed, and so on.

Securing a J2EE application is based on the specification of the application security requirements via the standard J2EE deployment descriptors. You secure access to EJBs and web components in an enterprise application by using the `ejb-jar.xml` and `web.xml` deployment descriptors. Figure 8.1 and Figure 8.2 illustrate the security-related elements in the EJB 2.0 and Servlet 2.2 deployment descriptors, respectively.

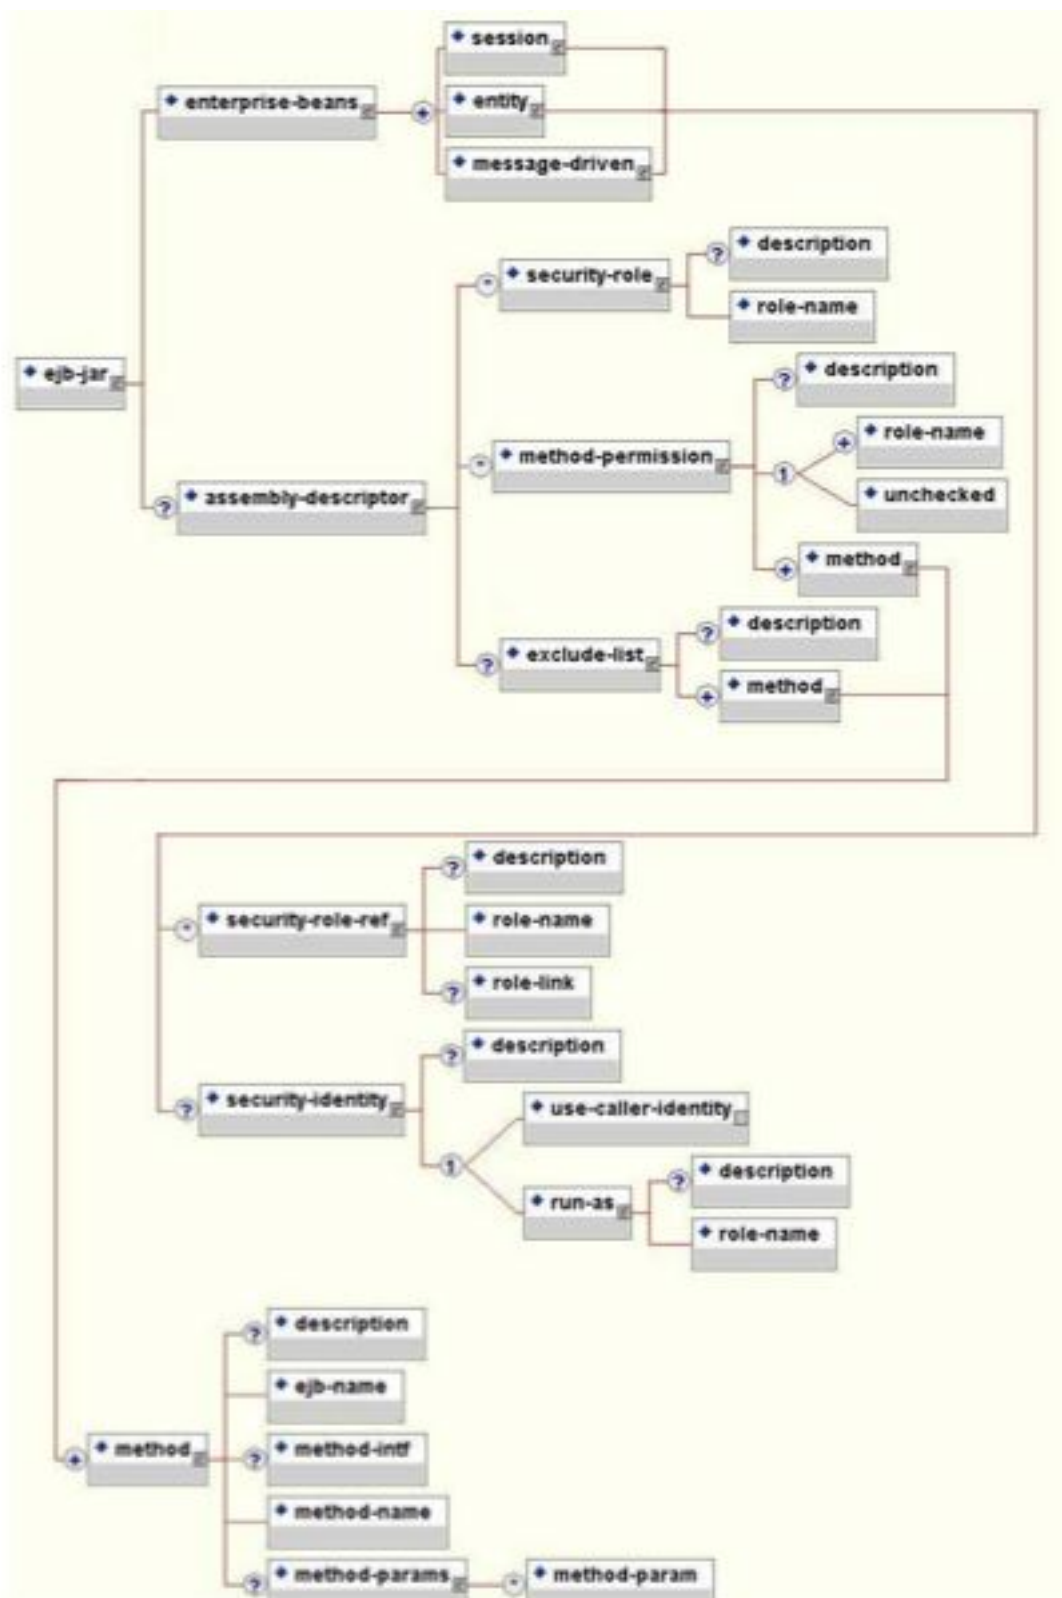


Figure 8.1. A subset of the EJB 2.0 deployment descriptor content model that shows the security related elements.

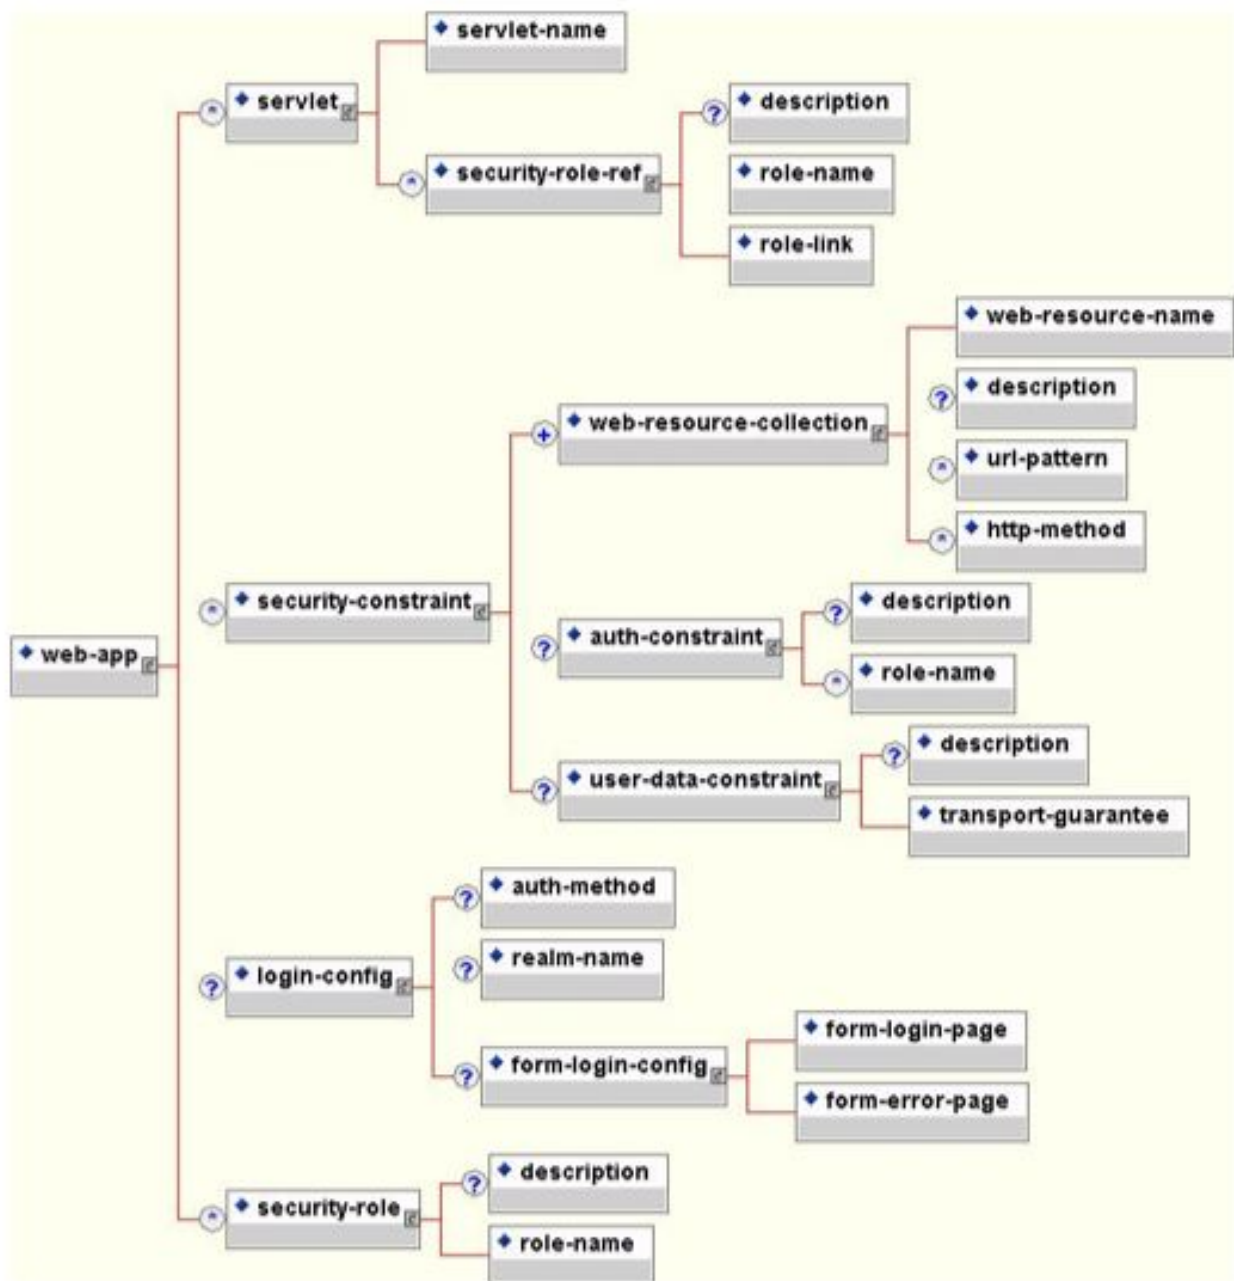


Figure 8.2. A subset of the Servlet 2.2 deployment descriptor content model that shows the security related elements.

The purpose and usage of the various security elements given in Figure 8.1 and Figure 8.2 is discussed in the following subsections.

8.1.1. Security References

Both EJBs and servlets may declare one or more `security-role-ref` elements. This element is used to declare that a component is using the `role-name` value as an argument to the `isCallerInRole(String)` method. Using the `isCallerInRole` method, a component can verify if the caller is in a role that has been declared with a `security-role-ref/role-name` element. The `role-name` element value must link to a `security-role` element through the `role-link` element. The typical use of `isCallerInRole` is to perform a security check that cannot

be defined using the role based `method-permissions` elements. However, use of `isCallerInRole` is discouraged because this results in security logic embedded inside of the component code. Example descriptor fragments that illustrate the `security-role-ref` element usage are presented in Example 8.4 and Example 8.5.

Example 8.1. An example `ejb-jar.xml` descriptor fragments which illustrate the `security-role-ref` element usage.

```
<!-- A sample ejb-jar.xml fragment -->
<ejb-jar>
  <enterprise-beans>
    <session>
      <ejb-name>ASessionBean</ejb-name>
      ...
      <security-role-ref>
        <role-name>TheRoleICheck</role-name>
        <role-link>TheApplicationRole</role-link>
      </security-role-ref>
    </session>
  </enterprise-beans>
  ...
</ejb-jar>
```

Example 8.2. An example `web.xml` descriptor fragments which illustrate the `security-role-ref` element usage.

```
<web-app>
  <servlet>
    <servlet-name>AServlet</servlet-name>
    ...
    <security-role-ref>
      <role-name>TheServletRole</role-name>
      <role-link>TheApplicationRole</role-link>
    </security-role-ref>
  </servlet>
  ...
</web-app>
```

8.1.2. Security Identity

EJBs can optionally declare a `security-identity` element. New to EJB 2.0 is the capability to specify what identity an EJB should use when it invokes methods on other components. The invocation identity can be that of the current caller, or a specific role. The application assembler uses the `security-identity` element with a `use-caller-identity` child element to indicate the current caller's identity should be propagated as the security identity for method invocations made by the EJB. Propagation of the caller's identity is the default used in the absence of an explicit `security-identity` element declaration.

Alternatively, the application assembler can use the `run-as/role-name` child element to specify that a specific security role given by the `role-name` value should be used as the security identity for method invocations made by the EJB. Note that this does not change the caller's identity as seen by `EJBContext.getCallerPrincipal()`. Rather, the caller's security roles are set to the single role specified by the `run-as/role-name` element value. One use case for the `run-as` element is to prevent external clients from accessing internal EJBs. This is accomplished by assigning the internal EJB `method-permission` elements that restrict access to a role never assigned to an external client. EJBs that need to use internal EJB are then configured with a `run-as/role-name` equal to the restricted role. An example descriptor fragment that illustrates `security-identity` element usage is presented in Example 8.3.

Example 8.3. An example `ejb-jar.xml` descriptor fragment which illustrates the `security-identity` element usage.

```
<!-- A sample ejb-jar.xml fragment -->
<ejb-jar>
  <enterprise-beans>
    <session>
      <ejb-name>ASessionBean</ejb-name>
      <!-- ... -->
      <security-identity>
        <use-caller-identity/>
      </security-identity>
    </session>
    <session>
      <ejb-name>RunAsBean</ejb-name>
      <!-- ... -->
      <security-identity>
        <run-as>
          <description>A private internal role</description>
          <role-name>InternalRole</role-name>
        </run-as>
      </security-identity>
    </session>
  </enterprise-beans>
  <!-- ... -->
</ejb-jar>
```

8.1.3. Security roles

The security role name referenced by either the `security-role-ref` or `security-identity` element needs to map to one of the application's declared roles. An application assembler defines logical security roles by declaring `security-role` elements. The `role-name` value is a logical application role name like `Administrator`, `Architect`, `SalesManager`, etc.

What is a role? The J2EE specifications note that it is important to keep in mind that the security roles in the deployment descriptor are used to define the logical security view of an application. Roles defined in the J2EE deployment descriptors should not be confused with the user groups, users, principals, and other concepts that exist in the target enterprise's operational environment. The deployment descriptor roles are application constructs with application domain specific names. For example, a banking application might use role names like `BankManager`, `Teller`, and `Customer`.

In JBoss, a `security-role` is only used to map `security-role-ref/role-name` values to the logical role that the component role referenced. The user's assigned roles are a dynamic function of the application's security manager, as you will see when we discuss the JBossSX implementation details. JBoss does not require the definition of `security-roles` in order to declare method permissions. Therefore, the specification of `security-role` elements is simply a good practice to ensure portability across application servers and for deployment descriptor maintenance. Example descriptor fragments that illustrate `security-role` usage are presented in Example 8.4 and Example 8.5.

Example 8.4. An example `ejb-jar.xml` descriptor fragments which illustrate the `security-role` element usage.

```
<!-- A sample ejb-jar.xml fragment -->
<ejb-jar>
  <!-- ... -->
  <assembly-descriptor>
```

```

    <security-role>
      <description>The single application role</description>
      <role-name>TheApplicationRole</role-name>
    </security-role>
  </assembly-descriptor>
</ejb-jar>

```

Example 8.5. An example web.xml descriptor fragment which illustrate the security-role element usage.

```

<!-- A sample web.xml fragment -->
<web-app>
  <!-- ... -->
  <security-role>
    <description>The single application role</description>
    <role-name>TheApplicationRole</role-name>
  </security-role>
</web-app>

```

8.1.4. EJB method permissions

An application assembler can set the roles that are allowed to invoke an EJB's home and remote interface methods through method-permission element declarations. Each `method-permission` element contains one or more `role-name` child elements that define the logical roles allowed access the EJB methods as identified by method child elements. As of EJB 2.0, you can now specify an `unchecked` element instead of the `role-name` element to declare that any authenticated user can access the methods identified by method child elements. In addition, you can declare that no one should have access to a method with the `exclude-list` element. If an EJB has methods that have not been declared as accessible by a role using a `method-permission` element, the EJB methods default to being excluded from use. This is equivalent to defaulting the methods into the `exclude-list`.

There are three supported styles of method element declarations.

- Style 1 is used for referring to all of the home and component interface methods of the named enterprise bean.

```

<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>*</method-name>
</method>

```

- Style 2 is used for referring to a specified method of the home or component interface of the named enterprise bean. If there are multiple methods with the same overloaded name, this style refers to all of the overloaded methods.

```

<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
</method>

```

- Style 3 is used to refer to a specified method within a set of methods with an overloaded name. The method must be defined in the specified enterprise bean's home or remote interface. The `method-param` element values are the fully qualified name of the corresponding method parameter type. If there are multiple methods

with the same overloaded signature, the permission applies to all of the matching overloaded methods.

```
<method>
  <ejb-name>EJBNAME</ejb-name>
  <method-name>METHOD</method-name>
  <method-params>
    <method-param>PARAMETER_1</method-param>
    <!-- ... -->
    <method-param>PARAMETER_N</method-param>
  </method-params>
</method>
```

The optional method-intf element can be used to differentiate methods with the same name and signature that are defined in both the home and remote interfaces of an enterprise bean. Example 8.6 provides examples of the method-permission element usage.

Example 8.6. An example ejb-jar.xml descriptor fragment which illustrates the method-permission element usage.

```
<ejb-jar>
  <assembly-descriptor>
    <method-permission>
      <description>The employee and temp-employee roles may access any
        method of the EmployeeService bean </description>
      <role-name>employee</role-name>
      <role-name>temp-employee</role-name>
      <method>
        <ejb-name>EmployeeService</ejb-name>
        <method-name>*</method-name>
      </method>
    </method-permission>
    <method-permission>
      <description>The employee role may access the findByPrimaryKey,
        getEmployeeInfo, and the updateEmployeeInfo(String) method of
        the AardvarkPayroll bean </description>
      <role-name>employee</role-name>
      <method>
        <ejb-name>AardvarkPayroll</ejb-name>
        <method-name>findByPrimaryKey</method-name>
      </method>
      <method>
        <ejb-name>AardvarkPayroll</ejb-name>
        <method-name>getEmployeeInfo</method-name>
      </method>
      <method>
        <ejb-name>AardvarkPayroll</ejb-name>
        <method-name>updateEmployeeInfo</method-name>
        <method-params>
          <method-param>java.lang.String</method-param>
        </method-params>
      </method>
    </method-permission>
    <method-permission>
      <description>The admin role may access any method of the
        EmployeeServiceAdmin bean </description>
      <role-name>admin</role-name>
      <method>
        <ejb-name>EmployeeServiceAdmin</ejb-name>
        <method-name>*</method-name>
      </method>
    </method-permission>
    <method-permission>
      <description>Any authenticated user may access any method of the
        EmployeeServiceHelp bean</description>
      <unchecked/>
      <method>
```

```

        <ejb-name>EmployeeServiceHelp</ejb-name>
        <method-name>*</method-name>
    </method>
</method-permission>
<exclude-list>
    <description>No fireTheCTO methods of the EmployeeFiring bean may be
        used in this deployment</description>
    <method>
        <ejb-name>EmployeeFiring</ejb-name>
        <method-name>fireTheCTO</method-name>
    </method>
</exclude-list>
</assembly-descriptor>
</ejb-jar>

```

8.1.5. Web Content Security Constraints

In a web application, security is defined by the roles allowed access to content by a URL pattern that identifies the protected content. This set of information is declared using the `web.xml` `security-constraint` element. The content to be secured is declared using one or more `web-resource-collection` elements. Each `web-resource-collection` element contains an optional series of `url-pattern` elements followed by an optional series of `http-method` elements. The `url-pattern` element value specifies a URL pattern against which a request URL must match for the request to correspond to an attempt to access secured content. The `http-method` element value specifies a type of HTTP request to allow.

The optional `user-data-constraint` element specifies the requirements for the transport layer of the client to server connection. The requirement may be for content integrity (preventing data tampering in the communication process) or for confidentiality (preventing reading while in transit). The `transport-guarantee` element value specifies the degree to which communication between client and server should be protected. Its values are `NONE`, `INTEGRAL`, or `CONFIDENTIAL`. A value of `NONE` means that the application does not require any transport guarantees. A value of `INTEGRAL` means that the application requires the data sent between the client and server be sent in such a way that it can't be changed in transit. A value of `CONFIDENTIAL` means that the application requires the data be transmitted in a fashion that prevents other entities from observing the contents of the transmission. In most cases, the presence of the `INTEGRAL` or `CONFIDENTIAL` flag indicates that the use of SSL is required.

The optional `login-config` is used to configure the authentication method that should be used, the realm name that should be used for this application, and the attributes that are needed by the form login mechanism. The `auth-method` child element specifies the authentication mechanism for the web application. As a prerequisite to gaining access to any web resources that are protected by an authorization constraint, a user must have authenticated using the configured mechanism. Legal values for `auth-method` are `BASIC`, `DIGEST`, `FORM`, or `CLIENT-CERT`. The `realm-name` child element specifies the realm name to use in HTTP basic and digest authorization. The `form-login-config` child element specifies the log in as well as error pages that should be used in form-based login. If the `auth-method` value is not `FORM`, `form-login-config` and its child elements are ignored.

As an example, the `web.xml` descriptor fragment given in Example 8.7 indicates that any URL lying under the web application `/restricted` path requires an `AuthorizedUser` role. There is no required transport guarantee and the authentication method used for obtaining the user identity is `BASIC` HTTP authentication.

Example 8.7. A `web.xml` descriptor fragment which illustrates the use of the `security-constraint` and related elements.

```

<web-app>
    <!-- ... -->
    <security-constraint>
        <web-resource-collection>

```

```
<web-resource-name>Secure Content</web-resource-name>
<url-pattern>/restricted/*</url-pattern>
</web-resource-collection>
<auth-constraint>
  <role-name>AuthorizedUser</role-name>
</auth-constraint>
<user-data-constraint>
  <transport-guarantee>NONE</transport-guarantee>
</user-data-constraint>
</security-constraint>
<!-- ... -->
<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>The Restricted Zone</realm-name>
</login-config>
<!-- ... -->
<security-role>
  <description>The role required to access restricted content </description>
  <role-name>AuthorizedUser</role-name>
</security-role>
</web-app>
```

8.1.6. Enabling Declarative Security in JBoss

The J2EE security elements that have been covered describe only the security requirements from the application's perspective. Since J2EE security elements declare logical roles, the application deployer maps the roles from the application domain onto the deployment environment. The J2EE specifications omit these application-server-specific details. In JBoss, mapping the application roles onto the deployment environment entails specifying a security manager that implements the J2EE security model using JBoss server specific deployment descriptors. We will avoid discussion the details of this step for now. The details behind the security configuration will be discussed when we describe the generic JBoss server security interfaces in Section 8.3.

8.2. An Introduction to JAAS

The default implementation of the JBossSX framework is based on the JAAS API. It is important that you understand the basic elements of the JAAS API to understand the implementation details of JBossSX. This section provides an introduction to JAAS to prepare you for the JBossSX architecture discussion. Additional details on the JAAS package can be found at the JAAS home page at: <http://java.sun.com/products/jaas/>.

8.2.1. What is JAAS?

The JAAS 1.0 API consists of a set of Java packages designed for user authentication and authorization. It implements a Java version of the standard Pluggable Authentication Module (PAM) framework and compatibly extends the Java 2 Platform's access control architecture to support user-based authorization. JAAS was first released as an extension package for JDK 1.3 and is bundled with JDK 1.4+. Because the JBossSX framework uses only the authentication capabilities of JAAS to implement the declarative role-based J2EE security model, this introduction focuses on only that topic.

Much of this section's material is derived from the JAAS 1.0 Developers Guide, so if you are familiar with its content you can skip ahead to the JBossSX architecture discussion in Section 8.4

JAAS authentication is performed in a pluggable fashion. This permits Java applications to remain independent from underlying authentication technologies and allows the JBossSX security manager to work in different security infrastructures. Integration with a security infrastructure can be achieved without changing the JBossSX

security manager implementation. All that needs to change is the configuration of the authentication stack that JAAS uses.

8.2.1.1. The JAAS Core Classes

The JAAS core classes can be broken down into three categories: common, authentication, and authorization. The following list presents only the common and authentication classes because these are the specific classes used to implement the functionality of JBossSX covered in this chapter.

Common classes:

- `Subject` (`javax.security.auth.Subject`)
- `Principal` (`java.security.Principal`)

Authentication classes:

- `Callback` (`javax.security.auth.callback.Callback`)
- `CallbackHandler` (`javax.security.auth.callback.CallbackHandler`)
- `Configuration` (`javax.security.auth.login.Configuration`)
- `LoginContext` (`javax.security.auth.login.LoginContext`)
- `LoginModule` (`javax.security.auth.spi.LoginModule`)

8.2.1.1.1. Subject and Principal

To authorize access to resources, applications first need to authenticate the request's source. The JAAS framework defines the term `subject` to represent a request's source. The `Subject` class is the central class in JAAS. A `Subject` represents information for a single entity, such as a person or service. It encompasses the entity's principals, public credentials, and private credentials. The JAAS APIs use the existing Java 2 `java.security.Principal` interface to represent a principal, which is essentially just a typed name.

During the authentication process, a subject is populated with associated identities, or principals. A subject may have many principals. For example, a person may have a name principal (John Doe), a social security number principal (123-45-6789), and a username principal (johnd), all of which help distinguish the subject from other subjects. To retrieve the principals associated with a `Subject`, two methods are available:

```
public Set getPrincipals() {...}
public Set getPrincipals(Class c) {...}
```

The first method returns all principals contained in the subject. The second method only returns those principals that are instances of `Class c` or one of its subclasses. An empty set will be returned if the subject has no matching principals. Note that the `java.security.acl.Group` interface is a subinterface of `java.security.Principal`, and so an instance in the principals set may represent a logical grouping of other principals or groups of principals.

8.2.1.1.2. Authentication of a Subject

Authentication of a subject requires a JAAS login. The login procedure consists of the following steps:

- An application instantiates a `LoginContext` passing in the name of the login configuration and a `CallbackHandler` to populate the `Callback` objects as required by the configuration `LoginModules`.
- The `LoginContext` consults a `Configuration` to load all of the `LoginModules` included in the named login configuration. If no such named configuration exists the other configuration is used as a default.

- The application invokes the `LoginContext.login` method.
- The login method invokes all the loaded `LoginModules`. As each `LoginModule` attempts to authenticate the subject, it invokes the handle method on the associated `CallbackHandler` to obtain the information required for the authentication process. The required information is passed to the handle method in the form of an array of `Callback` objects. Upon success, the `LoginModules` associate relevant principals and credentials with the subject.
- The `LoginContext` returns the authentication status to the application. Success is represented by a return from the login method. Failure is represented through a `LoginException` being thrown by the login method.
- If authentication succeeds, the application retrieves the authenticated subject using the `LoginContext.getSubject` method.
- After the scope of the subject authentication is complete, all principals and related information associated with the subject by the login method may be removed by invoking the `LoginContext.logout` method.

The `LoginContext` class provides the basic methods for authenticating subjects and offers a way to develop an application independent of the underlying authentication technology. The `LoginContext` consults a `Configuration` to determine the authentication services configured for a particular application. `LoginModule` classes represent the authentication services. Therefore, you can plug in different login modules into an application without changing the application itself. Example 8.8 provides code fragments that illustrate the steps required by an application to authenticate a subject.

Example 8.8. An illustration of the steps of the authentication process from the application perspective.

```
CallbackHandler handler = new MyHandler();
LoginContext lc = new LoginContext("some-config", handler);

try {
    lc.login();
    Subject subject = lc.getSubject();
} catch(LoginException e) {
    System.out.println("authentication failed");
    e.printStackTrace();
}

// Perform work as authenticated Subject
// ...

// Scope of work complete, logout to remove authentication info
try {
    lc.logout();
} catch(LoginException e) {
    System.out.println("logout failed");
    e.printStackTrace();
}

// A sample MyHandler class
class MyHandler
    implements CallbackHandler
{
    public void handle(Callback[] callbacks) throws
        IOException, UnsupportedCallbackException
    {
        for (int i = 0; i < callbacks.length; i++) {
            if (callbacks[i] instanceof NameCallback) {
                NameCallback nc = (NameCallback)callbacks[i];
                nc.setName(username);
            } else if (callbacks[i] instanceof PasswordCallback) {
                PasswordCallback pc = (PasswordCallback)callbacks[i];
                pc.setPassword(password);
            }
        }
    }
}
```

```

        } else {
            throw new UnsupportedOperationException(callbacks[i],
                                                "Unrecognized Callback");
        }
    }
}
}

```

Developers integrate with an authentication technology by creating an implementation of the `LoginModule` interface. This allows different authentication technologies to be plugged into an application by administrator. Multiple `LoginModules` can be chained together to allow for more than one authentication technology as part of the authentication process. For example, one `LoginModule` may perform username/password-based authentication, while another may interface to hardware devices such as smart card readers or biometric authenticators. The life cycle of a `LoginModule` is driven by the `LoginContext` object against which the client creates and issues the login method. The process consists of a two phases. The steps of the process are as follows:

- The `LoginContext` creates each configured `LoginModule` using its public no-arg constructor.
- Each `LoginModule` is initialized with a call to its `initialize` method. The `Subject` argument is guaranteed to be non-null. The signature of the `initialize` method is: `public void initialize(Subject subject, CallbackHandler callbackHandler, Map sharedState, Map options)`.
- The `login` method is then called to start the authentication process. An example method implementation might prompt the user for a username and password, and then verify the information against data stored in a naming service such as NIS or LDAP. Alternative implementations might interface to smart cards and biometric devices, or simply extract user information from the underlying operating system. The validation of user identity by each `LoginModule` is considered phase 1 of JAAS authentication. The signature of the `login` method is: `boolean login() throws LoginException`. Failure is indicated by throwing a `LoginException`. A return of `true` indicates that the method succeeded, while a return of `false` indicates that the login module should be ignored.
- If the `LoginContext`'s overall authentication succeeds, `commit` is invoked on each `LoginModule`. If phase 1 succeeded for a `LoginModule`, then the `commit` method continues with phase 2: associating relevant principals, public credentials, and/or private credentials with the subject. If phase 1 fails for a `LoginModule`, then `commit` removes any previously stored authentication state, such as usernames or passwords. The signature of the `commit` method is: `boolean commit() throws LoginException`. Failure to complete the `commit` phase is indicated by throwing a `LoginException`. A return of `true` indicates that the method succeeded, while a return of `false` indicates that the login module should be ignored.
- If the `LoginContext`'s overall authentication failed, then the `abort` method is invoked on each `LoginModule`. The `abort` method removes/destroys any authentication state created by the `login` or `initialize` methods. The signature of the `abort` method is: `boolean abort() throws LoginException`. Failure to complete the `abort` phase is indicated by throwing a `LoginException`. A return of `true` indicates that the method succeeded, while a return of `false` indicates that the login module should be ignored.
- Removal of the authentication state after a successful login is accomplished when the application invokes `logout` on the `LoginContext`. This in turn results in a `logout` method invocation on each `LoginModule`. The `logout` method removes the principals and credentials originally associated with the subject during the `commit` operation. Credentials should be destroyed upon removal. The signature of the `logout` method is: `boolean logout() throws LoginException`. Failure to complete the `logout` process is indicated by throwing a `LoginException`. A return of `true` indicates that the method succeeded, while a return of `false` indicates that the login module should be ignored.

When a `LoginModule` must communicate with the user to obtain authentication information, it uses a `CallbackHandler` object. Applications implement the `CallbackHandler` interface and pass it to the `LoginContext`, which forwards it directly to the underlying login modules. Login modules use the `CallbackHandler` both to gather input from users, such as a password or smart-card PIN number, and to supply information to users, such as status information. By allowing the application to specify the `CallbackHandler`, underlying `LoginModules` remain independent from the different ways applications interact with users. For example, a `CallbackHandler`'s implementation for a GUI application might display a window to solicit user input. On the other hand, a `callbackhandler`'s implementation for a non-GUI environment, such as an application server, might simply obtain credential information using an application server API. The `callbackhandler` interface has one method to implement:

```
void handle(Callback[] callbacks)
throws java.io.IOException, UnsupportedCallbackException;
```

The last authentication class to cover is the `Callback` interface. This is a tagging interface for which several default implementations are provided, including `NameCallback` and `PasswordCallback` that were used in Example 8.8. `LoginModule`s use a `Callback` to request information required by the authentication mechanism the `LoginModule` encapsulates. `LoginModule`s pass an array of `Callback`s directly to the `CallbackHandler.handle` method during the authentication's login phase. If a `callbackhandler` does not understand how to use a `Callback` object passed into the `handle` method, it throws an `UnsupportedCallbackException` to abort the login call.

8.3. The JBoss Security Model

Similar to the rest of the JBoss architecture, security at the lowest level is defined as a set of interfaces for which alternate implementations may be provided. There are three basic interfaces that define the JBoss server security layer: `org.jboss.security.AuthenticationManager`, `org.jboss.security.RealmMapping`, and `org.jboss.security.SecurityProxy`. Figure 8.3 shows a class diagram of the security interfaces and their relationship to the EJB container architecture.

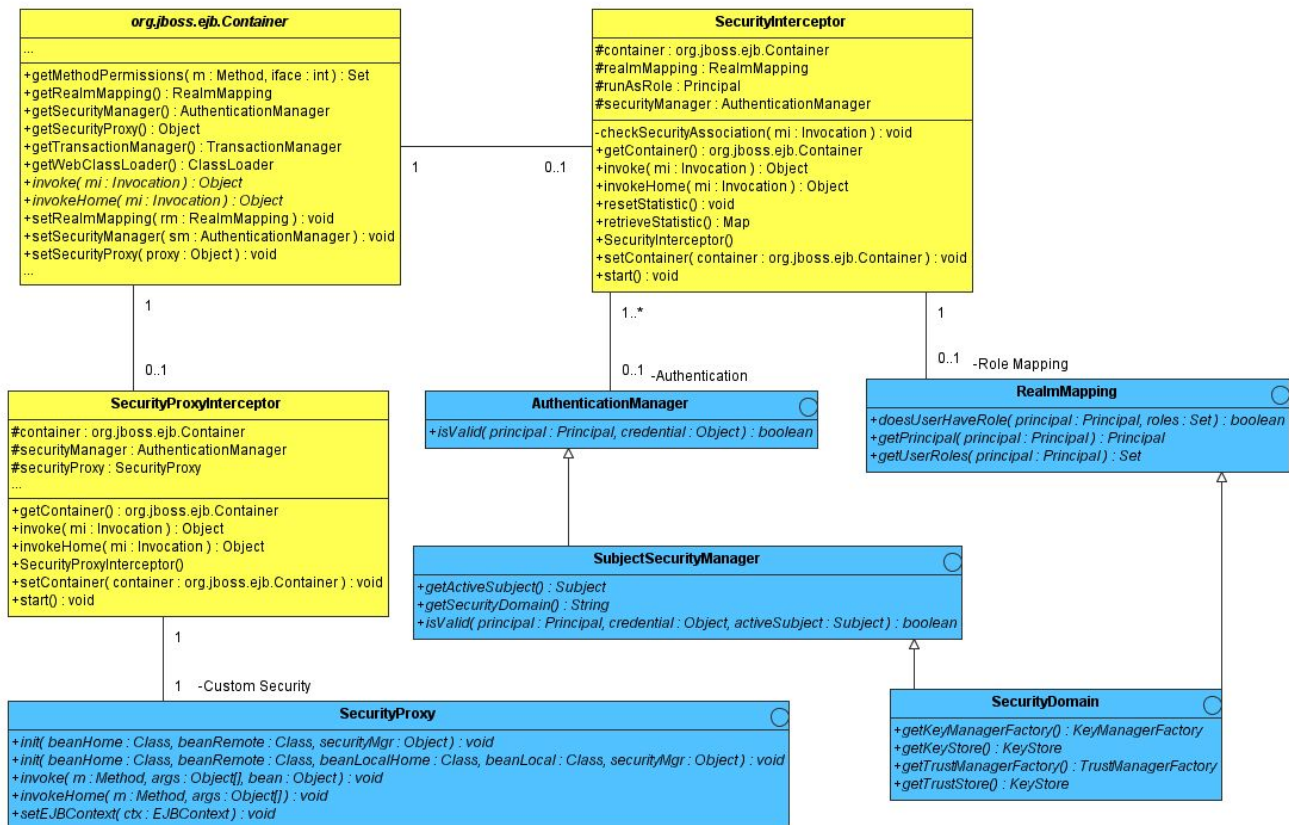


Figure 8.3. The key security model interfaces and their relationship to the JBoss server EJB container elements.

The light blue classes represent the security interfaces while the yellow classes represent the EJB container layer. The two interfaces required for the implementation of the J2EE security model are the `org.jboss.security.AuthenticationManager` and `org.jboss.security.RealmMapping`. The roles of the security interfaces presented in Figure 8.3 are summarized in the following list.

- `AuthenticationManager` is an interface responsible for validating credentials associated with principals. Principals are identities and examples include usernames, employee numbers, social security numbers, and so on. Credentials are proof of the identity and examples include passwords, session keys, digital signatures, and so on. The `isValid` method is invoked to see if a user identity and associated credentials as known in the operational environment are valid proof of the user identity.
- `RealmMapping` is an interface responsible for principal mapping and role mapping. The `getPrincipal` method takes a user identity as known in the operational environment and returns the application domain identity. The `doesUserHaveRole` method validates that the user identity in the operation environment has been assigned the indicated role from the application domain.
- `SecurityProxy` is an interface describing the requirements for a custom `SecurityProxyInterceptor` plugin. A `SecurityProxy` allows for the externalization of custom security checks on a per-method basis for both the EJB home and remote interface methods.
- `SubjectSecurityManager` is a subinterface of `AuthenticationManager` that simply adds accessor methods for obtaining the security domain name of the security manager and the current thread's authenticated Subject.
- `SecurityDomain` is an extension of the `AuthenticationManager`, `RealmMapping`, and `SubjectSecurity-`

Manager interfaces. It is a move to a comprehensive security interface based on the JAAS Subject, a `java.security.KeyStore`, and the JSSE `com.sun.net.ssl.KeyManagerFactory` and `com.sun.net.ssl.TrustManagerFactory` interfaces. This interface is still a work in progress that will be the basis of a multi-domain security architecture that will better support ASP style deployments of applications and resources.

Note that the `AuthenticationManager`, `RealmMapping` and `SecurityProxy` interfaces have no association to JAAS related classes. Although the JBossSX framework is heavily dependent on JAAS, the basic security interfaces required for implementation of the J2EE security model are not. The JBossSX framework is simply an implementation of the basic security plug-in interfaces that are based on JAAS. The component diagram presented in Figure 8.4 illustrates this fact. The implication of this plug-in architecture is that you are free to replace the JAAS-based JBossSX implementation classes with your own custom security manager implementation that does not make use of JAAS, if you so desire. You'll see how to do this when you look at the JBossSX MBeans available for the configuration of JBossSX in Figure 8.4.

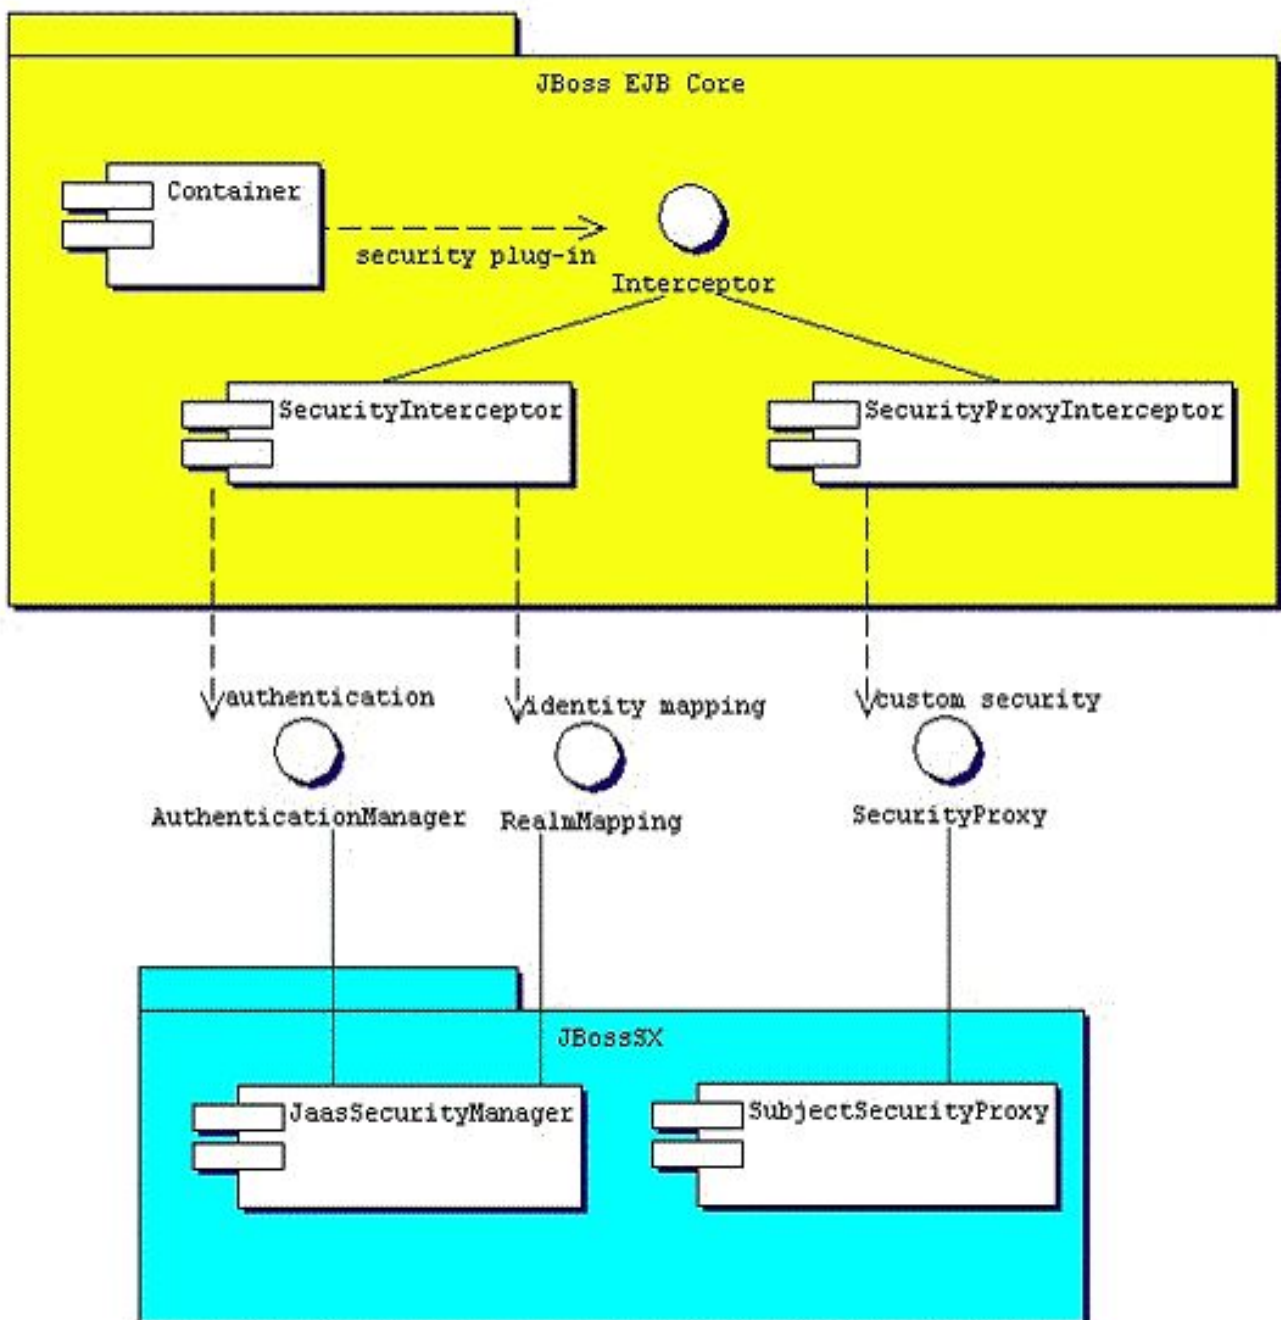


Figure 8.4. The relationship between the JBossSX framework implementation classes and the JBoss server EJB container layer.

8.3.1. Enabling Declarative Security in JBoss Revisited

Recall that our discussion of the J2EE standard security model ended with a requirement for the use of JBoss server specific deployment descriptor to enable security. The details of this configuration is presented here, as this is part of the generic JBoss security model. Figure 8.5 shows the JBoss-specific EJB and web application deployment descriptor's security-related elements.

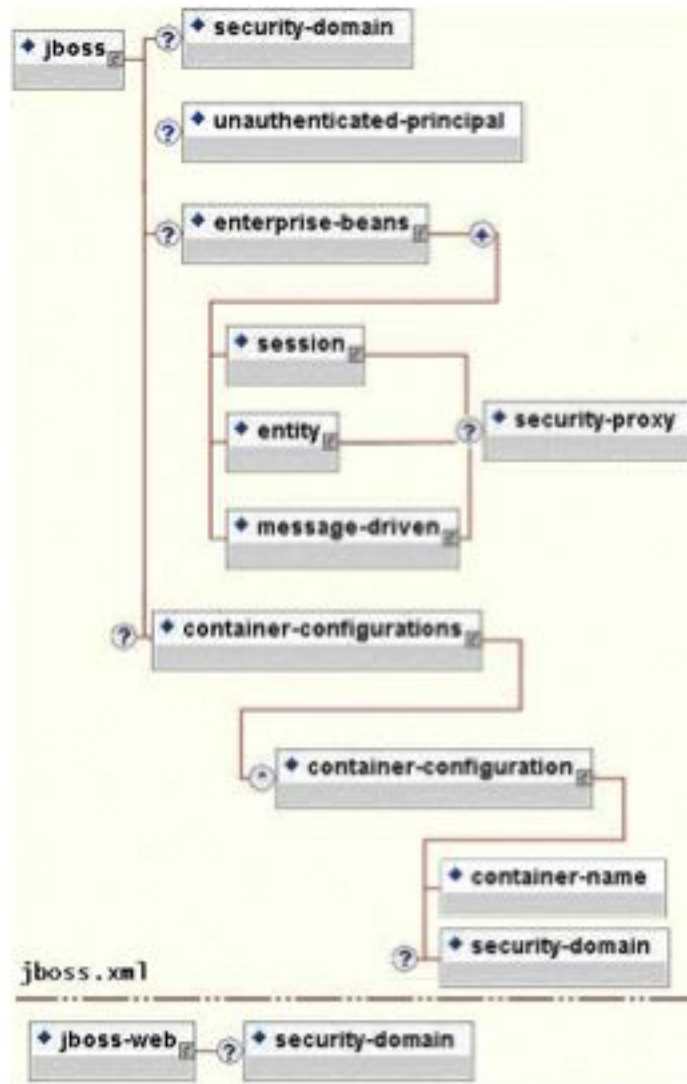


Figure 8.5. The security element subsets of the JBoss server `jboss.xml` and `jboss-web.xml` deployment descriptors.

The value of a `security-domain` element specifies the JNDI name of the security manager interface implementation that JBoss uses for the EJB and web containers. This is an object that implements both of the `AuthenticationManager` and `RealmMapping` interfaces. When specified as a top-level element it defines what security domain in effect for all EJBs in the deployment unit. This is the typical usage because mixing security managers within a deployment unit complicates inter-component operation and administration.

To specify the security domain for an individual EJB, you specify the `security-domain` at the container configuration level. This will override any top-level security-domain element.

The `unauthenticated-principal` element specifies the name to use for the `Principal` object returned by the `EJBContext.getUserPrincipal` method when an unauthenticated user invokes an EJB. Note that this conveys no special permissions to an unauthenticated caller. Its primary purpose is to allow unsecured servlets and JSP pages to invoke unsecured EJBs and allow the target EJB to obtain a non-null `Principal` for the caller using the `getUserPrincipal` method. This is a J2EE specification requirement.

The `security-proxy` element identifies a custom security proxy implementation that allows per-request security checks outside the scope of the EJB declarative security model without embedding security logic into the EJB implementation. This may be an implementation of the `org.jboss.security.SecurityProxy` interface, or just an object that implements methods in the home, remote, local home or local interfaces of the EJB to secure without implementing any common interface. If the given class does not implement the `SecurityProxy` interface, the instance must be wrapped in a `SecurityProxy` implementation that delegates the method invocations to the object. The `org.jboss.security.SubjectSecurityProxy` is an example `SecurityProxy` implementation used by the default JBossSX installation.

Take a look at a simple example of a custom `SecurityProxy` in the context of a trivial stateless session bean. The custom `SecurityProxy` validates that no one invokes the bean's `echo` method with a four-letter word as its argument. This is a check that is not possible with role-based security; you cannot define a `FourLetterEchoInvoker` role because the security context is the method argument, not a property of the caller. The code for the custom `SecurityProxy` is given in Example 8.9, and the full source code is available in the `src/main/org/jboss/chap8/ex1` directory of the book examples. The associated `jboss.xml` descriptor that installs the `EchoSecurityProxy` as the custom proxy for the `EchoBean` is given in Example 8.10.

Example 8.9. The example 1 custom `EchoSecurityProxy` implementation that enforces the echo argument-based security constraint.

```
package org.jboss.chap8.ex1;

import java.lang.reflect.Method;
import javax.ejb.EJBContext;

import org.apache.log4j.Category;

import org.jboss.security.SecurityProxy;

/** A simple example of a custom SecurityProxy implementation
 *  that demonstrates method argument based security checks.
 *  @author Scott.Stark@jboss.org
 *  @version $Revision: 1.15 $
 */
public class EchoSecurityProxy implements SecurityProxy
{
    Category log = Category.getInstance(EchoSecurityProxy.class);
    Method echo;

    public void init(Class beanHome, Class beanRemote,
                    Object securityMgr)
        throws InstantiationException
    {
        log.debug("init, beanHome="+beanHome
                  + ", beanRemote="+beanRemote
                  + ", securityMgr="+securityMgr);
        // Get the echo method for equality testing in invoke
        try {
            Class[] params = {String.class};
            echo = beanRemote.getDeclaredMethod("echo", params);
        } catch (Exception e) {
```



```

        String msg = "Failed to finde an echo(String) method";
        log.error(msg, e);
        throw new InstantiationException(msg);
    }
}

public void setEJBContext(EJBContext ctx)
{
    log.debug("setEJBContext, ctx="+ctx);
}

public void invokeHome(Method m, Object[] args)
    throws SecurityException
{
    // We don't validate access to home methods
}

public void invoke(Method m, Object[] args, Object bean)
    throws SecurityException
{
    log.debug("invoke, m="+m);
    // Check for the echo method
    if (m.equals(echo)) {
        // Validate that the msg arg is not 4 letter word
        String arg = (String) args[0];
        if (arg == null || arg.length() == 4)
            throw new SecurityException("No 4 letter words");
    }
    // We are not responsible for doing the invoke
}
}

```

Example 8.10. The jboss.xml descriptor which configures the EchoSecurityProxy as the custom security proxy for the EchoBean.

```

<jboss>
  <security-domain>java:/jaas/other</security-domain>

  <enterprise-beans>
    <session>
      <ejb-name>EchoBean</ejb-name>
      <security-proxy>org.jboss.chap8.ex1.EchoSecurityProxy</security-proxy>
    </session>
  </enterprise-beans>
</jboss>

```

The `EchoSecurityProxy` checks that the method to be invoked on the bean instance corresponds to the `echo(String)` method loaded the init method. If there is a match, the method argument is obtained and its length compared against 4 or null. Either case results in a `SecurityException` being thrown. Certainly this is a contrived example, but only in its application. It is a common requirement that applications must perform security checks based on the value of method arguments. The point of the example is to demonstrate how custom security beyond the scope of the standard declarative security model can be introduced independent of the bean implementation. This allows the specification and coding of the security requirements to be delegated to security experts. Since the security proxy layer can be done independent of the bean implementation, security can be changed to match the deployment environment requirements.

Now test the custom proxy by running a client that attempts to invoke the `EchoBean.echo` method with the arguments `Hello` and `Four` as illustrated in this fragment:

```
public class ExClient
```

```

{
    public static void main(String args[])
        throws Exception
    {
        Logger log = Logger.getLogger("ExClient");
        log.info("Looking up EchoBean");

        InitialContext iniCtx = new InitialContext();
        Object ref = iniCtx.lookup("EchoBean");
        EchoHome home = (EchoHome) ref;
        Echo echo = home.create();

        log.info("Created Echo");
        log.info("Echo.echo('Hello') = "+echo.echo("Hello"));
        log.info("Echo.echo('Four') = "+echo.echo("Four"));
    }
}

```

The first call should succeed, while the second should fail due to the fact that `Four` is a four-letter word. Run the client as follows using Ant from the examples directory:

```

[nr@toki examples]$ ant -Dchap=chap8 -Dex=1 run-example
run-example1:
[copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
[echo] Waiting for 5 seconds for deploy...
[java] [INFO,ExClient] Looking up EchoBean
[java] [INFO,ExClient] Created Echo
[java] [INFO,ExClient] Echo.echo('Hello') = Hello
[java] Exception in thread "main" java.rmi.ServerException: RemoteException occurred
in server thread; nested exception is:
[java]     java.rmi.AccessException: SecurityException; nested exception is:
[java]     java.lang.SecurityException: No 4 letter words
...
[java]     at org.jboss.chap8.ex1.ExClient.main(ExClient.java:28)
[java] Caused by: java.rmi.AccessException: SecurityException; nested exception is:
[java]     java.lang.SecurityException: No 4 letter words
...

```

The result is that the `echo('Hello')` method call succeeds as expected and the `echo('Four')` method call results in a rather messy looking exception, which is also expected. The above output has been truncated to fit in the book. The key part to the exception is that the `SecurityException("No 4 letter words")` generated by the `EchoSecurityProxy` was thrown to abort the attempted method invocation as desired.

8.4. The JBoss Security Extension Architecture

The preceding discussion of the general JBoss security layer has stated that the JBossSX security extension framework is an implementation of the security layer interfaces. This is the primary purpose of the JBossSX framework. The details of the implementation are interesting in that it offers a great deal of customization for integration into existing security infrastructures. A security infrastructure can be anything from a database or LDAP server to a sophisticated security software suite. The integration flexibility is achieved using the plugable authentication model available in the JAAS framework.

The heart of the JBossSX framework is `org.jboss.security.plugins.JaasSecurityManager`. This is the default implementation of the `AuthenticationManager` and `RealmMapping` interfaces. Figure 8.6 shows how the `JaasSecurityManager` integrates into the EJB and web container layers based on the `security-domain` element of the corresponding component deployment descriptor.

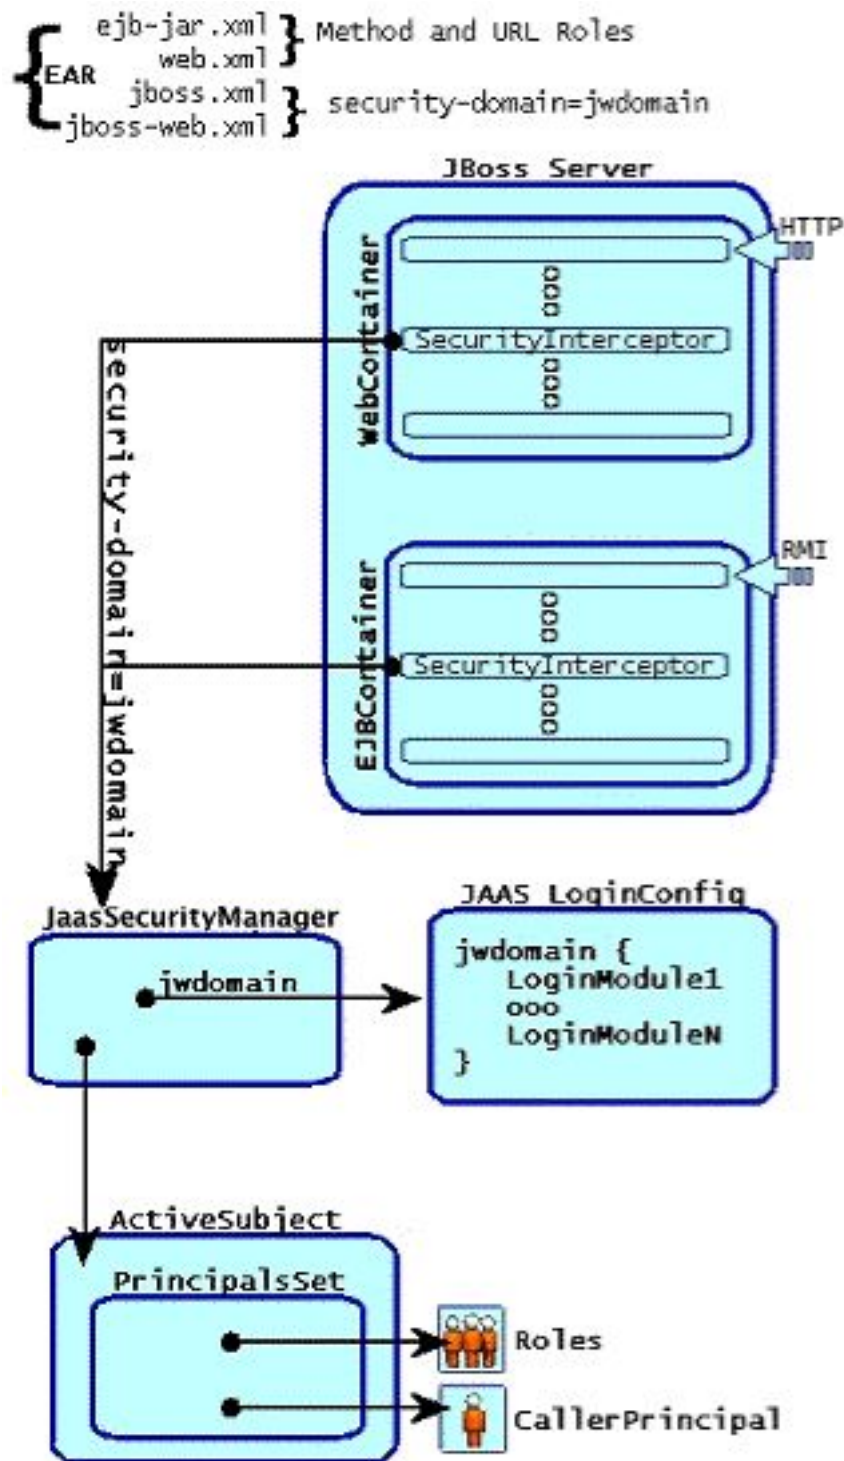


Figure 8.6. The relationship between the security-domain component deployment descriptor value, the component container and the `JaasSecurityManager`.

Figure 8.6 depicts an enterprise application that contains both EJBs and web content secured under the security domain `jwdomain`. The EJB and web containers have a request interceptor architecture that includes a security interceptor, which enforces the container security model. At deployment time, the `security-domain` element value in the `jboss.xml` and `jboss-web.xml` descriptors is used to obtain the security manager instance associated with the container. The security interceptor then uses the security manager to perform its role. When a secured component is requested, the security interceptor delegates security checks to the security manager instance associated with the container.

The JBossSX `JaasSecurityManager` implementation, shown in Figure 8.6 as the `JaasSecurityMgr` component, performs security checks based on the information associated with the `Subject` instance that results from executing the JAAS login modules configured under the name matching the security-domain element value. We will drill into the `JaasSecurityManager` implementation and its use of JAAS in the following section.

8.4.1. How the `JaasSecurityManager` Uses JAAS

The `JaasSecurityManager` uses the JAAS packages to implement the `AuthenticationManager` and `RealmMapping` interface behavior. In particular, its behavior derives from the execution of the login module instances that are configured under the name that matches the security domain to which the `JaasSecurityManager` has been assigned. The login modules implement the security domain's principal authentication and role-mapping behavior. Thus, you can use the `JaasSecurityManager` across different security domains simply by plugging in different login module configurations for the domains.

To illustrate the details of the `JaasSecurityManager`'s usage of the JAAS authentication process, you will walk through a client invocation of an EJB home method invocation. The prerequisite setting is that the EJB has been deployed in the JBoss server and its home interface methods have been secured using `method-permission` elements in the `ejb-jar.xml` descriptor, and it has been assigned a security domain named `jwdomain` using the `jboss.xml` descriptor `security-domain` element.

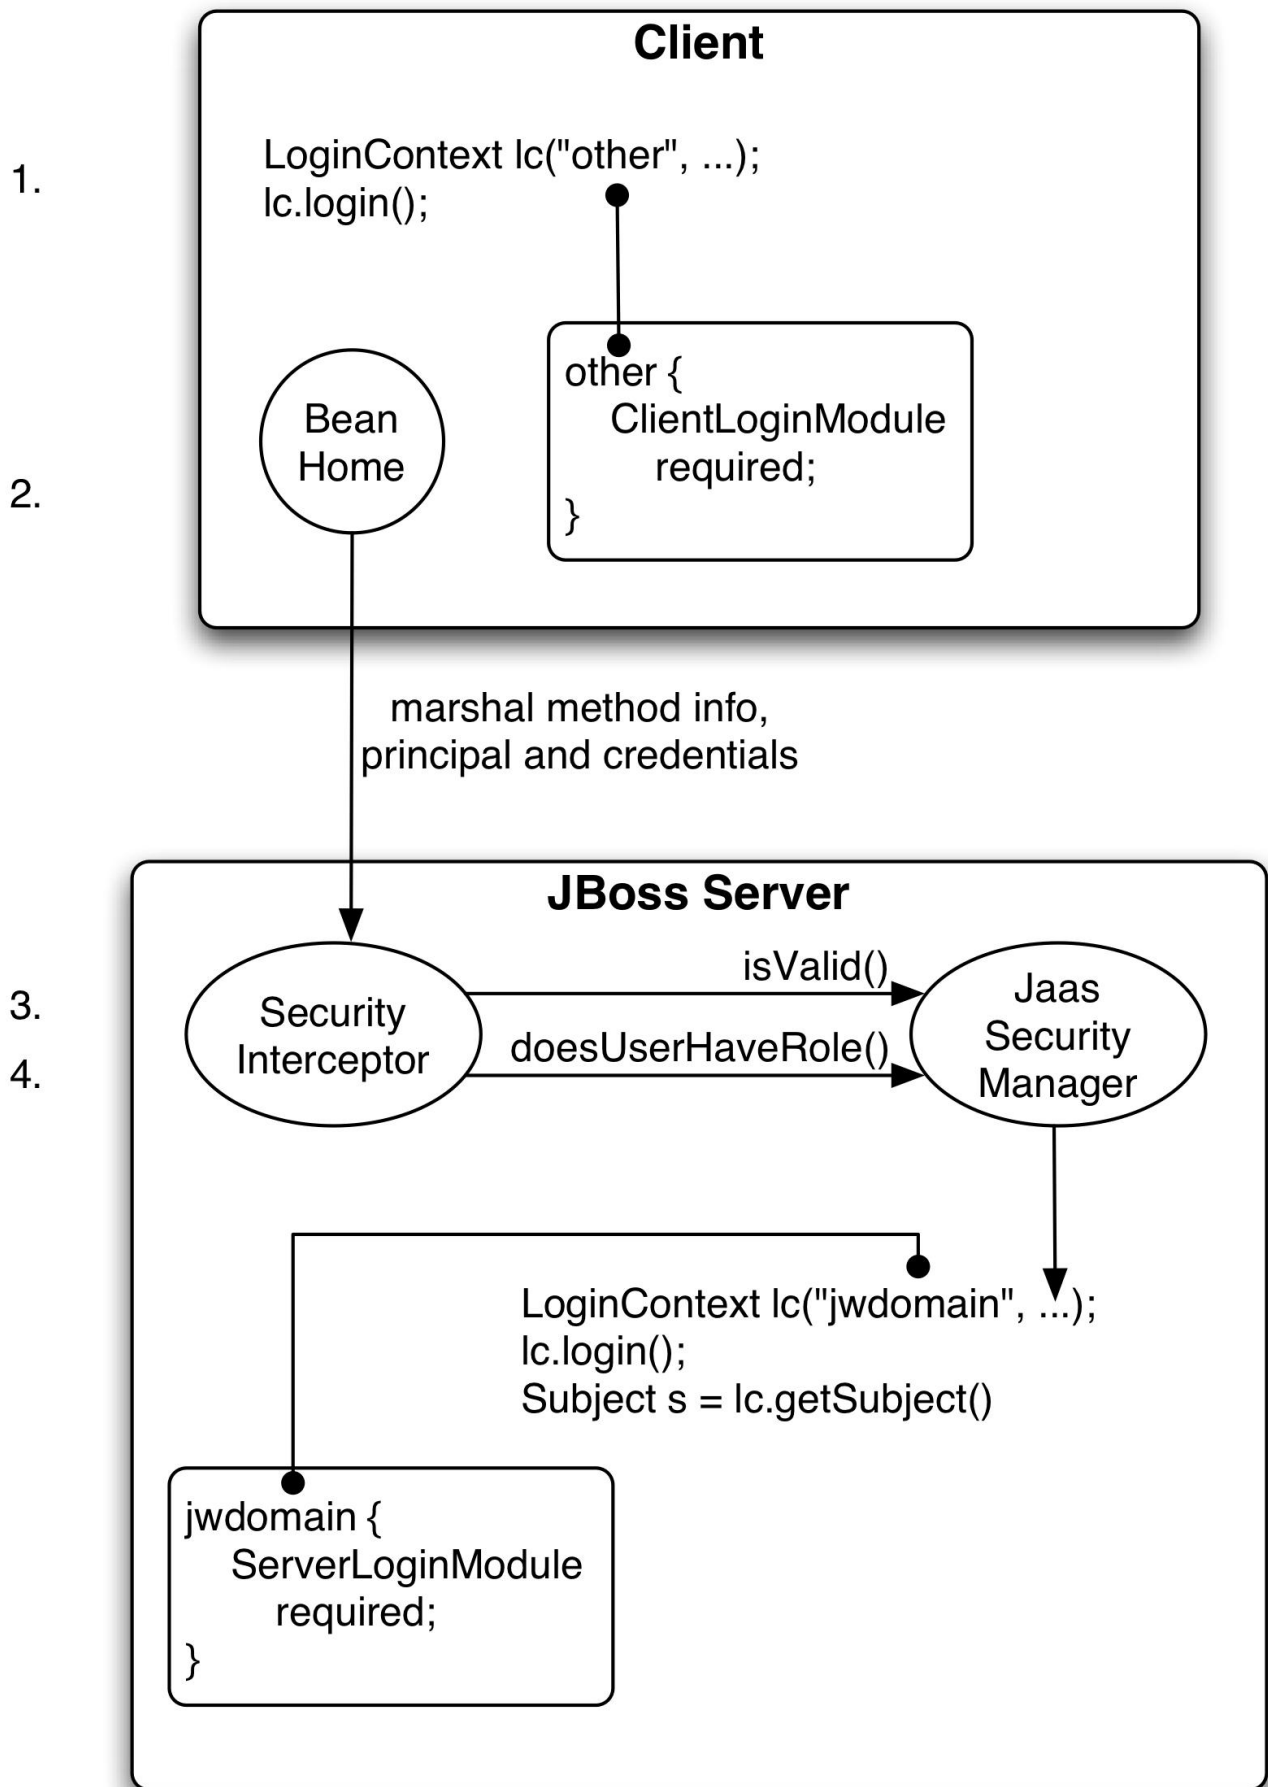


Figure 8.7. An illustration of the steps involved in the authentication and authorization of a secured EJB home method invocation.

Figure 8.7 provides a view of the client to server communication we will discuss. The numbered steps shown are:

1. The client first has to perform a JAAS login to establish the principal and credentials for authentication, and this is labeled *Client Side Login* in the figure. This is how clients establish their login identities in JBoss. Support for presenting the login information via JNDI `InitialContext` properties is provided via an alternate configuration. A JAAS login entails creating a `LoginContext` instance and passing the name of the configuration to use. The configuration name is `other`. This one-time login associates the login principal and credentials with all subsequent EJB method invocations. Note that the process might not authenticate the user. The nature of the client-side login depends on the login module configuration that the client uses. In this example, the `other` client-side login configuration entry is set up to use the `ClientLoginModule` module (an `org.jboss.security.ClientLoginModule`). This is the default client side module that simply binds the username and password to the JBoss EJB invocation layer for later authentication on the server. The identity of the client is not authenticated on the client.
2. Later, the client obtains the EJB home interface and attempts to create a bean. This event is labeled as *Home Method Invocation*. This results in a home interface method invocation being sent to the JBoss server. The invocation includes the method arguments passed by the client along with the user identity and credentials from the client-side JAAS login performed in step 1.
3. On the server side, the security interceptor first requires authentication of the user invoking the call, which, as on the client side, involves a JAAS login.
4. The security domain under which the EJB is secured determines the choice of login modules. The security domain name is used as the login configuration entry name passed to the `LoginContext` constructor. The EJB security domain is `jwdomain`. If the JAAS login authenticates the user, a JAAS `Subject` is created that contains the following in its `PrincipalsSet`:
 - A `java.security.Principal` that corresponds to the client identity as known in the deployment security environment.
 - A `java.security.acl.Group` named `Roles` that contains the role names from the application domain to which the user has been assigned. `org.jboss.security.SimplePrincipal` objects are used to represent the role names; `SimplePrincipal` is a simple string-based implementation of `Principal`. These roles are used to validate the roles assigned to methods in `ejb-jar.xml` and the `EJBContext.isCallerInRole(String)` method implementation.
 - An optional `java.security.acl.Group` named `CallerPrincipal`, which contains a single `org.jboss.security.SimplePrincipal` that corresponds to the identity of the application domain's caller. The `CallerPrincipal` sole group member will be the value returned by the `EJBContext.getCallerPrincipal()` method. The purpose of this mapping is to allow a `Principal` as known in the operational security environment to map to a `Principal` with a name known to the application. In the absence of a `CallerPrincipal` mapping the deployment security environment principal is used as the `getCallerPrincipal` method value. That is, the operational principal is the same as the application domain principal.
5. The final step of the security interceptor check is to verify that the authenticated user has permission to invoke the requested method. This is labeled as *Server Side Authorization* in Figure 8.7. Performing the authorization this entails the following steps:
 - Obtain the names of the roles allowed to access the EJB method from the EJB container. The role names are determined by `ejb-jar.xml` descriptor role-name elements of all `method-permission` ele-

ments containing the invoked method.

- If no roles have been assigned, or the method is specified in an `exclude-list` element, then access to the method is denied. Otherwise, the `doesUserHaveRole` method is invoked on the security manager by the security interceptor to see if the caller has one of the assigned role names. This method iterates through the role names and checks if the authenticated user's `Subject Roles` group contains a `SimplePrincipal` with the assigned role name. Access is allowed if any role name is a member of the `Roles` group. Access is denied if none of the role names are members.
- If the EJB was configured with a custom security proxy, the method invocation is delegated to it. If the security proxy wants to deny access to the caller, it will throw a `java.lang.SecurityException`. If no `SecurityException` is thrown, access to the EJB method is allowed and the method invocation passes to the next container interceptor. Note that the `SecurityProxyInterceptor` handles this check and this interceptor is not shown.

Every secured EJB method invocation, or secured web content access, requires the authentication and authorization of the caller because security information is handled as a stateless attribute of the request that must be presented and validated on each request. This can be an expensive operation if the JAAS login involves client-to-server communication. Because of this, the `JaasSecurityManager` supports the notion of an authentication cache that is used to store principal and credential information from previous successful logins. You can specify the authentication cache instance to use as part of the `JaasSecurityManager` configuration as you will see when the associated MBean service is discussed in following section. In the absence of any user-defined cache, a default cache that maintains credential information for a configurable period of time is used.

8.4.2. The `JaasSecurityManagerService` MBean

The `JaasSecurityManagerService` MBean service manages security managers. Although its name begins with *Jaas*, the security managers it handles need not use JAAS in their implementation. The name arose from the fact that the default security manager implementation is the `JaasSecurityManager`. The primary role of the `JaasSecurityManagerService` is to externalize the security manager implementation. You can change the security manager implementation by providing an alternate implementation of the `AuthenticationManager` and `RealmMapping` interfaces. Of course this is optional because, by default, the `JaasSecurityManager` implementation is used.

The second fundamental role of the `JaasSecurityManagerService` is to provide a JNDI `javax.naming.spi.ObjectFactory` implementation to allow for simple code-free management of the JNDI name to security manager implementation mapping. It has been mentioned that security is enabled by specifying the JNDI name of the security manager implementation via the `security-domain` deployment descriptor element. When you specify a JNDI name, there has to be an object-binding there to use. To simplify the setup of the JNDI name to security manager bindings, the `JaasSecurityManagerService` manages the association of security manager instances to names by binding a next naming system reference with itself as the JNDI Object-Factory under the name `java:/jaas`. This allows one to use a naming convention of the form `java:/jaas/XYZ` as the value for the `security-domain` element, and the security manager instance for the `XYZ` security domain will be created as needed for you. The security manager for the domain `XYZ` is created on the first lookup against the `java:/jaas/XYZ` binding by creating an instance of the class specified by the `SecurityManager-ClassName` attribute using a constructor that takes the name of the security domain. For example, consider the following container security configuration snippet:

```
<jboss>
  <!-- Configure all containers to be secured under the "hades" security domain -->
  <security-domain>java:/jaas/hades</security-domain>
  <!-- ... -->
</jboss>
```

Any lookup of the name `java:/jaas/hades` will return a security manager instance that has been associated with the security domain named `hades`. This security manager will implement the `AuthenticationManager` and `RealmMapping` security interfaces and will be of the type specified by the `JaasSecurityManagerService` `SecurityManagerClassName` attribute.

The `JaasSecurityManagerService` MBean is configured by default for use in the standard JBoss distribution, and you can often use the default configuration as is. The configurable attributes of the `JaasSecurityManagerService` include:

- **SecurityManagerClassName:** The name of the class that provides the security manager implementation. The implementation must support both the `org.jboss.security.AuthenticationManager` and `org.jboss.security.RealmMapping` interfaces. If not specified this defaults to the JAAS-based `org.jboss.security.plugins.JaasSecurityManager`.
- **CallbackHandlerClassName:** The name of the class that provides the `javax.security.auth.callback.CallbackHandler` implementation used by the `JaasSecurityManager`. You can override the handler used by the `JaasSecurityManager` if the default implementation (`org.jboss.security.auth.callback.SecurityAssociationHandler`) does not meet your needs. This is a rather deep configuration that generally should not be set unless you know what you are doing.
- **SecurityProxyFactoryClassName:** The name of the class that provides the `org.jboss.security.SecurityProxyFactory` implementation. If not specified this defaults to `org.jboss.security.SubjectSecurityProxyFactory`.
- **AuthenticationCacheJndiName:** Specifies the location of the security credential cache policy. This is first treated as an `ObjectFactory` location capable of returning `CachePolicy` instances on a per-security-domain basis. This is done by appending the name of the security domain to this name when looking up the `CachePolicy` for a domain. If this fails, the location is treated as a single `CachePolicy` for all security domains. As a default, a timed cache policy is used.
- **DefaultCacheTimeout:** Specifies the default timed cache policy timeout in seconds. The default value is 1800 seconds (30 minutes). The value you use for the timeout is a tradeoff between frequent authentication operations and how long credential information may be out of synch with respect to the security information store. If you want to disable caching of security credentials, set this to 0 to force authentication to occur every time. This has no affect if the `AuthenticationCacheJndiName` has been changed from the default value.
- **DefaultCacheResolution:** Specifies the default timed cache policy resolution in seconds. This controls the interval at which the cache current timestamp is updated and should be less than the `DefaultCacheTimeout` in order for the timeout to be meaningful. The default resolution is 60 seconds(1 minute). This has no affect if the `AuthenticationCacheJndiName` has been changed from the default value.

The `JaasSecurityManagerService` also supports a number of useful operations. These include flushing any security domain authentication cache at runtime, getting the list of active users in a security domain authentication cache, and any of the security manager interface methods.

Flushing a security domain authentication cache can be used to drop all cached credentials when the underlying store has been updated and you want the store state to be used immediately. The MBean operation signature is:
`public void flushAuthenticationCache(String securityDomain).`

This can be invoked programmatically using the following code snippet:

```
MBeanServer server = ...;
String jaasMgrName = "jboss.security:service=JaasSecurityManager";
```



```
ObjectName jaasMgr = new ObjectName(jaasMgrName);
Object[] params = {domainName};
String[] signature = {"java.lang.String"};
server.invoke(jaasMgr, "flushAuthenticationCache", params, signature);
```

Getting the list of active users provides a snapshot of the `Principals` keys in a security domain authentication cache that are not expired. The MBean operation signature is: `public List getAuthenticationCachePrincipals(String securityDomain)`.

This can be invoked programmatically using the following code snippet:

```
MBeanServer server = ...;
String jaasMgrName = "jboss.security:service=JaasSecurityManager";
ObjectName jaasMgr = new ObjectName(jaasMgrName);
Object[] params = {domainName};
String[] signature = {"java.lang.String"};
List users = (List) server.invoke(jaasMgr, "getAuthenticationCachePrincipals",
                                params, signature);
```

The security manager has a few additional access methods.

```
public boolean isValid(String securityDomain, Principal principal, Object credential);
public Principal getPrincipal(String securityDomain, Principal principal);
public boolean doesUserHaveRole(String securityDomain, Principal principal,
                                Object credential, Set roles);
public Set getUserRoles(String securityDomain, Principal principal, Object credential);
```

They provide access to the corresponding `AuthenticationManager` and `RealmMapping` interface method of the associated security domain named by the `securityDomain` argument.

8.4.3. The JaasSecurityDomain MBean

The `org.jboss.security.plugins.JaasSecurityDomain` is an extension of `JaasSecurityManager` that adds the notion of a `KeyStore`, a `JSSE KeyManagerFactory` and a `TrustManagerFactory` for supporting SSL and other cryptographic use cases. The additional configurable attributes of the `JaasSecurityDomain` include:

JKS

- **KeyStoreType:** The type of the `KeyStore` implementation. This is the type argument passed to the `java.security.KeyStore.getInstance(String type)` factory method.
- **KeyStoreURL:** A URL to the location of the `KeyStore` database. This is used to obtain an `InputStream` to initialize the `KeyStore`. If the string is not a value URL, it is treated as a file.
- **KeyStorePass:** The password associated with the `KeyStore` database contents. The `KeyStorePass` is also used in combination with the `Salt` and `IterationCount` attributes to create a PBE secret key used with the encode/decode operations. The `KeyStorePass` attribute value format is one of the following:
 - The plaintext password for the `KeyStore`. The `toCharArray()` value of the string is used without any manipulation.
 - A command to execute to obtain the plaintext password. The format is `{EXT}...` where the `...` is the exact command line that will be passed to the `Runtime.exec(String)` method to execute a platform-specific command. The first line of the command output is used as the password.
 - A class to create to obtain the plaintext password. The format is `{CLASS}classname[:ctorarg]` where

the `[:ctorarg]` is an optional string that will be passed to the constructor when instantiating the the `classname`. The password is obtained from `classname` by invoking a `toCharArray()` method if found, otherwise, the `toString()` method is used.

- **Salt:** The `PBEParameterSpec` salt value.
- **IterationCount:** The `PBEParameterSpec` iteration count value.
- **ManagerServiceName:** Sets the JMX object name string of the security manager service MBean. This is used to register the defaults to register the `JaasSecurityDomain` as a the security manager under `java:/jaas/<domain>` where `<domain>` is the name passed to the MBean constructor. The name defaults to `jboss.security:service=JaasSecurityManager`.
- **LoadSunJSSEProvider:** A flag indicating if the Sun `com.sun.net.ssl.internal.ssl.Provider` security provider should be loaded on startup. This is needed when using the Sun JSSE jars without them installed as an extension with JDK 1.3. This should be set to false with JDK 1.4 or when using an alternate JSSE provider. This flag currently defaults to true.

8.4.4. An XML JAAS Login Configuration MBean

JBoss uses a custom implementation of the `javax.security.auth.login.Configuration` class that is provided by the `org.jboss.security.auth.login.XMLLoginConfig` MBean. This configuration implementation uses an XML format that conforms to the DTD given by Figure 8.8.

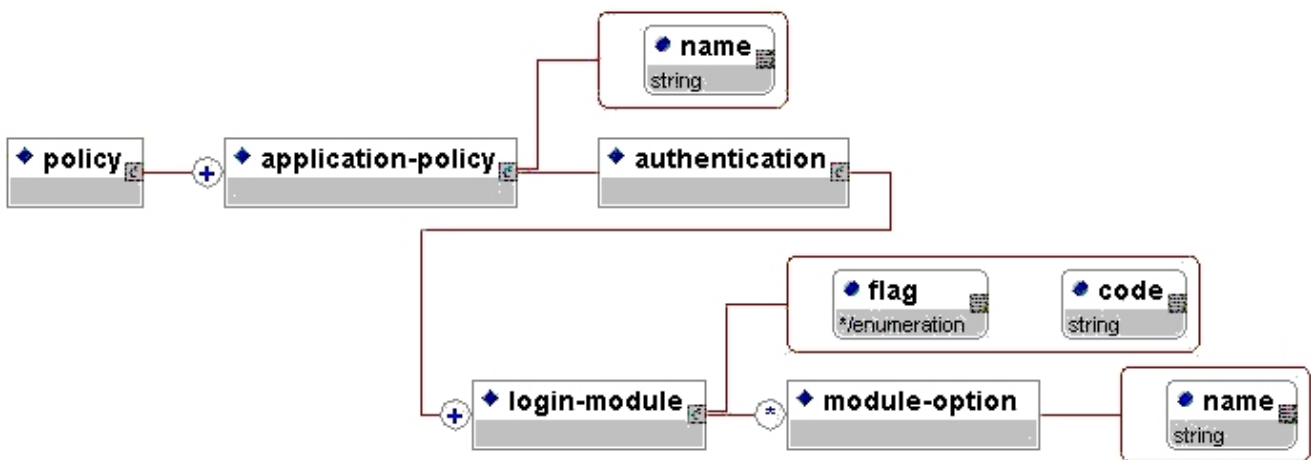


Figure 8.8. The XMLLoginConfig DTD

The `name` attribute of the `application-policy` is the login configuration name. This corresponds to the portion of the `jboss.xml` and `jboss-web.xml` `security-domain` element value after the `java:/jaas/` prefix. The `code` attribute of the `login-module` element specifies the class name of the login module implementation. The `flag` attribute controls the overall behavior of the authentication stack. The allowed values and meanings are:

- **required:** the `LoginModule` is required to succeed. If it succeeds or fails, authentication still continues to proceed down the `LoginModule` list.
- **requisite:** the `LoginModule` is required to succeed. If it succeeds, authentication continues down the `LoginModule` list. If it fails, control immediately returns to the application (authentication does not proceed down the `LoginModule` list).

- **sufficient:** the `LoginModule` is not required to succeed. If it does succeed, control immediately returns to the application (authentication does not proceed down the `LoginModule` list). If it fails, authentication continues down the `LoginModule` list.
- **optional:** the `LoginModule` is not required to succeed. If it succeeds or fails, authentication still continues to proceed down the `LoginModule` list.

Zero or more `module-option` elements may be specified as child elements of a `login-module`. These define name/value string pairs that are made available to the login module during initialization. The name attribute specifies the option name while the `module-option` body provides the value. An example login configuration is given in Example 8.11.

Example 8.11. A sample login module configuration suitable for use with XMLLoginConfig

```
<policy>
  <application-policy name="srp-test">
    <authentication>
      <login-module code="org.jboss.security.srp.jaas.SRPCacheLoginModule"
                    flag="required">
        <module-option name="cacheJndiName">srp-test/AuthenticationCache</module-option>
      </login-module>

      <login-module code="org.jboss.security.auth.spi.UsersRolesLoginModule"
                    flag="required">
        <module-option name="password-stacking">useFirstPass</module-option>
      </login-module>
    </authentication>
  </application-policy>
</policy>
```

The `XMLLoginConfig` MBean supports the following attributes:

- **ConfigURL:** specifies the URL of the XML login configuration file that should be loaded by this MBean on startup. This must be a valid URL string representation.
- **ConfigResource:** specifies the resource name of the XML login configuration file that should be loaded by this MBean on startup. The name is treated as a classpath resource for which a URL is located using the thread context class loader.
- **ValidateDTD:** a flag indicating if the XML configuration should be validated against its DTD. This defaults to true.

The MBean also supports the following operations that allow one to dynamically extend the login configurations at runtime. Note that any operation that attempts to alter login configuration requires a `javax.security.auth.AuthPermission("refreshLoginConfiguration")` when running with a security manager. The `org.jboss.chap8.service.SecurityConfig` service demonstrates how this can be used to add/remove a deployment specific security configuration dynamically.

- `void addAppConfig(String appName, AppConfigurationEntry[] entries):` this adds the given login module configuration stack to the current configuration under the given `appName`. This replaces any existing entry under that name.
- `void removeAppConfig(String appName):` this removes the login module configuration registered under the given `appName`.

- `String[] loadConfig(URL configURL)` throws `Exception`: this loads one or more login configurations from a URL representing either an XML or legacy Sun login configuration file. Note that all login configurations must be added or none will be added. It returns the names of the login configurations that were added.
- `void removeConfigs(String[] appNames)`: this removes the login configurations specified `appNames` array.
- `String displayAppConfig(String appName)`: this operation displays a simple string format of the named configuration if it exists.

8.4.5. The JAAS Login Configuration Management MBean

The installation of the custom `javax.security.auth.login.Configuration` is managed by the `org.jboss.security.plugins.SecurityConfig` MBean. There is one configurable attribute:

- **LoginConfig**: Specifies the JMX `ObjectName` string of the that provides the default JAAS login configuration. When the `SecurityConfig` is started, this mean is queried for its `javax.security.auth.login.Configuration` by calling its `getConfiguration(Configuration currentConfig)` operation. If the `LoginConfig` attribute is not specified then the default Sun Configuration implementation described in the `Configuration` class JavaDocs is used.

In addition to allowing for a custom JAAS login configuration implementation, this service allows configurations to be chained together in a stack at runtime. This allows one to push a login configuration onto the stack and latter pop it. This is a feature used by the security unit tests to install custom login configurations into a default JBoss installation. Pusing a new configuration is done using:

```
public void pushLoginConfig(String objectName) throws
    JMException, MalformedObjectNameException;
```

The `objectName` parameters specifies an MBean similar to the `LoginConfig` attribute. The current login configuration may be removed using:

```
public void popLoginConfig() throws JMException;
```

8.4.6. Using and Writing JBossSX Login Modules

The `JaasSecurityManager` implementation allows complete customization of the authentication mechanism using JAAS login module configurations. By defining the login module configuration entry that corresponds to the security domain name you have used to secure access to your J2EE components, you define the authentication mechanism and integration implementation.

The JBossSX framework includes a number of bundled login modules suitable for integration with standard security infrastructure store protocols such as LDAP and JDBC. It also includes standard base class implementations that help enforce the expected `LoginModule` to `Subject` usage pattern that was described in the Section 8.4.7. These implementations allow for easy integration of your own authentication protocol, if none of the bundled login modules prove suitable. In this section we will first describe the useful bundled login modules and their configuration, and then end with a discussion of how to create your own custom `LoginModule` implementations for use with JBoss.

8.4.6.1. `org.jboss.security.auth.spi.IdentityLoginModule`

The `IdentityLoginModule` is a simple login module that associates the principal specified in the module options with any subject authenticated against the module. It creates a `SimplePrincipal` instance using the name specified by the `principal` option. Although this is certainly not an appropriate login module for production strength authentication, it can be of use in development environments when you want to test the security associated with a given principal and associated roles.

The supported login module configuration options include:

- **principal=string:** The name to use for the `SimplePrincipal` all users are authenticated as. The principal name defaults to `guest` if no principal option is specified.
- **roles=string-list:** The names of the roles that will be assigned to the user principal. The value is a comma-delimited list of role names.
- **password-stacking=useFirstPass:** When password-stacking option is set, this module first looks for a shared username under the property name `javax.security.auth.login.name` in the login module shared state map. If found this is used as the principal name. If not found the principal name set by this login module is stored under the property name `javax.security.auth.login.name`.

A sample legacy Sun format login configuration entry that would authenticate all users as the principal named `jduke` and assign role names of `TheDuke`, and `AnimatedCharacter` is:

```
testIdentity {  
    org.jboss.security.auth.spi.IdentityLoginModule required  
    principal=jduke  
    roles=TheDuke,AnimatedCharater;  
};
```

The corresponding `XMLLoginConfig` format is:

```
<policy>  
  <application-policy name="testIdentity">  
    <authentication>  
      <login-module code="org.jboss.security.auth.spi.IdentityLoginModule"  
        flag="required">  
        <module-option name="principal">jduke</module-option>  
        <module-option name="roles">TheDuke,AnimatedCharater</module-option>  
      </login-module>  
    </authentication>  
  </application-policy>  
</policy>
```

To add this entry to a JBoss server login configuration found in the default configuration file set you would modify the `conf/default/auth.conf` file of the JBoss distribution.

8.4.6.2. `org.jboss.security.auth.spi.UsersRolesLoginModule`

The `UsersRolesLoginModule` is another simple login module that supports multiple users and user roles, and is based on two Java Properties formatted text files. The username-to-password mapping file is called `users.properties` and the username-to-roles mapping file is called `roles.properties`. The properties files are loaded during initialization using the `initialize` method thread context class loader. This means that these files can be placed into the J2EE deployment JAR, the JBoss configuration directory, or any directory on the JBoss server or system classpath. The primary purpose of this login module is to easily test the security settings of multiple users and roles using properties files deployed with the application.

The `users.properties` file uses a `username=password` format with each user entry on a separate line as show here:

```
username1=password1
username2=password2
...
```

The `roles.properties` file uses as `username=role1,role2,...` format with an optional group name value. For example:

```
username1=role1,role2,...
username1.RoleGroup1=role3,role4,...
username2=role1,role3,...
```

The `username.xxx` form of property name is used to assign the username roles to a particular named group of roles where the `xxx` portion of the property name is the group name. The `username=...` form is an abbreviation for `username.Roles=...`, where the `Roles` group name is the standard name the `JaasSecurityManager` expects to contain the roles which define the users permissions.

The following would be equivalent definitions for the `jduke` username:

```
jduke=TheDuke,AnimatedCharacter
jduke.Roles=TheDuke,AnimatedCharacter
```

The supported login module configuration options include the following:

- **unauthenticatedIdentity=name:** Defines the principal name that should be assigned to requests that contain no authentication information. This can be used to allow unprotected servlets to invoke methods on EJBs that do not require a specific role. Such a principal has no associated roles and so can only access either unsecured EJBs or EJB methods that are associated with the unchecked permission constraint.
- **password-stacking=useFirstPass:** When password-stacking option is set, this module first looks for a shared username and password under the property names `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively in the login module shared state map. If found these are used as the principal name and password. If not found the principal name and password are set by this login module and stored under the property names `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively.
- **hashAlgorithm=string:** The name of the `java.security.MessageDigest` algorithm to use to hash the password. There is no default so this option must be specified to enable hashing. When `hashAlgorithm` is specified, the clear text password obtained from the `callbackhandler` is hashed before it is passed to `UsernamePasswordLoginModule.validatePassword` as the `inputPassword` argument. The `expectedPassword` as stored in the `users.properties` file must be comparably hashed.
- **hashEncoding=base64|hex:** The string format for the hashed pass and must be either `base64` or `hex`. `Base64` is the default.
- **hashCharset=string:** The encoding used to convert the clear text password to a byte array. The platform default encoding is the default.
- **usersProperties=string:** The name of the properties resource containing the username to password mappings. This defaults to `users.properties`.
- **rolesProperties=string:** The name of the properties resource containing the username to roles mappings. This defaults to `roles.properties`.

A sample legacy Sun format login configuration entry that assigned unauthenticated users the principal name `nobody` and contains based64 encoded, MD5 hashes of the passwords in a `usersb64.properties` file is:

```
testUsersRoles {
    org.jboss.security.auth.spi.UsersRolesLoginModule required
    usersProperties=usersb64.properties
    hashAlgorithm=MD5
    hashEncoding=base64
    unauthenticatedIdentity=nobody
    ;
};
```

The corresponding XMLLoginConfig format is:

```
<policy>
  <application-policy name="testUsersRoles">
    <authentication>
      <login-module code="org.jboss.security.auth.spi.UsersRolesLoginModule"
        flag="required">
        <module-option name="usersProperties">usersb64.properties</module-option>
        <module-option name="hashAlgorithm">MD5</module-option>
        <module-option name="hashEncoding">base64</module-option>
        <module-option name="unauthenticatedIdentity">nobody</module-option>
      </login-module>
    </authentication>
  </application-policy>
</policy>
```

8.4.6.3. org.jboss.security.auth.spi.LdapLoginModule

The `LdapLoginModule` is a `LoginModule` implementation that authenticates against an LDAP server using JNDI login using the login module configuration options. You would use the `LdapLoginModule` if your username and credential information are store in an LDAP server that is accessible using a JNDI LDAP provider.

The LDAP connectivity information is provided as configuration options that are passed through to the environment object used to create JNDI initial context. The standard LDAP JNDI properties used include the following:

- **java.naming.factory.initial:** The classname of the `InitialContextFactory` implementation. This defaults to the Sun LDAP provider implementation `com.sun.jndi.ldap.LdapCtxFactory`.
- **java.naming.provider.url:** The LDAP URL for the LDAP server
- **java.naming.security.authentication:** The security level to use. This defaults to `simple`.
- **java.naming.security.protocol:** The transport protocol to use for secure access, such as, `ssl`.
- **java.naming.security.principal:** The principal for authenticating the caller to the service. This is built from other properties as described below.
- **java.naming.security.credentials:** The value of the property depends on the authentication scheme. For example, it could be a hashed password, clear-text password, key, certificate, and so on.

The supported login module configuration options include the following:

- **principalDNPrefix=string:** A prefix to add to the username to form the user distinguished name. See `principalDNSuffix` for more info.
- **principalDNSuffix=string:** A suffix to add to the username when forming the user distinguished name. This is useful if you prompt a user for a username and you don't want the user to have to enter the fully distinguished name. Using this property and `principalDNSuffix` the `userDN` will be formed as:

```
String userDN = principalDNPrefix + username + principalDNSuffix;
```

- **useObjectCredential=true|false:** Indicates that the credential should be obtained as an opaque Object using the `org.jboss.security.auth.callback.ObjectCallback` type of Callback rather than as a `char[]` password using a JAAS `PasswordCallback`. This allows for passing non-`char[]` credential information to the LDAP server.
- **rolesCtxDN=string:** The fixed distinguished name to the context to search for user roles.
- **userRolesCtxDNAttributeName=string:** The name of an attribute in the user object that contains the distinguished name to the context to search for user roles. This differs from `rolesCtxDN` in that the context to search for a user's roles can be unique for each user.
- **roleAttributeID=string:** The name of the attribute that contains the user roles. If not specified this defaults to `roles`.
- **roleAttributeIsDN=string:** A flag indicating whether the `roleAttributeID` contains the fully distinguished name of a role object, or the role name. If false, the role name is taken from the value of `roleAttributeID`. If true, the role attribute represents the distinguished name of a role object. The role name is taken from the value of the `roleNameAttributeID` attribute of the context name by the distinguished name. In certain directory schemas (e.g., MS ActiveDirectory), role attributes in the user object are stored as DNs to role objects instead of as simple names, in which case, this property should be set to true. The default is false.
- **roleNameAttributeID=string:** The name of the attribute of the in the context pointed to by the `roleCtxDN` distinguished name value which contains the role name. If the `roleAttributeIsDN` property is set to true, this property is used to find the role object's name attribute. The default is `group`.
- **uidAttributeID=string:** The name of the attribute in the object containing the user roles that corresponds to the `userid`. This is used to locate the user roles. If not specified this defaults to `uid`.
- **matchOnUserDN=true|false:** A flag indicating if the search for user roles should match on the user's fully distinguished name. If false, just the username is used as the match value against the `uidAttributeName` attribute. If true, the full `userDN` is used as the match value.
- **unauthenticatedIdentity=string:** The principal name that should be assigned to requests that contain no authentication information. This behavior is inherited from the `UsernamePasswordLoginModule` superclass.
- **password-stacking=useFirstPass:** When the password-stacking option is set, this module first looks for a shared username and password under the property names `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively in the login module shared state map. If found these are used as the principal name and password. If not found the principal name and password are set by this login module and stored under the property names `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively.
- **hashAlgorithm=string:** The name of the `java.security.MessageDigest` algorithm to use to hash the password. There is no default so this option must be specified to enable hashing. When `hashAlgorithm` is specified, the clear text password obtained from the `callbackhandler` is hashed before it is passed to `UsernamePasswordLoginModule.validatePassword` as the `inputPassword` argument. The `expectedPassword` as stored in the LDAP server must be comparably hashed.
- **hashEncoding=base64|hex:** The string format for the hashed pass and must be either `base64` or `hex`. Base64 is the default.

- **hashCharset=string:** The encoding used to convert the clear text password to a byte array. The platform default encoding is the default.
- **allowEmptyPasswords:** A flag indicating if empty (length 0) passwords should be passed to the LDAP server. An empty password is treated as an anonymous login by some LDAP servers and this may not be a desirable feature. Set this to false to reject empty passwords, true to have the LDAP server validate the empty password. The default is true.

The authentication of a user is performed by connecting to the LDAP server based on the login module configuration options. Connecting to the LDAP server is done by creating an `InitialLdapContext` with an environment composed of the LDAP JNDI properties described previously in this section. The `Context.SECURITY_PRINCIPAL` is set to the distinguished name of the user as obtained by the callback handler in combination with the `principalDNPrefix` and `principalDNSuffix` option values, and the `Context.SECURITY_CREDENTIALS` property is either set to the `String` password or the `Object` credential depending on the `useObjectCredential` option.

Once authentication has succeeded by virtue of being able to create an `InitialLdapContext` instance, the user's roles are queried by performing a search on the `rolesCtxDN` location with search attributes set to the `roleAttributeName` and `uidAttributeName` option values. The roles names are obtaining by invoking the `toString` method on the role attributes in the search result set.

A sample Sun legacy format login configuration entry is:

```
testLdap {
    org.jboss.security.auth.spi.LdapLoginModule required
    java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory
    java.naming.provider.url="ldap://ldaphost.jboss.org:1389/"
    java.naming.security.authentication=simple
    principalDNPrefix=uid=
    uidAttributeID=userid
    roleAttributeID=roleName
    principalDNSuffix=,ou=People,o=jboss.org
    rolesCtxDN=cn=JBossSX Tests,ou=Roles,o=jboss.org
};
```

The corresponding XMLLoginConfig format is:

```
<policy>
  <application-policy name="testLdap">
    <authentication>
      <login-module code="org.jboss.security.auth.spi.LdapLoginModule"
        flag="required">
        <module-option name="java.naming.factory.initial">
          com.sun.jndi.ldap.LdapCtxFactory
        </module-option>
        <module-option name="java.naming.provider.url">
          ldap://ldaphost.jboss.org:1389/
        </module-option>
        <module-option name="java.naming.security.authentication">
          simple
        </module-option>
        <module-option name="principalDNPrefix">uid=</module-option>
        <module-option name="uidAttributeID">userid</module-option>
        <module-option name="roleAttributeID">roleName</module-option>
        <module-option name="principalDNSuffix">,ou=People,o=jboss.org
        </module-option>
        <module-option name="rolesCtxDN">cn=JBossSX Tests,ou=Roles,o=jboss.org
        </module-option>
      </login-module>
    </authentication>
  </application-policy>
</policy>
```

To help you understand all of the options of the `LdapLoginModule`, consider the sample LDAP server data shown in Figure 8.9. This figure corresponds to the `testLdap` login configuration just shown.

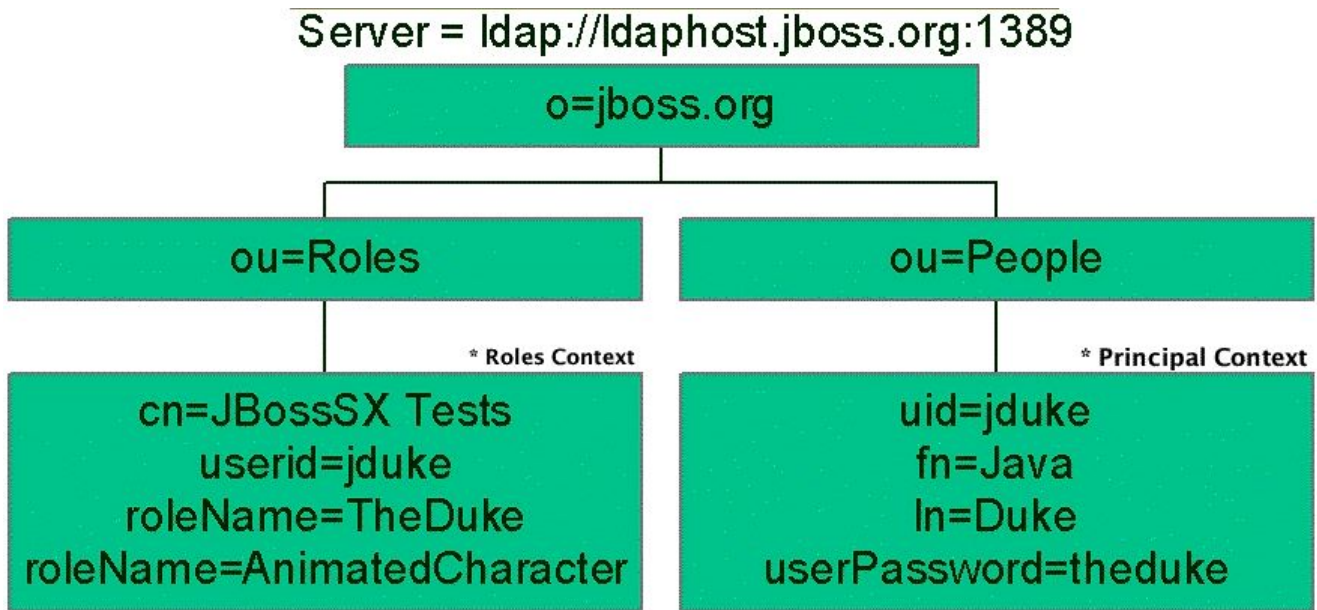


Figure 8.9. An LDAP server configuration compatible with the `testLdap` sample configuration.

Take a look at the `testLdap` login module configuration in comparison to the Figure 8.9 schema. The `java.naming.factory.initial`, `java.naming.factory.url` and `java.naming.security` options indicate the Sun LDAP JNDI provider implementation will be used, the LDAP server is located on host `ldaphost.jboss.org` on port 1389, and that simple username and password will be used to authenticate clients connecting to the LDAP server.

When the `LdapLoginModule` performs authentication of a user, it does so by connecting to the LDAP server specified by the `java.naming.factory.url`. The `java.naming.security.principal` property is built from the `principalDNPrefix`, passed in `username` and `principalDNSuffix` as described above. For the `testLdap` configuration example and a username of `jduke`, the `java.naming.security.principal` string would be `uid=jduke,ou=People,o=jboss.org`. This corresponds to the LDAP context on the lower right of Figure 8.9 labeled as *Principal Context*. The `java.naming.security.credentials` property would be set to the passed in password and it would have to match the `userPassword` attribute of the principal context. How a secured LDAP context stores the authentication credential information depends on the LDAP server, so your LDAP server may handle the validation of the `java.naming.security.credentials` property differently.

Once authentication succeeds, the roles on which authorization will be based are retrieved by performing a JNDI search of the LDAP context whose distinguished name is given by the `rolesCtxDN` option value. For the `testLdap` configuration this is `cn=JBossSX Tests,ou=Roles,o=jboss.org` and corresponds to the LDAP context on the lower left of Figure 8.9 labeled *Roles Context*. The search attempts to locate any subcontexts that contain an attribute whose name is given by the `uidAttributeID` option, and whose value matches the username passed to the login module. For any matching context, all values of the attribute whose name is given by the `roleAttributeID` option are obtained. For the `testLdap` configuration the attribute name that contains the roles is called `roleName`. The resulting `roleName` values are stored in the JAAS Subject associated with the `LdapLoginModule` as the Roles group principals that will be used for role-based authorization. For the LDAP schema shown in Figure 8.9, the roles that will be assigned to the user `jduke` are `TheDuke` and `AnimatedCharacter`.

8.4.6.4. `org.jboss.security.auth.spi.DatabaseServerLoginModule`

The `DatabaseServerLoginModule` is a JDBC based login module that supports authentication and role mapping. You would use this login module if you have your username, password and role information in a JDBC accessible database. The `DatabaseServerLoginModule` is based on two logical tables:

```
Table Principals(PrincipalID text, Password text)
Table Roles(PrincipalID text, Role text, RoleGroup text)
```

The `Principals` table associates the user `PrincipalID` with the valid password and the `Roles` table associates the user `PrincipalID` with its role sets. The roles used for user permissions must be contained in rows with a `RoleGroup` column value of `Roles`. The tables are logical in that you can specify the SQL query that the login module uses. All that is required is that the `java.sql.ResultSet` has the same logical structure as the `Principals` and `Roles` tables described previously. The actual names of the tables and columns are not relevant as the results are accessed based on the column index. To clarify this notion, consider a database with two tables, `Principals` and `Roles`, as already declared. The following statements build the tables to contain a `PrincipalID java` with a `Password` of `echoman` in the `Principals` table, a `PrincipalID java` with a role named `Echo` in the `Roles` `RoleGroup` in the `Roles` table, and a `PrincipalID java` with a role named `caller_java` in the `CallerPrincipal` `RoleGroup` in the `Roles` table:

```
INSERT INTO Principals VALUES('java', 'echoman')
INSERT INTO Roles VALUES('java', 'Echo', 'Roles')
INSERT INTO Roles VALUES('java', 'caller_java', 'CallerPrincipal')
```

The supported login module configuration options include the following:

- **dsJndiName:** The JNDI name for the `DataSource` of the database containing the logical `Principals` and `Roles` tables. If not specified this defaults to `java:/DefaultDS`.
- **principalsQuery:** The prepared statement query equivalent to: `select Password from Principals where PrincipalID=?`. If not specified this is the exact prepared statement that will be used.
- **rolesQuery:** The prepared statement query equivalent to: `select Role, RoleGroup from Roles where PrincipalID=?`. If not specified this is the exact prepared statement that will be used.
- **unauthenticatedIdentity=string:** The principal name that should be assigned to requests that contain no authentication information.
- **password-stacking=useFirstPass:** When `password-stacking` option is set, this module first looks for a shared username and password under the property names `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively in the login module shared state map. If found these are used as the principal name and password. If not found the principal name and password are set by this login module and stored under the property names `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively.
- **hashAlgorithm=string:** The name of the `java.security.MessageDigest` algorithm to use to hash the password. There is no default so this option must be specified to enable hashing. When `hashAlgorithm` is specified, the clear text password obtained from the `callbackhandler` is hashed before it is passed to `UserNamePasswordLoginModule.validatePassword` as the `inputPassword` argument. The `expectedPassword` as obtained from the database must be comparably hashed.
- **hashEncoding=base64|hex:** The string format for the hashed pass and must be either `base64` or `hex`. `Base64` is the default.
- **hashCharset=string:** The encoding used to convert the clear text password to a byte array. The platform default encoding is the default.

- **ignorePasswordCase=true|false:** A boolean flag indicating if the password comparison should ignore case. This can be useful for hashed password encoding where the case of the hashed password is not significant.
- **principalClass:** An option that specifies a Principal implementation class. This must support a constructor taking a string argument for the principal name.

As an example `DatabaseServerLoginModule` configuration, consider a custom table schema like the following:

```
CREATE TABLE Users(username VARCHAR(64) PRIMARY KEY, passwd VARCHAR(64))
CREATE TABLE UserRoles(username VARCHAR(64), userRoles VARCHAR(32))
```

A sample Sun legacy format corresponding `DatabaseServerLoginModule` configuration would be:

```
testDB {
    org.jboss.security.auth.spi.DatabaseServerLoginModule required
    dsJndiName="java:/MyDatabaseDS"
    principalsQuery="select passwd from Users username where username=?"
    rolesQuery="select userRoles, 'Roles' from UserRoles where username=?"
    ;
};
```

The corresponding `XMLLoginConfig` format is:

```
<policy>
  <application-policy name="testDB">
    <authentication>
      <login-module code="org.jboss.security.auth.spi.DatabaseServerLoginModule"
        flag="required">
        <module-option name="dsJndiName">java:/MyDatabaseDS</module-option>
        <module-option name="principalsQuery">
          select passwd from Users username where username=?</module-option>
        <module-option name="rolesQuery">
          select userRoles, 'Roles' from UserRoles where username=?</module-option>
        </login-module>
      </authentication>
    </application-policy>
  </policy>
```

8.4.6.5. BaseCertLoginModule

This is a login module which authenticates users based on X509 certificates. A typical usecase for this login module is *CLIENT-CERT* authentication in the web tier. This login module only performs authentication. You need to combine it with another login module capable of acquiring the authorization roles to completely define access to a secured web or EJB component. Two subclasses of this login module, *CertRolesLoginModule* and *DatabaseCertLoginModule* extend the behavior to obtain the authorization roles from either a properties file or database.

The *BaseCertLoginModule* needs a *KeyStore* to perform user validation. This is obtained through a *org.jboss.security.SecurityDomain* implementation. Typically, the *SecurityDomain* implementation is configured using the *org.jboss.security.plugins.JaasSecurityDomain* MBean as shown in this *jboss-service.xml* configuration fragment:

```
<mbean code="org.jboss.security.plugins.JaasSecurityDomain"
  name="jboss.web:service=SecurityDomain">
  <constructor>
    <arg type="java.lang.String" value="jmx-console"/>
  </constructor>
  <attribute name="KeyStoreURL">resource:localhost.keystore</attribute>
  <attribute name="KeyStorePass">unit-tests-server</attribute>
</mbean>
```

This creates a security domain with the name *jmx-console* whose *SecurityDomain* implementation is available via JNDI under the name *java:/jaas/jmx-console* following the JBossSX security domain naming pattern. To secure a web application such as the *jmx-console.war* using client certs and role based authorization, one would first modify the *web.xml* to declare the resources to be secured, along with the allowed roles and security domain to be used for authentication and authorization.

```
<?xml version="1.0"?>
<!DOCTYPE web-app PUBLIC
    "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
    "http://java.sun.com/dtd/web-app_2_3.dtd">

<web-app>
    ...
    <security-constraint>
        <web-resource-collection>
            <web-resource-name>HtmlAdaptor</web-resource-name>
            <description>An example security config that only allows users with
                the role JBossAdmin to access the HTML JMX console web
                application </description>
            <url-pattern>/*</url-pattern>
            <http-method>GET</http-method>
            <http-method>POST</http-method>
        </web-resource-collection>
        <auth-constraint>
            <role-name>JBossAdmin</role-name>
        </auth-constraint>
    </security-constraint>
    <login-config>
        <auth-method>CLIENT-CERT</auth-method>
        <realm-name>JBoss JMX Console</realm-name>
    </login-config>
    <security-role>
        <role-name>JBossAdmin</role-name>
    </security-role>
</web-app>
```

Next we, need to specify the JBoss security domain in *jboss-web.xml*:

```
<jboss-web>
    <security-domain>java:/jaas/jmx-console</security-domain>
</jboss-web>
```

Finally, you need to define the login module configuration for the *jmx-console* security domain you just specified. This is done in the *conf/login-config.xml* file.

```
<application-policy name="jmx-console">
    <authentication>
        <login-module code="org.jboss.security.auth.spi.BaseCertLoginModule"
            flag="required">
            <module-option name="password-stacking">useFirstPass</module-option>
            <module-option name="securityDomain">java:/jaas/jmx-console</module-option>
        </login-module>
        <login-module code="org.jboss.security.auth.spi.UsersRolesLoginModule"
            flag="required">
            <module-option name="password-stacking">useFirstPass</module-option>
            <module-option name="usersProperties">jmx-console-users.properties</module-option>
            <module-option name="rolesProperties">jmx-console-roles.properties</module-option>
        </login-module>
    </authentication>
</application-policy>
```

Here the *BaseCertLoginModule* is used for authentication of the client cert, and the *UsersRolesLoginModule* is only used for authorization due to the *password-stacking=useFirstPass* option. Both the *local-host.keystore* and the *jmx-console-roles.properties* need an entry that maps to the principal associated with the client cert. By default, the principal is created using the client certificate distinguished name. Consider

the following certificate:

```
[starksm@banshee9100 conf]$ keytool -printcert -file unit-tests-client.export
Owner: CN=unit-tests-client, OU=JBoss Inc., O=JBoss Inc., ST=Washington, C=US
Issuer: CN=jboss.com, C=US, ST=Washington, L=Snoqualmie Pass, EMAILADDRESS=admin
@jboss.com, OU=QA, O=JBoss Inc.
Serial number: 100103
Valid from: Wed May 26 07:34:34 PDT 2004 until: Thu May 26 07:34:34 PDT 2005
Certificate fingerprints:
    MD5:  4A:9C:2B:CD:1B:50:AA:85:DD:89:F6:1D:F5:AF:9E:AB
    SHA1: DE:DE:86:59:05:6C:00:E8:CC:C0:16:D3:C2:68:BF:95:B8:83:E9:58
```

The `localhost.keystore` would need this cert stored with an alias of `CN=unit-tests-client, OU=JBoss Inc., O=JBoss Inc., ST=Washington, C=US` and the `jmx-console-roles.properties` would also need an entry for the same entry. Since the DN contains many characters that are normally treated as delimiters, you will need to escape the problem characters using a backslash (`\`) as shown here:

```
# A sample roles.properties file for use with the UsersRolesLoginModule
CN\=unit-tests-client,\ OU\=JBoss\ Inc.,\ O\=JBoss\ Inc.,\ ST\=Washington,\ C\=US=JBossAdmin
admin=JBossAdmin
```

8.4.6.6. org.jboss.security.auth.spi.ProxyLoginModule

The `ProxyLoginModule` is a login module that loads a delegate `LoginModule` using the current thread context class loader. The purpose of this module is to work around the current JAAS 1.0 class loader limitation that requires `LoginModules` to be on the system classpath². Some custom `LoginModules` use classes that are loaded from the JBoss server `lib/ext` directory and these are not available if the `LoginModule` is placed on the system classpath. To work around this limitation you use the `ProxyLoginModule` to bootstrap the custom `LoginModule`. The `ProxyLoginModule` has one required configuration option called `moduleName`. It specifies the fully qualified class name of the `LoginModule` implementation that is to be bootstrapped. Any number of additional configuration options may be specified, and they will be passed to the bootstrapped login module.

As an example, consider a custom login module that makes use of some service that is loaded from the JBoss `lib/ext` directory. The class name of the custom login module is `com.biz.CustomServiceLoginModule`. A suitable Sun legacy format `ProxyLoginModule` configuration entry for bootstrapping this custom login module would be:

```
testProxy {
    org.jboss.security.auth.spi.ProxyLoginModule required
    moduleName=com.biz.CustomServiceLoginModule
    customOption1=value1
    customOption2=value2
    customOption3=value3;
};
```

The corresponding `XMLLoginConfig` format is:

```
<policy>
  <application-policy name="testProxy">
    <authentication>
      <login-module code="org.jboss.security.auth.spi.ProxyLoginModule"
        flag="required">
        <module-option name="moduleName">
          com.biz.CustomServiceLoginModule
        </module-option>
        <module-option name="customOption1">value1</module-option>
        <module-option name="customOption2">value2</module-option>
        <module-option name="customOption3">value3</module-option>
      </login-module>
    </authentication>
  </application-policy>
</policy>
```

² The `ProxyLoginModule` is generally not needed in JBoss since we have our own VFS implementation that solves this issue, and the JDK 1.4 JAAS implementation behaves in the same way. The `ProxyLoginModule` remains for backward compatibility.

```
</application-policy>
</policy>
```

8.4.6.7. org.jboss.security.auth.spi.RunAsLoginModule

JBoss has a helper login module called `RunAsLoginModule` that pushes a run as role for the duration of the login phase of authentication, and pops the run as role in either the commit or abort phase. The purpose of this login module is to provide a role for other login modules that need to access secured resources in order to perform their authentication. An example would be a login module that accesses an secured EJB. This login module must be configured ahead of the login module(s) that need a run as role established.

The only login module configuration option is:

- **roleName:** the name of the role to use as the run as role during login phase. If not specified a default of `nobody` is used.

8.4.6.8. org.jboss.security.ClientLoginModule

The `ClientLoginModule` is an implementation of `LoginModule` for use by JBoss clients for the establishment of the caller identity and credentials. This simply sets the `org.jboss.security.SecurityAssociation.principal` to the value of the `NameCallback` filled in by the `callbackhandler`, and the `org.jboss.security.SecurityAssociation.credential` to the value of the `PasswordCallback` filled in by the `callbackhandler`. This is the only supported mechanism for a client to establish the current thread's caller. Both stand-alone client applications and server environments, acting as JBoss EJB clients where the security environment has not been configured to use JBossSX transparently, need to use the `ClientLoginModule`. Of course, you could always set the `org.jboss.security.SecurityAssociation` information directly, but this is considered an internal API that is subject to change without notice.

Note that this login module does not perform any authentication. It merely copies the login information provided to it into the JBoss server EJB invocation layer for subsequent authentication on the server. If you need to perform client-side authentication of users you would need to configure another login module in addition to the `ClientLoginModule`.

The supported login module configuration options include the following:

- **multi-threaded=true|false:** When the multi-threaded option is set to true, each login thread has its own principal and credential storage. This is useful in client environments where multiple user identities are active in separate threads. When true, each separate thread must perform its own login. When set to false the login identity and credentials are global variables that apply to all threads in the VM. The default for this option is false.
- **password-stacking=useFirstPass:** When `password-stacking` option is set, this module first looks for a shared username and password using `javax.security.auth.login.name` and `javax.security.auth.login.password` respectively in the login module shared state map. This allows a module configured prior to this one to establish a valid username and password that should be passed to JBoss. You would use this option if you want to perform client-side authentication of clients using some other login module such as the `LdapLoginModule`.
- **restore-login-identity=[true|false]:** When `restore-login-identity` is true, the `SecurityAssociation` principal and credential seen on entry to the `login()` method are saved and restored on either abort or logout. When false (the default), the abort and logout simply clear the `SecurityAssociation`. A `restore-login-identity` of true is needed if one need to change identities and then restore the original caller identity.

A sample login configuration for `ClientLoginModule` is the default configuration entry found in the JBoss distribution `client/auth.conf` file. The configuration is:

```
other {
    // Put your login modules that work without jBoss here

    // jBoss LoginModule
    org.jboss.security.ClientLoginModule required;

    // Put your login modules that need jBoss here
};
```

8.4.7. Writing Custom Login Modules

If the login modules bundled with the JBossSX framework do not work with your security environment, you can write your own custom login module implementation that does.

Recall from the section on the `JaasSecurityManager` architecture that the `JaasSecurityManager` expected a particular usage pattern of the `Subject` principals set. You need to understand the JAAS `Subject` class's information storage features and the expected usage of these features to be able to write a login module that works with the `JaasSecurityManager`. This section examines this requirement and introduces two abstract base `LoginModule` implementations that can help you implement your own custom login modules.

You can obtain security information associated with a `Subject` in six ways using the following methods:

```
java.util.Set getPrincipals()
java.util.Set getPrincipals(java.lang.Class c)
java.util.Set getPrivateCredentials()
java.util.Set getPrivateCredentials(java.lang.Class c)
java.util.Set getPublicCredentials()
java.util.Set getPublicCredentials(java.lang.Class c)
```

For `Subject` identities and roles, JBossSX has selected the most natural choice: the principals sets obtained via `getPrincipals()` and `getPrincipals(java.lang.Class)`. The usage pattern is as follows:

- User identities (username, social security number, employee ID, and so on) are stored as `java.security.Principal` objects in the `SubjectPrincipals` set. The `Principal` implementation that represents the user identity must base comparisons and equality on the name of the principal. A suitable implementation is available as the `org.jboss.security.SimplePrincipal` class. Other `Principal` instances may be added to the `SubjectPrincipals` set as needed.
- The assigned user roles are also stored in the `Principals` set, but they are grouped in named role sets using `java.security.acl.Group` instances. The `Group` interface defines a collection of `Principals` and/or `Groups`, and is a subinterface of `java.security.Principal`. Any number of role sets can be assigned to a `Subject`. Currently, the JBossSX framework uses two well-known role sets with the names `Roles` and `CallerPrincipal`. The `Roles Group` is the collection of `Principals` for the named roles as known in the application domain under which the `Subject` has been authenticated. This role set is used by methods like the `EJBContext.isCallerInRole(String)`, which EJBs can use to see if the current caller belongs to the named application domain role. The security interceptor logic that performs method permission checks also uses this role set. The `CallerPrincipalGroup` consists of the single `Principal` identity assigned to the user in the application domain. The `EJBContext.getCallerPrincipal()` method uses the `CallerPrincipal` to allow the application domain to map from the operation environment identity to a user identity suitable for the application. If a `Subject` does not have a `CallerPrincipalGroup`, the application identity is the same as operational environment identity.

8.4.7.1. Support for the Subject Usage Pattern

To simplify correct implementation of the Subject usage patterns described in the preceding section, JBossSX includes two abstract login modules that handle the population of the authenticated Subject with a template pattern that enforces correct Subject usage. The most generic of the two is the `org.jboss.security.auth.spi.AbstractServerLoginModule` class. It provides a concrete implementation of the `javax.security.auth.spi.LoginModule` interface and offers abstract methods for the key tasks specific to an operation environment security infrastructure. The key details of the class are highlighted in the following class fragment. The JavaDoc comments detail the responsibilities of subclasses.

```
package org.jboss.security.auth.spi;
/**
 * This class implements the common functionality required for a JAAS
 * server-side LoginModule and implements the JBossSX standard
 * Subject usage pattern of storing identities and roles. Subclass
 * this module to create your own custom LoginModule and override the
 * login(), getRoleSets(), and getIdentity() methods.
 */
public abstract class AbstractServerLoginModule
    implements javax.security.auth.spi.LoginModule
{
    protected Subject subject;
    protected CallbackHandler callbackHandler;
    protected Map sharedState;
    protected Map options;
    protected Logger log;

    /** Flag indicating if the shared credential should be used */
    protected boolean useFirstPass;
    /**
     * Flag indicating if the login phase succeeded. Subclasses that
     * override the login method must set this to true on successful
     * completion of login
     */
    protected boolean loginOk;

    // ...
    /**
     * Initialize the login module. This stores the subject,
     * callbackHandler and sharedState and options for the login
     * session. Subclasses should override if they need to process
     * their own options. A call to super.initialize(...) must be
     * made in the case of an override.
     *
     * <p>
     * The options are checked for the <em>password-stacking</em> parameter.
     * If this is set to "useFirstPass", the login identity will be taken from the
     * <code>javax.security.auth.login.name</code> value of the sharedState map,
     * and the proof of identity from the
     * <code>javax.security.auth.login.password</code> value of the sharedState map.
     *
     * @param subject the Subject to update after a successful login.
     * @param callbackHandler the CallbackHandler that will be used to obtain the
     * the user identity and credentials.
     * @param sharedState a Map shared between all configured login module instances
     * @param options the parameters passed to the login module.
     */
    public void initialize(Subject subject,
                          CallbackHandler callbackHandler,
                          Map sharedState,
                          Map options)
    {
        // ...
    }

    /**
```

```

    * Looks for javax.security.auth.login.name and
    * javax.security.auth.login.password values in the sharedState
    * map if the useFirstPass option was true and returns true if
    * they exist. If they do not or are null this method returns
    * false.
    * Note that subclasses that override the login method
    * must set the loginOk ivar to true if the login succeeds in
    * order for the commit phase to populate the Subject. This
    * implementation sets loginOk to true if the login() method
    * returns true, otherwise, it sets loginOk to false.
    */
    public boolean login()
        throws LoginException
    {
        // ...
    }

    /**
     * Overridden by subclasses to return the Principal that
     * corresponds to the user primary identity.
     */
    abstract protected Principal getIdentity();

    /**
     * Overridden by subclasses to return the Groups that correspond
     * to the role sets assigned to the user. Subclasses should
     * create at least a Group named "Roles" that contains the roles
     * assigned to the user. A second common group is
     * "CallerPrincipal," which provides the application identity of
     * the user rather than the security domain identity.
     *
     * @return Group[] containing the sets of roles
     */
    abstract protected Group[] getRoleSets() throws LoginException;
}

```

You'll need to pay attention to the `loginOk` instance variable. This must be set to true if the login succeeds, false otherwise by any subclasses that override the login method. Failure to set this variable correctly will result in the commit method either not updating the `Subject` when it should, or updating the `Subject` when it should not. Tracking the outcome of the login phase was added to allow login module to be chained together with control flags that do not require that the login module succeed in order for the overall login to succeed.

The second abstract base login module suitable for custom login modules is the `org.jboss.security.auth.spi.UsernamePasswordLoginModule`. The login module further simplifies custom login module implementation by enforcing a string-based username as the user identity and a `char[]` password as the authentication credential. It also supports the mapping of anonymous users (indicated by a null username and password) to a `Principal` with no roles. The key details of the class are highlighted in the following class fragment. The JavaDoc comments detail the responsibilities of subclasses.

```

package org.jboss.security.auth.spi;

/**
 * An abstract subclass of AbstractServerLoginModule that imposes a
 * an identity == String username, credentials == String password
 * view on the login process. Subclasses override the
 * getUsersPassword() and getUsersRoles() methods to return the
 * expected password and roles for the user.
 */
public abstract class UsernamePasswordLoginModule
    extends AbstractServerLoginModule
{
    /** The login identity */
    private Principal identity;
    /** The proof of login identity */
    private char[] credential;
    /** The principal to use when a null username and password are seen */
}

```

```

private Principal unauthenticatedIdentity;

/**
 * The message digest algorithm used to hash passwords. If null then
 * plain passwords will be used. */
private String hashAlgorithm = null;

/**
 * The name of the charset/encoding to use when converting the
 * password String to a byte array. Default is the platform's
 * default encoding.
 */
private String hashCharset = null;

/** The string encoding format to use. Defaults to base64. */
private String hashEncoding = null;

// ...

/**
 * Override the superclass method to look for an
 * unauthenticatedIdentity property. This method first invokes
 * the super version.
 *
 * @param options,
 * @option unauthenticatedIdentity: the name of the principal to
 * assign and authenticate when a null username and password are
 * seen.
 */
public void initialize(Subject subject,
                      CallbackHandler callbackHandler,
                      Map sharedState,
                      Map options)
{
    super.initialize(subject, callbackHandler, sharedState,
                    options);
    // Check for unauthenticatedIdentity option.
    Object option = options.get("unauthenticatedIdentity");
    String name = (String) option;
    if (name != null) {
        unauthenticatedIdentity = new SimplePrincipal(name);
    }
}

// ...

/**
 * A hook that allows subclasses to change the validation of the
 * input password against the expected password. This version
 * checks that neither inputPassword or expectedPassword are null
 * and that inputPassword.equals(expectedPassword) is true;
 *
 * @return true if the inputPassword is valid, false otherwise.
 */
protected boolean validatePassword(String inputPassword,
                                   String expectedPassword)
{
    {
        if (inputPassword == null || expectedPassword == null) {
            return false;
        }
        return inputPassword.equals(expectedPassword);
    }
}

/**
 * Get the expected password for the current username available
 * via the getUsername() method. This is called from within the
 * login() method after the CallbackHandler has returned the
 * username and candidate password.
 *
 * @return the valid password String
 */

```

```
abstract protected String getUsersPassword()  
    throws LoginException;  
}
```

The choice of subclassing the `AbstractServerLoginModule` versus `UsernamePasswordLoginModule` is simply based on whether a `String` based username and `String` credential are usable for the authentication technology you are writing the login module for. If the string based semantic is valid, then subclass `UsernamePasswordLoginModule`, else subclass `AbstractServerLoginModule`.

The steps you are required to perform when writing a custom login module are summarized in the following depending on which base login module class you choose. When writing a custom login module that integrates with your security infrastructure, you should start by subclassing `AbstractServerLoginModule` or `UsernamePasswordLoginModule` to ensure that your login module provides the authenticated `Principal` information in the form expected by the JBossSX security manager.

When subclassing the `AbstractServerLoginModule`, you need to override the following:

- `void initialize(Subject, CallbackHandler, Map, Map):` if you have custom options to parse.
- `boolean login():` to perform the authentication activity. Be sure to set the `loginOk` instance variable to true if login succeeds, false if it fails.
- `Principal getIdentity():` to return the `Principal` object for the user authenticated by the `log()` step.
- `Group[] getRoleSets():` to return at least one `Group` named `Roles` that contains the roles assigned to the `Principal` authenticated during `login()`. A second common `Group` is named `CallerPrincipal` and provides the user's application identity rather than the security domain identity.

When subclassing the `UsernamePasswordLoginModule`, you need to override the following:

- `void initialize(Subject, CallbackHandler, Map, Map):` if you have custom options to parse.
- `Group[] getRoleSets():` to return at least one `Group` named `Roles` that contains the roles assigned to the `Principal` authenticated during `login()`. A second common `Group` is named `CallerPrincipal` and provides the user's application identity rather than the security domain identity.
- `String getUsersPassword():` to return the expected password for the current username available via the `getUsername()` method. The `getUsersPassword()` method is called from within `login()` after the callbackhandler returns the username and candidate password.

8.4.7.2. A Custom LoginModule Example

In this section we will develop a custom login module example. It will extend the `UsernamePasswordLoginModule` and obtains a user's password and role names from a JNDI lookup. The idea is that there is a JNDI context that will return a user's password if you perform a lookup on the context using a name of the form `password/<username>` where `<username>` is the current user being authenticated. Similarly, a lookup of the form `roles/<username>` returns the requested user's roles.

The source code for the example is located in the `src/main/org/jboss/chap8/ex2` directory of the book examples. Example 8.12 shows the source code for the `JndiUserAndPass` custom login module. Note that because this extends the JBoss `UsernamePasswordLoginModule`, all the `JndiUserAndPass` does is obtain the user's password and roles from the JNDI store. The `JndiUserAndPass` does not concern itself with the JAAS `LoginModule` operations.

Example 8.12. A JndiUserAndPass custom login module

```

package org.jboss.chap8.ex2;

import java.security.acl.Group;
import java.util.Map;
import javax.naming.InitialContext;
import javax.naming.NamingException;
import javax.security.auth.Subject;
import javax.security.auth.callback.CallbackHandler;
import javax.security.auth.login.LoginException;

import org.jboss.security.SimpleGroup;
import org.jboss.security.SimplePrincipal;
import org.jboss.security.auth.spi.UsernamePasswordLoginModule;

/**
 * An example custom login module that obtains passwords and roles
 * for a user from a JNDI lookup.
 *
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.15 $
 */
public class JndiUserAndPass
    extends UsernamePasswordLoginModule
{
    /** The JNDI name to the context that handles the password/username lookup */
    private String userPathPrefix;
    /** The JNDI name to the context that handles the roles/ username lookup */
    private String rolesPathPrefix;

    /**
     * Override to obtain the userPathPrefix and rolesPathPrefix options.
     */
    public void initialize(Subject subject, CallbackHandler callbackHandler,
        Map sharedState, Map options)
    {
        super.initialize(subject, callbackHandler, sharedState, options);
        userPathPrefix = (String) options.get("userPathPrefix");
        rolesPathPrefix = (String) options.get("rolesPathPrefix");
    }

    /**
     * Get the roles the current user belongs to by querying the
     * rolesPathPrefix + '/' + super.getUsername() JNDI location.
     */
    protected Group[] getRoleSets() throws LoginException
    {
        try {
            InitialContext ctx = new InitialContext();
            String rolesPath = rolesPathPrefix + '/' + super.getUsername();

            String[] roles = (String[]) ctx.lookup(rolesPath);
            Group[] groups = {new SimpleGroup("Roles")};
            log.info("Getting roles for user="+super.getUsername());
            for(int r = 0; r < roles.length; r++) {
                SimplePrincipal role = new SimplePrincipal(roles[r]);
                log.info("Found role="+roles[r]);
                groups[0].addMember(role);
            }
            return groups;
        } catch(NamingException e) {
            log.error("Failed to obtain groups for
                user="+super.getUsername(), e);
            throw new LoginException(e.toString(true));
        }
    }
}

```

```

    * Get the password of the current user by querying the
    * userPathPrefix + '/' + super.getUsername() JNDI location.
    */
protected String getUsersPassword()
    throws LoginException
{
    try {
        InitialContext ctx = new InitialContext();
        String userPath = userPathPrefix + '/' + super.getUsername();
        log.info("Getting password for user="+super.getUsername());
        String passwd = (String) ctx.lookup(userPath);
        log.info("Found password="+passwd);
        return passwd;
    } catch(NamingException e) {
        log.error("Failed to obtain password for
            user="+super.getUsername(), e);
        throw new LoginException(e.toString(true));
    }
}

```

The details of the JNDI store are found in the `org.jboss.chap8.ex2.service.JndiStore` MBean. This service binds an `ObjectFactory` that returns a `javax.naming.Context` proxy into JNDI. The proxy handles lookup operations done against it by checking the prefix of the lookup name against `password` and `roles`. When the name begins with `password`, a user's password is being requested. When the name begins with `roles` the user's roles are being requested. The example implementation always returns a password of `theduke` and an array of roles names equal to `{"TheDuke", "Echo"}` regardless of what the username is. You can experiment with other implementations as you wish.

The example code includes a simple session bean for testing the custom login module. To build, deploy and run the example, execute the following command from the examples directory. Make sure you have the JBoss server running. The key lines from the client are given in Example 8.13 while the server side operation of the is shown in Example 8.14.

Example 8.13. The chap8-ex2 secured client access output

```

[nr@toki examples]$ ant -Dchap=chap8 -Dex=2 run-example
Buildfile: build.xml
...
run-example2:
    [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
    [echo] Waiting for 5 seconds for deploy...
    [java] [INFO,ExClient] Login with username=jduke, password=theduke
    [java] [INFO,ExClient] Looking up EchoBean2
    [java] [INFO,ExClient] Created Echo
    [java] [INFO,ExClient] Echo.echo('Hello') = Hello

```

Example 8.14. The chap8-ex2 server side behavior of the JndiUserAndPass

```

17:48:11,458 INFO    [EjbModule] Deploying EchoBean2
17:48:11,890 INFO    [JndiStore] Start, bound security/store
17:48:11,896 INFO    [SecurityConfig] Using JAAS AuthConfig: jar:file:/private/tmp/jboss-3.2.6/s
server/default/tmp/deploy/tmp22821chap8-ex2.jar-contents/chap8-ex2.sar!/META-INF/login-config.x
ml
17:48:12,355 INFO    [EJBDeployer] Deployed: file:/private/tmp/jboss-3.2.6/server/default/deploy
/chap8-ex2.jar

```

The choice of using the `JndiUserAndPass` custom login module for the server side authentication of the user is determined by the login configuration for the example security domain. The EJB JAR `META-INF/jboss.xml` descriptor sets the security domain and the `sar META-INF/login-config.xml` descriptor defines the login module configuration. The contents of these descriptors are shown in Example 8.15.

Example 8.15. The chap8-ex2 security domain and login module configuration

```
<?xml version="1.0"?>
<jboss>
  <security-domain>java:/jaas/chap8-ex2</security-domain>
</jboss>
```

Example 8.16. The login-config.xml configuration fragment for the chap8-ex2 application

```
<application-policy name = "chap8-ex2">
  <authentication>
    <login-module code="org.jboss.chap8.ex2.JndiUserAndPass"
      flag="required">
      <module-option name = "userPathPrefix">/security/store/password</module-option>
      <module-option name = "rolesPathPrefix">/security/store/roles</module-option>
    </login-module>
  </authentication>
</application-policy>
```

8.4.8. The DynamicLoginConfig service

Security domains defined in the `login-config.xml` file are essentially static. They are read when JBoss starts up, but there is no easy way to add a new security domain or change the definition for an existing one. The `DynamicLoginConfig` service allows you to dynamically deploy security domains. This allows you to specify JAAS login configuration as part of a deployment (or just as a standalone service) rather than having to edit the static `login-config.xml` file.

The service supports the following attributes:

- **AuthConfig:** The resource path to the JAAS login configuration file to use. This defaults to `login-config.xml`
- **LoginConfigService:** the `XMLLoginConfig` service name to use for loading. This service must support a `String loadConfig(URL)` operation to load the configurations.
- **SecurityManagerService:** The `SecurityManagerService` name used to flush the registered security domains. This service must support a `flushAuthenticationCache(String)` operation to flush the case for the argument security domain. Setting this triggers the flush of the authentication caches when the service is stopped.

Here is an example MBean definition using the `DynamicLoginConfig` service.

```
<server>
  <mbean code="org.jboss.security.auth.login.DynamicLoginConfig" name="...">
    <attribute name="AuthConfig">login-config.xml</attribute>
    <!-- The service which supports dynamic processing of login-config.xml
    configurations.
    -->
```

```
<depends optional-attribute-name="LoginConfigService">
    jboss.security:service=XMLLoginConfig </depends>
<!-- Optionally specify the security mgr service to use when
this service is stopped to flush the auth caches of the domains
registered by this service.
-->
<depends optional-attribute-name="SecurityManagerService">
    jboss.security:service=JaasSecurityManager </depends>
</mbean>
</server>
```

This will load the specified `AuthConfig` resource using the specified `LoginConfigService` MBean by invoking `loadConfig` with the appropriate resource URL. When the service is stopped the configurations are removed. The resource specified may be either an XML file, or a Sun JAAS login configuration.

8.5. The Secure Remote Password (SRP) Protocol

The SRP protocol is an implementation of a public key exchange handshake described in the Internet standards working group request for comments 2945(RFC2945). The RFC2945 abstract states:

This document describes a cryptographically strong network authentication mechanism known as the Secure Remote Password (SRP) protocol. This mechanism is suitable for negotiating secure connections using a user-supplied password, while eliminating the security problems traditionally associated with reusable passwords. This system also performs a secure key exchange in the process of authentication, allowing security layers (privacy and/or integrity protection) to be enabled during the session. Trusted key servers and certificate infrastructures are not required, and clients are not required to store or manage any long-term keys. SRP offers both security and deployment advantages over existing challenge-response techniques, making it an ideal drop-in replacement where secure password authentication is needed.

Note: The complete RFC2945 specification can be obtained from <http://www.rfc-editor.org/rfc.html>. Additional information on the SRP algorithm and its history can be found at <http://www-cs-students.stanford.edu/~tjw/srp/>.

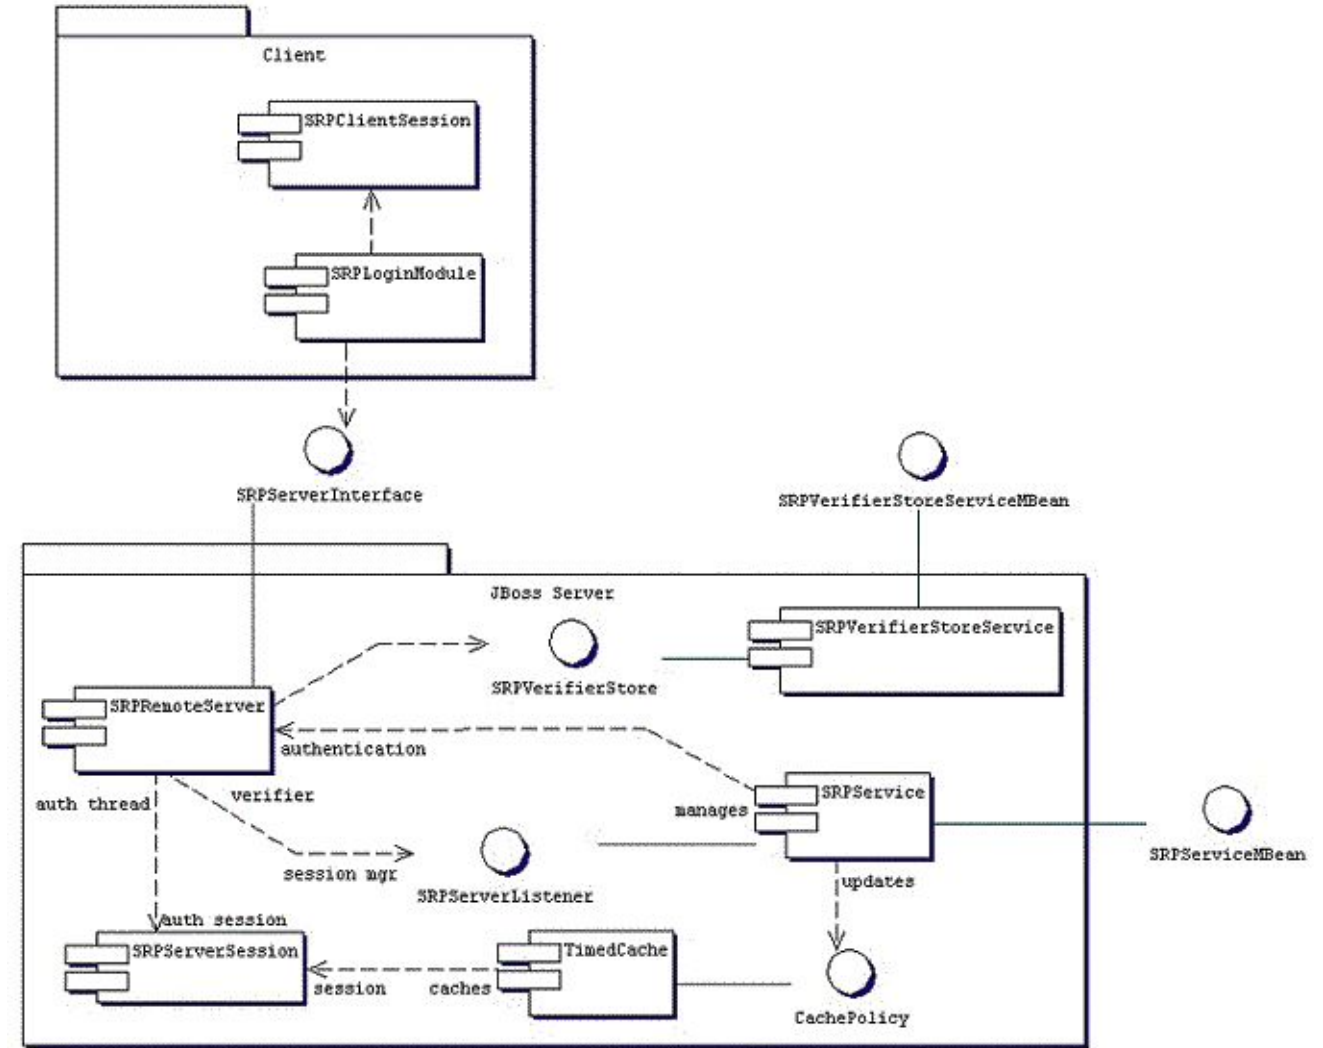
SRP is similar in concept and security to other public key exchange algorithms, such as Diffie-Hellman and RSA. SRP is based on simple string passwords in a way that does not require a clear text password to exist on the server. This is in contrast to other public key-based algorithms that require client certificates and the corresponding certificate management infrastructure.

Algorithms like Diffie-Hellman and RSA are known as public key exchange algorithms. The concept of public key algorithms is that you have two keys, one public that is available to everyone, and one that is private and known only to you. When someone wants to send encrypted information to you, then encrypt the information using your public key. Only you are able to decrypt the information using your private key. Contrast this with the more traditional shared password based encryption schemes that require the sender and receiver to know the shared password. Public key algorithms eliminate the need to share passwords. For more information on public key algorithms as well as numerous other cryptographic algorithms, see *Applied Cryptography, Second Edition* by Bruce Schneier, ISBN 0-471-11709-9.

The JBossSX framework includes an implementation of SRP that consists of the following elements:

- An implementation of the SRP handshake protocol that is independent of any particular client/server protocol
- An RMI implementation of the handshake protocol as the default client/server SRP implementation

- Figure 8.10 gives a diagram of the key components involved in the IBossSY implementation of the SPP client/



`org.jboss.security.srp.jaas.SRPPrincipal`.

- **srpServerJndiName:** The JNDI name of the `SRPServerInterface` object to use for communicating with the SRP authentication server. If both `srpServerJndiName` and `srpServerRmiUrl` options are specified, the `srpServerJndiName` is tried before `srpServerRmiUrl`.
- **srpServerRmiUrl:** The RMI protocol URL string for the location of the `SRPServerInterface` proxy to use for communicating with the SRP authentication server.
- **externalRandomA:** A true/false flag indicating if the random component of the client public key A should come from the user callback. This can be used to input a strong cryptographic random number coming from a hardware token for example.
- **hasAuxChallenge:** A true/false flag indicating an that a string will be sent to the server as an additional challenge for the server to validate. If the client session supports an encryption cipher then a temporary cipher will be created using the session private key and the challenge object sent as a `javax.crypto.SealedObject`.
- **multipleSessions:** a true/false flag indicating if a given client may have multiple SRP login sessions active simultaneously.

Any other options passed in that do not match one of the previous named options is treated as a JNDI property to use for the environment passed to the `IntialContext` constructor. This is useful if the SRP server interface is not available from the default `IntialContext`.

The `SRPLoginModule` needs to be configured along with the standard `ClientLoginModule` to allow the SRP authentication credentials to be used for validation of access to security J2EE components. An example login configuration entry that demonstrates such a setup is:

```
srp {  
    org.jboss.security.srp.jaas.SRPLoginModule required  
    srpServerJndiName="SRPServerInterface"  
    ;  
  
    org.jboss.security.ClientLoginModule required  
    password-stacking="useFirstPass"  
    ;  
};
```

On the JBoss server side, there are two MBeans that manage the objects that collectively make up the SRP server. The primary service is the `org.jboss.security.srp.SRPService` MBean, and it is responsible for exposing an RMI accessible version of the `SRPServerInterface` as well as updating the SRP authentication session cache. The configurable `SRPService` MBean attributes include the following:

- **JndiName:** The JNDI name from which the `SRPServerInterface` proxy should be available. This is the location where the `SRPService` binds the serializable dynamic proxy to the `SRPServerInterface`. If not specified it defaults to `srp/SRPServerInterface`.
- **VerifierSourceJndiName:** The JNDI name of the `SRPVerifierSource` implementation that should be used by the `SRPService`. If not set it defaults to `srp/DefaultVerifierSource`.
- **AuthenticationCacheJndiName:** The JNDI name under which the authentication `org.jboss.util.CachePolicy` implementation to be used for caching authentication information is bound. The SRP session cache is made available for use through this binding. If not specified it defaults to `srp/AuthenticationCache`.

- **ServerPort:** RMI port for the `SRPRemoteServerInterface`. If not specified it defaults to 10099.
- **ClientSocketFactory:** An optional custom `java.rmi.server.RMIClientSocketFactory` implementation class name used during the export of the `SRPServerInterface`. If not specified the default `RMIClientSocketFactory` is used.
- **ServerSocketFactory:** An optional custom `java.rmi.server.RMIServerSocketFactory` implementation class name used during the export of the `SRPServerInterface`. If not specified the default `RMIServerSocketFactory` is used.
- **AuthenticationCacheTimeout:** Specifies the timed cache policy timeout in seconds. If not specified this defaults to 1800 seconds(30 minutes).
- **AuthenticationCacheResolution:** Specifies the timed cache policy resolution in seconds. This controls the interval between checks for timeouts. If not specified this defaults to 60 seconds(1 minute).
- **RequireAuxChallenge:** Set if the client must supply an auxillary challenge as part of the verify phase. This gives control over whether the `SRPLoginModule` configuration used by the client must have the `useAuxChallenge` option enabled.
- **OverwriteSessions:** A flag indicating if a successful user auth for an existing session should overwrite the current session. This controls the behavior of the server SRP session cache when clients have not enabled the multiple session per user mode. The default is false meaning that the second attempt by a user to authentication will succeed, but the resulting SRP session will not overwrite the previous SRP session state.

The one input setting is the `VerifierSourceJndiName` attribute. This is the location of the SRP password information store implementation that must be provided and made available through JNDI. The `org.jboss.security.srp.SRPVerifierStoreService` is an example MBean service that binds an implementation of the `SRPVerifierStore` interface that uses a file of serialized objects as the persistent store. Although not realistic for a production environment, it does allow for testing of the SRP protocol and provides an example of the requirements for an `SRPVerifierStore` service. The configurable `SRPVerifierStoreService` MBean attributes include the following:

- **JndiName:** The JNDI name from which the `SRPVerifierStore` implementation should be available. If not specified it defaults to `srp/DefaultVerifierSource`.
- **StoreFile:** The location of the user password verifier serialized object store file. This can be either a URL or a resource name to be found in the classpath. If not specified it defaults to `SRPVerifierStore.ser`.

The `SRPVerifierStoreService` MBean also supports `addUser` and `delUser` operations for addition and deletion of users. The signatures are:

```
public void addUser(String username, String password) throws IOException;
public void delUser(String username) throws IOException;
```

An example configuration of these services is presented in Section 8.5.

8.5.1. Providing Password Information for SRP

The default implementation of the `SRPVerifierStore` interface is not likely to be suitable for your production security environment as it requires all password hash information to be available as a file of serialized objects. You need to provide an MBean service that provides an implementation of the `SRPVerifierStore` interface that integrates with your existing security information stores. The `SRPVerifierStore` interface is shown in.

Example 8.17. The SRPVerifierStore interface

```

package org.jboss.security.srp;

import java.io.IOException;
import java.io.Serializable;
import java.security.KeyException;

public interface SRPVerifierStore
{
    public static class VerifierInfo implements Serializable
    {
        /**
         * The username the information applies to. Perhaps redundant
         * but it makes the object self contained.
         */
        public String username;

        /** The SRP password verifier hash */
        public byte[] verifier;
        /** The random password salt originally used to verify the password */
        public byte[] salt;
        /** The SRP algorithm primitive generator */
        public byte[] g;
        /** The algorithm safe-prime modulus */
        public byte[] N;
    }

    /**
     * Get the indicated user's password verifier information.
     */
    public VerifierInfo getUserVerifier(String username)
        throws KeyException, IOException;

    /**
     * Set the indicated users' password verifier information. This
     * is equivalent to changing a user's password and should
     * generally invalidate any existing SRP sessions and caches.
     */
    public void setUserVerifier(String username, VerifierInfo info)
        throws IOException;

    /**
     * Verify an optional auxillary challenge sent from the client to
     * the server. The auxChallenge object will have been decrypted
     * if it was sent encrypted from the client. An example of a
     * auxillary challenge would be the validation of a hardware token
     * (SafeWord, SecureID, iButton) that the server validates to
     * further strengthen the SRP password exchange.
     */
    public void verifyUserChallenge(String username, Object auxChallenge)
        throws SecurityException;
}

```

The primary function of a `SRPVerifierStore` implementation is to provide access to the `SRPVerifierStore.VerifierInfo` object for a given username. The `getUserVerifier(String)` method is called by the `SRPService` at that start of a user SRP session to obtain the parameters needed by the SRP algorithm. The elements of the `VerifierInfo` objects are:

- **username:** The user's name or id used to login.
- **verifier:** This is the one-way hash of the password or PIN the user enters as proof of their identity. The `org.jboss.security.Util` class has a `calculateVerifier` method that performs that password hashing algorithm. The output password `H(salt | H(username | ':' | password))` as defined by RFC2945. Here `H` is the SHA secure hash function. The username is converted from a string to a `byte[]` using the UTF-8

encoding.

- **salt:** This is a random number used to increase the difficulty of a brute force dictionary attack on the verifier password database in the event that the database is compromised. It is a value that should be generated from a cryptographically strong random number algorithm when the user's existing clear-text password is hashed.
- **g:** The SRP algorithm primitive generator. In general this can be a well known fixed parameter rather than a per-user setting. The `org.jboss.security.srp.SRPConf` utility class provides several settings for `g` including a good default which can be obtained via `SRPConf.getDefaultParams().g()`.
- **N:** The SRP algorithm safe-prime modulus. In general this can be a well known fixed parameter rather than a per-user setting. The `org.jboss.security.srp.SRPConf` utility class provides several settings for `N` including a good default which can be obtained via `SRPConf.getDefaultParams().N()`.

So, step 1 of integrating your existing password store is the creation of a hashed version of the password information. If your passwords are already store in an irreversible hashed form, then this can only be done on a per-user basis as part of an upgrade procedure for example. Note that the `setUserVerifier(String, VerifierInfo)` method is not used by the current `SRPSerivce` and may be implemented as noop method, or even one that throws an exception stating that the store is read-only.

Step 2 is the creation of the custom `SRPVerifierStore` interface implementation that knows how to obtain the `VerifierInfo` from the store you created in step 1. The `verifyUserChallenge(String, Object)` method of the interface is only called if the client `SRPLoginModule` configuration specifies the `hasAuxChallenge` option. This can be used to integrate existing hardware token based schemes like `SafeWord` or `Radius` into the SRP algorithm.

Step 3 is the creation of an MBean that makes the step 2 implementation of the `SRPVerifierStore` interface available via JNDI, and exposes any configurable parameters you need. In addition to the default `org.jboss.security.srp.SRPVerifierStoreService` example, the SRP example presented later in this chapter provides a Java properties file based `SRPVerifierStore` implementation. Between the two examples you should have enough to integrate your security store.

8.5.2. Inside of the SRP algorithm

The appeal of the SRP algorithm is that it allows for mutual authentication of client and server using simple text passwords without a secure communication channel. You might be wondering how this is done. If you want the complete details and theory behind the algorithm, refer to the SRP references mentioned in a note earlier. There are six steps that are performed to complete authentication:

1. The client side `SRPLoginModule` retrieves the `SRPServerInterface` instance for the remote authentication server from the naming service.
2. The client side `SRPLoginModule` next requests the SRP parameters associated with the username attempting the login. There are a number of parameters involved in the SRP algorithm that must be chosen when the user password is first transformed into the verifier form used by the SRP algorithm. Rather than hard-coding the parameters (which could be done with minimal security risk), the JBossSX implementation allows a user to retrieve this information as part of the exchange protocol. The `getSRPParameters(username)` call retrieves the SRP parameters for the given username.
3. The client side `SRPLoginModule` begins an SRP session by creating an `SRPClientSession` object using the login username, clear-text password, and SRP parameters obtained from step 2. The client then creates a

random number `A` that will be used to build the private SRP session key. The client then initializes the server side of the SRP session by invoking the `SRPServerInterface.init` method and passes in the username and client generated random number `A`. The server returns its own random number `B`. This step corresponds to the exchange of public keys.

4. The client side `SRPLoginModule` obtains the private SRP session key that has been generated as a result of the previous messages exchanges. This is saved as a private credential in the login `Subject`. The server challenge response `M2` from step 4 is verified by invoking the `SRPClientSession.verify` method. If this succeeds, mutual authentication of the client to server, and server to client have been completed. The client side `SRPLoginModule` next creates a challenge `M1` to the server by invoking `SRPClientSession.response` method passing the server random number `B` as an argument. This challenge is sent to the server via the `SRPServerInterface.verify` method and server's response is saved as `M2`. This step corresponds to an exchange of challenges. At this point the server has verified that the user is who they say they are.
5. The client side `SRPLoginModule` saves the login username and `M1` challenge into the `LoginModule` shared-State map. This is used as the Principal name and credentials by the standard JBoss `ClientLoginModule`. The `M1` challenge is used in place of the password as proof of identity on any method invocations on J2EE components. The `M1` challenge is a cryptographically strong hash associated with the SRP session. Its interception via a third party cannot be used to obtain the user's password.
6. At the end of this authentication protocol, the `SRPServiceSession` has been placed into the `SRPService` authentication cache for subsequent use by the `SRPCacheLoginModule`.

Although SRP has many interesting properties, it is still an evolving component in the JBossSX framework and has some limitations of which you should be aware. Issues of note include the following:

- Because of how JBoss detaches the method transport protocol from the component container where authentication is performed, an unauthorized user could snoop the SRP `M1` challenge and effectively use the challenge to make requests as the associated username. Custom interceptors that encrypt the challenge using the SRP session key can be used to prevent this issue.
- The `SRPService` maintains a cache of SRP sessions that time out after a configurable period. Once they time out, any subsequent J2EE component access will fail because there is currently no mechanism for transparently renegotiating the SRP authentication credentials. You must either set the authentication cache timeout very long (up to 2,147,483,647 seconds, or approximately 68 years), or handle re-authentication in your code on failure.
- By default there can only be one SRP session for a given username. Because the negotiated SRP session produces a private session key that can be used for encryption/decryption between the client and server, the session is effectively a stateful one. JBoss supports for multiple SRP sessions per user, but you cannot encrypt data with one session key and then decrypt it with another.

To use end-to-end SRP authentication for J2EE component calls, you need to configure the security domain under which the components are secured to use the `org.jboss.security.srp.jaas.SRPCacheLoginModule`. The `SRPCacheLoginModule` has a single configuration option named `cacheJndiName` that sets the JNDI location of the SRP authentication `CachePolicy` instance. This must correspond to the `AuthenticationCacheJndiName` attribute value of the `SRPService` MBean. The `SRPCacheLoginModule` authenticates user credentials by obtaining the client challenge from the `SRPServiceSession` object in the authentication cache and comparing this to the challenge passed as the user credentials. Figure 8.11 illustrates the operation of the `SRPCacheLoginModule.login` method implementation.

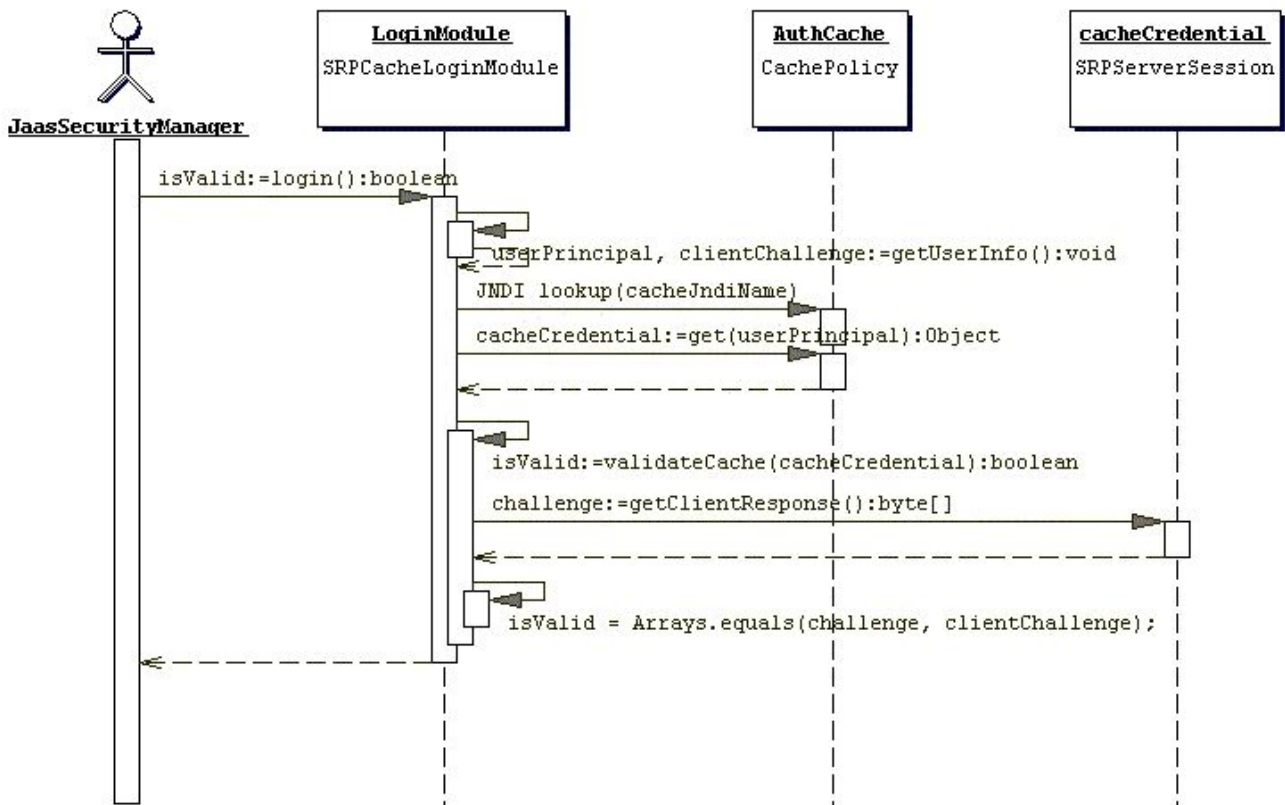


Figure 8.11. A sequence diagram illustrating the interaction of the SRPCacheLoginModule with the SRP session cache.

8.5.2.1. An SRP example

We have covered quite a bit of material on SRP and now its time to demonstrate SRP in practice with an example. The example demonstrates client side authentication of the user via SRP as well as subsequent secured access to a simple EJB using the SRP session challenge as the user credential. The test code deploys an EJB JAR that includes a sar for the configuration of the server side login module configuration and SRP services. As in the previous examples we will dynamically install the server side login module configuration using the `SecurityConfig` MBean. In this example we also use a custom implementation of the `SRPVerifierStore` interface that uses an in memory store that is seeded from a Java properties file rather than a serialized object store as used by the `SRPVerifierStoreService`. This custom service is `org.jboss.chap8.ex3.service.PropertiesVerifierStore`. The following shows the contents of the JAR that contains the example EJB and SRP services.

```
[orb@toki examples]$ java -cp output/classes ListJar output/chap8/chap8-ex3.jar
output/chap8/chap8-ex3.jar
+- META-INF/MANIFEST.MF
+- META-INF/ejb-jar.xml
+- META-INF/jboss.xml
+- org/jboss/chap8/ex3/Echo.class
+- org/jboss/chap8/ex3/EchoBean.class
+- org/jboss/chap8/ex3/EchoHome.class
+- roles.properties
+- users.properties
+- chap8-ex3.sar (archive)
| +- META-INF/MANIFEST.MF
| +- META-INF/jboss-service.xml
| +- META-INF/login-config.xml
| +- org/jboss/chap8/ex3/service/PropertiesVerifierStore$1.class
| +- org/jboss/chap8/ex3/service/PropertiesVerifierStore.class
| +- org/jboss/chap8/ex3/service/PropertiesVerifierStoreMBean.class
```

```
| +- org/jboss/chap8/service/SecurityConfig.class
| +- org/jboss/chap8/service/SecurityConfigMBean.class
```

The key SRP related items in this example are the SRP MBean services configuration, and the SRP login module configurations. The `jboss-service.xml` descriptor of the `chap8-ex3.sar` is given in Example 8.18, while Example 8.19 and Example 8.20 give the example client side and server side login module configurations.

Example 8.18. The `chap8-ex3.sar` `jboss-service.xml` descriptor for the SRP services

```
<server>
  <!-- The custom JAAS login configuration that installs
        a Configuration capable of dynamically updating the
        config settings -->

  <mbean code="org.jboss.chap8.service.SecurityConfig"
        name="jboss.docs.chap8:service=LoginConfig-EX3">
    <attribute name="AuthConfig">META-INF/login-config.xml</attribute>
    <attribute name="SecurityConfigName">jboss.security:name=SecurityConfig</attribute>
  </mbean>

  <!-- The SRP service that provides the SRP RMI server and server side
        authentication cache -->
  <mbean code="org.jboss.security.srp.SRPService"
        name="jboss.docs.chap8:service=SRPService">
    <attribute name="VerifierSourceJndiName">srp-test/chap8-ex3</attribute>
    <attribute name="JndiName">srp-test/SRPServerInterface</attribute>
    <attribute name="AuthenticationCacheJndiName">srp-test/AuthenticationCache</attribute>
    <attribute name="ServerPort">0</attribute>
    <depends>jboss.docs.chap8:service=PropertiesVerifierStore</depends>
  </mbean>

  <!-- The SRP store handler service that provides the user password verifier
        information -->
  <mbean code="org.jboss.chap8.ex3.service.PropertiesVerifierStore"
        name="jboss.docs.chap8:service=PropertiesVerifierStore">
    <attribute name="JndiName">srp-test/chap8-ex3</attribute>
  </mbean>
</server>
```

Example 8.19. The client side standard JAAS configuration

```
srp {
  org.jboss.security.srp.jaas.SRPLoginModule required
  srpServerJndiName="srp-test/SRPServerInterface"
  ;

  org.jboss.security.ClientLoginModule required
  password-stacking="useFirstPass"
  ;
};
```

Example 8.20. The server side `XMLLoginConfig` configuration

```
<application-policy name="chap8-ex3">
  <authentication>
    <login-module code="org.jboss.security.srp.jaas.SRPCacheLoginModule"
                  flag = "required">
      <module-option name="cacheJndiName">srp-test/AuthenticationCache</module-option>
    </login-module>
  </authentication>
</application-policy>
```



```
</login-module>
<login-module code="org.jboss.security.auth.spi.UsersRolesLoginModule"
    flag = "required">
    <module-option name="password-stacking">useFirstPass</module-option>
</login-module>
</authentication>
</application-policy>
```

The example services are the `ServiceConfig` and the `PropertiesVerifierStore` and `SRPService` MBeans. Note that the `JndiName` attribute of the `PropertiesVerifierStore` is equal to the `VerifierSourceJndiName` attribute of the `SRPService`, and that the `SRPService` depends on the `PropertiesVerifierStore`. This is required because the `SRPService` needs an implementation of the `SRPVerifierStore` interface for accessing user password verification information.

The client side login module configuration makes use of the `SRPLoginModule` with a `srpServerJndiName` option value that corresponds to the JBoss server component `SRPService` `JndiName` attribute value (`srp-test/SRPServiceInterface`). Also needed is the `ClientLoginModule` configured with the `password-stacking="useFirstPass"` value to propagate the user authentication credentials generated by the `SRPLoginModule` to the EJB invocation layer.

There are two issues to note about the server side login module configuration. First, note the `cacheJndiName=srp-test/AuthenticationCache` configuration option tells the `SRPCacheLoginModule` the location of the `CachePolicy` that contains the `SRPServerSession` for users who have authenticated against the `SRPService`. This value corresponds to the `SRPService` `AuthenticationCacheJndiName` attribute value. Second, the configuration includes a `UsersRolesLoginModule` with the `password-stacking=useFirstPass` configuration option. It is required to use a second login module with the `SRPCacheLoginModule` because SRP is only an authentication technology. A second login module needs to be configured that accepts the authentication credentials validated by the `SRPCacheLoginModule` to set the principal's roles that determines the principal's permissions. The `UsersRolesLoginModule` is augmenting the SRP authentication with properties file based authorization. The user's roles are coming from the `roles.properties` file included in the EJB JAR.

Now, run the example 3 client by executing the following command from the book examples directory:

```
[starksm@banshee examples]$ ant -Dchap=chap8 -Dex=3 run-example
Buildfile: build.xml
...
run-example3:
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
  [echo] Waiting for 5 seconds for deploy...
  [java] Logging in using the 'srp' configuration
  [java] Created Echo
  [java] Echo.echo()#1 = This is call 1
  [java] Echo.echo()#2 = This is call 2
```

In the `examples/logs` directory you will find a file called `ex3-trace.log`. This is a detailed trace of the client side of the SRP algorithm. The traces show step-by-step the construction of the public keys, challenges, session key and verification.

Note that the client has taken a long time to run relative to the other simple examples. The reason for this is the construction of the client's public key. This involves the creation of a cryptographically strong random number, and this process takes quite a bit of time the first time it occurs. If you were to log out and log in again within the same VM, the process would be much faster. Also note that `Echo.echo()#2` fails with an authentication exception. The client code sleeps for 15 seconds after making the first call to demonstrate the behavior of the `SRPService` cache expiration. The `SRPService` cache policy timeout has been set to a mere 10 seconds to force this issue. As stated earlier, you need to make the cache timeout very long, or handle re-authentication on failure.

8.6. Running JBoss with a Java 2 security manager

By default the JBoss server does not start with a Java 2 security manager. If you want to restrict privileges of code using Java 2 permissions you need to configure the JBoss server to run under a security manager. This is done by configuring the Java VM options in the `run.bat` or `run.sh` scripts in the JBoss server distribution bin directory. The two required VM options are as follows:

- **java.security.manager:** This is used without any value to specify that the default security manager should be used. This is the preferred security manager. You can also pass a value to the `java.security.manager` option to specify a custom security manager implementation. The value must be the fully qualified class name of a subclass of `java.lang.SecurityManager`. This form specifies that the policy file should augment the default security policy as configured by the VM installation.
- **java.security.policy:** This is used to specify the policy file that will augment the default security policy information for the VM. This option takes two forms: `java.security.policy=policyFileURL` and `java.security.policy==policyFileURL`. The first form specifies that the policy file should augment the default security policy as configured by the VM installation. The second form specifies that only the indicated policy file should be used. The `policyFileURL` value can be any URL for which a protocol handler exists, or a file path specification.

Example 8.21 illustrates a fragment of the standard `run.bat` start script for Win32 that shows the addition of these two options to the command line used to start JBoss.

Example 8.21. The modifications to the Win32 `run.bat` start script to run JBoss with a Java 2 security manager.

```
...

set CONFIG=%1
@if "%CONFIG%" == "" set CONFIG=default
set PF=../conf/%CONFIG%/server.policy
set OPTS=-Djava.security.manager
set OPTS=%OPTS% -Djava.security.policy=%PF%
echo JBOSS_CLASSPATH=%JBOSS_CLASSPATH%
java %JAXP% %OPTS% -classpath "%JBOSS_CLASSPATH%" org.jboss.Main %*
```

Example 8.22 shows a fragment of the standard `run.sh` start script for UNIX/Linux systems that shows the addition of these two options to the command line used to start JBoss.

Example 8.22. The modifications to the UNIX/Linux `run.sh` start script to run JBoss with a Java 2 security manager.

```
# ...

CONFIG=$1
if [ "$CONFIG" == "" ]; then CONFIG=default; fi
PF=../conf/$CONFIG/server.policy
OPTS=-Djava.security.manager
OPTS="$OPTS -Djava.security.policy=$PF"
echo JBOSS_CLASSPATH=$JBOSS_CLASSPATH
java $HOTSPOT $JAXP $OPTS -classpath $JBOSS_CLASSPATH org.jboss.Main $@
```

Both start scripts are setting the security policy file to the `server.policy` file located in the JBoss configuration file set directory that corresponds to the configuration name passed as the first argument to the script. This allows one maintain a security policy per configuration file set without having to modify the start script.

Enabling Java 2 security is the easy part. The difficult part of Java 2 security is establishing the allowed permissions. If you look at the `server.policy` file that is contained in the default configuration file set, you'll see that it contains the following permission grant statement:

```
grant {
    // Allow everything for now
    permission java.security.AllPermission;
};
```

This effectively disables security permission checking for all code as it says any code can do anything, which is not a reasonable default. What is a reasonable set of permissions is entirely up to you.

The current set of JBoss specific `java.lang.RuntimePermissions` that are required include:

TargetName	What the permission allows	Risks
<code>org.jboss.security.SecurityAssociation.getPrincipalInfo</code>	Access to the <code>org.jboss.security.SecurityAssociation.getPrincipal()</code> and <code>getCredentials()</code> methods.	The ability to see the current thread caller and credentials.
<code>org.jboss.security.SecurityAssociation.setPrincipalInfo</code>	Access to the <code>org.jboss.security.SecurityAssociation.setPrincipal()</code> and <code>setCredentials()</code> methods.	The ability to set the current thread caller and credentials.
<code>org.jboss.security.SecurityAssociation.setServer</code>	Access to the <code>org.jboss.security.SecurityAssociation.setServer</code> method.	The ability to enable or disable multithread storage of the caller principal and credential.
<code>org.jboss.security.SecurityAssociation.setRunAsRole</code>	Access to the <code>org.jboss.security.SecurityAssociation.pushRunAsRole</code> and <code>popRunAsRole</code> methods.	The ability to change the current caller run-as role principal.

To conclude this discussion, here is a little-known tidbit on debugging security policy settings. There are various debugging flag that you can set to determine how the security manager is using your security policy file as well as what policy files are contributing permissions. Running the VM as follows shows the possible debugging flag settings:

```
[nr@toki bin]$ java -Djava.security.debug=help

all          turn on all debugging
access       print all checkPermission results
combiner     SubjectDomainCombiner debugging
jar          jar verification
logincontext login context results
policy       loading and granting
provider     security provider debugging
scl          permissions SecureClassLoader assigns

The following can be used with access:

stack        include stack trace
```

```
domain    dumps all domains in context
failure   before throwing exception, dump stack
          and domain that didn't have permission
```

Note: Separate multiple options with a comma

Running with `-Djava.security.debug=all` provides the most output, but the output volume is torrential. This might be a good place to start if you don't understand a given security failure at all. A less verbose setting that helps debug permission failures is to use `-Djava.security.debug=access,failure`. This is still relatively verbose, but not nearly as bad as the all mode as the security domain information is only displayed on access failures.

8.7. Using SSL with JBoss using JSSE

JBoss uses JSSE the Java Secure Socket Extension (JSSE). JSSE is bundled with JBoss and it comes with JDK 1.4. For more information on JSSE see: <http://java.sun.com/products/jsse/index.html>. A simple test that you can use the JSSE as bundled with JBoss works is to run a program like the following:

```
import java.net.*;
import javax.net.ServerSocketFactory;
import javax.net.ssl.*;

public class JSSE_install_check
{
    public static void main(String[] args) throws Exception
    {
        Security.addProvider(new com.sun.net.ssl.internal.ssl.Provider());
        ServerSocketFactory factory =
            SSLServerSocketFactory.getDefault();
        SSLServerSocket sslSocket = (SSLServerSocket)
            factory.createServerSocket(12345);

        String [] cipherSuites = sslSocket.getEnabledCipherSuites();
        for(int i = 0; i < cipherSuites.length; i++) {
            System.out.println("Cipher Suite " + i + " = " + cipherSuites[i]);
        }
    }
}
```

The book examples includes a testcase for this which can be run using the following command. This will produce a lot of output as the `-Djavax.net.debug=all` option is passed to the VM.

```
[nr@toki examples]$ ant -Dchap=chap8 -Dex=4a run-example
...
run-example4a:
run-example4a:
    [echo] Testing JSSE availability
    [java] keyStore is :
    [java] keyStore type is : jks
    [java] init keystore
    [java] init keymanager of type SunX509
    [java] trustStore is: /System/Library/Frameworks/JavaVM.framework/Versions/1.4.2/Home
/lib/security/cacerts
    [java] trustStore type is : jks
    [java] init truststore
    ...
    [java] init context
    [java] trigger seeding of SecureRandom
    [java] done seeding SecureRandom
    [java] Cipher Suite 0 = SSL_RSA_WITH_RC4_128_MD5
    [java] Cipher Suite 1 = SSL_RSA_WITH_RC4_128_SHA
    [java] Cipher Suite 2 = TLS_RSA_WITH_AES_128_CBC_SHA
    [java] Cipher Suite 3 = TLS_DHE_RSA_WITH_AES_128_CBC_SHA
```

```
[java] Cipher Suite 4 = TLS_DHE_DSS_WITH_AES_128_CBC_SHA
[java] Cipher Suite 5 = SSL_RSA_WITH_3DES_EDE_CBC_SHA
[java] Cipher Suite 6 = SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
[java] Cipher Suite 7 = SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
[java] Cipher Suite 8 = SSL_RSA_WITH_DES_CBC_SHA
[java] Cipher Suite 9 = SSL_DHE_RSA_WITH_DES_CBC_SHA
[java] Cipher Suite 10 = SSL_DHE_DSS_WITH_DES_CBC_SHA
[java] Cipher Suite 11 = SSL_RSA_EXPORT_WITH_RC4_40_MD5
[java] Cipher Suite 12 = SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
[java] Cipher Suite 13 = SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA
[java] Cipher Suite 14 = SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
```

The JSSE jars include the `jcercert.jar`, `jnet.jar` and `jsse.jar` in the `JBOSS_DIST/client` directory.

Once you have tested that JSSE runs, you need a public key/private key pair in the form of an X509 certificate for use by the SSL server sockets. For the purpose of this example we have created a self-signed certificate using the JDK keytool and included the resulting keystore file in the `chap8` source directory as `chap8.keystore`. It was created using the following command and input:

```
[nr@toki examples]$ keytool -genkey -alias rmi+ssl -keyalg RSA -keystore chap8.keystore -v
validity 3650
[orb@toki examples]$ keytool -genkey -alias rmi+ssl -keyalg RSA -keystore chap8.keystore -v
validity 3650
Enter keystore password: rmi+ssl
What is your first and last name?
[Unknown]: Chapter 8 SSL Example
What is the name of your organizational unit?
[Unknown]: JBoss Book
What is the name of your organization?
[Unknown]: JBoss, Inc.
What is the name of your City or Locality?
[Unknown]: Issaquah
What is the name of your State or Province?
[Unknown]: WA
What is the two-letter country code for this unit?
[Unknown]: US
Is CN=Chapter 8 SSL Example, OU=JBoss Book, O="JBoss, Inc.", L=Issaquah, ST=WA, C=US correct?
[no]: yes

Enter key password for <rmi+ssl>
(RETURN if same as keystore password):
```

This produces a keystore file called `chap8.keystore`. A keystore is a database of security keys. There are two different types of entries in a keystore:

- **key entries:** each entry holds very sensitive cryptographic key information, which is stored in a protected format to prevent unauthorized access. Typically, a key stored in this type of entry is a secret key, or a private key accompanied by the certificate chain for the corresponding public key. The `keytool` and `jarsigner` tools only handle the later type of entry, that is private keys and their associated certificate chains.
- **trusted certificate entries:** each entry contains a single public key certificate belonging to another party. It is called a trusted certificate because the keystore owner trusts that the public key in the certificate indeed belongs to the identity identified by the subject (owner) of the certificate. The issuer of the certificate vouches for this, by signing the certificate.

Listing the `src/main/org/jboss/chap8/chap8.keystore` examples file contents using the `keytool` shows one self-signed certificate:

```
[nr@toki examples]$ keytool -list -v -keystore src/main/org/jboss/chap8/chap8.keystore
```

```

Enter keystore password: rmi+ssl

Keystore type: jks
Keystore provider: SUN

Your keystore contains 1 entry

Alias name: rmi+ssl
Creation date: Nov 8, 2001
Entry type: keyEntry
Certificate chain length: 1
Certificate[1]:
Owner: CN=Chapter8 SSL Example, OU=JBoss Book, O="JBoss Group, LLC", L=Issaquah, ST=WA, C=US
Issuer: CN=Chapter8 SSL Example, OU=JBoss Book, O="JBoss Group, LLC", L=Issaquah, ST=WA, C=US
Serial number: 3beb5271
Valid from: Thu Nov 08 21:50:09 CST 2001 until: Sun Nov 06 21:50:09 CST 2011
Certificate fingerprints:
MD5: F6:1B:2B:E9:A5:23:E7:22:B2:18:6F:3F:9F:E7:38:AE
SHA1: F2:20:50:36:97:86:52:89:71:48:A2:C3:06:C8:F9:2D:F7:79:00:36

*****
*****

```

With JSSE working and a keystore with the certificate you will use for the JBoss server, you are ready to configure JBoss to use SSL for EJB access. This is done by configuring the EJB invoker RMI socket factories. The JBossSX framework includes implementations of the `java.rmi.server.RMIServerSocketFactory` and `java.rmi.server.RMIClientSocketFactory` interfaces that enable the use of RMI over SSL encrypted sockets. The implementation classes are `org.jboss.security.ssl.RMISSLServerSocketFactory` and `org.jboss.security.ssl.RMISSLClientSocketFactory` respectively. There are two steps to enable the use of SSL for RMI access to EJBs. The first is to enable the use of a keystore as the database for the SSL server certificate, which is done by configuring an `org.jboss.security.plugins.JaasSecurityDomain` MBean. The `jboss-service.xml` descriptor in the `chap8/ex4` directory includes the `JaasSecurityDomain` definition shown in Example 8.23.

Example 8.23. A sample `JaasSecurityDomain` config for RMI/SSL

```

<!-- The SSL domain setup -->
<mbean code="org.jboss.security.plugins.JaasSecurityDomain"
      name="jboss.security:service=JaasSecurityDomain,domain=RMI+SSL">
  <constructor>
    <arg type="java.lang.String" value="RMI+SSL"/>
  </constructor>
  <attribute name="KeyStoreURL">chap8.keystore</attribute>
  <attribute name="KeyStorePass">rmi+ssl</attribute>
</mbean>

```

The `JaasSecurityDomain` is a subclass of the standard `JaasSecurityManager` class that adds the notions of a keystore as well JSSE `KeyManagerFactory` and `TrustManagerFactory` access. It extends the basic security manager to allow support for SSL and other cryptographic operations that require security keys. This configuration simply loads the `chap8.keystore` from the example 4 MBean sar using the indicated password.

The second step is to define an EJB invoker configuration that uses the JBossSX RMI socket factories that support SSL. To do this you need to define a custom configuration for the `JRMPInvoker` we saw in Chapter 5 as well as an EJB setup that makes use of this invoker. The configuration required to enable RMI over SSL access to stateless session bean is provided for you in Example 8.24 and Example 8.25 The top of the listing shows the `jboss-service.xml` descriptor that defines the custom `JRMPInovker`, and the bottom shows the example 4

EchoBean4 configuration needed to use the SSL invoker. You will use this configuration in a stateless session bean example.

Example 8.24. The `jboss-service.xml` configurations to enable SSL with the example 4 stateless session bean.

```
<mbean code="org.jboss.invocation.jrmp.server.JRMPInvoker
  name="jboss:service=invoker,type=jrmp,socketType=SSL">
  <attribute name="RMIObjectPort">14445</attribute>
  <attribute name="RMIClientSocketFactory">
    org.jboss.security.ssl.RMISSLClientSocketFactory
  </attribute>
  <attribute name="RMIServerSocketFactory">
    org.jboss.security.ssl.RMISSLServerSocketFactory
  </attribute>
  <attribute name="SecurityDomain">java:/jaas/RMI+SSL</attribute>
  <depends>jboss.security:service=JaasSecurityDomain, domain=RMI+SSL</depends>
</mbean>
```

Example 8.25. The `jboss.xml` configuration to enable SSL with the example 4 stateless session bean.

```
<?xml version="1.0"?>
<jboss>
  <enterprise-beans>
    <session>
      <ejb-name>EchoBean4</ejb-name>
      <configuration-name>Standard Stateless SessionBean</configuration-name>
      <home-invoker>jboss:service=invoker,type=jrmp,socketType=SSL</home-invoker>
      <bean-invoker>jboss:service=invoker,type=jrmp,socketType=SSL</bean-invoker>
    </session>
  </enterprise-beans>
</jboss>
```

The example 4 code is located under the `src/main/org/jboss/chap8/ex4` directory of the book examples. This is another simple stateless session bean with an echo method that returns its input argument. It is hard to tell when SSL is in use unless it fails, so we'll run the example 4 client in two different ways to demonstrate that the EJB deployment is in fact using SSL. Start the JBoss server using the default configuration and then run example 4b as follows:

```
[nr@toki examples]$ ant -Dchap=chap8 -Dex=4b run-example
...
run-example4b:
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
  [echo] Waiting for 15 seconds for deploy...
  [java] Exception in thread "main" java.rmi.ConnectIOException: error during JRMP connection establishment; nested exception is:
  [java]     javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: No trusted certificate found
  [java]     at sun.rmi.transport.tcp.TCPChannel.createConnection(TCPChannel.java:274)
  ...
  [java] Caused by: javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: No trusted certificate found
  [java]     at com.sun.net.ssl.internal.ssl.BaseSSLSocketImpl.a(DashoA12275)
  ...
  [java] Caused by: javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: No trusted certificate found
  [java]     at com.sun.net.ssl.internal.ssl.BaseSSLSocketImpl.a(DashoA12275)
  ...
```

The resulting exception is expected, and is the purpose of the 4b version of the example. Note that the excep-

tion stack trace has been edited to fit into the book format, so expect some difference. The key item to notice about the exception is it clearly shows you are using the Sun JSSE classes to communicate with the JBoss EJB container. The exception is saying that the self-signed certificate you are using as the JBoss server certificate cannot be validated as signed by any of the default certificate authorities. This is expected because the default certificate authority keystore that ships with the JSSE package only includes well known certificate authorities such as VeriSign, Thawte, and RSA Data Security. To get the EJB client to accept your self-signed certificate as valid, you need to tell the JSSE classes to use your `chap8.keystore` as its truststore. A truststore is just a keystore that contains public key certificates used to sign other certificates. To do this, run example 4 using `-Dex=4` rather than `-Dex=4b` to pass the location of the correct truststore using the `javax.net.ssl.trustStore` system property:

```
[nr@toki examples]$ ant -Dchap=chap8 -Dex=4 run-example
...
run-example4:
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
  [echo] Waiting for 5 seconds for deploy...
  [java] 0 [HandshakeCompletedNotify-Thread] DEBUG org.jboss.security.ssl.RMISSSLClientS
ocketFactory - SSL handshakeCompleted, cipher=SSL_RSA_WITH_RC4_128_MD5, peerHost=127.0.0.
1
    [java] Created Echo
    [java] Echo.echo()#1 = This is call 1
```

This time the only indication that an SSL socket is involved is because of the `SSL handshakeCompleted` message. This is coming from the `RMISSSLClientSocketFactory` class as a debug level log message. If you did not have the client configured to print out `log4j` debug level messages, there would be no direct indication that SSL was involved. If you note the run times and the load on your system CPU, there definitely is a difference. SSL, like SRP, involves the use of cryptographically strong random numbers that take time to seed the first time they are used. This shows up as high CPU utilization and start up times.

One consequence of this is that if you are running on a system that is slower than the one used to run the examples for the book, such as when running example 4b, you may see an exception similar to the following:

```
javax.naming.NameNotFoundException: EchoBean4 not bound
  at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer
  at sun.rmi.transport.StreamRemoteCall.executeCall
  at sun.rmi.server.UnicastRef.invoke
  at org.jnp.server.NamingServer_Stub.lookup
  at org.jnp.interfaces.NamingContext.lookup
  at org.jnp.interfaces.NamingContext.lookup
  at javax.naming.InitialContext.lookup
  at org.jboss.chap8.ex4.ExClient.main(ExClient.java:29)
```

The problem is that the JBoss server has not finished deploying the example EJB in the time the client allowed. This is due to the initial setup time of the secure random number generator used by the SSL server socket. If you see this issue, simply rerun the example again or increase the deployment wait time in the `chap8 build.xml` Ant script.

8.8. Configuring JBoss for use Behind a Firewall

JBoss comes with many socket based services that open listening ports. In this section we list the services that open ports that might need to be configured to work when accessing JBoss behind a firewall. The following table shows the ports, socket type, associated service for the services in the default configuration file set. Table 8.2 shows the same information for the additional ports that exist in the all configuration file set.

Table 8.1. The ports found in the default configuration

Port	Type	Service
1099	TCP	<code>org.jboss.naming.NamingService</code>
1098	TCP	<code>org.jboss.naming.NamingService</code>
1162	UDP	<code>org.jboss.jmx.adaptor.snmp.trapd.TrapdService</code>
4444	TCP	<code>org.jboss.invocation.jrmp.server.JRMPInvoker</code>
4445	TCP	<code>org.jboss.invocation.pooled.server.PooledInvoker</code>
8009	TCP	<code>org.jboss.web.tomcat.tc4.EmbeddedTomcatService</code>
8080	TCP	<code>org.jboss.web.tomcat.tc4.EmbeddedTomcatService</code>
8083	TCP	<code>org.jboss.web.WebService</code>
8090	TCP	<code>org.jboss.mq.il.oil.OILServerILService</code>
8092	TCP	<code>org.jboss.mq.il.oil2.OIL2ServerILService</code>
8093	TCP	<code>org.jboss.mq.il.uil2.UILServerILService</code>
0 ^a	TCP	<code>org.jboss.mq.il.rmi.RMIServerILService</code>
0 ^b	UDP	<code>org.jboss.jmx.adaptor.snmp.agent.SnmpAgentService</code>

^aThis service binds to an anonymous TCP port and does not support configuration of the port or bind interface.

^bThis service binds to an anonymous UDP port and does not support configuration of the port or bind interface.

Table 8.2. Additional ports in the all configuration

Port	Type	Service
1100	TCP	<code>org.jboss.ha.jndi.HANamingService</code>
0 ^a	TCP	<code>org.jboss.ha.jndi.HANamingService</code>
1102	UDP	<code>org.jboss.ha.jndi.HANamingService</code>
3528	TCP	<code>org.jboss.invocation.iiop.IIOPInvoker</code>
45566 ^b	UDP	<code>org.jboss.ha.framework.server.ClusterPartition</code>

^aCurrently anonymous but can be set via the `RmiPort` attribute.

^bPlus two additional anonymous UDP ports, one can be set using the `rcv_port`, and the other cannot be set.

8.9. How to Secure the JBoss Server

JBoss comes with several admin access points that need to be secured or removed to prevent unauthorized access to admin functions in a deployment. This section describes the various admin services and how to secure them.

8.9.1. The `jmx-console.war`

The `jmx-console.war` found in the `deploy` directory provides an html view into the JMX microkernel. As such, it provides access to arbitrary admin type access like shutting down the server, stopping services, deploying new services, etc. It should either be secured like any other web application, or removed.

8.9.2. The web-console.war

The `web-console.war` found in the `deploy/management` directory is another web application view into the JMX microkernel. This uses a combination of an applet and a HTML view and provides the same level of access to admin functionality as the `jmx-console.war`. As such, it should either be secured or removed. The `web-console.war` contains commented out templates for basic security in its `WEB-INF/web.xml` as well as commented out setup for a security domain in `WEB-INF/jboss-web.xml`.

8.9.3. The http-invoker.sar

The `http-invoker.sar` found in the `deploy` directory is a service that provides RMI/HTTP access for EJBs and the JNDI Naming service. This includes a servlet that processes posts of marshalled `org.jboss.invocation.Invocation` objects that represent invocations that should be dispatched onto the `MBeanServer`. Effectively this allows access to MBeans that support the detached invoker operation via HTTP since one could figure out how to format an appropriate HTTP post. To security this access point you would need to secure the `JMXInvokerServlet` servlet found in the `http-invoker.sar/invoker.war/WEB-INF/web.xml` descriptor. There is a secure mapping defined for the `/restricted/JMXInvokerServlet` path by default, one would simply have to remove the other paths and configure the `http-invoker` security domain setup in the `http-invoker.sar/invoker.war/WEB-INF/jboss-web.xml` descriptor.

8.9.4. The jmx-invoker-adaptor-server.sar

The `jmx-invoker-adaptor-server.sar` is a service that exposes the JMX `MBeanServer` interface via an RMI compatible interface using the RMI/JRMP detached invoker service. The only way for this service to be secured currently would be to switch the protocol to RMI/HTTP and secure the `http-invoker.sar` as described in the previous section. In the future this service will be deployed as an `XMBean` with a security interceptor that supports role based access checks. If your so inclined this is a configuration that can setup today following the procedure demonstrated in `XMBean` example: Section 2.4.3.2.3.

Integrating Servlet Containers

This chapter describes the steps for integrating a third party web container into the JBoss application server framework. A web container is a J2EE server component that enables access to servlets and JSP pages. The most widely used servlet container is Tomcat, and this is the default web container used by JBoss.

Integrating a servlet container into JBoss consists of mapping `web-app.xml` JNDI information into the JBoss JNDI namespace using an optional `jboss-web.xml` descriptor as well as delegating authentication and authorization to the JBoss security layer. The `org.jboss.web.AbstractWebContainer` class exists to simplify these tasks. The focus of the first part of this chapter is how to integrate a Web container using the `AbstractWebContainer` class. The chapter concludes with a discussion on configuration topics like the use of secure socket layer (SSL) encryption with the JBoss/Tomcat bundle, as well as how to configure Apache with the JBoss/Tomcat bundle.

9.1. The AbstractWebContainer Class

The `org.jboss.web.AbstractWebContainer` class is an implementation of a template pattern for web container integration into JBoss. Web container providers wishing to integrate their container into a JBoss server should create a subclass of `AbstractWebContainer` and provide the web container specific setup and WAR deployment steps. The `AbstractWebContainer` provides support for parsing the standard J2EE `web.xml` web application deployment descriptor JNDI and security elements as well as support for parsing the JBoss specific `jboss-web.xml` descriptor. Parsing of these deployment descriptors is performed to generate an integrated JNDI environment and security context. We have already seen the most of the elements of the `jboss-web.xml` descriptor in other chapters. Figure 9.1 provides an overview of the `jboss-web.xml` descriptor DTD for reference. The complete DTD with comments can be found in the `JBOSS_DIST/docs/dtd`.

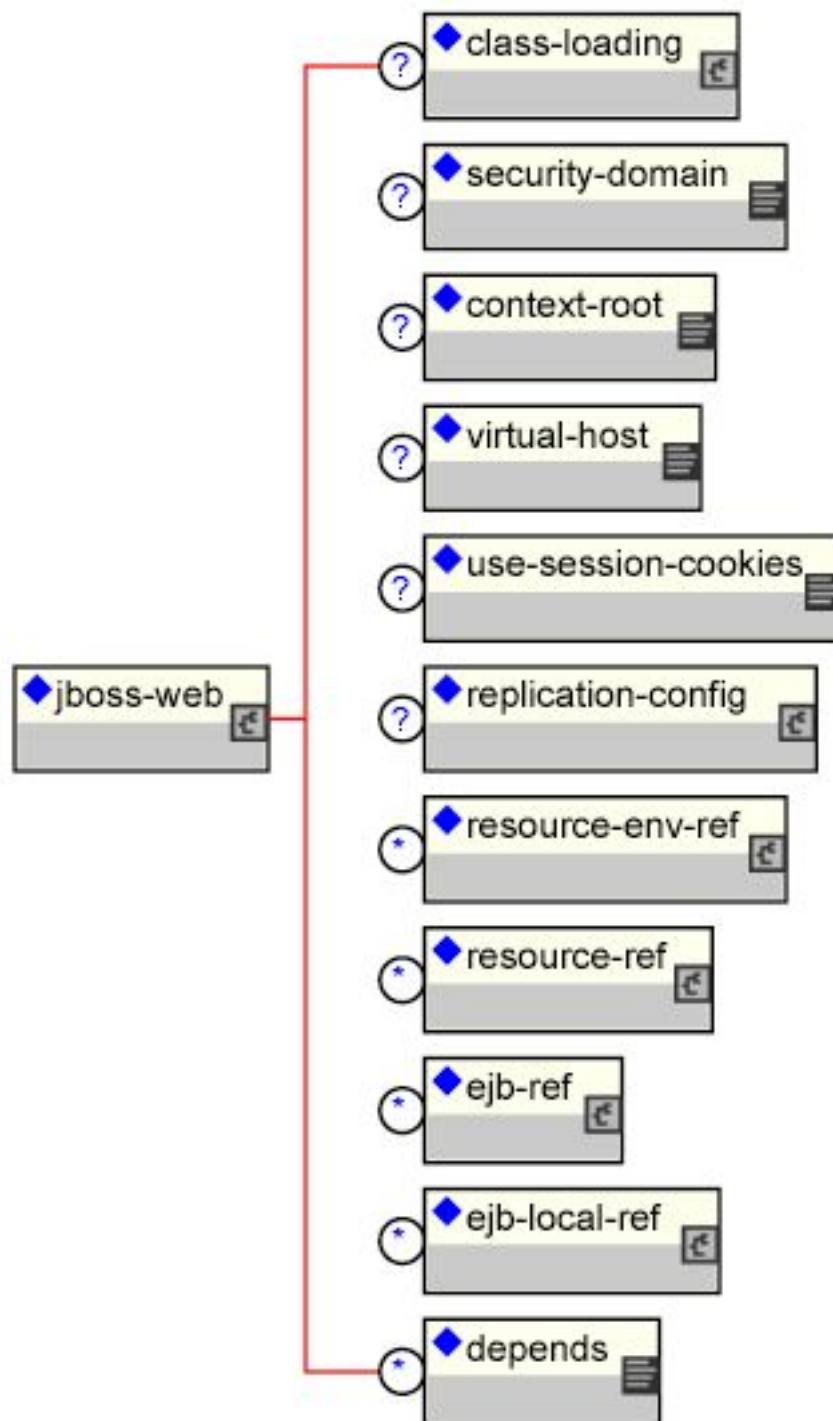


Figure 9.1. The complete jboss-web.xml descriptor DTD

The two elements that have not been discussed are the `context-root` and `virtual-host`. The `context-root` element allows one to specify the prefix under which web application is located. This is only applicable to stand-alone web application deployment as a WAR file. Web applications included as part of an EAR must set the root using the `context-root` element of the EAR `application.xml` descriptor. The sample `jboss-web.xml` descriptor shown in Example 9.1 illustrates mapping a war to the root context.

Example 9.1. A sample jboss-web.xml descriptor for mapping a war to the root context

```
<jboss-web>
```

```

<!-- An empty context root map the war to the root context,
     e.g., http://localhost:8080/ -->
<context-root />
</jboss-web>

```

The `virtual-host` element specifies the DNS name of the virtual host to which the web application should be deployed. The details of setting up virtual hosts for servlet contexts depends on the particular servlet container. We will look at examples of using the `virtual-host` element when we look at the Tomcat servlet containers later in this chapter.

9.1.1. The AbstractWebContainer Contract

The `AbstractWebContainer` is an abstract class that implements the `org.jboss.web.AbstractWebContainerMBean` interface used by the JBoss J2EE deployer to delegate the task of installing war files needing to be deployed. We'll look at some of the key methods of the `AbstractWebContainer` below.

```

public boolean accepts(DeploymentInfo sdi)
{
    String warFile = sdi.url.getFile();
    return warFile.endsWith("war") || warFile.endsWith("war/");
}

```

The `accepts` method is implemented by JBoss deployers to indicate which type of deployments they accepts. The `AbstractWebContainer` handles the deployments of WARs as JARs or unpacked directories.

```

public synchronized void start(DeploymentInfo di) throws DeploymentException
{
    Thread thread = Thread.currentThread();
    ClassLoader appClassLoader = thread.getContextClassLoader();

    try {
        // Create a classloader for the war to ensure a unique ENC
        URL[] empty = {};
        URLClassLoader warLoader = URLClassLoader.newInstance(empty, di.ucl);
        thread.setContextClassLoader(warLoader);
        WebDescriptorParser webAppParser = new DescriptorParser(di);

        String webContext = di.webContext;
        if (webContext != null) {
            if (webContext.length() > 0 && webContext.charAt(0) !=
                '/') {
                webContext = "/" + webContext;
            }
        }

        // Get the war URL
        URL warURL = di.localUrl != null ? di.localUrl : di.url;
        if (log.isDebugEnabled()) {
            log.debug("webContext: " + webContext);
            log.debug("warURL: " + warURL);
            log.debug("webAppParser: " + webAppParser);
        }

        // Parse the web.xml and jboss-web.xml descriptors
        WebMetaData metaData = (WebMetaData) di.metaData;
        parseMetaData(webContext, warURL, di.shortName, metaData);

        WebApplication warInfo = new WebApplication(metaData);
        warInfo.setDeploymentInfo(di);
        performDeploy(warInfo, warURL.toString(), webAppParser);
        deploymentMap.put(warURL.toString(), warInfo);
    }
}

```

```

        // Generate an event for the startup
        super.start(di);
    } catch (DeploymentException e) {
        throw e;
    } catch (Exception e) {
        throw new DeploymentException("Error during deploy", e);
    } finally {
        thread.setContextClassLoader(appClassLoader);
    }
}

```

This section corresponds to the `start` method. This method is a template pattern method implementation. The argument to the `deploy` method is the WAR deployment info object. This contains the URL to the WAR, the `UnifiedClassLoader` for the WAR, the parent archive such as an EAR, and the J2EE `application.xml` `context-root` if the WAR is part of an EAR.

The first step of the `start` method is to save the current thread context class loader and then create another `URLClassLoader` (`warLoader`) using the WAR `UnifiedClassLoader` as its parent. This `warLoader` is used to ensure a unique JNDI ENC (enterprise naming context) for the WAR will be created. Chapter 3 mentioned that the `java:comp` context's uniqueness was determined by the class loader that created the `java:comp` context. The `warLoader ClassLoader` is set as the current thread context class loader before the `performDeploy` call is made. Next, the `web.xml` and `jboss-web.xml` descriptors are parsed by calling `parseMetaData`. Next, the web container-specific subclass is asked to perform the actual deployment of the WAR through the `performDeploy` call. The `WebApplication` object for this deployment is stored in the deployed application map using the `warUrl` as the key. The final step is to restore the thread context class loader to the one that existed at the start of the method.

```

protected abstract void performDeploy(WebApplication webApp, String warUrl,
                                     WebDescriptorParser webAppParser)
    throws Exception;

```

This is the signature for the abstract `performDeploy` method. This method is called by the `start` method and must be overridden by subclasses to perform the web container specific deployment steps. A `WebApplication` is provided as an argument, and this contains the metadata from the `web.xml` descriptor, and the `jboss-web.xml` descriptor. The metadata contains the `context-root` value for the web module from the J2EE `application.xml` descriptor, or if this is a stand-alone deployment, the `jboss-web.xml` descriptor. The metadata also contains any `jboss-web.xml` descriptor `virtual-host` value. On return from `performDeploy`, the `WebApplication` must be populated with the class loader of the servlet context for the deployment. The `warUrl` argument is the string for the URL of the Web application WAR to deploy. The `webAppParser` argument is a callback handle the subclass must use to invoke the `parseWebAppDescriptors` method to set up the Web application JNDI environment. This callback provides a hook for the subclass to establish the web application JNDI environment before any servlets are created that are to be loaded on startup of the WAR. A subclass' `performDeploy` method implementation needs to be arranged so that it can call the `parseWebAppDescriptors` before starting any servlets that need to access JNDI for JBoss resources like EJBs, resource factories, and so on. One important setup detail that needs to be handled by a subclass implementation is to use the current thread context class loader as the parent class loader for any Web container-specific class loader created. Failure to do this results in problems for web applications that attempt to access EJBs or JBoss resources through the JNDI ENC.

```

public synchronized void stop(DeploymentInfo di)
    throws DeploymentException
{
    URL warURL = di.localUrl != null ? di.localUrl : di.url;
    String warUrl = warURL.toString();
    try {
        performUndeploy(warUrl);
        // Remove the web application ENC...
        deploymentMap.remove(warUrl);
        // Generate an event for the stop
    }
}

```

```

        super.stop(di);
    } catch(DeploymentException e) {
        throw e;
    } catch(Exception e) {
        throw new DeploymentException("Error during deploy", e);
    }
}

```

This is the `stop` method. It calls the subclass `performUndeploy` method to perform the container-specific undeployment steps. After undeploying the application, the `warUrl` is unregistered from the deployment map. The `warUrl` argument is the string URL of the WAR as originally passed to the `performDeploy` method.

```
protected abstract void performUndeploy(String warUrl) throws Exception;
```

This is the signature of the abstract `performUndeploy` method, which is called from the `stop` method. A call to `performUndeploy` asks the subclass to perform the Web container-specific undeployment steps.

```

public void setConfig(Element config)
{
}

```

The `setConfig` method is a stub method that subclasses can override if they want to support an arbitrary extended configuration beyond that which is possible through MBean attributes. The `config` argument is the parent DOM element for an arbitrary hierarchy given by the child element of the `Config` attribute in the `mbean` element specification of the `jboss-service.xml` descriptor of the web container service. You'll see an example use of this method and `config` value when you look at the MBean that supports embedding Tomcat into JBoss.

```

protected void parseWebAppDescriptors(DeploymentInfo di,
                                      ClassLoader loader,
                                      WebMetaData metaData)
    throws Exception
{
    log.debug("AbstractWebContainer.parseWebAppDescriptors, Begin");
    InitialContext iniCtx = new InitialContext();
    Context envCtx = null;
    Thread currentThread = Thread.currentThread();
    ClassLoader currentLoader = currentThread.getContextClassLoader();
    try {
        // Create a java:comp/env environment unique for the web application
        log.debug("Creating ENC using ClassLoader: "+loader);
        ClassLoader parent = loader.getParent();
        while (parent != null ) {
            log.debug(".." +parent);
            parent = parent.getParent();
        }
        currentThread.setContextClassLoader(loader);
        metaData.setENCLoader(loader);
        envCtx = (Context) iniCtx.lookup("java:comp");
        // Add a link to the global transaction manager
        envCtx.bind("UserTransaction", new LinkRef("UserTransaction"));
        log.debug("Linked java:comp/UserTransaction to JNDI name: UserTransaction");
        envCtx = envCtx.createSubcontext("env");
    } finally {
        currentThread.setContextClassLoader(currentLoader);
    }

    Iterator envEntries = metaData.getEnvironmentEntries();
    log.debug("addEnvEntries");
    addEnvEntries(envEntries, envCtx);

    Iterator resourceEnvRefs = metaData.getResourceEnvReferences();
    log.debug("linkResourceEnvRefs");
    linkResourceEnvRefs(resourceEnvRefs, envCtx);

    Iterator resourceRefs = metaData.getResourceReferences();

```

```
log.debug("linkResourceRefs");
linkResourceRefs(resourceRefs, envCtx);

Iterator ejbRefs = metaData.getEjbReferences();
log.debug("linkEjbRefs");
linkEjbRefs(ejbRefs, envCtx, di);

Iterator ejbLocalRefs = metaData.getEjbLocalReferences();
log.debug("linkEjbLocalRefs");
linkEjbLocalRefs(ejbLocalRefs, envCtx, di);

String securityDomain = metaData.getSecurityDomain();
log.debug("linkSecurityDomain");
linkSecurityDomain(securityDomain, envCtx);

log.debug("AbstractWebContainer.parseWebAppDescriptors, End");
}
```

The `parseWebAppDescriptors` method is invoked from within the subclass `performDeploy` method when it invokes the `webAppParser.parseWebAppDescriptors` callback to setup the web application ENC (`java:comp/env`) `env-entry`, `resource-env-ref`, `resource-ref`, `local-ejb-ref` and `ejb-ref` values declared in the `web.xml` descriptor. The creation of the `env-entry` values does not require a `jboss-web.xml` descriptor. The creation of the `resource-env-ref`, `resource-ref`, and `ejb-ref` elements does require a `jboss-web.xml` descriptor for the JNDI name of the deployed resources/EJBs. Because the ENC context is private to the web application, the web application class loader is used to identify the ENC. The loader argument is the class loader for the web application, and may not be null. The `metaData` argument is the `WebMetaData` argument passed to the subclass `performDeploy` method. The implementation of the `parseWebAppDescriptors` uses the metadata information from the WAR deployment descriptors and then creates the JNDI ENC bindings.

```
protected void addEnvEntries(Iterator envEntries, Context envCtx)
    throws ClassNotFoundException, NamingException
{
}
```

The `addEnvEntries` method creates the `java:comp/env` web application `env-entry` bindings that were specified in the `web.xml` descriptor.

```
protected void linkResourceEnvRefs(Iterator resourceEnvRefs, Context envCtx)
    throws NamingException
{
}
```

The `linkResourceEnvRefs` method maps the `java:comp/env/xxx` web application JNDI ENC `resource-env-ref` `web.xml` descriptor elements onto the deployed JNDI names using the mappings specified in the `jboss-web.xml` descriptor.

```
protected void linkResourceRefs(Iterator resourceRefs, Context envCtx)
    throws NamingException
{
}
```

The `linkResourceRefs` method maps the `java:comp/env/xxx` web application JNDI ENC `resource-ref` `web.xml` descriptor elements onto the deployed JNDI names using the mappings specified in the `jboss-web.xml` descriptor.

```
protected void linkEjbRefs(Iterator ejbRefs, Context envCtx, DeploymentInfo di)
    throws NamingException
{
}
```


The `linkEjbRefs` method maps the `java:comp/env/ejb` web application JNDI ENC `ejb-ref` `web.xml` descriptor elements onto the deployed JNDI names using the mappings specified in the `jboss-web.xml` descriptor.

```
protected void linkEjbLocalRefs(Iterator ejbRefs, Context envCtx,
                               DeploymentInfo di)
    throws NamingException
{
}
```

The `linkEjbLocalRefs` method maps the `java:comp/env/ejb` Web application JNDI ENC `ejb-local-ref` `web.xml` descriptor elements onto the deployed JNDI names using the `ejb-link` mappings specified in the `web.xml` descriptor.

```
protected void linkSecurityDomain(String securityDomain, Context envCtx)
    throws NamingException
{
}
```

The `linkSecurityDomain` method creates a `java:comp/env/security` context that contains a `securityMgr` binding pointing to the `AuthenticationManager` implementation and a `realmMapping` binding pointing to the `RealmMapping` implementation that is associated with the security domain for the web application. Also creates is a subject binding that provides dynamic access to the authenticated `Subject` associated with the request thread. If the `jboss-web.xml` descriptor contained a `security-domain` element, the bindings are `javax.naming.LinkRefs` to the JNDI name specified by the `security-domain` element, or subcontexts of this name. If there was no `security-domain` element, the bindings are to `org.jboss.security.plugins.NullSecurityManager` instance that simply allows all authentication and authorization checks.

```
public String[] getCompileClasspath(ClassLoader loader)
{
}
```

The `getCompileClasspath` method is a utility method available for web containers to generate a classpath that walks up the class loader chain starting at the given loader and queries each class loader for the URLs it serves to build a complete classpath of URL strings. This is needed by some JSP compiler implementations (Jasper for one) that expect to be given a complete classpath for compilation.

9.1.2. Creating an AbstractWebContainer Subclass

To integrate a web container into JBoss you need to create a subclass of `AbstractWebContainer` and implement the required `performDeploy(WebApplication, String, WebDescriptorParser)` and `performUndeploy(String)` methods as described in the preceding section. The following additional integration points should be considered as well.

9.1.2.1. Use the Thread Context Class Loader

Although this issue was noted in the `performDeploy` method description, we'll repeat it here since it is such a critical detail. During the setup of a WAR container, the current thread context class loader must be used as the parent class loader for any web container specific class loader that is created. Failure to do this will result in problems for web applications that attempt to access EJBs or JBoss resources through the JNDI ENC.

9.1.2.2. Integrate Logging Using log4j

JBoss uses the Apache log4j logging API as its internal logging API. For a web container to integrate well with JBoss it needs to provide a mapping between the web container logging abstraction to the log4j API. As a subclass of `AbstractWebContainer`, your integration class has access to the log4j interface via the `super.log` instance variable or equivalently, the superclass `getLog()` method. This is an instance of the `org.jboss.logging.Logger` class that wraps the log4j category. The name of the log4j category is the name of the container subclass.

9.1.2.3. Delegate web container authentication and authorization to JBossSX

Ideally both web application and EJB authentication and authorization are handled by the same security manager. To enable this for your web container you must hook into the JBoss security layer. This typically requires a request interceptor that maps from the web container security callouts to the JBoss security API. Integration with the JBossSX security framework is based on the establishment of a `java:comp/env/security` context as described in the `linkSecurityDomain` method comments in the previous section. The security context provides access to the JBossSX security manager interface implementations associated with the web application for use by subclass request interceptors. An outline of the steps for authenticating a user using the security context is presented in Example 9.2 in quasi pseudo-code. Example 9.3 provides the equivalent process for the authorization of a user.

Example 9.2. A pseudo-code description of authenticating a user via the JBossSX API and the `java:comp/env/security` JNDI context.

```
// Get the username and password from the request context...
HttpServletRequest request = ...;
String username = getUsername(request);
String password = getPassword(request);

// Get the JBoss security manager from the ENC context
InitialContext iniCtx = new InitialContext();
AuthenticationManager securityMgr = (AuthenticationManager)
    iniCtx.lookup("java:comp/env/security/securityMgr");

SimplePrincipal principal = new SimplePrincipal(username);
if (securityMgr.isValid(principal, password)) {
    // Indicate the user is allowed access to the web content...
    // Propagate the user info to JBoss for any calls into made by the servlet
    SecurityAssociation.setPrincipal(principal);
    SecurityAssociation.setCredential(password.toCharArray());
} else {
    // Deny access...
}
```

Example 9.3. A pseudo-code description of authorization a user via the JBossSX API and the `java:comp/env/security` JNDI context.

```
// Get the username & required roles from the request context...
HttpServletRequest request = ...;
String username = getUsername(request);
String[] roles = getContentRoles(request);

// Get the JBoss security manager from the ENC context
InitialContext iniCtx = new InitialContext();
RealmMapping securityMgr = (RealmMapping)
    iniCtx.lookup("java:comp/env/security/realmMapping");

SimplePrincipal principal = new SimplePrincipal(username);
Set requiredRoles = new HashSet(java.util.Arrays.asList(roles));
```

```
if (securityMgr.doesUserHaveRole(principal, requiredRoles)) {  
    // Indicate user has the required roles for the web content...  
} else {  
    // Deny access...  
}
```

9.2. JBoss/Tomcat-5 bundle notes

In this section we'll discuss configuration issues specific to the JBoss/Tomcat 5 integration bundle. Tomcat 5 is the latest release of the Apache Java servlet container. It supports the Servlet 2.4 and JSP 2.0 specifications. The JBoss/Tomcat integration layer is controlled by the JBoss MBean service configuration. The MBean used to embed the Tomcat-4.1.x series of web containers is `org.jboss.web.tomcat.tc5.Tomcat5`, and it is a subclass of the `AbstractWebContainer` class. Its configurable attributes include:

- **Java2ClassLoadingCompliance:** enables the standard Java2 parent delegation class loading model rather than the servlet model which loads from the WAR. This is true by default as loading from WARs that include client JARs with classes used by EJBs causes class loading conflicts. If you enable the servlet class loading model by setting this flag to false, you will need to organize your deployment package to avoid duplicate classes in the deployment.
- **UseJBossWebLoader:** A flag indicating if the class loader used by Tomcat as the web application class loader is a JBoss unified class loader. The default is true, which means that the classes available in the WAR inside of the `WEB-INF/classes` and `WEB-INF/lib` are incorporated into the default shared class loader repository described in Chapter 2. This may not be what you want as its contrary to the default servlet class loading model and can result in sharing of classes/resources between web applications. You can disable this by setting this attribute to false.
- **ManagerClass:** This is the class to use as the session manager for replicating the state of web applications marked as distributable. The only provided implementation session manager is `org.jboss.web.tomcat.tc5.session.JBossCacheManager`, which uses `JBossCache` to track the distributed state.
- **SnapshotMode:** Set the snapshot mode in a clustered environment. This must be one of `instant` or `interval`. In instant mode changes to a clustered session are instantly propagated whenever a modification is made. In interval mode all modifications are periodically propagated according to the `SnapshotInterval`.
- **SnapshotInterval:** Set the snapshot interval in ms for the `interval` snapshot mode. The default is 1000 ms, which is 1 second.

9.2.1. The Tomcat server.xml file

While the `jboss-service.xml` file controls the JBoss/Tomcat integration, Tomcat has its own configuration file which guides its operation. This is the `server.xml` descriptor that you will find in the `deploy/jboss-web-tomcat50.sar` directory.

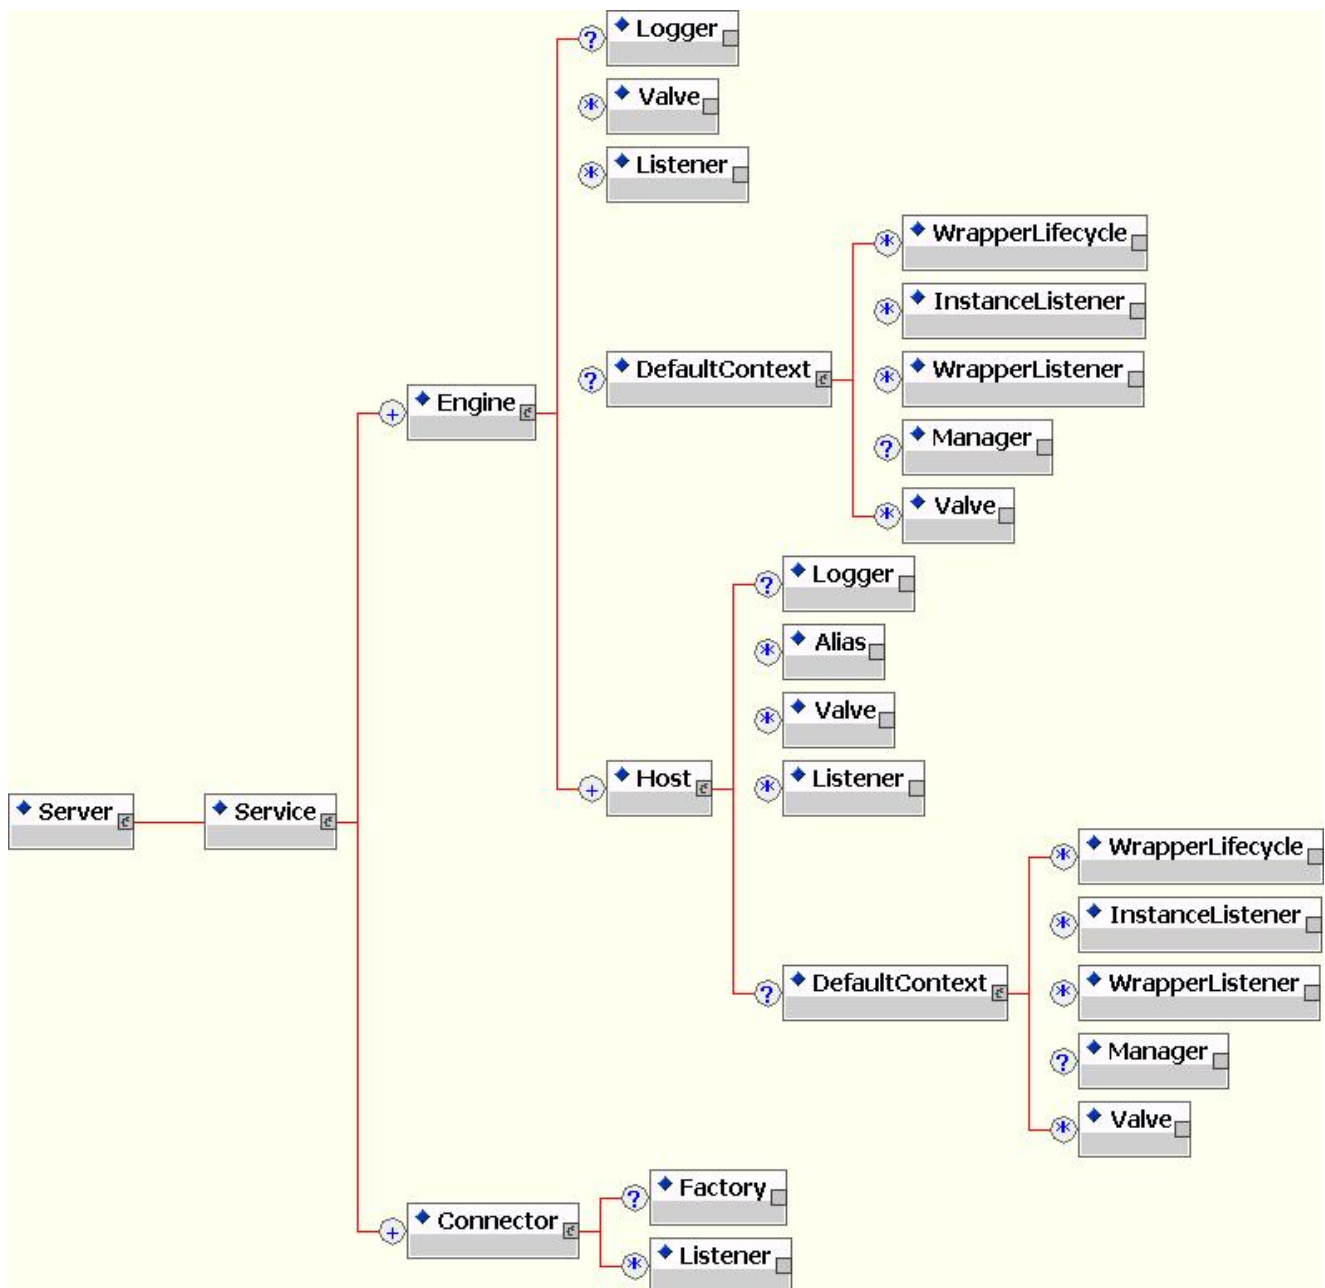


Figure 9.2. An overview of the Tomcat 5 configuration DTD supported by the `server.xml` file.

We'll now look at some of the configuration options available in the `server.xml` file. The top level element is the `Server` element, which should contain a `Service` element representing the entire web subsystem. The only supported attribute is:

- **name:** a unique name by which the service is known.

9.2.1.1. Connector

A `Connector` element configures a transport mechanism that allows clients to send requests and receive responses from the `Service` it is associated with. Connectors forward requests to the engine and return the results to the requesting client. Connectors support these attributes:

- **enableLookups:** a flag that enables DNS resolution of the client hostname as accessed via the `ServletRe-`

`quest.getRemoteHost` method. This flag defaults to false.

- **redirectPort:** the port to which non-SSL requests will be redirected when a request for content secured under a transport confidentiality or integrity constraint is received. This defaults to the standard HTTPS port of 443.
- **secure:** sets the `ServletRequest.isSecure` method value flag to indicate whether or not the transport channel is secure. This flag defaults to false.
- **scheme:** sets the protocol name as accessed by the `ServletRequest.getScheme` method. The scheme defaults to `http`.
- **acceptCount:** The maximum queue length for incoming connection requests when all possible request processing threads are in use. Any requests received when the queue is full will be refused. The default value is 10.
- **address.** For servers with more than one IP address, this attribute specifies which address will be used for listening on the specified port. By default, this port will be used on all IP addresses associated with the server.
- **bufferSize:** The size (in bytes) of the buffer to be provided for input streams created by this connector. By default, buffers of 2048 bytes will be provided.
- **connectionTimeout:** The number of milliseconds this connector will wait, after accepting a connection, for the request URI line to be presented. The default value is 60000 (i.e. 60 seconds).
- **debug:** The debugging detail level of log messages generated by this component, with higher numbers creating more detailed output. If not specified, this attribute is set to zero (0). Whether or not this shows up in the log further depends on the log4j category `org.jboss.web.tomcat.tc5.Tomcat5` threshold.
- **maxThreads:** The maximum number of request processing threads to be created by this connector, which therefore determines the maximum number of simultaneous requests that can be handled. If not specified, this attribute is set to 200.
- **maxSpareThreads:** The maximum number of unused request processing threads that will be allowed to exist until the thread pool starts stopping the unnecessary threads. The default value is 50.
- **minSpareThreads:** The number of request processing threads that will be created when this connector is first started. The connector will also make sure it has the specified number of idle processing threads available. This attribute should be set to a value smaller than that set for `maxThreads`. The default value is 4.
- **port:** The TCP port number on which this connector will create a server socket and await incoming connections. Your operating system will allow only one server application to listen to a particular port number on a particular IP address.
- **proxyName:** If this connector is being used in a proxy configuration, configure this attribute to specify the server name to be returned for calls to `request.getServerName()`.
- **proxyPort:** If this connector is being used in a proxy configuration, configure this attribute to specify the server port to be returned for calls to `request.getServerPort()`.
- **tcpNoDelay:** If set to true, the `TCP_NO_DELAY` option will be set on the server socket, which improves performance under most circumstances. This is set to true by default.

Additional attribute descriptions may be found in the Tomcat website document: <http://tomcat.apache.org/tomcat-5.0-doc/config/connector.html>

[tp://jakarta.apache.org/tomcat/tomcat-5.0-doc/config/http11.html](http://jakarta.apache.org/tomcat/tomcat-5.0-doc/config/http11.html)

9.2.1.2. Engine

Each `Service` must have a single `Engine` configuration. An engine handles the requests submitted to a service via the configured connectors. The child elements supported by the embedded service include `Host`, `Logger`, `DefaultContext`, `Valve` and `Listener`. The supported attributes include:

- **className:** the fully qualified class name of the `org.apache.catalina.Engine` interface implementation to use. If not specified this defaults to `org.apache.catalina.core.StandardEngine`.
- **defaultHost:** the name of a `Host` configured under the `Engine` that will handle requests with host names that do not match a `Host` configuration.
- **name:** a logical name to assign the `Engine`. It will be used in log messages produced by the `Engine`.

Additional information on the `Engine` element may be found in the Tomcat website document <http://jakarta.apache.org/tomcat/tomcat-5.1-doc/config/engine.html>.

9.2.1.3. Host

A `Host` element represents a virtual host configuration. It is a container for web applications with a specified DNS hostname. The child elements supported by the embedded service include `Alias`, `Logger`, `DefaultContext`, `Valve` and `Listener`. The supported attributes include:

- **className:** the fully qualified class name of the `org.apache.catalina.Host` interface implementation to use. If not specified this defaults to `org.apache.catalina.core.StandardHost`.
- **name:** the DNS name of the virtual host. At least one `Host` element must be configured with a name that corresponds to the `defaultHost` value of the containing `Engine`.

The `Alias` element is an optional child element of the `Host` element. Each `Alias` content specifies an alternate DNS name for the enclosing `Host`.

Additional information on the `Host` element may be found in the Tomcat website document <http://jakarta.apache.org/tomcat/tomcat-4.1-doc/config/host.html>.

9.2.1.4. DefaultContext

The `DefaultContext` element is a configuration template for web application contexts. It may be defined at the `Engine` or `Host` level. The child elements supported by the embedded service include `WrapperLifecycle`, `InstanceListener`, `WrapperListener`, and `Manager`. The supported attributes include:

- **className:** the fully qualified class name of the `org.apache.catalina.core.DefaultContext` implementation. This defaults to `org.apache.catalina.core.DefaultContext` and if overridden must be a subclass of `DefaultContext`.
- **cookies:** a flag indicating if sessions will be tracked using cookies. The default is true.
- **crossContext:** A flag indicating if the `ServletContext.getContext(String path)` method should return contexts for other web applications deployed in the calling web application's virtual host. The default is false.

9.2.1.5. Logger

The `Logger` element specifies a logging configuration the Tomcat instance. The supported attributes include:

- **className**: The fully qualified class name of the `org.apache.catalina.Logger` interface implementation. For integration with JBoss logging this should be set to `org.jboss.web.tomcat.Log4jLogger`.
- **verbosity**: The default log level.
- **category**: The default log category.

9.2.1.6. Valve

A `Valve` element configures a hook into the request processing pipeline for the web container. Valves must implement the `org.apache.catalina.Valve` interface. There is only one required configuration attribute:

- **className**: The fully qualified class name of the `org.apache.catalina.Valve` interface implementation.

The most commonly used valve is the `AccessLogValve`, which keeps a standard HTTP access log of incoming requests. The `className` for the access log value is `org.jboss.web.catalina.valves.AccessLogValue`. The addition Valve attributes supported by it include:

- **directory**: The directory path into which the access log files will be created.
- **pattern**: A pattern specifier that defines the format of the log messages. This defaults to `common`.
- **prefix**: The prefix to add to each log file name. This defaults to `access_log`.
- **suffix**: The suffix to add to each log file name. This default to an empty string, meaning that no suffix will be added.

Additional information on the `Valve` element and the available valve implementations may be found in the Tomcat website document <http://jakarta.apache.org/tomcat/tomcat-5.0-doc/config/valve.html>.

9.2.2. Using SSL with the JBoss/Tomcat bundle

There are a few ways one can configure HTTP over SSL for the embedded Tomcat servlet container. The main difference is whether or not you use the JBoss specific connector socket factory that allows one to obtain the JSSE server certificate information from a `JBossSX SecurityDomain`. This requires establishing a `SecurityDomain` using the `org.jboss.security.plugins.JaasSecurityDomain` MBean. These two steps are similar to the procedure we used in Chapter 8 to enable RMI with SSL encryption. A `server.xml` configuration file that illustrates the setup of only an SSL connector via this approach is given in Example 9.4. This configuration includes the same `JaasSecurityDomain` setup as Chapter 8, but since the descriptor is not being deployed as part of a SAR that includes the `chap8.keystore`, you need to copy the `chap8.keystore` to the `server/default/conf` directory.

Example 9.4. The `JaasSecurityDomain` and `EmbeddedCatalinaSX` MBean configurations for setting up Tomcat 5 to use SSL as its primary connector protocol.

```
<Server>
  <Service name="jboss.web" className="org.jboss.web.tomcat.tc5.StandardService">
```

```
<Connector port="8080" address="${jboss.bind.address}" maxThreads="150"
  minSpareThreads="25" maxSpareThreads="75" enableLookups="false"
  redirectPort="443" acceptCount="100" connectionTimeout="20000"
  disableUploadTimeout="true"/>

<Connector port="443" address="${jboss.bind.address}" maxThreads="100"
  minSpareThreads="5" maxSpareThreads="15" scheme="https"
  secure="true" clientAuth="false"
  keystoreFile="${jboss.server.home.dir}/conf/chap8.keystore"
  keystorePass="rmi+ssl" sslProtocol="TLS"/>

<Engine name="jboss.web" defaultHost="localhost">
  <Realm
    className="org.jboss.web.tomcat.security.JBossSecurityMgrRealm"
    certificatePrincipal="org.jboss.security.auth.certs.SubjectDNMapping"/>
  <Logger className="org.jboss.web.tomcat.Log4jLogger"
    verbosityLevel="WARNING" category="org.jboss.web.localhost.Engine"/>
  <Host name="localhost" autoDeploy="false" deployOnStartup="false"
    deployXML="false">
    <DefaultContext cookies="true" crossContext="true" override="true"/>
  </Host>
</Engine>
</Service>
</Server>
```

A quick test of this config can be made by accessing the JMX console web application using this URL <https://localhost/jmx-console/index.jsp>.

Note: if your running on a *nix system (Linux, Solaris, OS X) that only allows root to open ports below 1024 you will need to change the port number above to something like 8443.

Factory configuration attributes:

- **algorithm:** The certificate encoding algorithm to be used. If not specified, the default value is SunX509.
- **className:** The fully qualified class name of the SSL server socket factory implementation class. You must specify `org.apache.coyote.tomcat4.CoyoteServerSocketFactory` here. Using any other socket factory will not cause an error, but the server socket will not be using SSL.
- **clientAuth:** Set to true if you want the SSL stack to require a valid certificate chain from the client before accepting a connection. A false value (which is the default) will not require a certificate chain unless the client requests a resource protected by a security constraint that uses CLIENT-CERT authentication.
- **keystoreFile:** The pathname of the keystore file where you have stored the server certificate to be loaded. By default, the pathname is the file ".keystore" in the operating system home directory of the user that is running Tomcat.
- **keystorePass:** The password used to access the server certificate from the specified keystore file. The default value is "changeit".
- **keystoreType:** The type of keystore file to be used for the server certificate. If not specified, the default value is "JKS".
- **protocol:** The version of the SSL protocol to use. If not specified, the default is "TLS".

Note that if you try to test this configuration using the self-signed certificate from the Chapter 8 `chap8.keystore` and attempt to access content over an HTTPS connection, your browser should display a warning dialog indicating that it does not trust the certificate authority that signed the certificate of the server

you are connecting to. For example, when the first configuration example was tested, IE 5.5 showed the initial security alert dialog listed in Figure 9.3. Figure 9.4 shows the server certificate details. This warning is important as anyone can generate a self-signed certificate with any information they want. Your only way to verify that the system on the other side really represents the party it claim to is by verifying that it is signed by a trusted 3rd party.



Figure 9.3. The Internet Explorer 5.5 security alert dialog.

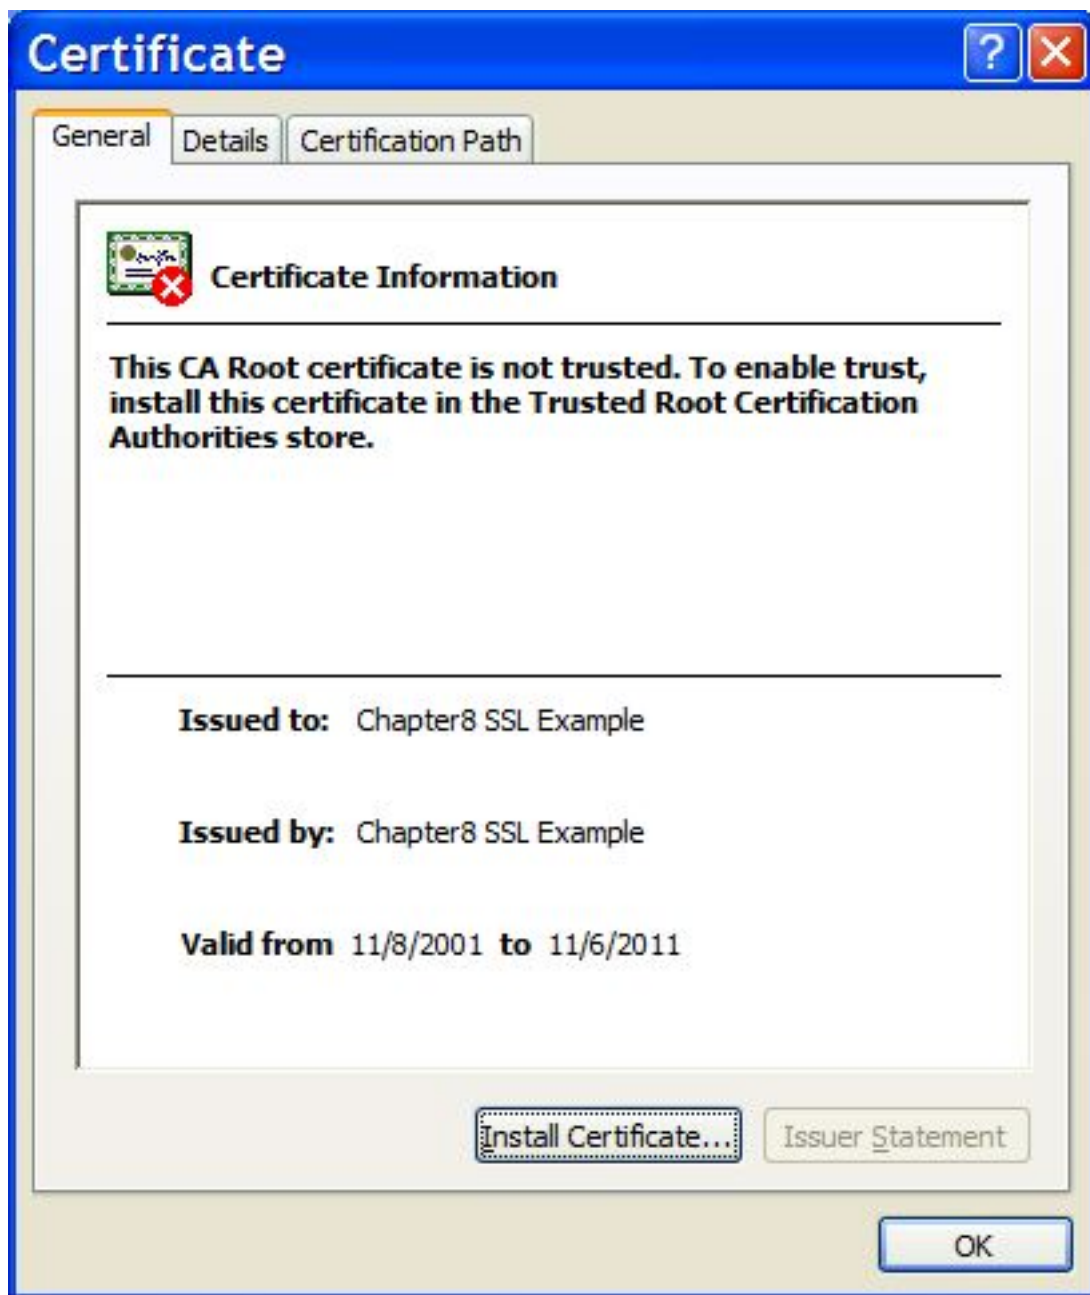


Figure 9.4. The Internet Explorer 5.5 SSL certificate details dialog.

9.2.3. Setting up Virtual Hosts

Virtual hosts allow you to group web applications according to the various DNS names by which the machine running JBoss is known. As an example, consider the `server.xml` configuration file given in Example 9.5. This configuration defines a default host named `vhost1.mydot.com` and a second host named `vhost2.mydot.com`, which also has the alias `www.mydot.com` associated with it.

Example 9.5. An example virtual host configuration.

```
<Server>
  <Service name="jboss.web"
    className="org.jboss.web.tomcat.tc5.StandardService">

    <!-- A HTTP/1.1 Connector on port 8080 -->
```

```

<Connector port="8080" address="${jboss.bind.address}"
    maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
    enableLookups="false" redirectPort="8443" acceptCount="100"
    connectionTimeout="20000" disableUploadTimeout="true"/>

<Engine name="jboss.web" defaultHost="vhost1">
    <Realm className="org.jboss.web.tomcat.security.JBossSecurityMgrRealm"
        certificatePrincipal="org.jboss.security.auth.certs.SubjectDNMapping"
    />
    <Logger className="org.jboss.web.tomcat.Log4jLogger"
        verbosityLevel="WARNING"
        category="org.jboss.web.localhost.Engine"/>

    <Host name="vhost1" autoDeploy="false"
        deployOnStartup="false" deployXML="false">
        <Alias>vhost1.mydot.com</Alias>
        <Valve className="org.apache.catalina.valves.AccessLogValve"
            prefix="vhost1" suffix=".log" pattern="common"
            directory="${jboss.server.home.dir}/log"/>

        <DefaultContext cookies="true" crossContext="true" override="true"/>
    </Host>
    <Host name="vhost2" autoDeploy="false"
        deployOnStartup="false" deployXML="false">
        <Alias>vhost2.mydot.com</Alias>
        <Alias>www.mydot.com</Alias>

        <Valve className="org.apache.catalina.valves.AccessLogValve"
            prefix="vhost2" suffix=".log" pattern="common"
            directory="${jboss.server.home.dir}/log"/>

        <DefaultContext cookies="true" crossContext="true" override="true"/>
    </Host>
</Engine>
</Service>
</Server>

```

When a WAR file is deployed, it will be by default associated with the virtual host whose name matches the `defaultHost` attribute of the containing `Engine`. To deploy a WAR to a specific virtual host you need to specify an appropriate `virtual-host` definition in your `e jboss-web.xml` descriptor. The following `jboss-web.xml` descriptor demonstrates how to deploy a WAR to the virtual host `www.mydot.com`. Note that we can use either the virtual host name in the config file and the actual host name.

```

<jboss-web>
    <context-root>/</context-root>
    <virtual-host>www.mydot.com</virtual-host>
</jboss-web>

```

9.2.4. Serving Static Content

JBoss provides a default application that serves content for the `root` application context. This default context is the `ROOT.war` application in the `jbossweb-tomcat50.sar` directory. You can serve static files not associated with any other application by adding that content to the `ROOT.war` directory. For example, if you want to have a shared image directory you could create an `image` subdirectory inside of `ROOT.war` and place the images there. You could then access an image named `myimage.jpg` at `http://localhost:8080/images/myimage.jpg`.

9.2.5. Using Apache with the Tomcat

In some architectures, it is useful to put an Apache web server in front of the JBoss server. External web clients

talk to an Apache instance, which in turn speaks to the Tomcat instance on behalf of the clients. Apache needs to be configured to use the `mod_jk` module which speaks the AJP protocol to an AJP connector running in Tomcat. The provided `server.xml` file comes with this AJP connector enabled.

```
<Connector port="8009" address="${jboss.bind.address}"
  enableLookups="false" redirectPort="8443" debug="0"
  protocol="AJP/1.3" />
```

You'll need to consult the Apache and `mod_jk` documentation for complete installation instructions. Assuming you have a properly configured Apache instance, the following configuration fragment shows an example of how to connect with a WAR deployed with a context root of `/jbosstest`.

```
...
LoadModule jk_module libexec/mod_jk.so
AddModule mod_jk.c

<IfModule mod_jk.c>
  JkWorkersFile /tmp/workers.properties
  JkLogFile /tmp/mod_jk.log
  JkLogLevel debug
  JkMount /jbosstest/* ajp13
</IfModule>
```

The `workers.properties` file contains the details of how to contact the JBoss instance.

9.2.6. Using Clustering

JBoss supports clustering in the embedded Tomcat service. The steps to setup clustering of Tomcat embedded containers is:

- If you are using a load balancer, make sure that your setup uses sticky sessions. This means that if a user that starts a session on node A, all subsequent requests are forwarded to node A as long node A is up and running. For configuration of the Apache webserver sticky sessions see <http://www.ubbeans.com/tomcat/> for details.
- If you aren't using the `all` configuration, make sure that `cluster-service.xml` is in your deploy directory. If it isn't, copy `cluster-service.xml` from `server/all/deploy` into your deploy directory. You also need the `jgroups.jar` in your `lib` directory. This can be found in the `server/all/lib` directory.
- Start JBoss to check if your setup works. Look at the JMX management console (<http://localhost:8080/jmx-console/>). Find the `jboss.cache:service=TomcatClusteringCache` MBean. The `StateString` must be `Started`. If it is `Stopped` look in the server's log file.
- To enable clustering of your web applications you must mark them as distributable in the `web.xml` descriptor. For example:

```
<?xml version="1.0"?>
<!DOCTYPE web-app PUBLIC
  "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
  "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>
  <distributable/>
  <!-- ... -->
</web-app>
```

- Deploy your WAR as usual and it should now be clustered.

If you have deployed and accessed your application, go back to the `jboss.cache:service=TomcatClusteringCache` MBean and invoke the `printDetails` operation. You should see output resembling the following.

```
/JSESSION  
  
/n6HywRwITbY-xvzaZ0LS5Q**  
n6HywRwITbY-xvzaZ0LS5Q**: org.jboss.invocation.MarshalledValue@9c1dddab  
  
/R1T4Dapn7c8T-+Ynd9v9MA**  
R1T4Dapn7c8T-+Ynd9v9MA**: org.jboss.invocation.MarshalledValue@8c0f60b6
```

This output shows two separate web sessions that are being shared via JBossCache. If you don't see any output, either the application was not correctly marked as `distributable` or you haven't accessed the a part of application that places values in the HTTP session.

MBean Services Miscellany

This chapter discusses useful MBean services that are not discussed elsewhere either because they are utility services not necessary for running JBoss, or they don't fit into a current section of the book.

10.1. System Properties Management

The management of system properties can be done using the `org.jboss.varia.property.SystemPropertiesService` MBean. It supports setting of the VM global property values just as `java.lang.System.setProperty` method and the VM command line arguments do.

Its configurable attributes include:

- **Properties:** a specification of multiple property name=value pairs using the `java.util.Properties.load(java.io.InputStream)` method format. Each property=value statement is given on a separate line within the body of the `Properties` attribute element.
- **URLList:** a comma separated list of URL strings from which to load properties file formatted content. If a component in the list is a relative path rather than a URL it will be treated as a file path relative to the `<jboss-dist>/server/<config>` directory. For example, a component of `conf/local.properties` would be treated as a file URL that points to the `<jboss-dist>/server/default/conf/local.properties` file when running with the default configuration file set.

Both attributes are illustrated in Example 10.1.

Example 10.1. An example `SystemPropertiesService` jboss-service descriptor

```
<server>
  <mbean code="org.jboss.varia.property.SystemPropertiesService"
        name="jboss.util:type=Service,name=SystemProperties">

    <!-- Load properties from each of the given comma seperated URLs -->
    <attribute name="URLList">
      http://somehost/some-location.properties,
      ./conf/somelocal.properties
    </attribute>

    <!-- Set propertuies using the properties file style. -->
    <attribute name="Properties">
      property1=This is the value of my property
      property2=This is the value of my other property
    </attribute>

  </mbean>
</server>
```

10.2. Property Editor Management

Support for managing `java.bean.PropertyEditor` instances is available through the `org.jboss.varia.property.PropertyEditorManagerService` MBean. This is a simple service that help define property editors using the `java.bean.PropertyEditorManager` class. This service is used in the main `jboss-service.xml` file to preload the custom JBoss `PropertyEditor` implementations. This is necessary for some JDK1.3.0 VMs that will only load property editors from the system classpath.

Its supported attributes include:

- **BootstrapEditors:** This is a listing of `property_editor_class=editor_value_type_class` pairs defining the property editor to type mappings that should be preloaded into the property `PropertyEditorManager` class using its `registerEditor(Class targetType, Class editorClass)` method. The value type of this attribute is a string so that it may be set from a string without requiring a custom property editor.
- **Editors:** This serves the same function as the `BootstrapEditors` attribute, but its type is a `java.util.Properties` class, and so setting this from a string value requires a custom property editor for `Properties`. In situations where custom property editors can be loaded from the thread context class loader, this may be used instead of the `BootstrapEditors` attribute.
- **EditorSearchPath:** This attribute allows one to set the `PropertyEditorManager` editor packages search path.

10.3. Services Binding Management

With all of the independently deployed services available in JBoss, running multiple instances on a given machine can be a tedious exercise in configuration file editing. The binding service, `org.jboss.services.binding.ServiceBindingManager`, allows one to map service attribute values from a central location. After a service's descriptor file is parsed and the initial attribute values have been applied to the service, the `ServiceConfigurator` queries the `ServiceBindingManager` to apply any overrides that may exist for the service. The `ServicesBindingManager` acts a coordinator between the `ServiceConfigurator`, a store of configuration overrides, the service configuration, and a configuration delegate that knows how to apply a configuration to a service. The classes in this act are shown in Figure 10.1.

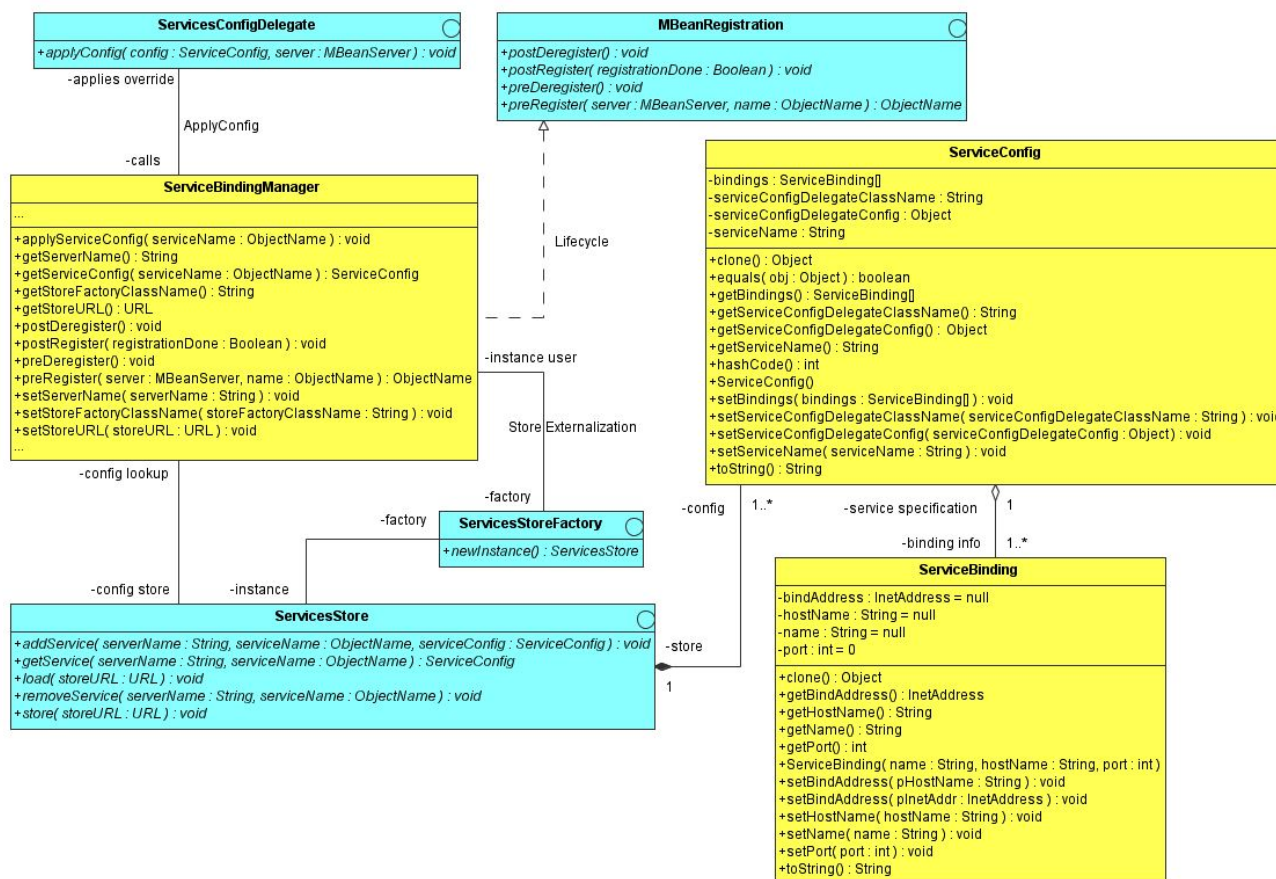


Figure 10.1. Class diagram for the org.jboss.services.binding package of the ServiceBindingManager

The first thing to note about the `ServiceBindingManager` is that it implements the `JMX MBeanRegistration` interface methods as its life cycle notification interface rather than the JBoss service interface. This is necessary because the `ServiceBindingManager` operates on other services attribute values. Attributes are set before any JBoss service life cycle methods are called, and so the `ServiceBindingManager` must be active as soon as it is registered with the MBean server. The setup of the `ServiceBindingManager` occurs in the `postRegister(Boolean)` callback method.

The `ServiceBindingManager` is associated with a `ServicesStore` through a `ServicesStoreFactory`. The `ServicesStoreFactory` is set through an attribute of the `ServiceBindingManager`. The set of configurable attributes of the `ServiceBindingManager` include:

- **ServerName:** The name of the server this manager is associated with. This is a logical name used to lookup `ServiceConfigs` from the `ServicesStore`.
- **StoreFactoryClassName:** The name of the class that implements the `ServicesStoreFactory` interface. You may provide your own implementation, or use the default XML based store `org.jboss.services.binding.XMLServicesStoreFactory`.
- **StoreURL:** The URL of the configuration store contents. This is passed to the `load(URL)` method of the `ServicesStore` instance obtained from the `ServicesStoreFactory`.

A `ServicesStore` is just a collection of `ServiceConfig` objects keyed by a JBoss instance name and the JMX `ObjectName` of the service. A `ServiceConfig` is a collection of `ServiceBinding` objects and a `ServicesConfigDelegate` that knows how to map a `ServiceBinding` onto a target MBean. The `ServiceConfig` may also

contain an arbitrary configuration for the delegate. A `ServiceBinding` is a named (interface, port) pair.

So what happens when the `ServiceBindingManager` is asked to override a service's configuration? The sequence of events is illustrated by Figure 10.2.



Figure 10.2. How the `ServiceConfigurator` queries the `ServiceBindingManager`

- The `ServiceConfigurator` queries the `ServiceBindingManager` to apply any configuration overrides for the MBean given by the `applyServiceConfig` method JMX `ObjectName`.
- The `ServiceBindingManager` queries the `ServicesStore` for the `ServiceConfig` for the named service, specifying the identity of the JBoss server instance in which it is operating. This is an attribute of the `ServiceBindingManager`, and can be taken from a system property as we will see in an example. If the `ServicesStore` contains a configuration override for the indicated `<serverName, serviceName>` pair, it returns the `ServiceConfig`.
- If there was a `ServiceConfig`, the `ServiceBindingManager` queries it for the name of the class implementing the `ServicesConfigDelegate` interface.
- The `ServicesConfigDelegate` class is loaded using the thread context class loader and an instance is created.
- The `ServicesConfigDelegate` instance is then asked to apply the `ServiceConfig` using the provided `MBeanServer`. The delegate would use the delegate configuration information along with the binding(s) to override the indicated attributes of the service by invoking attribute setters, or even operations on the service using the `MBeanServer`. The target service name is available in the `ServiceConfig`.

That is the generic overview of the `ServiceBindingManager`. Let's take a look at how you can use this service to bring up two JBoss instances of the default configuration set of services on the same machine to make this more concrete.

10.3.1. Running Two JBoss Instances

JBoss ships with a service configuration `ServiceBindingManager` for the along with a sample `ServicesStore` XML file for starting two JBoss instances on the same host. Here we will walk through the steps to bring up the two instances and look at the sample configuration. Start by making two server configuration file sets called `jboss0` and `jboss1` by running the following command from the book examples directory:

```
[nr@toki examples]$ ant -Dchap=chap10 -Dex=1 run-example
Buildfile: build.xml
```

```

...
[echo] Preparing jboss0 configuration fileset
[mkdir] Created dir: /tmp/jboss-3.2.6/server/jboss0
[copy] Copying 259 files to /tmp/jboss-3.2.6/server/jboss0
[copy] Copying 1 file to /tmp/jboss-3.2.6/server/jboss0/conf
[copy] Copying 1 file to /tmp/jboss-3.2.6/server
[echo] Preparing jboss1 configuration fileset
[mkdir] Created dir: /tmp/jboss-3.2.6/server/jboss1
[copy] Copying 259 files to /tmp/jboss-3.2.6/server/jboss1

```

BUILD SUCCESSFUL

This creates duplicates of the `server/default` configuration file sets as `server/jboss0` and `server/jboss1`, and then replaces the `conf/jboss-service.xml` descriptor with one that has the `ServiceBindingManager` configuration enabled as follows:

```

<mbean code="org.jboss.services.binding.ServiceBindingManager"
  name="jboss.system:service=ServiceBindingManager">
  <attribute name="ServerName">${jboss.server.name}</attribute>
  <attribute name="StoreURL">${jboss.server.base.dir}/chap10ex1-bindings.xml</attribute>
  <attribute name="StoreFactoryClassName">
    org.jboss.services.binding.XMLServicesStoreFactory
  </attribute>
</mbean>

```

The attribute values are:

- **ServerName:** This is the unique name for the JBoss server instance that will be used to distinguish what configuration overrides to apply. Here the `${jboss.server.name}` variable reference is the configuration file set directory name, either `jboss0` or `jboss1` in this example.
- **StoreURL:** This is the location of the `ServicesStore` configuration data that defines the overrides for the `jboss0` and `jboss1` instances. The `${jboss.server.base.dir}` variable reference is the URL to the root of the JBoss server directory. We are using the `chap10ex1-bindings.xml` which was installed as part of the example 1 setup.
- **StoreFactoryClassName:** This the default XML based `ServicesStore` implement ion.

The `chap10ex1-bindings.xml` file contains two server configurations named `jboss0` and `jboss1`. The `jboss0` configuration uses the default settings for the ports, while the `jboss1` configuration adds 100 to each port number. The bindings file is a duplicate of the `docs/examples.binding-service.sample-bindings.xml` with `jboss0` and `jboss1` as the server names.

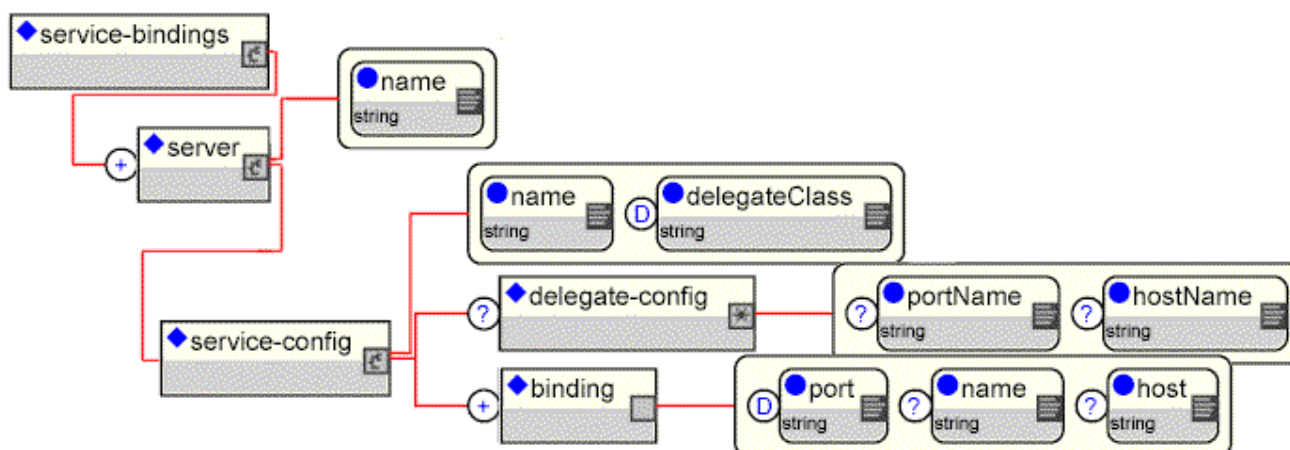


Figure 10.3. The binding service XMLServicesStoreFactory DTD

The DTD shown in Figure 10.3 is the one supported by the `XMLServicesStoreFactory` class. The elements are:

- **service-bindings**: the root element of the configuration file. It contains one or more server elements.
- **server**: This is the base of a JBoss server instance configuration. It has a required name attribute that defines the JBoss instance name to which it applies. This is the name that correlates with the `ServiceBindingManagerServerName` attribute value. The server element content consists of one or more `service-config` elements.
- **service-config**: This element represents a configuration override for an MBean service. It has a required name attribute that is the JMX `ObjectName` string of the MBean service the configuration applies to. It also has a required `delegateClass` name attribute that specifies the class name of the `ServicesConfigDelegate` implementation that knows how to handle bindings for the target service. Its contents consists of an optional `delegate-config` element and one or more binding elements.
- **binding**: A binding element specifies a named port, address pair. It has an optional name that can be used to provide multiple binding for a service. An example would be multiple virtual hosts for a web container. The port and address are specified via the optional port and host attributes respectively. If the port is not specified it defaults to 0 meaning choose an anonymous port. If the host is not specified it defaults to null meaning any address.
- **delegate-config**: The `delegate-config` element is an arbitrary XML fragment for use by the `ServicesConfigDelegate` implementation. The `hostName` and `portName` attributes only apply to the `AttributeMappingDelegate` of the example and are there to prevent DTD aware editors from complaining about their existence in the `AttributeMappingDelegate` configurations. Generally both the attributes and content of the `delegate-config` is arbitrary, but there is no way to specify and an element can have any number of attributes with a DTD.

The two `ServicesConfigDelegate` implementations are `AttributeMappingDelegate` and `XSLTConfigDelegate`. The `AttributeMappingDelegate` class is an implementation of the `ServicesConfigDelegate` that expects a `delegate-config` element of the form:

```
<delegate-config portName="portAttrName" hostName="hostAttrName">
  <attribute name="someAttrName">someHostPortExpr</attribute>
  <!-- ... -->
</delegate-config>
```

The `portAttrName` is the attribute name of the MBean service to which the binding port value should be applied, and the `hostAttrName` is the attribute name of the MBean service to which the binding host value should be applied. If the `portName` attribute is not specified then the binding port is not applied. Likewise, if the `hostName` attribute is not specified then the binding host is not applied. The optional attribute element(s) specify arbitrary MBean attribute names whose values are a function of the host and/or port settings. Any reference to `${host}` in the attribute content is replaced with the host binding and any `${port}` reference is replaced with the port binding. The `portName`, `hostName` attribute values and attribute element content may reference system properties using the `${x}` syntax that is supported by the JBoss services descriptor.

The sample listing illustrates the usage of `AttributeMappingDelegate`.

```
<service-config name="jboss:service=Naming"
  delegateClass="org.jboss.services.binding.AttributeMappingDelegate">
  <delegate-config portName="Port"/>
```

```
<binding port="1099" />
</service-config>
```

Here the `jboss:service=Naming` MBean service has its `Port` attribute value overridden to 1099. The corresponding setting from the `jboss1` server configuration overrides the port to 1199.

The `XSLTConfigDelegate` class is an implementation of the `ServicesConfigDelegate` that expects a `delegate-config` element of the form:

```
<delegate-config>
  <xslt-config configName="ConfigurationElement"><![CDATA[
    Any XSL document contents...
  ]]>
</xslt-config>
<xslt-param name="param-name">param-value</xslt-param>
<!-- ... -->
</delegate-config>
```

The `xslt-config` child element content specifies an arbitrary XSL script fragment that is to be applied to the MBean service attribute named by the `configName` attribute. The named attribute must be of type `org.w3c.dom.Element`. The optional `xslt-param` elements specify XSL script parameter values for parameters used in the script. There are two XSL parameters defined by default called `host` and `port`, and their values are set to the configuration host and port bindings.

The `XSLTConfigDelegate` is used to transform services whose `port/interface` configuration is specified using a nested XML fragment. The following illustrates an example from the `jboss1` server section which maps the Tomcat servlet container listening port to 8180 and maps the AJP listening port to 8109:

```
<!-- jbossweb-tomcat50.sar -->
<service-config name="jboss.web:service=WebServer"
  delegateClass="org.jboss.services.binding.XSLTFileDelegate"
>
  <delegate-config>
    <xslt-config configName="ConfigFile"><![CDATA[
<xsl:stylesheet
  xmlns:xsl='http://www.w3.org/1999/XSL/Transform' version='1.0'>

  <xsl:output method="xml" />
  <xsl:param name="port"/>

  <xsl:variable name="portAJP" select="$port - 71"/>
  <xsl:variable name="portHttps" select="$port + 363"/>

  <xsl:template match="/">
    <xsl:apply-templates/>
  </xsl:template>

  <xsl:template match = "Connector">
    <Connector>
      <xsl:for-each select="@*">
        <xsl:choose>
          <xsl:when test="(name() = 'port' and . = '8080')">
            <xsl:attribute name="port"><xsl:value-of select="$port" />
          </xsl:when>
          <xsl:when test="(name() = 'port' and . = '8009')">
            <xsl:attribute name="port"><xsl:value-of select="$portAJP" />
          </xsl:when>
          <xsl:when test="(name() = 'redirectPort')">
            <xsl:attribute name="redirectPort"><xsl:value-of select="$portHttps" />
          </xsl:when>
          <xsl:when test="(name() = 'port' and . = '8443')">
            <xsl:attribute name="port"><xsl:value-of select="$portHttps" />
          </xsl:when>
        </xsl:choose>
      </xsl:for-each>
    </Connector>
  </xsl:template>
</xsl:stylesheet>
```

```

        </xsl:attribute>
    </xsl:when>
    <xsl:otherwise>
        <xsl:attribute name="{name()}"><xsl:value-of select="." />
    </xsl:attribute>
    </xsl:otherwise>
</xsl:choose>
</xsl:for-each>
<xsl:apply-templates/>
</Connector>
</xsl:template>

<xsl:template match="*|@">
    <xsl:copy>
        <xsl:apply-templates select="@*|node()" />
    </xsl:copy>
</xsl:template>
</xsl:stylesheet>
]]>
    </xslt-config>
</delegate-config>
<binding port="8180"/>
</service-config>

```

To test the sample configuration, start two JBoss instances using the `jboss0` and `jboss1` configuration file sets created previously by running the `chap10 example1` build. Looking at the console for the service port numbers you should see the overridden mappings. For the `jboss1` server for example, here are some of the non-standard ports that show up:

```

16:04:39,246 INFO  [WebService] Using RMI server codebase: http://toki.local:8183/
16:04:40,442 INFO  [NamingService] Started jnpPort=1199, rmiPort=1198, backlog=50, bindAdd
ress=/0.0.0.0, Client SocketFactory=null, Server SocketFactory=org.jboss.net.sockets.Defau
ltSocketFactory@ad093076
16:05:28,596 INFO  [Http11Protocol] Initializing Coyote HTTP/1.1 on http-0.0.0.0-8180
16:05:55,165 INFO  [Http11Protocol] Starting Coyote HTTP/1.1 on http-0.0.0.0-8180
16:05:56,061 INFO  [ChannelSocket] JK2: ajp13 listening on /0.0.0.0:8109

```

10.4. Scheduling Tasks

Java includes a simple timer based capability through the `java.util.Timer` and `java.util.TimerTask` utility classes. JMX also includes a mechanism for scheduling JMX notifications at a given time with an optional repeat interval as the `javax.management.timer.TimerMBean` agent service.

JBoss includes two variations of the JMX timer service in the `org.jboss.varia.scheduler.Scheduler` and `org.jboss.varia.scheduler.ScheduleManager` MBeans. Both MBeans rely on the JMX timer service for the basic scheduling. They extend the behavior of the timer service as described in the following sections.

10.4.1. org.jboss.varia.scheduler.Scheduler

The `Scheduler` differs from the `TimerMBean` in that the `Scheduler` directly invokes a callback on an instance of a user defined class, or an operation of a user specified MBean.

- **InitialStartDate:** Date when the initial call is scheduled. It can be either:
 - `NOW`: date will be the current time plus 1 seconds
 - A number representing the milliseconds since 1/1/1970

- Date as String able to be parsed by `SimpleDateFormat` with default format pattern "M/d/yy h:mm a". If the date is in the past the `Scheduler` will search a start date in the future with respect to the initial repetitions and the period between calls. This means that when you restart the MBean (restarting JBoss etc.) it will start at the next scheduled time. When no start date is available in the future the `Scheduler` will not start.

For example, if you start your `Schedulable` everyday at Noon and you restart your JBoss server then it will start at the next Noon (the same if started before Noon or the next day if start after Noon).

- **InitialRepetitions:** The number of times the scheduler will invoke the target's callback. If -1 then the callback will be repeated until the server is stopped.
- **StartAtStartup:** A flag that determines if the `Scheduler` will start when it receives its `startService` life cycle notification. If true the `Scheduler` starts on its startup. If false, an explicit `startSchedule` operation must be invoked on the `Scheduler` to begin.
- **SchedulePeriod:** The interval between scheduled calls in milliseconds. This value must be bigger than 0.
- **SchedulableClass:** The fully qualified class name of the `org.jboss.varia.scheduler.Schedulable` interface implementation that is to be used by the `Scheduler`. The `SchedulableArguments` and `SchedulableArgumentTypes` must be populated to correspond to the constructor of the `Schedulable` implementation.
- **SchedulableArguments:** A comma separated list of arguments for the `Schedulable` implementation class constructor. Only primitive data types, `String` and classes with a constructor that accepts a `String` as its sole argument are supported.
- **SchedulableArgumentTypes:** A comma separated list of argument types for the `Schedulable` implementation class constructor. This will be used to find the correct constructor via reflection. Only primitive data types, `String` and classes with a constructor that accepts a `String` as its sole argument are supported.
- **SchedulableMBean:** Specifies the fully qualified JMX `ObjectName` name of the schedulable MBean to be called. If the MBean is not available it will not be called but the remaining repetitions will be decremented. When using `SchedulableMBean` the `SchedulableMBeanMethod` must also be specified.
- **SchedulableMBeanMethod:** Specifies the operation name to be called on the schedulable MBean. It can optionally be followed by an opening bracket, a comma separated list of parameter keywords, and a closing bracket. The supported parameter keywords include:
 - `NOTIFICATION` which will be replaced by the `timers` notification instance (`javax.management.Notification`)
 - `DATE` which will be replaced by the date of the notification call (`java.util.Date`)
 - `REPETITIONS` which will be replaced by the number of remaining repetitions (long)
 - `SCHEDULER_NAME` which will be replaced by the `ObjectName` of the `Scheduler`
 - Any fully qualified class name which the `Scheduler` will set to null.

A given `Scheduler` instance only support a single schedulable instance. If you need to configure multiple scheduled events you would use multiple `Scheduler` instances, each with a unique `ObjectName`. The following is an example of configuring a `Scheduler` to call a `Schedulable` implementation as well as a configuration for calling a MBean.

Example 10.2. An example Scheduler jboss-service descriptor

```

<server>

  <mbean code="org.jboss.varia.scheduler.Scheduler"
        name="jboss.docs.chap10:service=Scheduler">
    <attribute name="StartAtStartup">true</attribute>
    <attribute name="SchedulableClass">org.jboss.chap10.ex2.ExSchedulable</attribute>
    <attribute name="SchedulableArguments">TheName,123456789</attribute>
    <attribute name="SchedulableArgumentTypes">java.lang.String,long</attribute>

    <attribute name="InitialStartDate">NOW</attribute>
    <attribute name="SchedulePeriod">60000</attribute>
    <attribute name="InitialRepetitions">-1</attribute>
  </mbean>

</server>

```

The SchedulableClass `org.jboss.chap10.ex2.ExSchedulable` example class is given in Example 10.3.

Example 10.3. The ExSchedulable class code

```

package org.jboss.chap10.ex2;

import java.util.Date;
import org.jboss.varia.scheduler.Schedulable;

import org.apache.log4j.Logger;

/**
 * A simple Schedulable example.
 * @author Scott.Stark@jboss.org
 * @version $Revision: 1.7 $
 */
public class ExSchedulable implements Schedulable
{
    private static final Logger log = Logger.getLogger(ExSchedulable.class);

    private String name;
    private long value;

    public ExSchedulable(String name, long value)
    {
        this.name = name;
        this.value = value;
        log.info("ctor, name: " + name + ", value: " + value);
    }

    public void perform(Date now, long remainingRepetitions)
    {
        log.info("perform, now: " + now +
            ", remainingRepetitions: " + remainingRepetitions +
            ", name: " + name + ", value: " + value);
    }
}

```

Deploy the timer sar by running:

```

[nr@toki examples]$ ant -Dchap=chap10 -Dex=2 run-example
...
run-example2:
    [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy

```

The server console shows the following which includes the first two timer invocations, separated by 60 seconds:

```
16:44:49,275 INFO [ExSchedulable] ctor, name: TheName, value: 123456789
16:44:50,365 INFO [ExSchedulable] perform, now: Fri Oct 15 16:44:50 CDT 2004, remainingRe
petitions: -1, name: TheName, value: 123456789
16:45:50,317 INFO [ExSchedulable] perform, now: Fri Oct 15 16:45:50 CDT 2004, remainingRe
petitions: -1, name: TheName, value: 123456789
16:46:50,319 INFO [ExSchedulable] perform, now: Fri Oct 15 16:46:50 CDT 2004, remainingRe
petitions: -1, name: TheName, value: 123456789
```

10.5. The JBoss Logging Framework

In 3.2 the logging framework has been generalized to allow for any particular framework implementation. The JBoss classes themselves use the `org.jboss.logging.Logger` as the factory and logging interface. This class is essentially identical to the Log4j `org.apache.log4j.Logger` class, with the addition of support for a trace level priority. The `Logger` class delegates to a `LoggerPlugin` instance. The class diagram for `Logger` and `LoggerPlugin` are shown in Figure 10.4.

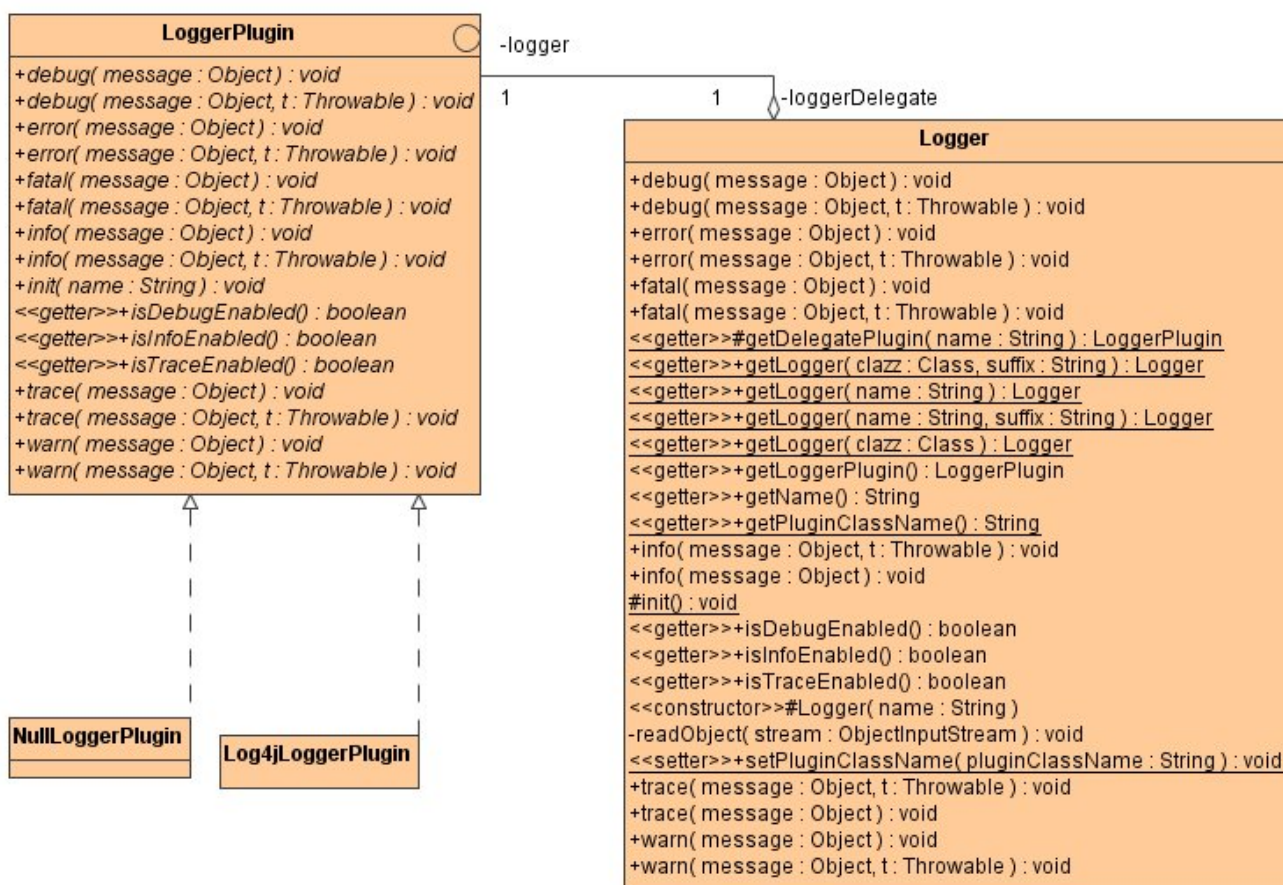


Figure 10.4. The JBoss logging framework classes.

By default we continue to use the Log4j framework as the underlying logging implementation, and this is what the `org.jboss.logging.Log4jLoggerPlugin` provides. To integrate an alternate logging implementation, you would provide an implementation of the `LoggerPlugin` interface and specify that it should be used by setting the `org.jboss.logging.Logger.pluginClass` system property to implementation class name. To disable all log-

ging, you can use the `org.jboss.logging.NullLoggerPlugin`. This implementation simply provides empty versions of the `LoggerPlugin` methods.

10.5.1. `org.jboss.logging.Log4jService`

The `Log4jService` MBean configures the Apache log4j system. JBoss uses the log4j framework as its internal logging API.

- **ConfigurationURL:** The URL for the log4j configuration file. This can refer to either a XML document parsed by the `org.apache.log4j.xml.DOMConfigurator` or a Java properties file parsed by the `org.apache.log4j.PropertyConfigurator`. The type of the file is determined by the URL content type, or if this is null, the file extension. The default setting of `resource:log4j.xml` refers to the `conf/log4j.xml` file of the active server configuration file set.
- **RefreshPeriod:** The time in seconds between checks for changes in the log4 configuration specified by the `ConfigurationURL` attribute. The default value is 60 seconds.
- **CatchSystemErr:** This boolean flag if true, indicates if the `System.err` stream should be redirected onto a log4j category called `STDERR`. The default is true.
- **CatchSystemOut:** This boolean flag if true, indicates if the `System.out` stream should be redirected onto a log4j category called `STDOUT`. The default is true.
- **Log4jQuietMode:** This boolean flag if true, sets the `org.apache.log4j.helpers.LogLog.setQuietMode`. As of log4j1.2.8 this needs to be set to avoid a possible deadlock on exception at the appender level. See bug#696819.

10.6. RMI Dynamic Class Loading

10.6.1. `org.jboss.web.WebService`

The `WebService` MBean provides dynamic class loading for RMI access to the server EJBs. The configurable attributes for the `WebService` are as follows:

- **Port:** the `WebService` listening port number. A port of 0 will use any available port.
- **Host:** Set the name of the public interface to use for the host portion of the RMI codebase URL.
- **BindAddress:** the specific address the `WebService` listens on. This can be used on a multi-homed host for a `java.net.ServerSocket` that will only accept connect requests to one of its addresses.
- **Backlog:** The maximum queue length for incoming connection indications (a request to connect) is set to the `backlog` parameter. If a connection indication arrives when the queue is full, the connection is refused.
- **DownloadServerClasses:** A flag indicating if the server should attempt to download classes from thread context class loader when a request arrives that does not have a class loader key prefix.

The CMP Engine

This chapter details the operation of the JBoss CMP2 engine. It does not provide an introduction to the EJB 2.0 container managed persistence (CMP2.0) model. To get started with CMP2.0, see the J2EE tutorial (<http://java.sun.com/j2ee/tutorial/index.html>), or *Enterprise Java Beans* - 3rd edition along with the companion JBoss workbook (<http://www.oreilly.com/catalog/entjbeans3/workbooks/index.html>).

11.1. Getting Started

JBossCMP is the default persistence manager for EJB 2.0 applications. Because JBossCMP is a core feature of JBoss, no action beyond the basic JBoss installation is required to use CMP 2.0, but there are some details to note when creating a new EJB 2.0 application or when upgrading an EJB 1.1 application.

When JBoss deploys an EJB JAR file, it uses the DOCTYPE of the `ejb-jar.xml` deployment descriptor to determine the version of the EJB jar. The correct DOCTYPE for EJB 2.0 is given below.

```
<!DOCTYPE ejb-jar PUBLIC
"-//Sun Microsystems, Inc.//DTD Enterprise JavaBeans 2.0//EN"
"http://java.sun.com/dtd/ejb-jar_2_0.dtd">
```

If the public identifier of the DOCTYPE is `"-//Sun Microsystems, Inc.//DTD Enterprise JavaBeans 2.0//EN"` JBossCMP will use the *Standard CMP 2.x EntityBean* configuration in the `standardjboss.xml` file. If you have an application that uses a custom entity bean configuration, and you are upgrading to EJB 2.0, you must change the `persistence-manager` and add the new interceptors (see the *Standard CMP 2.x EntityBean* configuration in the `standardjboss.xml` file for details). No further configuration is necessary to deploy and run your EJB 2.0 application successfully.

11.1.1. Example Code

The full source code for all of the examples presented in this documentation is available in the `examples/src/main/org/jboss/cmp2` directory. The code represents a Crime Portal, which models criminal organizations. A diagram of the portions of the Criminal Portal data model used in the example code is shown in Figure 11.1.

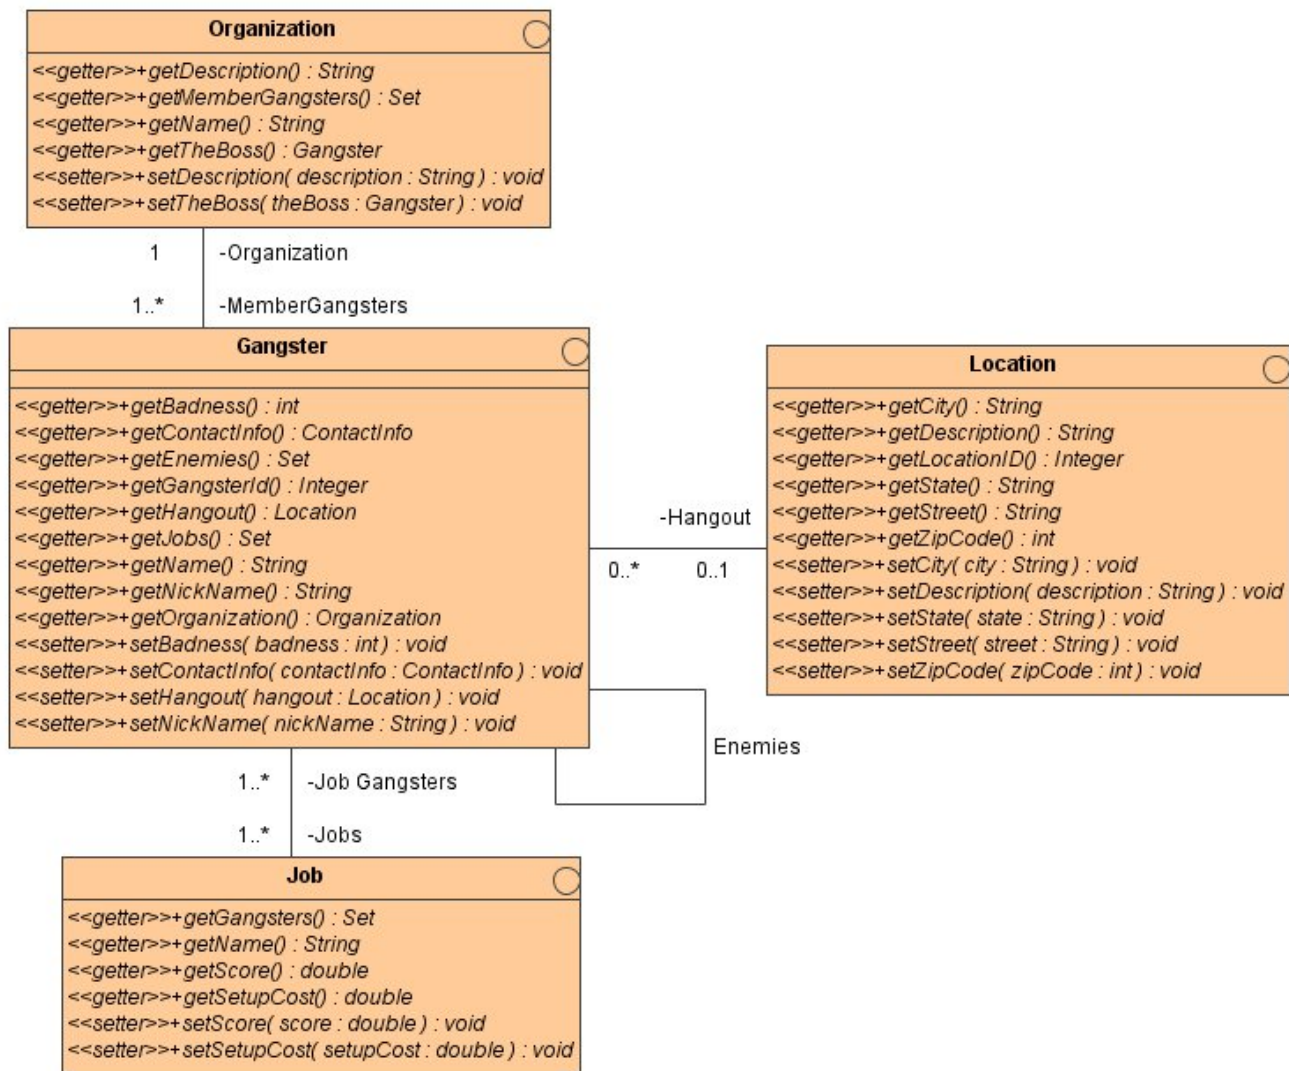


Figure 11.1. The main CMP2 example classes

To build the example code, execute ant with the following arguments:

```
[nr@toki examples]$ ant -Dchap=cmp2 config
...
config:
  [copy] Copying 1 file to /tmp/jboss-3.2.6/server/default/deploy
  [echo] Waiting for 5 seconds for deploy...
  [junit] .
  [junit] Time: 3.474

  [junit] OK (1 test)
```

This command builds and deploys the application to the JBoss server. When you start your JBoss server, or if it is already running, you should see the following deployment messages:

```
09:54:42,018 INFO [EjbModule] Deploying OrganizationEJB
09:54:42,399 INFO [EjbModule] Deploying GangsterEJB
09:54:42,438 INFO [EjbModule] Deploying JobEJB
09:54:42,468 INFO [EjbModule] Deploying LocationEJB
09:54:42,507 INFO [EjbModule] Deploying EJBTTestRunnerEJB
09:54:42,587 INFO [EjbModule] Deploying ReadAheadEJB
09:54:46,300 WARN [JDBCTypeFactory] Type not mapped: int
09:54:47,223 INFO [EJBDeployer] Deployed: file:/private/tmp/jboss-3.2.6/server/default/deploy/cmp2-ex1.jar
```

```

09:54:48,841 INFO [OrganizationBean$Proxy] Creating organization Yakuza, Japanese Gangsters
09:54:48,918 INFO [OrganizationBean$Proxy] Creating organization Mafia, Italian Bad Guys
09:54:48,925 INFO [OrganizationBean$Proxy] Creating organization Triads, Kung Fu Movie Extras
09:54:48,931 INFO [GangsterBean$Proxy] Creating Gangster 0 'Bodyguard' Yojimbo
09:54:49,068 INFO [GangsterBean$Proxy] Creating Gangster 1 'Master' Takeshi
09:54:49,106 INFO [GangsterBean$Proxy] Creating Gangster 2 'Four finger' Yuriko
09:54:49,117 INFO [GangsterBean$Proxy] Creating Gangster 3 'Killer' Chow
09:54:49,133 INFO [GangsterBean$Proxy] Creating Gangster 4 'Lightning' Shogi
09:54:49,143 INFO [GangsterBean$Proxy] Creating Gangster 5 'Pizza-Face' Valentino
09:54:49,184 INFO [GangsterBean$Proxy] Creating Gangster 6 'Toohless' Toni
09:54:49,202 INFO [GangsterBean$Proxy] Creating Gangster 7 'Godfather' Corleone
09:54:49,215 INFO [JobBean$Proxy] Creating Job 10th Street Jeweler Heist
09:54:49,224 INFO [JobBean$Proxy] Creating Job The Great Train Robbery
09:54:49,258 INFO [JobBean$Proxy] Creating Job Cheap Liquor Snatch and Grab

```

Before the chapter tests can be run, the log level of JBossCMP must be increased. To enable debug logging for the CMP subsystem, add the following category to your `log4j.xml` file:

```

<category name="org.jboss.ejb.plugins.cmp">
  <priority value="DEBUG"/>
</category>

```

In addition to this, it is necessary to decrease the threshold on the `CONSOLE` appender to allow debug level messages to be logged to the console. The following changes also need to be applied to the `log4j.xml` file.

```

<!-- ===== -->
<!-- Append messages to the console -->
<!-- ===== -->
<appender name="CONSOLE" class="org.apache.log4j.ConsoleAppender">
  <param name="Threshold" value="DEBUG"/>
  <param name="Target" value="System.out"/>
  <layout class="org.apache.log4j.PatternLayout">
    <!-- The default pattern: Date Priority [Category] Message\n -->
    <param name="ConversionPattern" value="%d{ABSOLUTE} %-5p [%c{1}] %m%n"/>
  </layout>
</appender>

```

To see the full workings of the CMP engine you would need to enable the custom `TRACE` level priority on the `org.jboss.ejb.plugins.cmp` category as shown here:

```

<category name="org.jboss.ejb.plugins.cmp">
  <priority value="TRACE" class="org.jboss.logging.XLevel"/>
</category>

```

One final note before moving on to look at how to run the chapter examples. Since the beans in the examples are configured to remove their tables on undeployment, anytime you restart the JBoss server you need to rerun the config target to reload the example data. Also, if you make changes to the examples and want to redeploy the example EJB JAR, this also should be done using the config target so that the example data is reloaded.

11.1.2. Tests

The first test target illustrates a number of the customization features that will be discussed throughout this chapter. To run these tests execute the following ant target:

```

[nr@toki examples]ant -Dchap=cmp2 -Dex=test run-example
14:03:27,920 DEBUG [OrganizationEJB#findByPrimaryKey] Executing SQL: SELECT name FROM ORGANIZATION WHERE name=?
14:03:28,011 DEBUG [OrganizationEJB] Executing SQL: SELECT desc, the_boss FROM ORGANIZATION

```

```

N WHERE (name=?)
14:03:28,020 DEBUG [OrganizationEJB] Executing SQL: SELECT id FROM GANGSTER WHERE (organization=?)
14:03:28,044 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,052 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,070 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,229 DEBUG [GangsterEJB#findBadDudes_ejbql] Executing SQL: SELECT t0_g.id FROM GANGSTER t0_g WHERE (t0_g.badness > ?)
14:03:28,256 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,264 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,270 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,276 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,281 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,395 DEBUG [GangsterEJB#findBadDudes_jbossql] Executing SQL: SELECT t0_g.id, t0_g.badness FROM GANGSTER t0_g WHERE (t0_g.badness > ?) ORDER BY t0_g.badness DESC
14:03:28,417 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,423 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,429 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,439 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,446 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,605 DEBUG [GangsterEJB#findBadDudes_declaredsql] Executing SQL: SELECT id FROM GANGSTER WHERE badness > ? ORDER BY badness DESC
14:03:28,613 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,631 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,641 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,647 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,655 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,783 DEBUG [GangsterEJB#ejbSelectBoss_ejbql] Executing SQL: SELECT DISTINCT t0_underling_organization_theBos.id FROM GANGSTER t1_underling, ORGANIZATION t4_underling_organization, GANGSTER t0_underling_organization_theBos WHERE ((t1_underling.name = ?) OR (t1_underling.nick_name = ?)) AND t1_underling.organization=t4_underling_organization.name AND t4_underling_organization.the_boss=t0_underling_organization_theBos.id
14:03:28,815 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
14:03:28,822 DEBUG [GangsterEJB#ejbSelectBoss_ejbql] Executing SQL: SELECT DISTINCT t0_underling_organization_theBos.id FROM GANGSTER t1_underling, ORGANIZATION t4_underling_organization, GANGSTER t0_underling_organization_theBos WHERE ((t1_underling.name = ?) OR (t1_underling.nick_name = ?)) AND t1_underling.organization=t4_underling_organization.name AND t4_underling_organization.the_boss=t0_underling_organization_theBos.id
14:03:28,829 DEBUG [GangsterEJB#findByPrimaryKey] Executing SQL: SELECT id FROM GANGSTER WHERE id=?
...
14:03:29,970 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout, organization FROM GANGSTER WHERE (id=?)
14:03:29,980 DEBUG [GangsterEJB] Executing SQL: SELECT cell_area, cell_exch, cell_ext, page_area, page_exch, page_ext, email FROM GANGSTER WHERE (id=?)
14:03:29,987 DEBUG [GangsterEJB] Executing SQL: UPDATE GANGSTER SET cell_area=?, cell_exch=?, cell_ext=?, page_area=?, page_exch=?, page_ext=?, email=? WHERE id=?
14:03:29,995 DEBUG [GangsterEJB] Rows affected = 1

```

These tests exercise various finders, selectors and object to table mapping issues. We will refer to the tests

throughout the chapter.

11.1.3. Read-ahead

The other main target runs a set of tests to demonstrate the optimized loading configurations presented in Section 11.7. Now that the logging is setup correctly, the read-ahead tests will display useful information about the queries performed. Note that you do not have to restart the JBoss server for it to recognize the changes to the `log4j.xml` file, but it may take a minute or so. The following shows the actual execution of the readahead client:

```
[starksm@banshee examples]$ ant -Dchap=cmp2 -Dex=readahead run-example
Buildfile: build.xml
...
run-example:

run-examplereadahead:
[junit] .
[junit] Time: 0.561

[junit] OK (1 test)
```

When the readahead client is executed, all of the SQL queries executed during the test are displayed in the JBoss server console. The important items of note when analyzing the output are the number of queries executed, the columns selected, and the number of rows loaded. The following shows the read-ahead none portion of the JBoss server console output from readahead:

```
#####
### read-ahead none
###
08:31:15,892 DEBUG [findAll_none] Executing SQL: SELECT t0_g.id, t0_g.id FROM GANGSTER
t0_g ORDER BY t0_g.id ASC
08:31:15,902 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,912 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,912 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,912 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,922 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,922 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,932 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,932 DEBUG [GangsterEJB] Executing SQL: SELECT name, nick_name, badness, hangout,
organization FROM GANGSTER WHERE (id=?)
08:31:15,942 INFO [ReadAheadTest]
###
#####
```

We will revisit this example and explore the output when we discuss the settings for optimized loading.

11.2. The jbosscomp-jdbc Structure

The `jbosscomp-jdbc.xml` descriptor is used to control the behavior of the JBossCMP engine. This can be done globally through the `conf/standardjbosscomp-jdbc.xml` descriptor found in the server configuration file set, or per EJB JAR deployment via a `META-INF/jbosscomp-jdbc.xml` descriptor. We will touch on the elements of the as we go through the following sections which describe the capabilities of the JBossCMP engine. The top level

elements are shown in Figure 11.2.

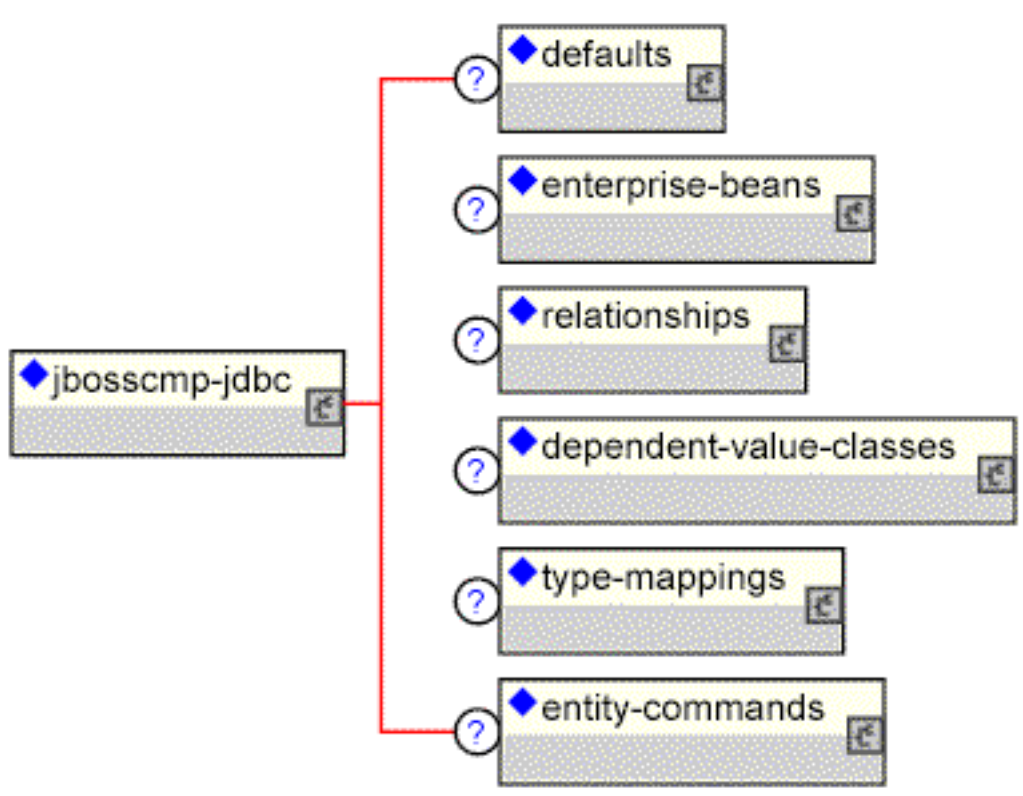


Figure 11.2. The jbosscmp-jdbc top level content model.

- **defaults:** The defaults section allows for the specification of default behavior/settings for behavior that controls entity beans. Use of this section simplifies the amount of information needed for the common behaviors found in the entity beans section. See Section 11.12 for the details of the defaults content.
- **enterprise-beans:** The `enterprise-beans` element allows for customization of entity beans defined in the `ejb-jar.xml` `enterprise-beans` descriptor. This is described in detail in Section 11.3.
- **relationships:** The `relationships` element allows for the customization of tables and the loading behavior of entity relationships. This is described in detail in Section 11.5.
- **dependent-value-classes:** The `dependent-value-classes` element allows for the customization of the mapping of dependent value classes to tables. This is described in detail in Section 11.4.6 (DVCs).
- **type-mappings:** The `type-mappings` element defines the Java to SQL type mappings for a database, along with SQL templates, and function mappings. This is described in detail in Section 11.13.
- **entity-commands:** The `entity-commands` element allows for the definition of the entity creation command instances that know how to create an entity instance in a persistent store. This is described in detail in Section 11.11.
- **user-type-mappings:** The `user-type-mappings` elements defines a mapping of a user types to a column using a mapper class. A mapper is like a mediator: when storing, it takes an instance of the user type and translates it to a column value. When loading, it takes a column value and translates it to an instance of the user type. Details of the user type mappings are described in Section 11.13.3.
- **reserved-words:** The `reserved-words` element defines one or more reserved words that should be escaped when generating tables. Each reserved word is specified as the content of a word element.

Example 11.1. DTD for jbosscmp-jdbc.xml

The DTD for the jbosscmp-jdbc.xml descriptor can be found in JBOSS_DIST/docs/dtd/jbosscmp-jdbc_3_2.dtd. The public doctype for this DTD is:

```
<!DOCTYPE jbosscmp-jdbc PUBLIC
    "-//JBoss//DTD JBOSSCMP-JDBC 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jbosscmp-jdbc_3_2.dtd">
```

11.3. Entity Beans

Although several new features have been added, and there have been major changes to cmp-fields and finders, the basic entity bean structure has not changed much in CMP 2.0. A new feature of EJB 2.0 is the addition of local interfaces. A local interface is composed of two interfaces, the local interface and the local home interface³. These interfaces are conceptually the same thing as the remote interface and home interface (sometimes referred to as the remote home), except that local interfaces are only accessible within the same Java VM. This allows local interfaces to use pass-by-reference semantics, removing the overhead associated with serializing and deserializing every method parameter⁴. Local interfaces are not unique to CMP and are not discussed in this documentation. The simplified code for the Gangster entity follows:

Example 11.2. Entity Local Home Interface

```
// Gangster Local Home Interface
public interface GangsterHome extends EJBLocalHome {
    Gangster create(Integer id, String name, String nickName)
        throws CreateException;
    Gangster findByPrimaryKey(Integer id)
        throws FinderException;
}
```

Example 11.3. Entity Local Interface

```
// Gangster Local Interface
public interface Gangster extends EJBLocalObject {
    Integer getGangsterId();
    String getName();
    String getNickName();
    void setNickName(String nickName);
}
```

Example 11.4. Entity Implementation Class

```
// Gangster Implementation Class
public abstract class GangsterBean
    implements EntityBean
{
    private EntityContext ctx;
```

³The term local interface is used to refer to the EJBLocalObject alone, as well as to refer to the EJBLocalObject/EJBLocalHome combination. Although this is confusing, it is the current usage of the term in the EJB community.

⁴Most J2EE servers, including JBoss, can optimize in-VM calls over a remote interface by using pass-by-reference semantics.


```

private Category log = Category.getInstance(getClass());
public Integer ejbCreate(Integer id, String name, String nickName)
    throws CreateException
{
    log.info("Creating Gangster " + id + " '" + nickName + "' " + name);
    setGangsterId(id);
    setName(name);
    setNickName(nickName);
    return null;
}

public void ejbPostCreate(Integer id, String name, String nickName) {
}

// CMP field accessors -----
public abstract Integer getGangsterId();
public abstract void setGangsterId(Integer gangsterId);
public abstract String getName();
public abstract void setName(String name);
public abstract String getNickName();
public abstract void setNickName(String nickName);
public abstract int getBadness();
public abstract void setBadness(int badness);
public abstract ContactInfo getContactInfo();
public abstract void setContactInfo(ContactInfo contactInfo);
//...

// EJB callbacks -----
public void setEntityContext(EntityContext context) { ctx = context; }
public void unsetEntityContext() { ctx = null; }
public void ejbActivate() { }
public void ejbPassivate() { }
public void ejbRemove() { log.info("Removing " + getName()); }
public void ejbStore() { }
public void ejbLoad() { }
}

```

The base declaration of an entity in the `ejb-jar.xml` file has not changed much in CMP 2.0. The declaration of the `GangsterEJB` interfaces and `cmp` fields is shown below.

Example 11.5. The `ejb-jar.xml` Entity Declaration

```

<?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>CMP 2.0 Lab Jar</display-name>
  <enterprise-beans>
    <entity>
      <display-name>Gangster Entity Bean</display-name>
      <ejb-name>GangsterEJB</ejb-name>
      <local-home>org.jboss.cmp2.crimeportal.GangsterHome</local-home>
      <local>org.jboss.cmp2.crimeportal.Gangster</local>
      <ejb-class>org.jboss.cmp2.crimeportal.GangsterBean</ejb-class>
      <persistence-type>Container</persistence-type>
      <prim-key-class>java.lang.Integer</prim-key-class>
      <reentrant>False</reentrant>
      <cmp-version>2.x</cmp-version>
      <abstract-schema-name>gangster</abstract-schema-name>
      <cmp-field>
        <field-name>gangsterId</field-name>
      </cmp-field>
      <cmp-field>
        <field-name>name</field-name>
      </cmp-field>
    </entity>
  </enterprise-beans>
</ejb-jar>

```

```

        </cmp-field>
        <cmp-field>
            <field-name>nickName</field-name>
        </cmp-field>
        <cmp-field>
            <field-name>badness</field-name>
        </cmp-field>
        <cmp-field>
            <field-name>contactInfo</field-name>
        </cmp-field>
        <primkey-field>gangsterId</primkey-field>
        <!-- ... -->
    </entity>
</enterprise-beans>
</ejb-jar>

```

The new `local-home` and `local` elements are equivalent to the `home` and `remote` elements. The `cmp-version` element is new and can be either 1.x or the default 2.x. This element was added so 1.x and 2.x entities could be mixed in the same application. The `abstract-schema-name` element is also new and is used to identify this entity type in EJB-QL queries, which are discussed in Section 11.6.

11.3.1. Entity Mapping

The JBossCMP configuration for the entity is declared with an entity element in the `jbosscomp-jdbc.xml` file. This file is located in the `META-INF` directory of the EJB JAR and contains all of the optional configuration information for JBossCMP. The entity elements are grouped together in the `enterprise-beans` element under the top level `jbosscomp-jdbc` element. An example entity configuration is shown below.

Example 11.6. A sample `jbosscomp-jdbc.xml` Entity Mapping

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE jbosscomp-jdbc PUBLIC
    "-//JBoss//DTD JBOSSCMP-JDBC 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jbosscomp-jdbc_3_2.dtd">
<jbosscomp-jdbc>
    <!-- ... -->
    <enterprise-beans>
        <entity>
            <ejb-name>GangsterEJB</ejb-name>
            <table-name>gangster</table-name>
            <!-- CMP Fields (see CMP-Fields) -->
            <!-- Load Groups (see Load Groups) -->
            <!-- Queries (see Queries) -->
        </entity>
    </enterprise-beans>
</jbosscomp-jdbc>

```

In this case the DOCTYPE declaration is optional, but will reduce configuration errors. In addition, all of the elements are optional except for `ejb-name`, which is used to match the configuration to an entity declared in the `ejb-jar.xml` file. Unless noted otherwise, the default values come from the defaults section of either the `jbosscomp-jdbc.xml` file, or the `conf/standardjbosscomp-jdbc.xml` file for the current server configuration file set. The defaults section is discussed in Section 11.12.

A detailed description of each entity element follows:

- **ejb-name:** This required element is the name of the EJB to which this configuration applies. This element

must match an `ejb-name` of an entity in the `ejb-jar.xml` file.

- **datasource:** This optional element is the `jndi-name` used to look up the datasource. All database connections used by an entity or relation-table are obtained from the datasource. Having different datasources for entities is not recommended, as it vastly constrains the domain over which finders and `ejbSelects` can query. The default is `java:/DefaultDS` unless overridden in the defaults section.
- **datasource-mapping:** This optional element specifies the name of the `type-mapping`, which determines how Java types are mapped to SQL types, and how EJB-QL functions are mapped to database specific functions. Type mappings are discussed in Section 11.13.2. The default is `Hypersonic SQL` unless overridden in the defaults section.
- **create-table:** This optional element when true, specifies that JBossCMP should attempt to create a table for the entity. When the application is deployed, JBossCMP checks if a table already exists before creating the table. If a table is found, it is logged, and the table is not created. This option is very useful during the early stages of development when the table structure changes often. The default is false unless overridden in the defaults section.
- **alter-table:** If `create-table` is used to automatically create the schema, `alter-table` can be used to keep the schema current with changes to the entity bean. Alter table will perform the following specific tasks:
 - new fields will be created
 - fields which are no longer used will be removed
 - string fields which are shorter than the declared length will have their length increased to the declared length. (not supported by all databases)
- **remove-table:** This optional element when true, JBossCMP will attempt to drop the table for each entity and each relation table mapped relationship. When the application is undeployed, JBossCMP will attempt to drop the table. This option is very useful during the early stages of development when the table structure changes often. The default is false unless overridden in the defaults section.
- **post-table-create:** This optional element specifies an arbitrary SQL statement that should be executed immediately after the database table is created. This command is only executed if `create-table` is true and the table did not previously exist.
- **read-only:** This optional element when true specifies that the bean provider will not be allowed to change the value of any fields. A field that is read-only will not be stored in, or inserted into, the database. If a primary key field is read-only, the create method will throw a `CreateException`. If a set accessor is called on a read-only field, it throws an `EJBException`. Read-only fields are useful for fields that are filled in by database triggers, such as last update. The `read-only` option can be overridden on a per `cmp-field` basis, and is discussed in Section 11.4.4. The default is false unless overridden in the defaults section.
- **read-time-out:** This optional element is the amount of time in milliseconds that a read on a read-only field is valid. A value of 0 means that the value is always reloaded at the start of a transaction, and a value of -1 means that the value never times out. This option can also be overridden on a per `cmp-field` basis. If `read-only` is false, this value is ignored. The default is -1 unless overridden in the defaults section.
- **row-locking:** This optional element if true specifies that JBossCMP will lock all rows loaded in a transaction. Most databases implement this by using the `SELECT FOR UPDATE` syntax when loading the entity, but the actual syntax is determined by the `row-locking-template` in the datasource-mapping used by this entity. The default is false unless overridden in the defaults section.
- **pk-constraint:** This optional element if true specifies that JBossCMP will add a primary key constraint when creating tables. The default is true unless overridden in the defaults section.

- **read-ahead:** This optional element controls caching of query results and `cmr-fields` for the entity. This option is discussed in Section 11.7.3.
- **list-cache-max:** This optional element specifies the number of read-lists that can be tracked by this entity. This option is discussed in `on-load`. The default is 1000 unless overridden in the defaults section.
- **fetch-size:** This optional element specifies the number of entities to read in one round-trip to the underlying datastore. The default is 0 unless overridden in the defaults section.
- **clean-read-ahead-on-load:** When an entity is loaded from the read ahead cache, JBoss can remove the data used from the read ahead cache. The default is `false`.
- **table-name:** This optional element is the name of the table that will hold data for this entity. Each entity instance will be stored in one row of this table. The default is the `ejb-name`.
- **cmp-field:** The optional element allows one to define how the `ejb-jar.xml` `cmp-field` is mapped onto the persistence store. This is discussed in Section 11.4.
- **load-groups:** This optional element specifies one or more groupings of CMP fields to declare load groupings of fields. This is discussed in Section 11.7.2.
- **eager-load-groups:** This optional element defines one or more load grouping as eager load groups. This is discussed in Section 11.8.2.
- **lazy-load-groups:** This optional element defines one or more load grouping as lazy load groups. This is discussed in Section 11.8.3.
- **query:** This optional element specifies the definition of finders and selectors. This is discussed in Section 11.6.
- **unknown-pk:** This optional element allows one to define how an unknown primary key type of `java.lang.Object` maps to the persistent store.
- **entity-command:** This optional element allows one to define the entity creation command instance. Typically this is used to define a custom command instance to allow for primary key generation. This is described in detail in Section 11.11.
- **optimistic-locking:** This optional element defines the strategy to use for optimistic locking. This is described in detail in Section 11.10.
- **audit:** This optional element defines the CMP fields that will be audited. This is described in detail in Section 11.4.5.

11.4. CMP-Fields

11.4.1. CMP-Field Abstract Accessors

Although CMP fields have not changed in CMP 2.0 with regards to functionality, they are no longer declared using fields in the bean implementation class. In CMP 2.0, CMP fields are not directly accessible; rather each CMP field is declared in the bean implementation class of the entity with a set of abstract accessor methods. Abstract accessors are similar to JavaBean property accessors, except no implementation is given. For example, the following listing declares the `gangsterId`, `name`, `nickName`, and `badness` CMP field accessors in the `gangster` entity:

Example 11.7. Sample cmp-field abstract accessor declaration

```
public abstract class GangsterBean implements EntityBean {
    public abstract Integer getGangsterId();
    public abstract void setGangsterId(Integer gangsterId);
    public abstract String getName();
    public abstract void setName(String param);
    public abstract String getNickName();
    public abstract void setNickName(String param);
    public abstract int getBadness();
    public abstract void setBadness(int param);
}
```

Each CMP field is required to have both a getter and a setter method, and each accessor method must be declared public abstract.

11.4.2. CMP-Field Declaration

The declaration of a `cmp-field` in the `ejb-jar.xml` file has not changed at all in EJB 2.0. For example, to declare the `gangsterId`, `name`, `nickName` and `badness` fields defined in Example 11.7 you would add the following to the `ejb-jar.xml` file:

Example 11.8. The ejb-jar.xml cmp-field Declaration

```
<ejb-jar>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <cmp-field><field-name>gangsterId</field-name></cmp-field>
      <cmp-field><field-name>name</field-name></cmp-field>
      <cmp-field><field-name>nickName</field-name></cmp-field>
      <cmp-field><field-name>badness</field-name></cmp-field>
    </entity>
  </enterprise-beans>
</ejb-jar>
```

11.4.3. CMP-Field Column Mapping

The mapping of an `ejb-jar.xml` `cmp-field` is declared in a `jbossCMP-jdbc.xml` `cmp-field` element within the entity. The content model of the `cmp-field` element of the `jbossCMP-jdbc.xml` is shown below.

The following is an example usage of `cmp-field` mapping.

```
<jbossCMP-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <table-name>gangster</table-name>

      <cmp-field>
        <field-name>gangsterId</field-name>
        <column-name>id</column-name>
      </cmp-field>
      <cmp-field>
        <field-name>name</field-name>
        <column-name>name</column-name>
        <not-null/>
      </cmp-field>
    </entity>
  </enterprise-beans>
</jbossCMP-jdbc>
```

```

    </cmp-field>
    <cmp-field>
      <field-name>nickName</field-name>
      <column-name>nick_name</column-name>
      <jdbc-type>VARCHAR</jdbc-type>
      <sql-type>VARCHAR(64)</sql-type>
    </cmp-field>
    <cmp-field>
      <field-name>badness</field-name>
      <column-name>badness</column-name>
    </cmp-field>
  </entity>
</enterprise-beans>
</jboss-cmp-jdbc>

```

In the `cmp-field` element, you can control the name and datatype of the column. A detailed description of each element follows:

- **field-name:** This required element is the name of the `cmp-field` that is being configured. It must match the `field-name` element of a `cmp-field` declared for this entity in the `ejb-jar.xml` file.
- **column-name:** This optional element is the name of the column to which the `cmp-field` is mapped. The default is to use the `field-name` value.
- **not-null:** This optional element indicates that JBossCMP should add a NOT NULL to the end of the column declaration when automatically creating the table for this entity. The default for primary key fields and primitives not null.
- **jdbc-type:** This is the JDBC type that is used when setting parameters in a JDBC PreparedStatement or loading data from a JDBC ResultSet. The valid types are defined in `java.sql.Types`. Only required if `sql-type` is specified, default is based on `datasource-mapping`.
- **sql-type:** This is the SQL type that is used in create table statements for this field. Valid `sql-types` are only limited by your database vendor. Only required if `jdbc-type` is specified, default is based on `datasource-mapping`.
- **property:** This optional element allows one to define how the properties of a dependent value class `cmp-field` should be mapped to the persistent store. This is discussed further in Dependent Value Classes (DVCs).
- **auto-increment:** The presence of this optional field indicates that it is auto-incremented by the database layer. This is used to map a field to a generated column as well as an externally manipulated column.
- **dbindex:** The presence of this optional field indicates that the server should create an index on the corresponding column in the database, and the index name will be `fieldname_index`.
- **check-dirty-after-get:** This value defaults to false for primitive types and the basic `java.lang` immutable wrappers (`Integer`, `String`, etc...). For potentially mutable objects, JBoss will mark they field as potentially dirty after a get operation. If the dirty check on an object is too expensive, you can optimize it away by setting `check-dirty-after-get` to false.
- **state-factory:** This specifies class name of a state factory object which can perform dirty checking for this field. State factory classes must implement the `CMPFieldStateFactory` interface.

11.4.4. Read-only Fields

Another benefit of abstract accessors for cmp-fields is the ability to have read-only fields. The 1.x CMP engine, JAWS, supported read-only with read-time-out for entities. However, the problem with CMP 1.x was the bean provider could always change the value of a field on a read-only entity, and there was nothing the container could do. With CMP 2.x, the container provides the implementation for the accessor, and therefore can throw an exception when the bean provider attempts to set the value of a read-only bean.

In JBossCMP this feature has been extended to the field level with the addition of the `read-only` and `read-time-out` elements to the `cmp-field` element. These elements work the same way as they do at the entity level. If a field is read-only, it will never be used in an `INSERT` or `UPDATE` statement. If a primary key field is read-only, the create method will throw a `CreateException`. If a set accessor is called for a read-only field, it throws an `EJBException`. Read-only fields are useful for fields that are filled in by database triggers, such as last update. A read-only `cmp-field` declaration example follows:

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <cmp-field>
        <field-name>lastUpdated</field-name>
        <read-only>true</read-only>
        <read-time-out>1000</read-time-out>
      </cmp-field>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

11.4.5. Auditing Entity Access

The audit element of the entity section allows one to specify how access to an entity bean is audited. This is only allowed when an entity bean is accessed under a security domain so that this is a caller identity established. The content model of the audit element is given in Figure 11.3.

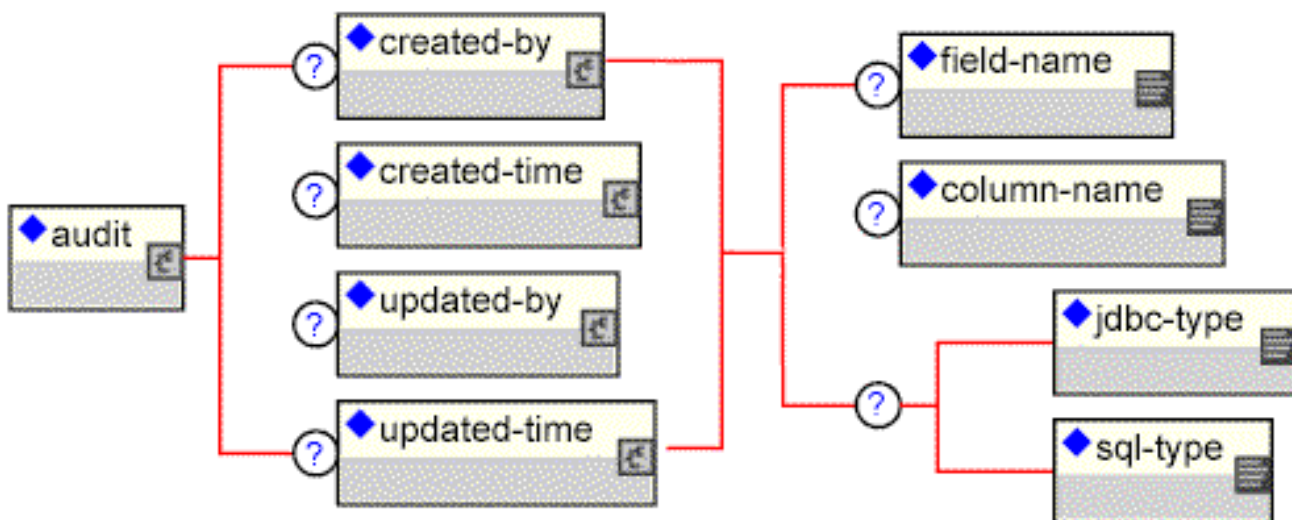


Figure 11.3. The `jbosscmp-jdbc.xml` audit element content model

- **created-by:** This optional element indicates that the caller who created the entity should be saved to either the indicated `column-name` or `cmp field-name`.
- **created-time:** This optional element indicates that the time of entity creation should be saved to either the

indicated `column-name` or `cmp field-name`.

- **updated-by:** This optional element indicates that the caller who last modified the entity should be saved to either the indicated `column-name` or `CMP field-name`.
- **updated-time:** This optional element indicates that the last time of entity modification should be saved to either the indicated `column-name` or `CMP field-name`.
- ***/field-name:** This element indicates that the corresponding audit information should be stored in the indicated `cmp-field` of the accessed entity bean. Note that there does not have to be an actual `CMP field` match in the entity. In case there are matching field names, you will be able to access audit fields in the application using the corresponding `CMP field` abstract getters and setters. Otherwise, audit fields will be created and added to entity internally, and you will be able to access audit information in `EJB-QL` queries using the audit field names, but not directly through the entity accessors.
- ***/column-name:** This element indicates that the corresponding audit information should be stored in the indicated column of the entity table. If `JBossCMP` is creating the table the `jdbc-type` and `sql-type` element can be used to define the storage type.

Example 11.9. A sample audit element declaration

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>AuditChangedNamesEJB</ejb-name>
      <table-name>cmp2_audit_changednames</table-name>
      <audit>
        <created-by>
          <column-name>createdby</column-name>
        </created-by>
        <created-time>
          <column-name>createdtime</column-name>
        </created-time>
        <updated-by>
          <column-name>updatedby</column-name></updated-by>
        <updated-time>
          <column-name>updatedtime</column-name>
        </updated-time>
      </audit>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

11.4.6. Dependent Value Classes (DVCs)

A Dependent Value Class (DVC) is a fancy term used to identify any Java class that is the type of a `cmp-field`, other than the automatically recognized types. See the *Enterprise JavaBeans Specification* for further requirements. By default, a DVC is serialized, and the serialized form is stored in a single database column. Although not discussed here, there are several known issues with the long-term storage of classes in serialized form. `JBossCMP` supports the storage of the internal data of a DVC into one or more columns. This is useful for supporting legacy `JavaBeans` and database structures. It is not uncommon to find a database with a highly flattened structure (e.g., a `PURCHASE_ORDER` table with the fields `SHIP_LINE1`, `SHIP_LINE2`, `SHIP_CITY`, etc. and an additional set of fields for the billing address). Other common database structures include telephone numbers with separate fields for area code, exchange, and extension, or a person's name spread across several fields. With a DVC, multiple columns can be mapped to one logical `JavaBean`.

JBossCMP requires that a DVC to be mapped must follow the JavaBeans naming specification for simple properties, and that each property to be stored in the database must have both a getter and a setter method⁵. Furthermore, the bean must be serializable and must have a no argument constructor. A property can be any simple type, an unmapped DVC or a mapped DVC, but cannot be an EJB⁶. A DVC mapping is specified within the `dependent-value-classes` element.

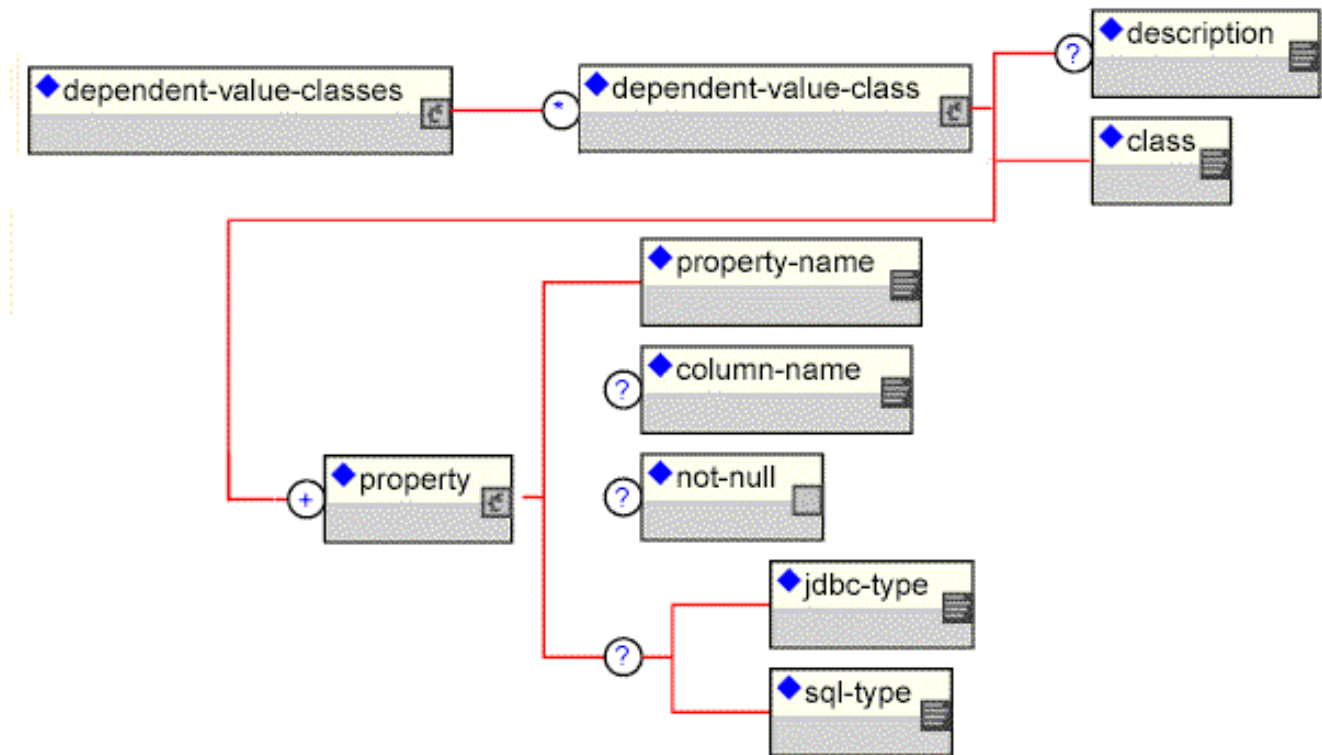


Figure 11.4. The `jbosscmp-jdbc` dependent-value-classes element model.

Here is an example of a simple `ContactInfo` DVC class.

```

public class ContactInfo
    implements Serializable
{
    /** The cell phone number. */
    private PhoneNumber cell;

    /** The pager number. */
    private PhoneNumber pager;

    /** The email address */
    private String email;

    // ...
}

```

The contact info includes a phone number, which is represented by another DVC class.

```

public class PhoneNumber
    implements Serializable
{
    /** The first three digits of the phone number. */
    private short areaCode;
}

```

⁵The requirement that a DVC use the JavaBeans naming convention will be removed in a future release of JBossCMP.

⁶This restriction will also be removed in a future release. The current proposal is to allow the value to be retrieved from any no argument method and to be set with any single argument method or with a constructor.

```

    /** The middle three digits of the phone number. */
    private short exchange;

    /** The last four digits of the phone number. */
    private short extension;

    // ...
}

```

The dependent-value-classes mapping for these two classes is relatively straight forward.

```

<jbosscmp-jdbc>
  <dependent-value-classes>
    <dependent-value-class>
      <description>A phone number</description>
      <class>org.jboss.cmp2.crimeportal.PhoneNumber</class>
      <property>
        <property-name>areaCode</property-name>
        <column-name>area_code</column-name>
      </property>
      <property>
        <property-name>exchange</property-name>
        <column-name>exchange</column-name>
      </property>
      <property>
        <property-name>extension</property-name>
        <column-name>extension</column-name>
      </property>
    </dependent-value-class>

    <dependent-value-class>
      <description>General contact info</description>
      <class>org.jboss.cmp2.crimeportal.ContactInfo</class>
      <property>
        <property-name>cell</property-name>
        <column-name>cell</column-name>
      </property>
      <property>
        <property-name>pager</property-name>
        <column-name>pager</column-name>
      </property>
      <property>
        <property-name>email</property-name>
        <column-name>email</column-name>
        <jdbc-type>VARCHAR</jdbc-type>
        <sql-type>VARCHAR(128)</sql-type>
      </property>
    </dependent-value-class>
  </dependent-value-classes>
</jbosscmp-jdbc>

```

Each DVC is declared with a `dependent-value-class` element. A DVC is identified by the Java class type declared in the `class` element. Each property to be persisted is declared with a `property` element. This specification is based on the `cmp-field` element, so it should be self-explanatory. This restriction will also be removed in a future release. The current proposal involves storing the primary key fields in the case of a local entity and the entity handle in the case of a remote entity.

The `dependent-value-classes` section defines the internal structure and default mapping of the classes. When JBossCMP encounters a field that has an unknown type, it searches the list of registered DVCs, and if a DVC is found, it persists this field into a set of columns, otherwise the field is stored in serialized form in a single column. JBossCMP does not support inheritance of DVCs; therefore, this search is only based on the declared type of the field. A DVC can be constructed from other DVCs, so when JBossCMP runs into a DVC, it flattens the DVC tree structure into a set of columns. If JBossCMP finds a DVC circuit during startup, it will throw an `EJBException`. The default column name of a property is the column name of the base `cmp-field` followed by

an underscore and then the property column name. If the property is a DVC, the process is repeated. For example, a cmp-field named `info` that uses the `ContactInfo` DVC would have the following columns:

```
info_cell_area_code
info_cell_exchange
info_cell_extension
info_pager_area_code
info_pager_exchange
info_pager_extension
info_email
```

The automatically generated column names can quickly become excessively long and awkward. The default mappings of columns can be overridden in the entity element as follows:

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <cmp-field>
        <field-name>contactInfo</field-name>
        <property>
          <property-name>cell.areaCode</property-name>
          <column-name>cell_area</column-name>
        </property>
        <property>
          <property-name>cell.exchange</property-name>
          <column-name>cell_exch</column-name>
        </property>
        <property>
          <property-name>cell.extension</property-name>
          <column-name>cell_ext</column-name>
        </property>

        <property>
          <property-name>pager.areaCode</property-name>
          <column-name>page_area</column-name>
        </property>
        <property>
          <property-name>pager.exchange</property-name>
          <column-name>page_exch</column-name>
        </property>
        <property>
          <property-name>pager.extension</property-name>
          <column-name>page_ext</column-name>
        </property>

        <property>
          <property-name>email</property-name>
          <column-name>email</column-name>
          <jdbc-type>VARCHAR</jdbc-type>
          <sql-type>VARCHAR(128)</sql-type>
        </property>
      </cmp-field>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

When overriding property info for the entity, you need to refer to the property from a flat perspective as in `cell.areaCode`.

11.5. Container Managed Relationships

Container Managed Relationships (CMRs) are a powerful new feature of CMP 2.0. Programmers have been creating relationships between entity objects since EJB 1.0 was introduced (not to mention since the introduc-

tion of databases), but before CMP 2.0 the programmer had to write a lot of code for each relationship in order to extract the primary key of the related entity and store it in a pseudo foreign key field. The simplest relationships were tedious to code, and complex relationships with referential integrity required many hours to code. With CMP 2.0 there is no need to code relationships by hand. The container can manage one-to-one, one-to-many and many-to-many relationships, with referential integrity. One restriction with CMRs is that they are only defined between local interfaces. This means that a relationship cannot be created between two entities in different virtual machines⁷.

There are two basic steps to create a container managed relationship: create the `cmr-field` abstract accessors and declare the relationship in the `ejb-jar.xml` file. The following two sections describe these steps.

11.5.1. CMR-Field Abstract Accessors

CMR-Field abstract accessors have the same signatures as `cmp-fields`, except that single-valued relationships must return the local interface of the related entity, and multi-valued relationships can only return a `java.util.Collection` (or `java.util.Set`) object. As with `cmp-fields`, at least one of the two entities in a relationship must have `cmr-field` abstract accessors. For example, to declare a one-to-many relationship between `Organization` and `Gangster`, first add the following to the `OrganizationBean` class:

```
public abstract class OrganizationBean
    implements EntityBean
{
    public abstract Set getMemberGangsters();
    public abstract void setMemberGangsters(Set gangsters);
}
```

Second, add the following to the `GangsterBean` class:

```
public abstract class GangsterBean
    implements EntityBean
{
    public abstract Organization getOrganization();
    public abstract void setOrganization(Organization org);
}
```

Although each bean declared a `cmr-field`, only one of the two beans in a relationship must have a set of accessors. As with `cmp-fields`, a `cmr-field` is required to have both a getter and a setter method.

11.5.2. Relationship Declaration

The declaration of relationships in the `ejb-jar.xml` file is complicated and error prone. The XML used to declare relationships is as inconsistent as Visual Basic syntax. The best way to configure a relationship is to use a tool, such as XDoclet, or cut and paste a working relationship. The declaration of the `Organization-Gangster` relationship follows:

Example 11.10. The `ejb-jar.xml` relationship Declaration

```
<ejb-jar>
  <relationships>
    <ejb-relation>
      <ejb-relation-name>Organization-Gangster</ejb-relation-name>
      <ejb-relationship-role>
        <ejb-relationship-role-name>org-has-gangsters </ejb-relationship-role-name>
        <multiplicity>One</multiplicity>
        <relationship-role-source>
```

⁷The EJB specification does not even allow for relationships between entities in different applications within the same VM.

```

        <ejb-name>OrganizationEJB</ejb-name>
    </relationship-role-source>
    <cmr-field>
        <cmr-field-name>memberGangsters</cmr-field-name>
        <cmr-field-type>java.util.Set</cmr-field-type>
    </cmr-field>
</ejb-relationship-role>
<ejb-relationship-role>
    <ejb-relationship-role-name>gangster-belongs-to-org </ejb-relationship-role-name>
    <multiplicity>Many</multiplicity>
    <cascade-delete/>
    <relationship-role-source>
        <ejb-name>GangsterEJB</ejb-name>
    </relationship-role-source>
    <cmr-field>
        <cmr-field-name>organization</cmr-field-name>
    </cmr-field>
</ejb-relationship-role>
</ejb-relation>
</relationships>
</ejb-jar>

```

As you can see, each relationship is declared with an `ejb-relation` element within the top level `relationships`⁸ element, and each `ejb-relation` contains two `ejb-relationship-role` elements (one for each entity in the relationship). The `ejb-relationship-role` tags are as follows:

- **ejb-relationshiprole-name:** This optional element is used to identify the role and match the database mapping the `jbosscmp-jdbc.xml` file. The name cannot be the same as the related role.
- **multiplicity:** This required element must be `One` or `Many`. Note, as with all XML elements, this element is case sensitive.
- **cascade-delete:** When this optional element is present, JBossCMP will delete the child entity when the parent entity is deleted. Cascade deletion is only allowed for a role where the other side of the relationship has a multiplicity of one. The default is to not cascade delete.
- **relationship-role-source/ ejb-name:** This required element gives the name of the entity that has the role.
- **cmr-field/ cmr-field-name:** This is the name of the cmr-field of the entity has one, if entity has a cmrfield abstract accessor.
- **cmr-field/ cmr-field-type:** This is the type of the cmr-field. Must be `java.util.Collection` or `java.util.Set`. Only required if cmr-field abstract accessor is collection valued

After adding the `cmr-field` abstract accessors and declaring the relationship, the relationship should be functional. For more information on relationships, see section 10.3 of the EJB 2.0 specification. The next section discusses the database mapping of the relationship.

11.5.3. Relationship Mapping

Relationships can be mapped using either a foreign key or a separate `relation-table`. One-to-one and one-to-many relationships use the foreign key mapping style by default, and many-to-many relationships use only the `relation-table` mapping style. The mapping of a relationship is declared in the `relationships` section of the `jbosscmp-jdbc.xml` descriptor via `ejb-relation` elements. Relationships are identified by the `ejb-relation-name` from the `ejb-jar.xml` file. The `jbosscmp-jdbc.xml` `ejb-relation` element content model is

⁸This is the first place where the specification is inconsistent. It would be much easier if the specification used the following tags: `relationships`, `relationship`, and `relationship-name`.

shown in Figure 11.5.

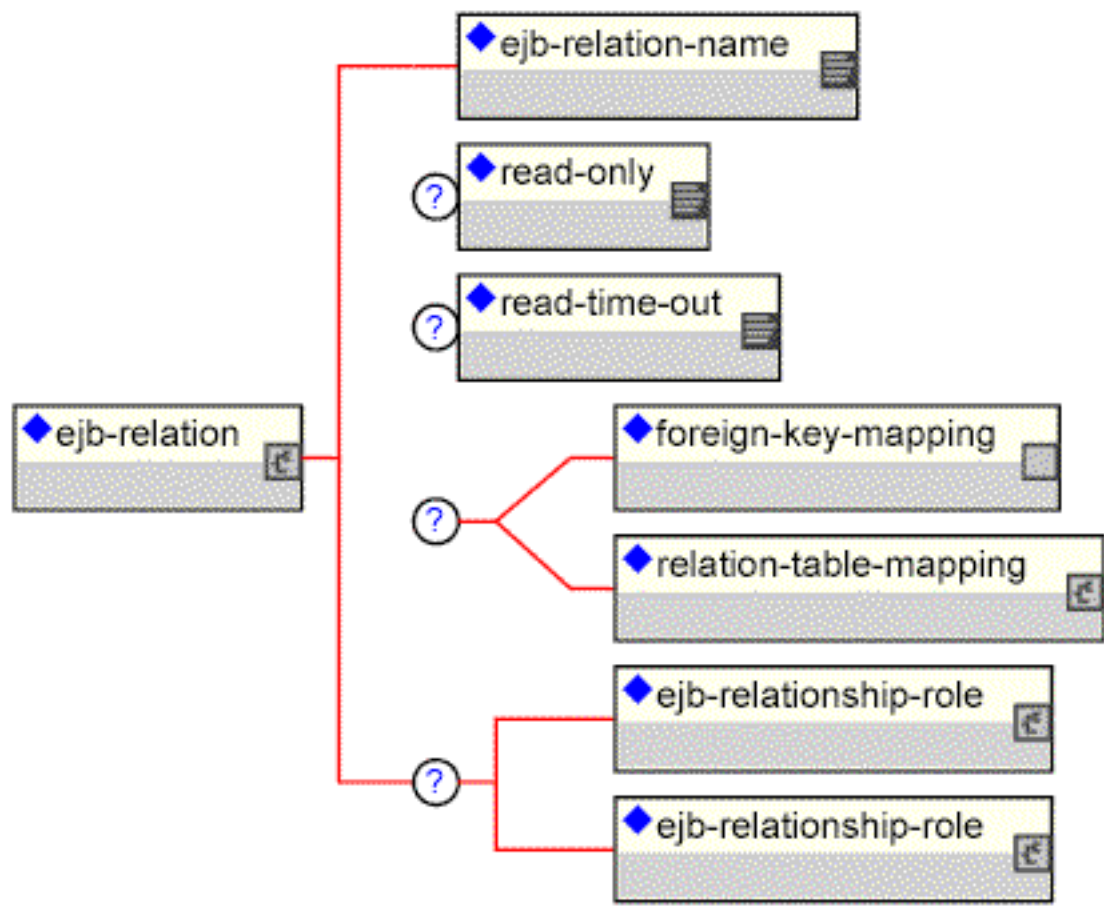


Figure 11.5. The jbosscmp-jdbc.xml ejb-relation element content model

The basic template of the relationship mapping declaration for Organization-Gangster follows:

```

<jbosscmp-jdbc>
  <relationships>
    <ejb-relation>
      <ejb-relation-name>Organization-Gangster</ejb-relation-name>
      <foreign-key-mapping/>
      <ejb-relationship-role>
        <ejb-relationship-role-name>org-has-gangsters</ejb-relationship-role-name>
        <key-fields>
          <key-field>
            <field-name>name</field-name>
            <column-name>organization</column-name>
          </key-field>
        </key-fields>
      </ejb-relationship-role>
      <ejb-relationship-role>
        <ejb-relationship-role-name>gangster-belongs-to-org</ejb-relationship-role-name>
        <key-fields/>
      </ejb-relationship-role>
    </ejb-relation>
  </relationships>
</jbosscmp-jdbc>

```

After the `ejb-relation-name` of the relationship being mapped is declared, the relationship can be declared as read only using the `read-only` and `read-time-out` elements. They have the same semantics as their counterparts in the entity element.

The `ejb-relation` element must contain either a `foreign-key-mapping` element or a `relation-table-mapping` element, which are described in the foreign key mapping and relation-table mapping sections respectively. This element may also contain a pair of `ejb-relationship-role` elements as described in the following section.

11.5.3.1. Relationship Role Mapping

Each of the two `ejb-relationship-role` elements contains mapping information specific to an entity in the relationship, and the content model of the `ejb-relationship-role` element is shown in Figure 11.6 .

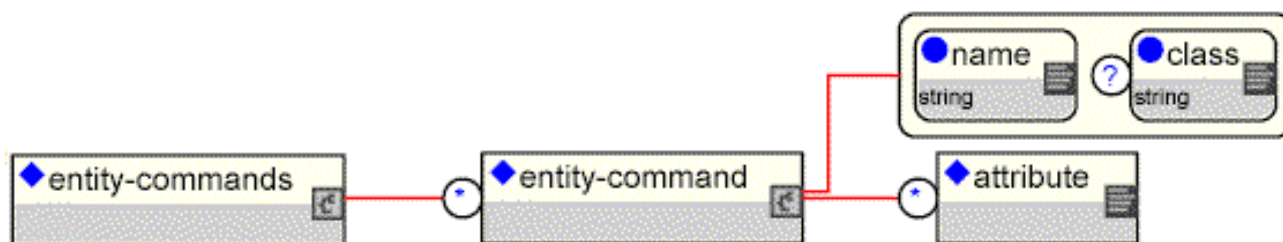


Figure 11.6. The `jbosscmp-jdbc` `ejb-relationship-role` element content model

A detailed description of the main elements follows:

- **ejb-relationship-role-name:** This required element gives the name of the role to which this configuration applies. This element must match the name of one of the roles declared for this query in the `ejb-jar.xml` file.
- **fk-constraint:** This optional element if true indicates that JBossCMP should add a foreign key constraint to the tables. JBossCMP will only add the constraint if both the primary table and the related table were created by JBossCMP during deployment.
- **key-fields:** This optional element specifies the mapping of the primary key fields of the current entity. This element is only necessary if exact field mapping is desired. Otherwise, the `key-fields` element must⁹ contain a `key-field` element for each primary key field of the current entity. The details of this element are described below.
- **read-ahead:** This optional element controls the caching of this relationship. This option is discussed in Section 11.8.3.1.
- **batch-cascade-delete:** When a relationship is marked as `batch-delete` in the `ejb-jar.xml`, the corresponding relationship can be marked with `batch-cascade-delete`. In this case, the cascade delete will be performed with a single SQL statement.

As noted above, the `key-fields` element contains a `key-field` for each primary key field of the current entity. The `key-field` element uses the same syntax as the `cmp-field` element of the entity, except that `key-field` does not support the `not-null` option. Key fields of a `relation-table` are automatically not null, because they are the primary key of the table. On the other hand, foreign key fields must be nullable by default. This is because the current implementation of JBossCMP inserts a row into the database for a new entity between `ejb-Create` and `ejbPostCreate`. Since the EJB specification does not allow a relationship to be modified until `ejbPostCreate`, a foreign key will be initially set to null. There is a similar problem with removal. You can change this insert behavior using the `jboss.xml` `insert-after-ejb-post-create` container configuration flag. The following example illustrates the use of `insert-after-ejb-post-create`.

Note that with foreign key mapping this element can be empty; meaning that there will be not be a foreign key for the current entity. This is required for the many side of a one-to-many relationship, such as Gangster in the Organization-Gangster example.

```

<jboss>
  <!-- ... -->
  <container-configurations>
    <container-configuration extends="Standard CMP 2.x EntityBean">
      <container-name>INSERT after ejbPostCreate Container</container-name>
      <insert-after-ejb-post-create>true</insert-after-ejb-post-create>
    </container-configuration>
  </container-configurations>
</jboss>

```

An alternate means of working around the non-null foreign key issue is to map the foreign key elements onto non-null CMP fields. In this case you simply populate the foreign key fields in `ejbCreate` using the associated CMP field setters.

The content model of the key-fields element is Figure 11.7.

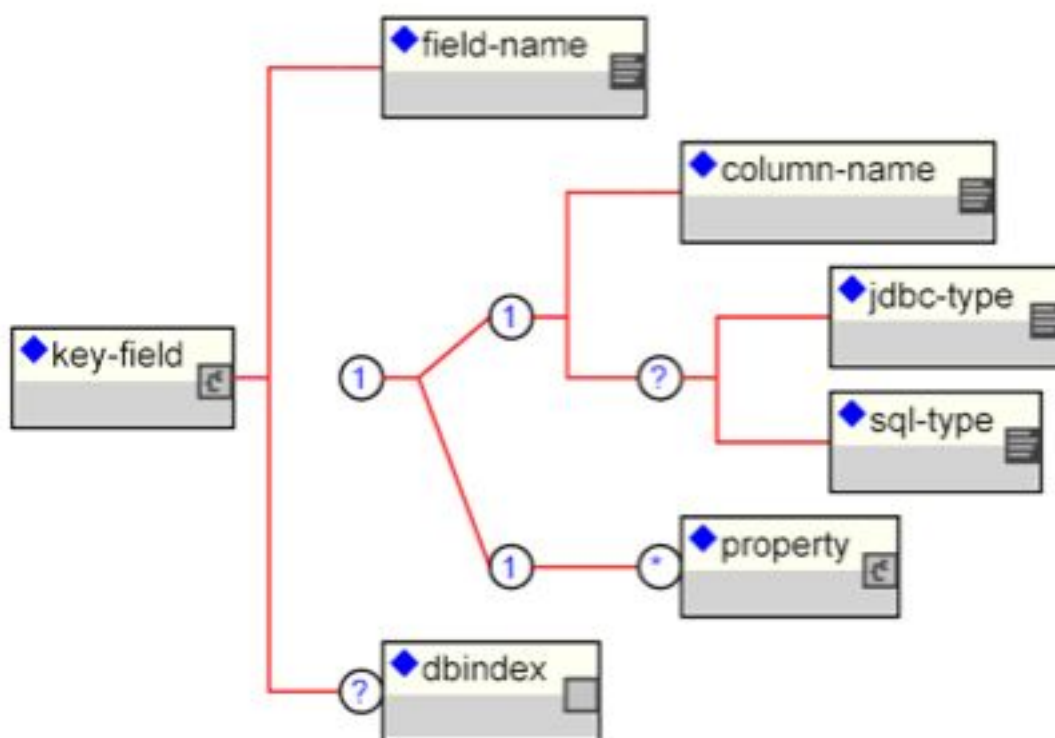


Figure 11.7. The `jbossCMP-jdbc` key-fields element content model

A detailed description of the elements contained in the `key-field` element follows:

- **field-name:** This required element identifies the field to which this mapping applies. This name must match a primary key field of the current entity.
- **column-name:** Use this element to specify the column name in which this primary key field will be stored. If this is relationship uses `foreign-key-mapping`, this column will be added to the table for the related entity. If this relationship uses `relation-table-mapping`, this column is added to the `relation-table`. This element is not allowed for mapped dependent value class; instead use the `property` element.
- **jdbc-type:** This is the JDBC type that is used when setting parameters in a JDBC `PreparedStatement` or loading data from a JDBC `ResultSet`. The valid types are defined in `java.sql.Types`.
- **sql-type:** This is the SQL type that is used in create table statements for this field. Valid types are only lim-

ited by your database vendor.

- **property:** Use this element for to specify the mapping of a primary key field which is a dependent value class.
- **dbindex:** The presence of this optional field indicates that the server should create an index on the corresponding column in the database, and the index name will be `fieldname_index`.

11.5.3.2. Foreign Key Mapping

Foreign key mapping is the most common mapping style for one-to-one and one-to-many relationships, but is not allowed for many-to many relationships. The foreign key mapping element is simply declared by adding an empty `foreign-key-mapping` element to the `ejb-relation` element.

As noted in the previous section, with a foreign key mapping the `key-fields` declared in the `ejb-relationship-role` are added to the table of the related entity. If the `key-fields` element is empty, a foreign key will not be created for the entity. In a one-to-many relationship, the many side (`Gangster` in the example) must have an empty `key-fields` element, and the one side (`Organization` in the example) must have a `key-fields` mapping. In one-to-one relationships, one or both roles can have foreign keys.

The foreign key mapping is not dependent on the direction of the relationship. This means that in a one-to-one unidirectional relationship (only one side has an accessor) one or both roles can still have foreign keys. The complete foreign key mapping for the `Organization-Gangster` relationship is shown below with the foreign key elements highlighted in bold:

```
<jbosscmp-jdbc>
  <relationships>
    <ejb-relation>
      <ejb-relation-name>Organization-Gangster</ejb-relation-name>
      <foreign-key-mapping/>
      <ejb-relationship-role>
        <ejb-relationship-role-name>org-has-gangsters</ejb-relationship-role-name>
        <key-fields>
          <key-field>
            <field-name>name</field-name>
            <column-name>organization</column-name>
          </key-field>
        </key-fields>
      </ejb-relationship-role>
      <ejb-relationship-role>
        <ejb-relationship-role-name>gangster-belongs-to-org</ejb-relationship-role-name>
        <key-fields/>
      </ejb-relationship-role>
    </ejb-relation>
  </relationships>
</jbosscmp-jdbc>
```

11.5.3.3. Relation-table Mapping

Relation table mapping is less common for one-to-one and one-to-many relationships, but is the only mapping style allowed for many-to-many relationships. Relation table mapping is defined using the `relation-table-mapping` element, the content model of which is shown below.

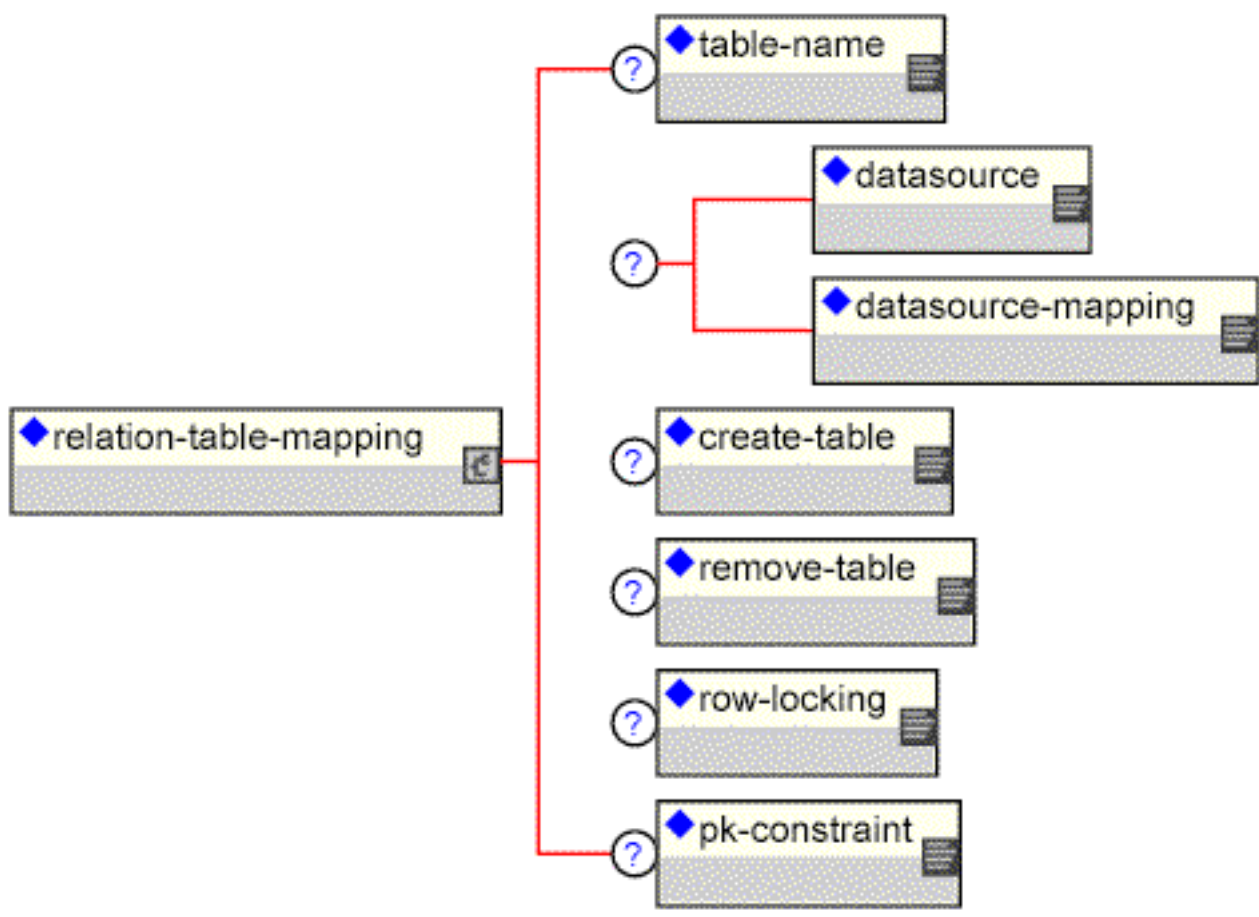


Figure 11.8. The jbossCMP-jdbc relation-table-mapping element content model

The relation-table-mapping for the `Gangster-Job` relationship is shown in with table mapping elements highlighted in bold:

Example 11.11. The jbossCMP-jdbc.xml Relation-table Mapping

```

<jbossCMP-jdbc>
  <relationships>
    <ejb-relation>
      <ejb-relation-name>Gangster-Jobs</ejb-relation-name>
      <relation-table-mapping>
        <table-name>gangster_job</table-name>
      </relation-table-mapping>
      <ejb-relationship-role>
        <ejb-relationship-role-name>gangster-has-jobs</ejb-relationship-role-name>
        <key-fields>
          <key-field>
            <field-name>gangsterId</field-name>
            <column-name>gangster</column-name>
          </key-field>
        </key-fields>
      </ejb-relationship-role>
      <ejb-relationship-role>
        <ejb-relationship-role-name>job-has-gangsters</ejb-relationship-role-name>
        <key-fields>
          <key-field>
            <field-name>name</field-name>
            <column-name>job</column-name>
          </key-field>
        </key-fields>
      </ejb-relationship-role>
    </ejb-relation>
  </relationships>
</jbossCMP-jdbc>
  
```

```
        </ejb-relationship-role>
    </ejb-relation>
</relationships>
</jbosscomp-jdbc>
```

The `relation-table-mapping` element contains a subset of the options available in the `entity` element. A detailed description of these elements is reproduced here for convenience:

- **table-name:** This optional element gives the name of the table that will hold data for this relationship. The default is based on the entity and `cmr-field` names.
- **datasource:** This optional element gives the `jndi-name` used to look up the datasource. All database connections are obtained from the datasource. Having different datasources for entities is not recommended, as it vastly constrains the domain over which finders and `ejbSelects` can query.
- **datasourcemapping:** This optional element allows one to specify the name of the `type-mapping` to use.
- **create-table:** This optional element if true indicates JBossCMP should attempt to create a table for the relationship. When the application is deployed, JBossCMP checks if a table already exists before creating the table. If a table is found, it is logged, and the table is not created. This option is very useful during the early stages of development when the table structure changes often.
- **post-table-create:** This optional element specifies an arbitrary SQL statement that should be executed immediately after the database table is created. This command is only executed if `create-table` is true and the table did not previously exist.
- **remove-table:** This optional element if true indicates JBossCMP should attempt to drop the `relation-table` when the application is undeployed. This option is very useful during the early stages of development when the table structure changes often.
- **row-locking:** This optional element if true indicates JBossCMP should lock all rows loaded in a transaction. Most databases implement this by using the `SELECT FOR UPDATE` syntax when loading the entity, but the actual syntax is determined by the `row-locking-template` in the `datasource-mapping` used by this entity.
- **pk-constraint:** This optional element if true indicates JBossCMP should add a primary key constraint when creating tables.

11.6. Queries

Another powerful new feature of CMP 2.0 is the introduction of the EJB Query Language (EJB-QL) and `ejbSelect` methods. In CMP 1.1, every EJB container had a different way to specify finders, and this was a serious threat to J2EE portability. In CMP 2.0, EJB-QL was created to specify finders and `ejbSelect` methods in a platform independent way. The `ejbSelect` method is designed to provide private query statements to an entity implementation. Unlike finders, which are restricted to only return entities of the same type as the home interface on which they are defined, `ejbSelect` methods can return any entity type or just one field of the entity.

EJB-QL is beyond the scope of this documentation, so only the basic method coding and query declaration will be covered here. For more information, see the Enterprise JavaBeans Specification.

11.6.1. Finder and ejbSelect Declaration

The declaration of finders has not changed in CMP 2.0. Finders are still declared in the home interface (local or remote) of the entity. Finders defined on the local home interface do not throw a `RemoteException`. The following code declares the `findBadDudes_ejbql`¹⁰ finder on the `GangsterHome` interface:

Example 11.12. Finder Declaration

```
public interface GangsterHome
    extends EJBLocalHome
{
    Collection findBadDudes_ejbql(int badness) throws FinderException;
}
```

The `ejbSelect` methods are declared in the entity implementation class, and must be public abstract just like `cmp-field` and `cmr-field` abstract accessors. Select methods must be declared to throw a `FinderException`, but not a `RemoteException`. The following code declares an `ejbSelect` method:

Example 11.13. ejbSelect Declaration

```
public abstract class GangsterBean
    implements EntityBean
{
    public abstract Set ejbSelectBoss_ejbql(String name)
        throws FinderException;
}
```

11.6.2. EJB-QL Declaration

The EJB 2.0 specification requires that every `ejbSelect` or finder method (except `findByPrimaryKey`) have an EJB-QL query defined in the `ejb-jar.xml` file. The EJB-QL query is declared in a query element, which is contained in the entity element. The following are the declarations for `findBadDudes_ejbql` and `ejbSelectBoss_ejbql` queries.

```
<ejb-jar>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!-- ... -->
      <query>
        <query-method>
          <method-name>findBadDudes_ejbql</method-name>
          <method-params>
            <method-param>int</method-param>
          </method-params>
        </query-method>
        <ejb-ql><![CDATA[
          SELECT OBJECT(g)
          FROM gangster g
          WHERE g.badness > ?1
        ]]></ejb-ql>
      </query>
      <query>
        <query-method>
          <method-name>ejbSelectBoss_ejbql</method-name>
```

¹⁰ Ignore the `_ejbql` suffix, it is not required. Later this query will be implemented using `JBossQL` and declared SQL, and the suffix is used to separate the different query specifications in the `jbosscmp-jdbc.xml` file.

```

        <method-params>
            <method-param>java.lang.String</method-param>
        </method-params>
    </query-method>
    <ejb-ql><![CDATA[
        SELECT DISTINCT underling.organization.theBoss
        FROM gangster underling
        WHERE underling.name = ?1 OR underling.nickName = ?1
    ]]></ejb-ql>
</query>
</entity>
</enterprise-beans>
</ejb-jar>

```

EJB-QL is similar to SQL but has some surprising differences. The following are some important things to note about EJB-QL:

- EJB-QL is a typed language, meaning that it only allows comparison of like types (i.e., strings can only be compared with strings).
- In an equals comparison a variable (single valued path) must be on the left hand side. Some examples follow¹¹:

```

g.hangout.state = 'CA' Legal
'CA' = g.shippingAddress.state NOT Legal
'CA' = 'CA' NOT Legal
(r.amountPaid * .01) > 300 NOT Legal
r.amountPaid > (300 / .01) Legal

```

- Parameters use a base 1 index like `java.sql.PreparedStatement`.
- Parameters are only allowed on the right hand side of a comparison. For example:

```

gangster.hangout.state = ?1 Legal
?1 = gangster.hangout.state NOT Legal

```

11.6.3. Overriding the EJB-QL to SQL Mapping

The EJB-QL to SQL mapping can be overridden in the `jbosscmp-jdbc.xml` file. The finder or `ejbSelect` is still required to have an EJB-QL declaration in the `ejb-jar.xml` file, but the `ejb-ql` element can be left empty. Currently the SQL can be overridden with JBossQL, DynamicQL, DeclaredSQL or a BMP style custom `ejbFind` method. All EJB-QL overrides are non-standard extensions to the EJB 2.0 specification, so use of these extensions will limit portability of your application. All of the EJB-QL overrides, except for BMP custom finders, are declared using the `entity/query` element, and the content model is shown in Figure 11.9.

¹¹The example "`(r.amountPaid * .01) > 300`" is presented on page 244 of "Enterprise JavaBeans 3rd Edition" by Richard Monson-Haefel to demonstrate the use of arithmetic operators in a WHERE clause, and is included here to highlight the fact that it is not legal EJB-QL syntax

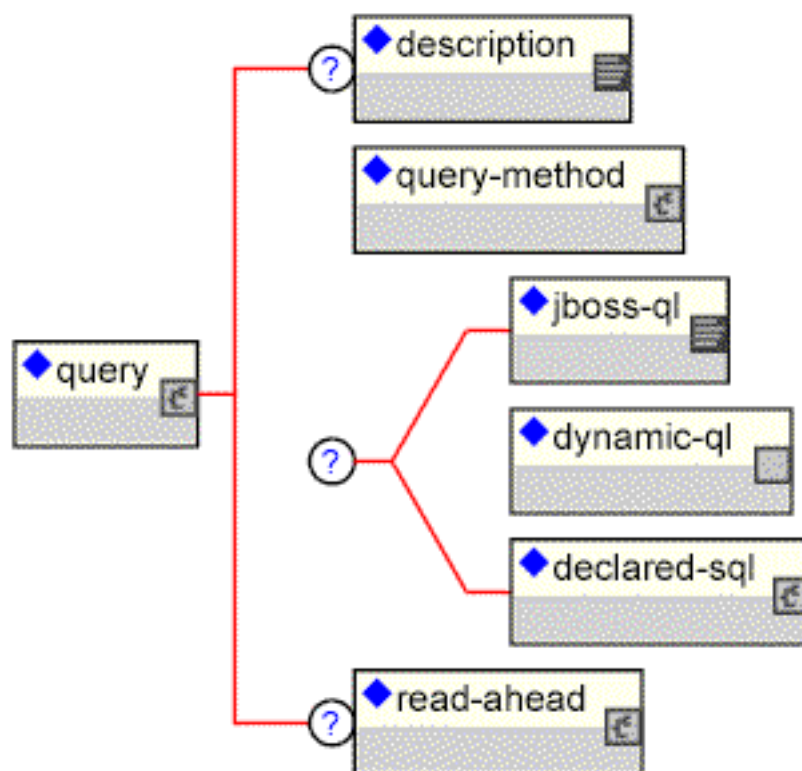


Figure 11.9. The jbosscmp-jdbc query element content model

- **description:** An optional description for the query.
- **query-method:** This required element specifies the query method that being configured. This must match a query-method declared for this entity in the ejb-jar.xml file.
- **jboss-ql, dynamic-ql, declared-sql:** These elements are alternate ways to specify the query method and each is discussed in its own section.
- **read-ahead:** This optional element allows one to optimize the loading of additional fields for use with the entities referenced by the query. This is discussed in detail in Section 11.7.

11.6.4. JBossQL

JBossQL is a superset of EJB-QL that is designed to address some of the inadequacies of EJB-QL. In addition to a more flexible syntax, new functions, key words, and clauses have been added to JBossQL. At the time of this writing, JBossQL includes support for an `ORDER BY`, `OFFSET` and `LIMIT` clauses, parameters in the `IN` and `LIKE` operators, the `COUNT`, `MAX`, `MIN`, `AVG`, `SUM`, `UCASE` and `LCASE` functions, and queries can also include functions in the `SELECT` clause for `ejbSelect` methods.

JBossQL is declared in the `jbosscmp-jdbc.xml` file with a `query/jboss-ql` element containing the JBossQL query. The following example provides an example JBossQL declaration.

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <query>
        <query-method>
          <method-name>findBadDudes_jbossql</method-name>
```

```

        <method-params>
            <method-param>int</method-param>
        </method-params>
    </query-method>
    <jboss-ql><![CDATA[
        SELECT OBJECT(g)
        FROM gangster g
        WHERE g.badness > ?1
        ORDER BY g.badness DESC
    ]]></jboss-ql>
</query>
</entity>
</enterprise-beans>
</jbosscomp-jdbc>

```

The corresponding generated SQL is straightforward.

```

SELECT t0_g.id
FROM gangster t0_g
WHERE t0_g.badness > ?
ORDER BY t0_g.badness DESC

```

Another capability of JBossQL is the ability to retrieve finder results in blocks using the LIMIT and OFFSET functions. For example, to iterate through the large number of jobs performed, the following `findManyJobs_jbossql` finder may be defined.

```

<jbosscomp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <query>
        <query-method>
          <method-name>findManyJobs_jbossql</method-name>
          <method-params>
            <method-param>int</method-param>
          </method-params>
          <method-params>
            <method-param>int</method-param>
          </method-params>
        </query-method>
        <jboss-ql><![CDATA[
          SELECT OBJECT(j)
          FROM jobs j
          OFFSET ?1 LIMIT ?2
        ]]></jboss-ql>
      </query>
    </entity>
  </enterprise-beans>
</jbosscomp-jdbc>

```

11.6.5. DynamicQL

DynamicQL allows the runtime generation and execution of JBossQL queries. A DynamicQL query method is an abstract method that takes a JBossQL query and the query arguments as parameters. JBossCMP compiles the JBossQL and executes the generated SQL. The following generates a JBossQL query that selects all the gangsters that have a hangout in any state in the states set:

```

public abstract class GangsterBean
    implements EntityBean
{
    public Set ejbHomeSelectInStates(Set states)
        throws FinderException
    {
        // generate JBossQL query
    }
}

```

```

StringBuffer jbossQl = new StringBuffer();
jbossQl.append("SELECT OBJECT(g) ");
jbossQl.append("FROM gangster g ");
jbossQl.append("WHERE g.hangout.state IN (");
for(int i = 0; i < states.size(); i++) {
    if(i > 0) {
        jbossQl.append(", ");
    }

    jbossQl.append("?").append(i+1);
}

jbossQl.append(") ORDER BY g.name");

// pack arguments into an Object[]
Object[] args = states.toArray(new Object[states.size()]);

// call dynamic-ql query
return ejbSelectGeneric(jbossQl.toString(), args);
}
}

```

The `DynamicQL` `ejbSelect` method may have any valid `ejbSelect` method name, but the method must always take a `String` and `Object` array as parameters. `DynamicQL` is declared in the `jbosscmp-jdbc.xml` file with an empty `query/dynamic-ql` element. The following is the declaration for `ejbSelectGeneric`.

```

<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <query>
        <query-method>
          <method-name>ejbSelectGeneric</method-name>
          <method-params>
            <method-param>java.lang.String</method-param>
            <method-param>java.lang.Object[]</method-param>
          </method-params>
        </query-method>
        <dynamic-ql/>
      </query>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>

```

11.6.6. DeclaredSQL

`DeclaredSQL` is based on the legacy JAWS CMP 1.1 engine finder declaration, but has been updated for CMP 2.0. Commonly this declaration is used to limit a query with a `WHERE` clause that cannot be represented in `EJB-QL` or `JBossQL`. The content model for the `declared-sql` element is given in Figure 11.10.

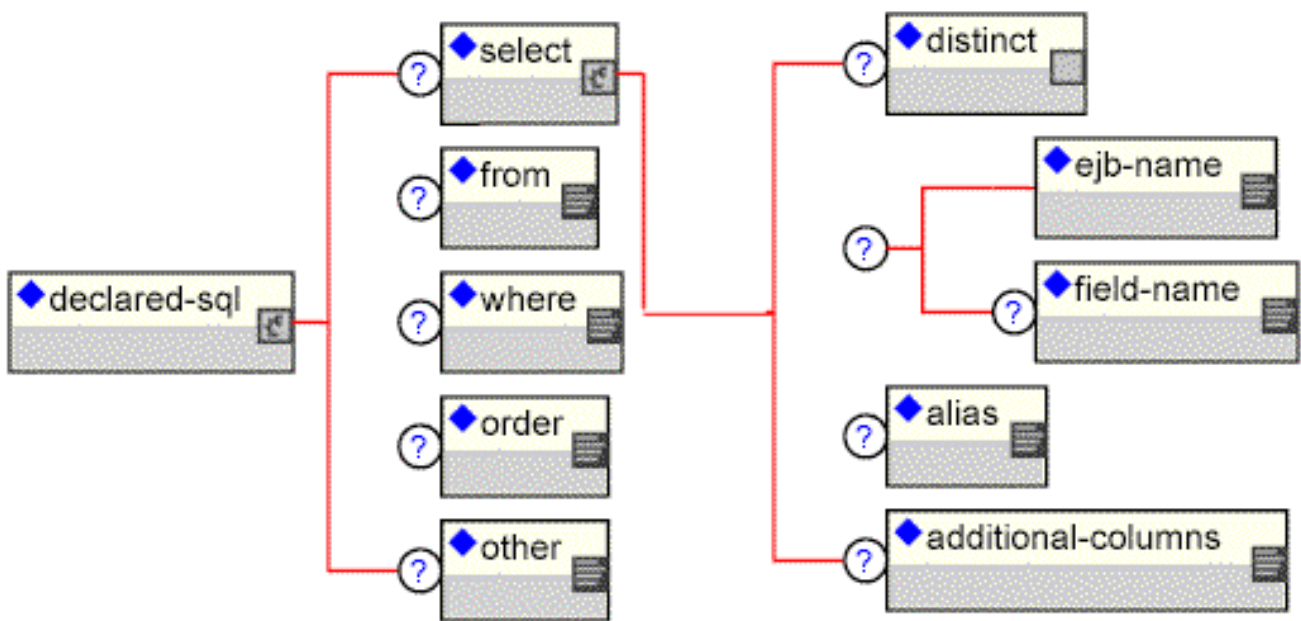


Figure 11.10. The jbosscmp-jdbc declared-sql element content model.>

- **select:** Specifies what is to be selected and consists of the following elements:
 - **distinct:** If this empty element is present, JBossCMP will add the `DISTINCT` keyword to the generated `SELECT` clause. The default is to use `DISTINCT` if method returns a `java.util.Set`
 - **ejb-name:** This is the `ejb-name` of the entity that will be selected. Only required if the query is for an `ejbSelect` method.
 - **field-name:** This is the name of the `cmp-field` that will be selected from the specified entity. The default is to select entire entity.
 - **alias:** This specifies the alias that will be used for the main select table. The default is to use the `ejb-name`.
 - **additional-columns:** Declares other columns to be selected to satisfy ordering by arbitrary columns with finders or to facilitate aggregate functions in selects.
- **from:** Declares additional SQL to append to the generated `from` clause.
- **where:** Declares the `where` clause for the query.
- **order:** Declares the `order` clause for the query.
- **other:** Declares additional SQL that is appended to the end of a query.

The following is an example DeclaredSQL declaration.

```

<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <query>
        <query-method>
          <method-name>findBadDudes_declaredsql</method-name>
        </query-method>
      </query>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>

```

```

        <method-params>
            <method-param>int</method-param>
        </method-params>
    </query-method>
    <declared-sql>
        <where><![CDATA[ badness > {0} ]]></where>
        <order><![CDATA[ badness DESC ]]></order>
    </declared-sql>
</query>
</entity>
</enterprise-beans>
</jbosscomp-jdbc>

```

The generated SQL would be:

```

SELECT id
FROM gangster
WHERE badness > ?
ORDER BY badness DESC

```

As you can see, JBossCMP generates the SELECT and FROM clauses necessary to select the primary key for this entity. If desired an additional FROM clause can be specified that is appended to the end of the automatically generated FROM clause. The following is example DeclaredSQL declaration with an additional FROM clause.

```

<jbosscomp-jdbc>
    <enterprise-beans>
        <entity>
            <ejb-name>GangsterEJB</ejb-name>
            <query>
                <query-method>
                    <method-name>ejbSelectBoss_declaredsql</method-name>
                    <method-params>
                        <method-param>java.lang.String</method-param>
                    </method-params>
                </query-method>
                <declared-sql>
                    <select>
                        <distinct/>
                        <ejb-name>GangsterEJB</ejb-name>
                        <alias>boss</alias>
                    </select>
                    <from><![CDATA[, gangster g, organization o]]></from>
                    <where><![CDATA[
                        (LCASE(g.name) = {0} OR LCASE(g.nick_name) = {0}) AND
                        g.organization = o.name AND o.the_boss = boss.id
                    ]]></where>
                </declared-sql>
            </query>
        </entity>
    </enterprise-beans>
</jbosscomp-jdbc>

```

The generated SQL would be:

```

SELECT DISTINCT boss.id
FROM gangster boss, gangster g, organization o
WHERE (LCASE(g.name) = ? OR LCASE(g.nick_name) = ?) AND
      g.organization = o.name AND o.the_boss = boss.id

```

Notice that the FROM clause starts with a comma. This is because the container appends the declared FROM clause to the end of the generated FROM clause. It is also possible for the FROM clause to start with a SQL JOIN statement. Since this is an ejbSelect method, it must have a select element to declare the entity that will be selected. Note that an alias is also declared for the query. If an alias is not declared, the table-name is used as the

alias, resulting in a `SELECT` clause with the `table_name.field_name` style column declarations. Not all database vendors support the that syntax, so the declaration of an alias is preferred. The optional empty `distinct` element causes the `SELECT` clause to use the `SELECT DISTINCT` declaration. The `DeclaredSQL` declaration can also be used in `ejbSelect` methods to select a `cmp-field`.

Now we will see an example which overrides an `ejbSelect` to select all of the zip codes in which an `Organization` operates.

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>OrganizationEJB</ejb-name>
      <query>
        <query-method>
          <method-name>ejbSelectOperatingZipCodes_declaredsql</method-name>
          <method-params>
            <method-param>java.lang.String</method-param>
          </method-params>
        </query-method>
        <declared-sql>
          <select>
            <distinct/>
            <ejb-name>LocationEJB</ejb-name>
            <field-name>zipCode</field-name>
            <alias>hangout</alias>
          </select>
          <from><![CDATA[ , organization o, gangster g ]]></from>
          <where><![CDATA[
            LCASE(o.name) = {0} AND o.name = g.organization AND
            g.hangout = hangout.id
          ]]></where>
          <order><![CDATA[ hangout.zip ]]></order>
        </declared-sql>
      </query>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

The corresponding SQL would be:

```
SELECT DISTINCT hangout.zip
FROM location hangout, organization o, gangster g
WHERE LCASE(o.name) = ? AND o.name = g.organization AND g.hangout = hangout.id
ORDER BY hangout.zip
```

11.6.6.1. Parameters

JBossCMP `DeclaredSQL` uses a completely new parameter handling system, which supports entity and DVC parameters. Parameters are enclosed in curly brackets and use a base zero index, which is different from the base one EJB-QL parameters. There are three categories of parameters: simple, DVC, and entity:

- **simple:** A simple parameter can be of any type except for a known (mapped) DVC or an entity. A simple parameter only contains the argument number, such as `{0}`. When a simple parameter is set, the JDBC type used to set the parameter is determined by the `datasourcemapping` for the entity. An unknown DVC is serialized and then set as a parameter. Note that most databases do not support the use of a BLOB value in a `WHERE` clause.
- **DVC:** A DVC parameter can be any known (mapped) DVC. A DVC parameter must be dereferenced down to a simple property (one that is not another DVC). For example, if we had a CVS property of type `ContactInfo`, valid parameter declarations would be `{0.email}` and `{0.cell.areaCode}` but not `{0.cell}`. The JDBC type used to set a parameter is based on the class type of the property and the `datasourcemapping`.

ping of the entity. The JDBC type used to set the parameter is the JDBC type that is declared for that property in the `dependent-value-class` element.

- **entity:** An entity parameter can be any entity in the application. An entity parameter must be dereferenced down to a simple primary key field or simple property of a DVC primary key field. For example, if we had a parameter of type `Gangster`, a valid parameter declaration would be `{0.gangsterId}`. If we had some entity with a primary key field named `info` of type `ContactInfo`, a valid parameter declaration would be `{0.info.cell.areaCode}`. Only fields that are members of the primary key of the entity can be dereferenced (this restriction may be removed in later versions). The JDBC type used to set the parameter is the JDBC type that is declared for that field in the entity declaration.

11.6.7. EJBQL 2.1 and SQL92 queries

The default query compiler doesn't fully support EJB-QL 2.1 or the SQL92 standard. If you need either of these functions, you can replace the query compiler. The default compiler is specified in `standardjbosscmp-jdbc.xml`.

```
<defaults>
...
<ql-compiler>org.jboss.ejb.plugins.cmp.jdbc.JDBCEJBQLCompiler</ql-compiler>
...
</defaults>
```

To use the SQL92 compiler, simply specify the SQL92 compiler in `ql-compiler` element.

```
<defaults>
...
<ql-compiler>org.jboss.ejb.plugins.cmp.jdbc.EJBQLToSQL92Compiler</ql-compiler>
...
</defaults>
```

This changes the query compiler for all beans in the entire system. You can also specify the `ql-compiler` for each element in `jbosscmp-jdbc.xml`. Here is an example using one of our earlier queries.

```
<query>
  <query-method>
    <method-name>findBadDudes_ejbql</method-name>
    <method-params>
      <method-param>int</method-param>
    </method-params>
  </query-method>
  <ejb-ql><![CDATA[
    SELECT OBJECT(g)
    FROM gangster g
    WHERE g.badness > ?1]]>
  </ejb-ql>
  <ql-compiler>org.jboss.ejb.plugins.cmp.jdbc.EJBQLToSQL92Compiler</ql-compiler>
</query>
```

One important limitation of SQL92 query compiler is that it always selects all the fields of an entity regardless the read-ahead strategy in use. For example, if a query is configured with the `on-load` read-ahead strategy, the first query will include all the fields, not just primary key fields but only the primary key fields will be read from the `ResultSet`. Then, on load, other fields will be actually loaded into the read-ahead cache. The `on-find` read-ahead with the default load group * works as expected.

11.6.8. BMP Custom Finders

JBossCMP continues the tradition of JAWS in supporting bean managed persistence custom finders. If a custom finder matches a finder declared in the home or local home interface, JBossCMP will always call the custom finder over any other implementation declared in the ejb-jar.xml or jbosscomp-jdbc.xml files. The following simple example finds the entities by a collection of primary keys¹²:

Example 11.14. Custom Finder Example Code

```
public abstract class GangsterBean
    implements EntityBean
{
    public Collection ejbFindByPrimaryKeys(Collection keys)
    {
        return keys;
    }
}
```

11.7. Optimized Loading

The goal of optimized loading is to load the smallest amount of data required to complete a transaction in the least number of queries. The tuning of JBossCMP depends on a detailed knowledge of the loading process. This section describes the internals of the JBossCMP loading process and its configuration. Tuning of the loading process really requires a holistic understanding of the loading system, so this chapter may have to be read more than once.

11.7.1. Loading Scenario

The easiest way to investigate the loading process is to look at a usage scenario. The most common scenario is to locate a collection of entities and iterate over the results performing some operation. The following example generates an html table containing all of the gangsters:

Example 11.15. Loading Scenario Example Code

```
public String createGangsterHtmlTable_none()
    throws FinderException
{
    StringBuffer table = new StringBuffer();
    table.append("<table>");

    Collection gangsters = gangsterHome.findAll_none();
    for(Iterator iter = gangsters.iterator(); iter.hasNext(); ) {
        Gangster gangster = (Gangster)iter.next();
        table.append("<tr>");
        table.append("<td>").append(gangster.getName());
        table.append("</td>");
        table.append("<td>").append(gangster.getNickName());
        table.append("</td>");
        table.append("<td>").append(gangster.getBadness());
    }
}
```

¹²This is a very useful finder because it quickly converts primary keys into real Entity objects without contacting the database. One drawback is that it can create an Entity object with a primary key that does not exist in the database. If any method is invoked on the bad Entity, a NoSuchEntityException will be thrown. Another drawback is that the resulting entity bean violates the EJB specification in that it implements a finder, and the JBoss EJB verifier will fail the deployment of such an entity unless the StrictVerifier attribute is set to false.

```

        table.append("</td>");
        table.append("</tr>");
    }

    return table.toString();
}

```

Assume this code is called within a single transaction and all optimized loading has been disabled. At line 5, JBossCMP will execute the following query:

Example 11.16. Unoptimized findAll Query

```

SELECT t0_g.id
  FROM gangster t0_g
 ORDER BY t0_g.id ASC

```

Then at line 8, in order to load the eight Gangsters in the sample database, JBossCMP executes the following eight queries:

Example 11.17. Unoptimized Load Queries

```

SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=0)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=1)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=2)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=3)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=4)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=5)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=6)
SELECT name, nick_name, badness, hangout, organization
  FROM gangster WHERE (id=7)

```

There are two problems with this scenario. First, an excessive number of queries are executed because JBossCMP executes one query for findAll and one query for each element found. This is known as the "n+1" problem¹³ and is addressed with the read-ahead strategies described in the following sections. Second, values of unused fields are loaded because JBossCMP loads the hangout and organization fields¹⁴, which are never accessed. Configuration of eager loading is described in Section 11.8.2. The following table shows the execution of the queries:

Table 11.1. Unoptimized Query Execution

¹³The reason for this behavior has to do with the handling of query results inside the JBoss container. Although it appears that the actual entity beans selected are returned when a query is executed, JBoss really only returns the primary keys of the matching entities, and does not load the entity until a method is invoked on it.

¹⁴Normally JBossCMP would also load the contactInfo field, but for the sake of readability, it has been disabled in this example because contact info maps to seven columns. The actual configuration used to disable the default loading of the contactInfo field is presented in Listing 6-12.

id	name	nick_name	badness	hangout	organization
0	Yojimbo	Bodyguard	7	0	Yakuza
1	Takeshi	Master	10	1	Yakuza
2	Yuriko	Four finger	4	2	Yakuza
3	Chow	Killer	9	3	Triads
4	Shogi	Lightning	8	4	Triads
5	Valentino	Pizza-Face	4	5	Mafia
6	Toni	Toothless	2	6	Mafia
7	Corleone	Godfather	6	7	Mafia

11.7.2. Load Groups

The configuration and optimization of the loading system begins with the declaration of named load groups in the entity. A load group contains the names of `cmp-fields` and `cmr-fields` with a foreign key (e.g., `Gangster` in the `Organization-Gangster` example) that will be loaded in a single operation. An example configuration is shown below:

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!-- ... -->
      <load-groups>
        <load-group>
          <load-group-name>basic</load-group-name>
          <field-name>name</field-name>
          <field-name>nickName</field-name>
          <field-name>badness</field-name>
        </load-group>
        <load-group>
          <load-group-name>contact info</load-group-name>
          <field-name>nickName</field-name>
          <field-name>contactInfo</field-name>
          <field-name>hangout</field-name>
        </load-group>
      </load-groups>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

In this example, two load groups are declared: `basic` and `contact info`. Note that the load groups do not need to be mutually exclusive. For example, both of the load groups contain the `nickName` field. In addition to the declared load groups, JBossCMP automatically adds a group named `*` (the star group) that contains every `cmp-field` and `cmr-field` with a foreign key in the entity.

11.7.3. Read-ahead

Optimized loading in JBossCMP is called read-ahead. This term was inherited from JAWS, and refer to the technique of reading the row for an entity being loaded, as well as the next several rows; hence the term read-ahead. JBossCMP implements two main strategies (`on-find` and `on-load`) to optimize the loading problem identified in the previous section. The extra data loaded during read-ahead is not immediately associated with

an entity object in memory, as entities are not materialized in JBoss until actually accessed. Instead, it is stored in the preload cache where it remains until it is loaded into an entity or the end of the transaction occurs. The following sections describe the read-ahead strategies.

11.7.3.1. on-find

The `on-find` strategy reads additional columns when the query is invoked. If the query in the scenario detailed Example 11.15 is `on-find` optimized, JBossCMP will execute the following query at line 5:

```
SELECT t0_g.id, t0_g.name, t0_g.nick_name, t0_g.badness
FROM gangster t0_g
ORDER BY t0_g.id ASC
```

Then at line 8, all of the required data would be in the preload cache, so no additional queries would be executed. This strategy is effective for queries that return a small amount of data, but becomes very inefficient when trying to load a large result set into memory¹⁵. The following table shows the execution of this query:

Table 11.2. on-find Optimized Query Execution

id	name	nick_name	badness	hangout	organization
0	Yojimbo	Bodyguard	7	0	Yakuza
1	Takeshi	Master	10	1	Yakuza
2	Yuriko	Four finger	4	2	Yakuza
3	Chow	Killer	9	3	Triads
4	Shogi	Lightning	8	4	Triads
5	Valentino	Pizza-Face	4	5	Mafia
6	Toni	Toothless	2	6	Mafia
7	Corleone	Godfather	6	7	Mafia

The `read-ahead` strategy and `load-group` for a query is defined in the `query` element. If a `read-ahead` strategy is not declared in the `query` element, the strategy declared in the `entity` element or `defaults` element is used. The `on-find` configuration follows:

```
<jbossCMP-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!--...-->
      <query>
        <query-method>
          <method-name>findAll_onfind</method-name>
          <method-params/>
        </query-method>
        <jboss-ql><![CDATA[
          SELECT OBJECT(g)
          FROM gangster g
          ORDER BY g.gangsterId
        ]]></jboss-ql>
        <read-ahead>
          <strategy>on-find</strategy>
          <page-size>4</page-size>
          <eager-load-group>basic</eager-load-group>
        </read-ahead>
      </query>
    </entity>
  </enterprise-beans>
</jbossCMP-jdbc>
```

¹⁵ JBossCMP uses soft references in the read-ahead cache implementation, so data will be cached and then immediately released.


```

        </read-ahead>
    </query>
</entity>
</enterprise-beans>
</jboss-cmp-jdbc>

```

One problem with the `on-find` strategy is that it must load additional data for every entity selected. Commonly in web applications only a fixed number of results are rendered on a page. Since the preloaded data is only valid for the length of the transaction, and a transaction is limited to a single web HTTP hit, most of the preloaded data is not used. The `on-load` strategy discussed in the next section does not suffer from this problem.

11.7.3.1.1. Left join read ahead

Left join read ahead is an enhanced `on-find read-ahead` strategy. It allows you to preload in one SQL query not only fields from the base instance but also related instances which can be reached from the base instance by CMR navigation. There are no limitation for the depth of CMR navigations. There are also no limitations for cardinality of CMR fields used in navigation and relationship type mapping, i.e. both foreign key and relation-table mapping styles are supported. Let's look at some examples. Entity and relationship declarations can be found below.

11.7.3.1.2. D#findByPrimaryKey

Suppose we have an entity `D`. A typical SQL query generated for the `findByPrimaryKey` would look like this:

```
SELECT t0_D.id, t0_D.name FROM D t0_D WHERE t0_D.id=?
```

Suppose that while executing `findByPrimaryKey` we also want to preload two collection-valued CMR fields `bs` and `cs`.

```

<query>
  <query-method>
    <method-name>findByPrimaryKey</method-name>
    <method-params>
      <method-param>java.lang.Long</method-param>
    </method-params>
  </query-method>
  <jboss-ql><![CDATA[SELECT OBJECT(o) FROM D AS o WHERE o.id = ?1]]></jboss-ql>
  <read-ahead>
    <strategy>on-find</strategy>
    <page-size>4</page-size>
    <eager-load-group>basic</eager-load-group>
    <left-join cmr-field="bs" eager-load-group="basic"/>
    <left-join cmr-field="cs" eager-load-group="basic"/>
  </read-ahead>
</query>

```

The `left-join` declares the relations to be eager loaded. The generated SQL would look like this:

```

SELECT t0_D.id, t0_D.name,
       t1_D_bs.id, t1_D_bs.name,
       t2_D_cs.id, t2_D_cs.name
FROM D t0_D
     LEFT OUTER JOIN B t1_D_bs ON t0_D.id=t1_D_bs.D_FK
     LEFT OUTER JOIN C t2_D_cs ON t0_D.id=t2_D_cs.D_FK
WHERE t0_D.id=?

```

For the `D` with the specific `id` we preload all its related `B`'s and `C`'s and can access those instance loading them from the read ahead cache, not from the database.

11.7.3.1.3. D#findAll

In the same way, we could optimize the `findAll` method on `D` selects all the `D`'s. A normal `findAll` query would look like this:

```
SELECT DISTINCT t0_o.id, t0_o.name FROM D t0_o ORDER BY t0_o.id DESC
```

To preload the relations, we simply need to add the `left-join` elements to the query.

```
<query>
  <query-method>
    <method-name>findAll</method-name>
  </query-method>
  <jboss-ql><![CDATA[SELECT DISTINCT OBJECT(o) FROM D AS o ORDER BY o.id DESC]]></jboss-ql>
  <read-ahead>
    <strategy>on-find</strategy>
    <page-size>4</page-size>
    <eager-load-group>basic</eager-load-group>
    <left-join cmr-field="bs" eager-load-group="basic"/>
    <left-join cmr-field="cs" eager-load-group="basic"/>
  </read-ahead>
</query>
```

And here is the generated SQL:

```
SELECT DISTINCT t0_o.id, t0_o.name,
                t1_o_bs.id, t1_o_bs.name,
                t2_o_cs.id, t2_o_cs.name
FROM D t0_o
     LEFT OUTER JOIN B t1_o_bs ON t0_o.id=t1_o_bs.D_FK
     LEFT OUTER JOIN C t2_o_cs ON t0_o.id=t2_o_cs.D_FK
ORDER BY t0_o.id DESC
```

Now the simple `findAll` query now preloads the related `B` and `C` objects for each `D` object.

11.7.3.1.4. A#findAll

Now let's look at a more complex configuration. Here we want to preload instance `A` along with several relations.

- its parent (self-relation) reached from `A` with CMR field `parent`
- the `B` reached from `A` with CMR field `b`, and the related `C` reached from `B` with CMR field `c`
- `B` reached from `A` but this time with CMR field `b2` and related to it `C` reached from `B` with CMR field `c`.

For reference, the standard query would be:

```
SELECT t0_o.id, t0_o.name FROM A t0_o ORDER BY t0_o.id DESC FOR UPDATE
```

The following metadata describes our preloading plan.

```
<query>
  <query-method>
    <method-name>findAll</method-name>
  </query-method>
  <jboss-ql><![CDATA[SELECT OBJECT(o) FROM A AS o ORDER BY o.id DESC]]></jboss-ql>
  <read-ahead>
    <strategy>on-find</strategy>
    <page-size>4</page-size>
```

```

<eager-load-group>basic</eager-load-group>
<left-join cmr-field="parent" eager-load-group="basic"/>
<left-join cmr-field="b" eager-load-group="basic">
  <left-join cmr-field="c" eager-load-group="basic"/>
</left-join>
<left-join cmr-field="b2" eager-load-group="basic">
  <left-join cmr-field="c" eager-load-group="basic"/>
</left-join>
</read-ahead>
</query>

```

The SQL query generated would be:

```

SELECT t0_o.id, t0_o.name,
       t1_o_parent.id, t1_o_parent.name,
       t2_o_b.id, t2_o_b.name,
       t3_o_b_c.id, t3_o_b_c.name,
       t4_o_b2.id, t4_o_b2.name,
       t5_o_b2_c.id, t5_o_b2_c.name
FROM A t0_o
LEFT OUTER JOIN A t1_o_parent ON t0_o.PARENT=t1_o_parent.id
LEFT OUTER JOIN B t2_o_b ON t0_o.B_FK=t2_o_b.id
LEFT OUTER JOIN C t3_o_b_c ON t2_o_b.C_FK=t3_o_b_c.id
LEFT OUTER JOIN B t4_o_b2 ON t0_o.B2_FK=t4_o_b2.id
LEFT OUTER JOIN C t5_o_b2_c ON t4_o_b2.C_FK=t5_o_b2_c.id
ORDER BY t0_o.id DESC FOR UPDATE

```

With this configuration, you can navigate CMRs from any found instance of A without an additional database load.

11.7.3.1.5. A#findMeParentGrandParent

Here is some more example of self-relation. Suppose, we want to write a method that would preload an instance, its parent, grand-parent and its grand-grand-parent in one query. To do this, we would used nested left-join declaration.

```

<query>
  <query-method>
    <method-name>findMeParentGrandParent</method-name>
    <method-params>
      <method-param>java.lang.Long</method-param>
    </method-params>
  </query-method>
  <jboss-ql><![CDATA[SELECT OBJECT(o) FROM A AS o WHERE o.id = ?1]]></jboss-ql>
  <read-ahead>
    <strategy>on-find</strategy>
    <page-size>4</page-size>
    <eager-load-group>*</eager-load-group>
    <left-join cmr-field="parent" eager-load-group="basic">
      <left-join cmr-field="parent" eager-load-group="basic">
        <left-join cmr-field="parent" eager-load-group="basic"/>
      </left-join>
    </left-join>
  </read-ahead>
</query>

```

The generated SQL would be:

```

SELECT t0_o.id, t0_o.name, t0_o.secondName, t0_o.B_FK, t0_o.B2_FK, t0_o.PARENT,
       t1_o_parent.id, t1_o_parent.name,
       t2_o_parent_parent.id, t2_o_parent_parent.name,
       t3_o_parent_parent_parent.id, t3_o_parent_parent_parent.name
FROM A t0_o
LEFT OUTER JOIN A t1_o_parent ON t0_o.PARENT=t1_o_parent.id
LEFT OUTER JOIN A t2_o_parent_parent ON t1_o_parent.PARENT=t2_o_parent_parent.id

```

```
LEFT OUTER JOIN A t3_o_parent_parent ON t2_o_parent_parent.PARENT=t3_o_parent_parent
WHERE (t0_o.id = ?) FOR UPDATE
```

Note, if we remove left-join metadata we will have only

```
SELECT t0_o.id, t0_o.name, t0_o.secondName, t0_o.B2_FK, t0_o.PARENT FOR UPDATE
```

11.7.3.2. on-load

The on-load strategy block-loads additional data for several entities when an entity is loaded, starting with the requested entity and the next several entities in the order they were selected. This strategy is based on the theory that the results of a find or select will be accessed in forward order. When a query is executed, JBossCMP stores the order of the entities found in the list cache. Later, when one of the entities is loaded, JBossCMP uses this list to determine the block of entities to load. The number of lists stored in the cache is specified with the list-cachemax element of the entity. This strategy is also used when faulting in data not loaded in the on-find strategy. With this strategy, the query executed at line 5 of Example 11.15 remains unchanged.

Example 11.18. on-load (Unoptimized) findAll Query

```
SELECT t0_g.id
FROM gangster t0_g
ORDER BY t0_g.id ASC
```

If, for example, the on-load/page-size is set to four, JBossCMP will execute the following two queries to load the name, nickName and badness fields for the entities:

Example 11.19. on-load Optimized Load Queries

```
SELECT id, name, nick_name, badness
FROM gangster
WHERE (id=0) OR (id=1) OR (id=2) OR (id=3)
SELECT id, name, nick_name, badness
FROM gangster
WHERE (id=4) OR (id=5) OR (id=6) OR (id=7)
```

The following table shows the execution of these queries:

Table 11.3. on-load Optimized Query Execution

id	name	nick_name	badness	hangout	organization
0	Yojimbo	Bodyguard	7	0	Yakuza
1	Takeshi	Master	10	1	Yakuza
2	Yuriko	Four finger	4	2	Yakuza
3	Chow	Killer	9	3	Triads
4	Shogi	Lightning	8	4	Triads
5	Valentino	Pizza-Face	4	5	Mafia
6	Toni	Toothless	2	6	Mafia

id	name	nick_name	badness	hangout	organization
7	Corleone	Godfather	6	7	Mafia

As with the on-find strategy, on-load is declared in the read-ahead element. The on-load configuration for this example is shown below.

Example 11.20. The jbosscmp-jdbc.xml on-load Declaration

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!-- ... -->
      <query>
        <query-method>
          <method-name>findAll_onload</method-name>
          <method-params/>
        </query-method>
        <jboss-ql><![CDATA[
          SELECT OBJECT(g)
          FROM gangster g
          ORDER BY g.gangsterId
        ]]></jboss-ql>
        <read-ahead>
          <strategy>on-load</strategy>
          <page-size>4</page-size>
          <eager-load-group>basic</eager-load-group>
        </read-ahead>
      </query>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

11.7.3.3. none

The none strategy is really an anti-strategy. This strategy causes the system to fall back to the default lazy-load code, and specifically does not read-ahead any data or remember the order of the found entities. This results in the queries and performance shown at the beginning of this chapter. The none strategy is declared with a read-ahead element. If the read-ahead element contains a page-size element or eager-load-group, it is ignored. The none strategy is declared the following example.

Example 11.21. The jbosscmp-jdbc.xml none Declaration

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!-- ... -->
      <query>
        <query-method>
          <method-name>findAll_none</method-name>
          <method-params/>
        </query-method>
        <jboss-ql><![CDATA[
          SELECT OBJECT(g)
          FROM gangster g
          ORDER BY g.gangsterId
        ]]></jboss-ql>
        <read-ahead>
```

```
        <strategy>none</strategy>
      </read-ahead>
    </query>
  </entity>
</enterprise-beans>
</jboss-cmp-jdbc>
```

11.8. Loading Process

In the previous section several steps use the phrase "when the entity is loaded." This was intentionally left vague because the commit option specified for the entity and the current state of the transaction determine when an entity is loaded. The following section describes the commit options and the loading processes.

11.8.1. Commit Options

Central to the loading process are the commit options, which control when the data for an entity expires. JBoss supports four commit options A, B, C and D. The first three are described in the Enterprise JavaBeans Specification, but the last one is specific to JBoss. A detailed description of each commit option follows:

- **A:** JBossCMP assumes it is the sole user of the database; therefore, JBossCMP can cache the current value of an entity between transactions, which can result in substantial performance gains. As a result of this assumption, no data managed by JBossCMP can be changed outside of JBossCMP. For example, changing data in another program or with the use of direct JDBC (even within JBoss) will result in an inconsistent database state.
- **B:** JBossCMP assumes that there is more than one user of the database but keeps the context information about entities between transactions. This context information is used for optimizing loading of the entity. This is the default commit option.
- **C:** JBossCMP discards all entity context information at the end of the transaction.
- **D:** This is a JBoss specific commit option. This option is similar to commit option A, except that the data only remains valid for a specified amount of time.

The commit option is declared in the `jboss.xml` file. For a detailed description of this file see Chapter 5. The following example changes the commit option to A for all entity beans in the application:

Example 11.22. The `jboss.xml` Commit Option Declaration

```
<jboss>
  <container-configurations>
    <container-configuration>
      <container-name>Standard CMP 2.x EntityBean</container-name>
      <commit-option>A</commit-option>
    </container-configuration>
  </container-configurations>
</jboss>
```

11.8.2. Eager-loading Process

One of the most important changes in CMP 2.0 is the change from using class fields for CMP fields to abstract accessor methods. In CMP 1.x, the container could not know which fields were required in a transaction, so the container had to eager load every field when loading the bean¹⁶. In CMP 2.x, the container creates the implementation for the abstract accessors, so the container can know when the data for a field is required. JBossCMP can be configured to eager load only some of the fields when loading an entity, and later lazy load the remaining fields as needed.

When an entity is loaded, JBossCMP must determine the fields that need to be loaded. By default, JBossCMP will use the `eager-load-group` of the last query that selected this entity. If the entity has not been selected in a query, or the last query used the `none` read-ahead strategy, JBossCMP will use the default `eager-load-group` declared for the entity. In the following example configuration, the `basic` load group is set as the default `eager-load-group` for the `GangsterEJB` entity:

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!-- ... -->
      <load-groups>
        <load-group>
          <load-group-name>most</load-group-name>
          <field-name>name</field-name>
          <field-name>nickName</field-name>
          <field-name>badness</field-name>
          <field-name>hangout</field-name>
          <field-name>organization</field-name>
        </load-group>
      </load-groups>
      <eager-load-group>most</eager-load-group>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

The eager loading process is initiated the first time a method is called on an entity in a transaction. A detailed description of the load process follows:

- If the entity context is still valid, no loading is necessary, and therefore the loading process is done. The entity context will be valid when using commit option `A`, or when using commit option `D`, and the data has not timed out.
- Any residual data in the entity context is flushed. This assures that old data does not bleed into the new load.
- The primary key value is injected back into the primary key fields. The primary key object is actually independent of the fields and needs to be reloaded after the flush in step 2.
- All data in the preload cache for this entity is loaded into the fields.
- JBossCMP determines the additional fields that still need to be loaded. Normally the fields to load are determined by the `eager-load-group` of the entity, but can be overridden if the entity was located using a query or CMR field with an `on-find` or `on-load` read ahead strategy. If all of the fields have already been loaded, the load process skips to step 7.
- A query is executed to select the necessary column. If this entity is using the `on-load` strategy, a page of data is loaded as described in Section 11.7.3.2. The data for the current entity is stored in the context and

¹⁶In a future version, JBossCMP will be able to keep the current data of a commit option `B` entity between transactions and validate that the data is still current using `last-update` optimistic locking. For entities that contain a large amount of data, this will result in a significant performance enhancement.

the data for the other entities is stored in the preload cache.

- The `ejbLoad` method of the entity is called.

11.8.3. Lazy loading Process

Lazy loading is the other half of eager loading. If a field is not eager loaded, it must be lazy loaded. When the bean accesses an unloaded field, JBossCMP loads the field and any field in a `lazy-load-group` of which the unloaded field is a member. JBossCMP performs a set join and then removes any field that is already loaded. An example configuration is shown below.

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>GangsterEJB</ejb-name>
      <!-- ... -->
      <load-groups>
        <load-group>
          <load-group-name>basic</load-group-name>
          <field-name>name</field-name>
          <field-name>nickName</field-name>
          <field-name>badness</field-name>
        </load-group>
        <load-group>
          <load-group-name>contact info</load-group-name>
          <field-name>nickName</field-name>
          <field-name>contactInfo</field-name>
          <field-name>hangout</field-name>
        </load-group>
      </load-groups>
      <!-- ... -->
      <lazy-load-groups>
        <load-group-name>basic</load-group-name>
        <load-group-name>contact info</load-group-name>
      </lazy-load-groups>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

When the bean provider calls `getName()` with this configuration, JBossCMP loads `name`, `nickName` and `badness` (assuming they are not already loaded). When the bean provider calls `getNickName()`, the `name`, `nickName`, `badness`, `contactInfo`, and `hangout` are loaded. A detailed description of the lazy loading process follows:

- All data in the preload cache for this entity is loaded into the fields.
- If the field value was loaded by the preload cache the lazy load process is finished.
- JBossCMP finds all of the lazy load groups that contain this field, performs a set join on the groups, and removes any field that has already been loaded.
- A query is executed to select the necessary columns. As in the basic load process, JBossCMP may load a block of entities. The data for the current entity is stored in the context and the data for the other entities is stored in the preload cache.

11.8.3.1. Relationships

Relationships are a special case in lazy loading because a CMR field is both a field and query. As a field it can be on-load block loaded, meaning the value of the currently sought entity and the values of the CMR field for the next several entities are loaded. As a query, the field values of the related entity can be preloaded using on-

find.

Again, the easiest way to investigate the loading is to look at a usage scenario. In this example, an HTML table is generated containing each gangster and their hangout. The example code follows:

Example 11.23. Relationship Lazy Loading Example Code

```
public String createGangsterHangoutHtmlTable()
    throws FinderException
{
    StringBuffer table = new StringBuffer();
    table.append("<table>");
    Collection gangsters = gangsterHome.findAll_onfind();
    for (Iterator iter = gangsters.iterator(); iter.hasNext(); ) {
        Gangster gangster = (Gangster)iter.next();

        Location hangout = gangster.getHangout();
        table.append("<tr>");
        table.append("<td>").append(gangster.getName());
        table.append("</td>");
        table.append("<td>").append(gangster.getNickName());
        table.append("</td>");
        table.append("<td>").append(gangster.getBadness());
        table.append("</td>");
        table.append("<td>").append(hangout.getCity());
        table.append("</td>");
        table.append("<td>").append(hangout.getState());
        table.append("</td>");
        table.append("<td>").append(hangout.getZipCode());
        table.append("</td>");
        table.append("</tr>");
    }

    table.append("</table>");return table.toString();
}
```

For this example, the configuration of the Gangster findAll_onfind query is unchanged from the on-find section. The configuration of the Location entity and Gangster-Hangout relationship follows:

Example 11.24. The jbosscmp-jdbc.xml Relationship Lazy Loading Configuration

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>LocationEJB</ejb-name>
      <load-groups>
        <load-group>
          <load-group-name>quick info</load-group-name>
          <field-name>city</field-name>
          <field-name>state</field-name>
          <field-name>zipCode</field-name>
        </load-group>
      </load-groups>
      <eager-load-group/>
    </entity>
  </enterprise-beans>
  <relationships>
    <ejb-relation>
      <ejb-relation-name>Gangster-Hangout</ejb-relation-name>
      <foreign-key-mapping/>
      <ejb-relationship-role>
        <ejb-relationship-role-name>
          gangster-has-a-hangout
        </ejb-relationship-role-name>
      </ejb-relationship-role>
    </ejb-relation>
  </relationships>
</jbosscmp-jdbc>
```

```

        <key-fields/>
        <read-ahead>
            <strategy>on-find</strategy>
            <page-size>4</page-size>
            <eager-load-group>quick info</eager-load-group>
        </read-ahead>
    </ejb-relationship-role>
</ejb-relationship-role>
    <ejb-relationship-role-name>
        hangout-for-a-gangster
    </ejb-relationship-role-name>
    <key-fields>
        <key-field>
            <field-name>locationID</field-name>
            <column-name>hangout</column-name>
        </key-field>
    </key-fields>
</ejb-relationship-role>
</ejb-relation>
</relationships>
</jbosscmp-jdbc>

```

At line 25, JBossCMP will execute the following query:

```

SELECT t0_g.id, t0_g.name, t0_g.nick_name, t0_g.badness
FROM gangster t0_g
ORDER BY t0_g.id ASC

```

Then at line 29, JBossCMP executes the following two queries to load the city, state, and zip fields of the hideout:

```

SELECT gangster.id, gangster.hangout,
    location.city, location.st, location.zip
FROM gangster, location
WHERE (gangster.hangout=location.id) AND
    ((gangster.id=0) OR (gangster.id=1) OR
    (gangster.id=2) OR (gangster.id=3))
SELECT gangster.id, gangster.hangout,
    location.city, location.st, location.zip
FROM gangster, location
WHERE (gangster.hangout=location.id) AND
    ((gangster.id=4) OR (gangster.id=5) OR
    (gangster.id=6) OR (gangster.id=7))

```

The following table shows the execution of the queries:

id	name	nick_name	badness	hangout	id	city	st	zip
0	Yojimbo	Bodyguard	7	0	0	San Fran	CA	94108
1	Takeshi	Master	10	1	1	San Fran	CA	94133
2	Yuriko	Four finger	4	2	2	San Fran	CA	94133
3	Chow	Killer	9	3	3	San Fran	CA	94133
4	Shogi	Lightning	8	4	4	San Fran	CA	94133
5	Valentino	Pizza-Face	4	5	5	New York	NY	10017
6	Toni	Toothless	2	6	6	Chicago	IL	60661
7	Corleone	Godfather	6	7	7	Las Vegas	NV	89109

11.8.4. Lazy loading result sets

By default, when a multiobject finder or `ejbSelect` method is executed the `ResultSet` is read to the end immediately. The client receives a collection of `EJBLocalObject` or CMP field values which it can then iterate through. For big result sets this approach is not efficient. In some cases it is better to delay reading the next row in the `ResultSet` until the client tries to read the corresponding value from the collection. You can get this behaviour for a query using the `lazy-resultset-loading` element.

```
<query>
  <query-method>
    <method-name>findAll</method-name>
  </query-method>
  <jboss-ql><![CDATA[select object(o) from A o]]></jboss-ql>
  <lazy-resultset-loading>true</lazy-resultset-loading>
</query>
```

There are some issues you should be aware of when using lazy result set loading. Special care should be taken when working with a `Collection` associated with a lazily loaded result set. The first call to `iterator()` returns a special `Iterator` that reads from the `ResultSet`. Until this `Iterator` has been exhausted, subsequent calls to `iterator()` or calls to the `add()` method will result in an exception. The `remove()` and `size()` methods work as would be expected.

11.9. Transactions

All of the examples presented in this chapter have been defined to run in a transaction. Transaction granularity is a dominating factor in optimized loading because transactions define the lifetime of preloaded data. If the transaction completes, commits, or rolls back, the data in the preload cache is lost. This can result in a severe negative performance impact.

The performance impact of running without a transaction will be demonstrated with an example similar to Example 11.15. This example uses an `on-find` optimized query that selects the first four gangsters (to keep the result set small), and it is executed without a wrapper transaction. The example code follows:

Example 11.25. No Transaction Loading Example Code

```
public String createGangsterHtmlTable_no_tx() throws FinderException
{
    StringBuffer table = new StringBuffer();
    table.append("<table>");

    Collection gangsters = gangsterHome.findFour();
    for(Iterator iter = gangsters.iterator(); iter.hasNext(); ) {
        Gangster gangster = (Gangster)iter.next();
        table.append("<tr>");
        table.append("<td>").append(gangster.getName());
        table.append("</td>");
        table.append("<td>").append(gangster.getNickName());
        table.append("</td>");
        table.append("<td>").append(gangster.getBadness());
        table.append("</td>");
        table.append("</tr>");
    }

    table.append("</table>");
    return table.toString();
}
```

The following is the query executed at line 53.

Example 11.26. No Transaction on-find Optimized findAll Query

```
SELECT t0_g.id, t0_g.name, t0_g.nick_name, t0_g.badness
FROM gangster t0_g
WHERE t0_g.id < 4followi
ORDER BY t0_g.id ASC
```

Normally this would be the only query executed, but since this code is not running in a transaction, all of the preloaded data is thrown away as soon as `findAll` returns. Then at line 56 JBossCMP executes the following four queries¹⁷ (one for each loop):

Example 11.27. No Transaction on-load Optimized Load Queries

```
SELECT id, name, nick_name, badness
FROM gangster
WHERE (id=0) OR (id=1) OR (id=2) OR (id=3)
SELECT id, name, nick_name, badness
FROM gangster
WHERE (id=1) OR (id=2) OR (id=3)
SELECT id, name, nick_name, badness
FROM gangster
WHERE (id=2) OR (id=3)
SELECT name, nick_name, badness
FROM gangster
WHERE (id=3)
```

The following figure shows the execution of the queries:

id§	name§	nick name§	badness§
0§	Yojimbo	Bodyguard	7
1§	Takeshi	Master	10
2§	Yuriko	Four finger	4
3§	Chow	Killer	9

Figure 11.11. No Transaction on-find optimized query execution

This performance is much worse than read ahead none because of the amount of data loaded from the database. The number of rows loaded is determined by the following equation:

$$n + n - 1 + n - 2 + \dots + 1 + = \frac{n * (n + 1)}{2} = O(n^2)$$

¹⁷It's actually worse than this. JBossCMP executes each of these queries three times; once for each CMP field that is accessed. This is because the preloaded values are discarded between the CMP field accessor calls.

This all happens because the transaction in the example is bounded by a single call on the entity. This brings up the important question "How do I run my code in a transaction?" The answer depends on where the code runs. If it runs in an EJB (session, entity, or message driven), the method must be marked with the `Required` or `RequiresNew` trans-attribute in the assembly-descriptor. If the code is not running in an EJB, a user transaction is necessary. The following code wraps a call to the declared method with a user transaction:

Example 11.28. User Transaction Example Code

```
public String createGangsterHtmlTable_with_tx()
    throws FinderException
{
    UserTransaction tx = null;
    try {
        InitialContext ctx = new InitialContext();
        tx = (UserTransaction) ctx.lookup("UserTransaction");
        tx.begin();

        String table = createGangsterHtmlTable_no_tx();

        if (tx.getStatus() == Status.STATUS_ACTIVE) {
            tx.commit();
        }
        return table;
    } catch (Exception e) {
        try {
            if (tx != null) tx.rollback();
        } catch (SystemException unused) {
            // eat the exception we are exceptioning out anyway
        }
        if (e instanceof FinderException) {
            throw (FinderException) e;
        }
        if (e instanceof RuntimeException) {
            throw (RuntimeException) e;
        }
        throw new EJBException(e);
    }
}
```

11.10. Optimistic Locking

JBoss has supports for optimistic locking of entity beans. Optimistic locking allows multiple instances of the same entity bean to be active simultaneously. Consistency is enforced based on the optimistic locking policy choice. The optimistic locking policy choice defines the set of fields that are used in the commit time write of modified data to the database. The optimistic consistency check asserts that the values of the chosen set of fields has the same values in the database as existed when the current transaction was started. This is done using a `select for UPDATE WHERE ...` statement that contains the value assertions.

You specify the optimistic locking policy choice using an `entity/optimistic-locking` element in the `jbosscmp-jdbc.xml` descriptor. The content model of the `optimistic-locking` element is shown below and the description of the elements follows.

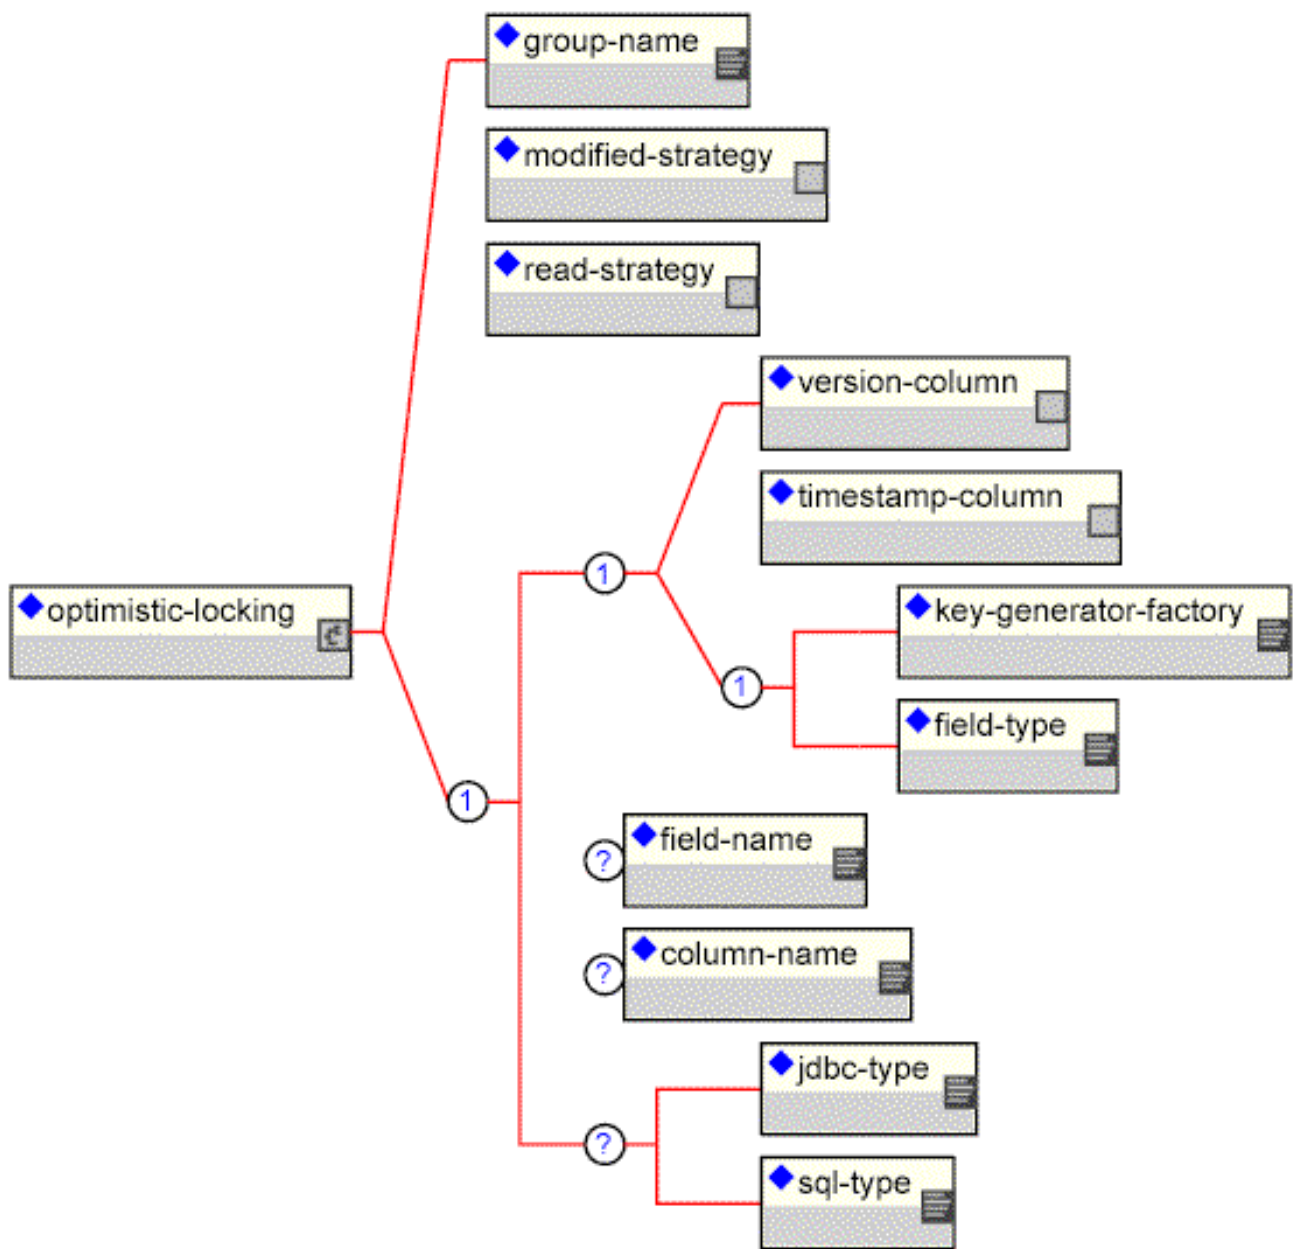


Figure 11.12. The jbossCMP-jdbc optimistic-locking element content model

- **group-name:** This element specifies that optimistic locking is based on the fields of a `load-group`. This value of this element must match one of the entity's `load-group-name`. The fields in this group will be used for optimistic locking.
- **modified-strategy:** This element specifies that optimistic locking is based on the modified fields. This strategy implies that the fields that were modified during transaction will be used for optimistic locking.
- **read-strategy:** This element specifies that optimistic locking is based on the fields read. This strategy implies that the fields that were read/changed in the transaction will be used for optimistic locking.
- **version-column:** This element specifies that optimistic locking is based on a version column strategy. Specifying this element will add an additional version field of type `java.lang.Long` to the entity bean for optimistic locking. Each update of the entity will increase the value of this field. The `field-name` element allows for the specification of the name of the CMP field while the `column-name` element allows for the spe-

cification of the corresponding table column.

- **timestamp-column:** This element specifies that optimistic locking is based on a timestamp column strategy. Specifying this element will add an additional version field of type `java.util.Date` to the entity bean for optimistic locking. Each update of the entity will set the value of this field to the current time. The `field-name` element allows for the specification of the name of the CMP field while the `column-name` element allows for the specification of the corresponding table column.
- **key-generator-factory:** This element specifies that optimistic locking is based on key generation. The value of the element is the JNDI name of a `org.jboss.ejb.plugins.keygenerator.KeyGeneratorFactory` implementation. Specifying this element will add an additional version field to the entity bean for optimistic locking. The type of the field must be specified via the `field-type` element. Each update of the entity will update the key field by obtaining a new value from the key generator. The `field-name` element allows for the specification of the name of the CMP field while the `column-name` element allows for the specification of the corresponding table column.

A sample `jbosscmp-jdbc.xml` descriptor illustrating all of the optimistic locking strategies is given below.

Example 11.29. A sample `jbosscmp-jdbc.xml` descriptor illustrating the optimistic locking strategies

```
<!DOCTYPE jbosscmp-jdbc PUBLIC
    "-//JBoss//DTD JBOSSCMP-JDBC 3.2//EN"
    "http://www.jboss.org/j2ee/dtd/jbosscmp-jdbc_3_2.dtd">
<jbosscmp-jdbc>
  <defaults>
    <datasource>java:/DefaultDS</datasource>
    <datasource-mapping>Hypersonic SQL</datasource-mapping>
  </defaults>
  <enterprise-beans>
    <entity>
      <ejb-name>EntityGroupLocking</ejb-name>
      <create-table>true</create-table>
      <remove-table>true</remove-table>
      <table-name>entitygrouplocking</table-name>
      <cmp-field>
        <field-name>dateField</field-name>
      </cmp-field>
      <cmp-field>
        <field-name>integerField</field-name>
      </cmp-field>
      <cmp-field>
        <field-name>stringField</field-name>
      </cmp-field>
      <load-groups>
        <load-group>
          <load-group-name>string</load-group-name>
          <field-name>stringField</field-name>
        </load-group>
        <load-group>
          <load-group-name>all</load-group-name>
          <field-name>stringField</field-name>
          <field-name>dateField</field-name>
        </load-group>
      </load-groups>
      <optimistic-locking>
        <group-name>string</group-name>
      </optimistic-locking>
    </entity>
    <entity>
      <ejb-name>EntityModifiedLocking</ejb-name>
      <create-table>true</create-table>
      <remove-table>true</remove-table>
      <table-name>entitymodifiedlocking</table-name>
```

```

    <cmp-field>
      <field-name>dateField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>integerField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>stringField</field-name>
    </cmp-field>
    <optimistic-locking>
      <modified-strategy/>
    </optimistic-locking>
  </entity>
  <entity>
    <ejb-name>EntityReadLocking</ejb-name>
    <create-table>true</create-table>
    <remove-table>true</remove-table>
    <table-name>entityreadlocking</table-name>
    <cmp-field>
      <field-name>dateField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>integerField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>stringField</field-name>
    </cmp-field>
    <optimistic-locking>
      <read-strategy/>
    </optimistic-locking>
  </entity>
  <entity>
    <ejb-name>EntityVersionLocking</ejb-name>
    <create-table>true</create-table>
    <remove-table>true</remove-table>
    <table-name>entityversionlocking</table-name>
    <cmp-field>
      <field-name>dateField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>integerField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>stringField</field-name>
    </cmp-field>
    <optimistic-locking>
      <version-column/>
      <field-name>versionField</field-name>
      <column-name>ol_version</column-name>
      <jdbc-type>INTEGER</jdbc-type>
      <sql-type>INTEGER(5)</sql-type>
    </optimistic-locking>
  </entity>
  <entity>
    <ejb-name>EntityTimestampLocking</ejb-name>
    <create-table>true</create-table>
    <remove-table>true</remove-table>
    <table-name>entitytimestamplocking</table-name>
    <cmp-field>
      <field-name>dateField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>integerField</field-name>
    </cmp-field>
    <cmp-field>
      <field-name>stringField</field-name>
    </cmp-field>
    <optimistic-locking>
      <timestamp-column/>
      <field-name>versionField</field-name>
      <column-name>ol_timestamp</column-name>

```




Figure 11.13. The jbosscmp-jdbc.xml entity-command element model

vendor specific commands typically subclass the `org.jboss.ejb.plugins.cmp.jdbc.JDBCIdentityColumnCreateCommand` if the database generates the primary key as a side effect of doing an insert, or the `org.jboss.ejb.plugins.cmp.jdbc.JDBCInsertPKCreateCommand` if the command must insert the generated key.

- **entity-command/attribute:** The optional `attribute` element(s) allows for the specification of arbitrary name/value property pairs that will be available to the entity command implementation class. The `attribute` element has a required `name` attribute that specifies the name property, and the `attribute` element content is the value of the property. The attribute values are accessible through the `org.jboss.ejb.plugins.cmp.jdbc.metadata.JDBCEntityCommandMetaData.getAttribute(String)` method.

11.11.1. Existing Entity Commands

The following are the current `entity-command` definitions found in the `standardjbosscmp-jdbc.xml` descriptor:

- **default** (`org.jboss.ejb.plugins.cmp.jdbc.JDBCCreateEntityCommand`) The `JDBCCreateEntityCommand` is the default entity creation as it is the `entity-command` referenced in the `standardjbosscmp-jdbc.xml` defaults element. This entity-command executes an `INSERT INTO` query using the assigned primary key value.
- **no-select-before-insert:** (`org.jboss.ejb.plugins.cmp.jdbc.JDBCCreateEntityCommand`) This is a variation on default that skips select before insert by specifying an attribute `name="SQLExceptionProcessor"` that points to the `jbosscmp-jdbc:service=SQLExceptionProcessor` service. The `SQLExceptionProcessor` service provides a boolean `isDuplicateKey(SQLException e)` operation that allows for determination of any unique constraint violation.
- **pk-sql** (`org.jboss.ejb.plugins.cmp.jdbc.JDBCPkSqlCreateCommand`) The `JDBCPkSqlCreateCommand` executes an `INSERT INTO` query statement provided by the `pk-sql` attribute to obtain the next primary key value. Its primary target usage are databases with sequence support.
- **mysql-get-generated-keys:** (`org.jboss.ejb.plugins.cmp.jdbc.mysql.JDBCMySQLCreateCommand`) The `JDBCMySQLCreateCommand` executes an `INSERT INTO` query using the `getGeneratedKeys` method from MySQL native `java.sql.Statement` interface implementation to fetch the generated key.
- **oracle-sequence:** (`org.jboss.ejb.plugins.cmp.jdbc.keygen.JDBCOracleCreateCommand`) The `JDBCOracleCreateCommand` is a create command for use with Oracle that uses a sequence in conjunction with a `RETURNING` clause to generate keys in a single statement. It has a required `sequence` element that specifies the name of the sequence column.
- **hsqldb-fetch-key:** (`org.jboss.ejb.plugins.cmp.jdbc.hsqldb.JDBCHsqldbCreateCommand`) The `JDBCHsqldbCreateCommand` executes an `INSERT INTO` query after executing a `CALL IDENTITY()` statement to fetch the generated key.
- **sybase-fetch-key:** (`org.jboss.ejb.plugins.cmp.jdbc.sybase.JDBCSybaseCreateCommand`) The `JDBCSybaseCreateCommand` executes an `INSERT INTO` query after executing a `SELECT @@IDENTITY` statement to fetch the generated key.
- **mssql-fetch-key:** (`org.jboss.ejb.plugins.cmp.jdbc.keygen.JDBCSQLServerCreateCommand`) The `JDBCSQLServerCreateCommand` for Microsoft SQL Server that uses the value from an `IDENTITY` columns. By default uses `SELECT SCOPE_IDENTITY()` to reduce the impact of triggers; can be overridden with `pk-sql` attribute e.g. for V7.

- **informix-serial:** (org.jboss.ejb.plugins.cmp.jdbc.informix.JDBCInformixCreateCommand) The JDBCInformixCreateCommand executes an INSERT INTO query after using the getSerial method from Informix native java.sql.Statement interface implementation to fetch the generated key.
- **postgresql-fetch-seq** (org.jboss.ejb.plugins.cmp.jdbc.keygen.JDBCPostgreSQLCreateCommand) The JDBCPostgreSQLCreateCommand for PostgreSQL that fetches the currval of the sequence. The optional sequence attribute can be used to change the name of the sequence, with the default being table_pkColumn_seq.
- **key-generator:** (org.jboss.ejb.plugins.cmp.jdbc.JDBCKeyGeneratorCreateCommand) The JDBCKeyGeneratorCreateCommand executes an INSERT INTO query after obtaining a value for the primary key from the key generator referenced by the key-generator-factory. The key-generator-factory attribute must provide the name of a JNDI binding of the org.jboss.ejb.plugins.keygenerator.KeyGeneratorFactory implementation.
- **get-generated-keys:** (org.jboss.ejb.plugins.cmp.jdbc.jdbc3.JDBCGetGeneratedKeysCreateCommand) The JDBCGetGeneratedKeysCreateCommand executes an INSERT INTO query using a statement built using the JDBC3 prepareStatement(String, Statement.RETURN_GENERATED_KEYS) that has the capability to retrieve the auto-generated key. The generated key is obtained by calling the PreparedStatement.getGeneratedKeys method. Since this requires JDBC3 support it is only available in JDK1.4.1+ with a supporting JDBC driver.

An example configuration using the hsqldb-fetch-key entity-command with the generated key mapped to a known primary key cmp-field is shown below.

Example 11.30. A sample autogenerated key config for a known pk cmp-field

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>LocationEJB</ejb-name>
      <pk-constraint>false</pk-constraint>
      <table-name>location</table-name>

      <cmp-field>
        <field-name>locationID</field-name>
        <column-name>id</column-name>
        <auto-increment/>
      </cmp-field>
      <!-- ... -->
      <entity-command name="hsqldb-fetch-key"/>

    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

An alternate example using an unknown primary key without an explicit cmp-field is shown below.

```
<jbosscmp-jdbc>
  <enterprise-beans>
    <entity>
      <ejb-name>LocationEJB</ejb-name>
      <pk-constraint>false</pk-constraint>
      <table-name>location</table-name>
      <unknown-pk>
        <unknown-pk-class>java.lang.Integer</unknown-pk-class>
        <field-name>locationID</field-name>
        <column-name>id</column-name>
        <jdbc-type>INTEGER</jdbc-type>
      </unknown-pk>
    </entity>
  </enterprise-beans>
</jbosscmp-jdbc>
```

```
        <sql-type>INTEGER</sql-type>
        <auto-increment/>
    </unknown-pk>
    <!--...-->
    <entity-command name="hsqldb-fetch-key"/>
</entity>
</enterprise-beans>
</jbosscomp-jdbc>
```

11.12. Defaults

JBossCMP global defaults are defined in the `standardjbosscomp-jdbc.xml` file of the `server/<server-name>/conf/` directory. Each application can override the global defaults in the `jbosscomp-jdbc.xml` file. The default options are contained in a `defaults` element of the configuration file, and the content model is shown below.

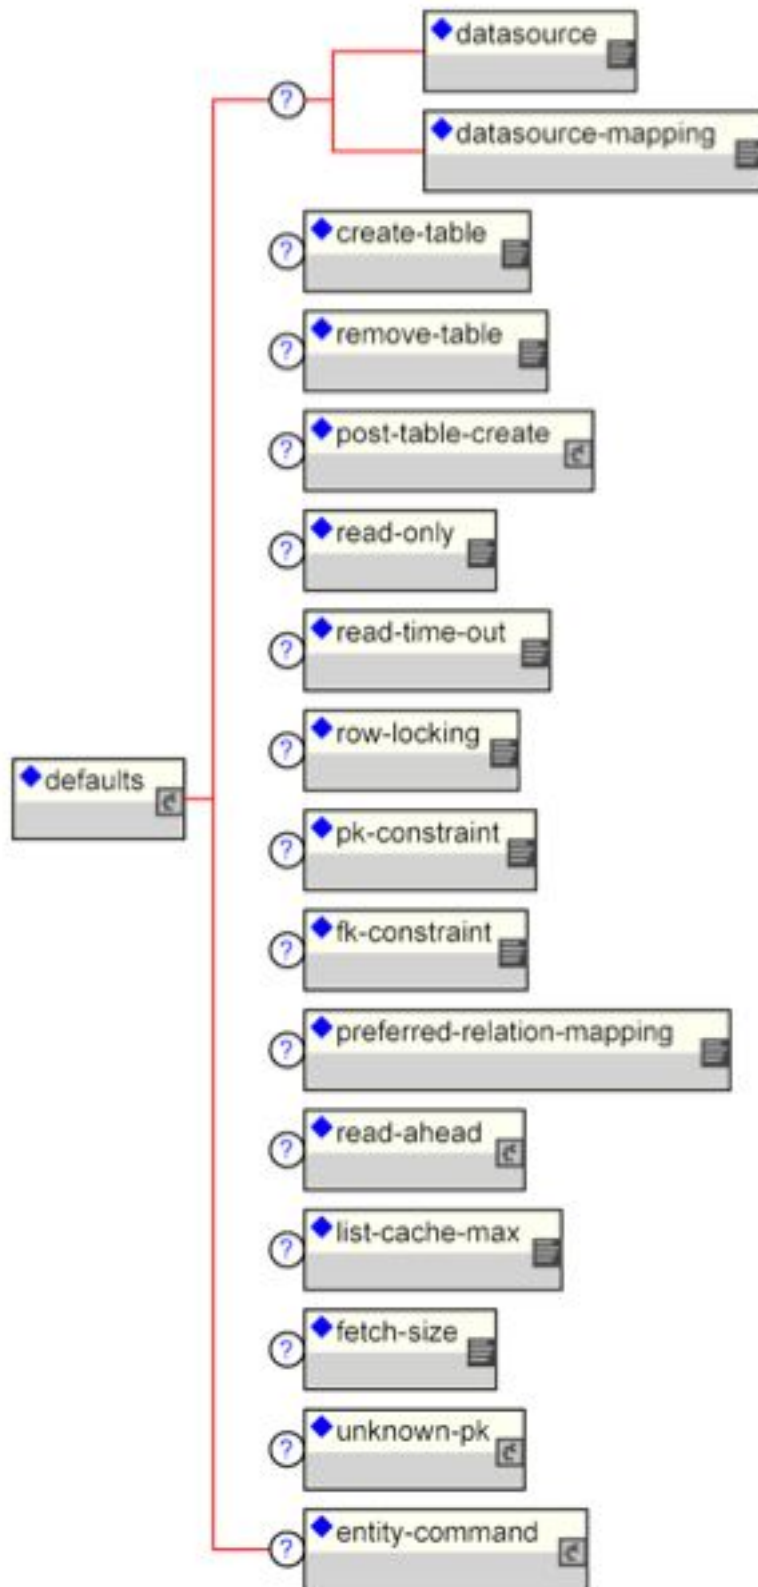


Figure 11.14. The `jbosscmp-jdbc/defaults` content model

An example of the defaults section follows:

```
<jbosscmp-jdbc>
  <defaults>
    <datasource>java:/DefaultDS</datasource>
    <datasource-mapping>Hypersonic SQL</datasource-mapping>
```

```

<create-table>true</create-table>
<remove-table>false</remove-table>
<read-only>false</read-only>
<read-time-out>300000</read-time-out>
<pk-constraint>true</pk-constraint>
<fk-constraint>false</fk-constraint>
<row-locking>false</row-locking>
<preferred-relation-mapping>foreign-key</preferred-relation-mapping>
<read-ahead>
  <strategy>on-load</strategy>
  <page-size>1000</page-size>
  <eager-load-group>*</eager-load-group>
</read-ahead>
<list-cache-max>1000</list-cache-max>
</defaults>
</jbosscmp-jdbc>

```

11.12.1. A sample jbosscmp-jdbc.xml defaults declaration

Each option can apply to entities, relationships, or both, and can be overridden in the specific entity or relationship. A detailed description of each option follows:

- **datasource:** This optional element is the `jndi-name` used to look up the datasource. All database connections used by an entity or `relation-table` are obtained from the datasource. Having different datasources for entities is not recommended, as it vastly constrains the domain over which finders and `ejbSelects` can query.
- **datasource-mapping:** This optional element specifies the name of the `type-mapping`, which determines how Java types are mapped to SQL types, and how EJB-QL functions are mapped to database specific functions. Type mappings are discussed in Section 11.13.2.
- **create-table:** This optional element when true, specifies that JBossCMP should attempt to create a table for the entity. When the application is deployed, JBossCMP checks if a table already exists before creating the table. If a table is found, it is logged, and the table is not created. This option is very useful during the early stages of development when the table structure changes often. The default is false.
- **alter-table:** If `create-table` is used to automatically create the schema, `alter-table` can be used to keep the schema current with changes to the entity bean. Alter table will perform the following specific tasks:
 - new fields will be created
 - fields which are no longer used will be removed
 - string fields which are shorter than the declared length will have their length increased to the declared length. (not supported by all databases)
- **remove-table:** This optional element when true, JBossCMP will attempt to drop the table for each entity and each relation table mapped relationship. When the application is undeployed, JBossCMP will attempt to drop the table. This option is very useful during the early stages of development when the table structure changes often. The default is false.
- **read-only:** This optional element when true specifies that the bean provider will not be allowed to change the value of any fields. A field that is read-only will not be stored in, or inserted into, the database. If a primary key field is read-only, the create method will throw a `CreateException`. If a set accessor is called on a read-only field, it throws an `EJBException`. Read only fields are useful for fields that are filled in by database triggers, such as last update. The `read-only` option can be overridden on a per field basis. The default is false.

- **read-time-out:** This optional element is the amount of time in milliseconds that a read on a read only field is valid. A value of 0 means that the value is always reloaded at the start of a transaction, and a value of -1 means that the value never times out. This option can also be overridden on a per CMP field basis. If `read-only` is false, this value is ignored. The default is -1.
- **row-locking:** This optional element if true specifies that JBossCMP will lock all rows loaded in a transaction. Most databases implement this by using the `SELECT FOR UPDATE` syntax when loading the entity, but the actual syntax is determined by the `row-locking-template` in the `datasource-mapping` used by this entity. The default is false.
- **pk-constraint:** This optional element if true specifies that JBossCMP will add a primary key constraint when creating tables. The default is true.
- **preferred-relation-mapping:** This optional element specifies the preferred mapping style for relationships. The `preferred-relation-mapping` element must be either `foreign-key` or `relation-table`.
- **read-ahead:** This optional element controls caching of query results and CMR fields for the entity. This option is discussed in Section 11.7.3.
- **list-cache-max:** This optional element specifies the number of `read-lists` that can be tracked by this entity. This option is discussed in Section 11.7.3.2. The default is 1000.
- **clean-read-ahead-on-load:** When an entity is loaded from the read ahead cache, JBoss can remove the data used from the read ahead cache. The default is false.
- **fetch-size:** This optional element specifies the number of entities to read in one round-trip to the underlying datastore. The default is 0.
- **unknown-pk:** This optional element allows one to define the default mapping of an unknown primary key type of `java.lang.Object` maps to the persistent store.
- **entity-command:** This optional element allows one to define the default command for entity creation. This is described in detail in Section 11.11.

11.13. Datasource Customization

JBossCMP includes predefined type-mappings for many databases including: Cloudscape, DB2, DB2/400, Hypersonic SQL, InformixDB, InterBase, MS SQLSERVER, MS SQLSERVER2000, MySQL, Oracle7, Oracle8, Oracle9i, PointBase, PostgreSQL, PostgreSQL 7.2, SapDB, SOLID, and Sybase. If you do not like the supplied mapping, or a mapping is not supplied for your database, you will have to define a new mapping. If you find an error in one of the supplied mappings, or if you create a new mapping for a new database, please consider posting a patch at the JBoss project page on SourceForge.

Customization of a database is done through the `type-mapping` section of the `jbosscmp-jdbc.xml` descriptor. The content model for the type-mapping element is given in Figure 11.15. The elements are:

- **name:** This required element provides the name identifying the database customization. It is used to refer to the mapping by the `datasource-mapping` elements found in defaults and entity.
- **row-locking-template:** This required element gives the `PreparedStatement` template used to create a row lock on the selected rows. The template must support three arguments:

1. the select clause
2. the from clause. The order of the tables is currently not guaranteed
3. the where clause

If row locking is not supported in select statement this element should be empty. The most common form of row locking is select for update as in: `SELECT ?1 FROM ?2 WHERE ?3 FOR UPDATE.`

- **pk-constraint-template:** This required element gives the `PreparedStatement` template used to create a primary key constraint in the create table statement. The template must support two arguments
 1. Primary key constraint name; which is always `pk_{table-name}`
 2. Comma separated list of primary key column names

If a primary key constraint clause is not supported in a create table statement this element should be empty. The most common form of a primary key constraint is: `CONSTRAINT ?1 PRIMARY KEY (?2)`

- **fk-constraint-template:** This is the template used to create a foreign key constraint in separate statement. The template must support five arguments:
 1. Table name
 2. Foreign key constraint name; which is always `fk_{table-name}_{cmr-field-name}`
 3. Comma separated list of foreign key column names
 4. References table name
 5. Comma separated list of the referenced primary key column names

If the datasource does not support foreign key constraints this element should be empty. The most common form of a foreign key constraint is: `ALTER TABLE ?1 ADD CONSTRAINT ?2 FOREIGN KEY (?3) REFERENCES ?4 (?5).`

- **auto-increment-template:** This declares the SQL template for specifying auto increment columns.
- **add-column-template:** When `alter-table` is true, this SQL template specifies the syntax for adding a column to an existing table. The default value is `ALTER TABLE ?1 ADD ?2 ?3`. The parameters are:
 1. the table name
 2. the column name
 3. the column type
- **drop-column-template:** When `alter-table` is true, this SQL template specifies the syntax for dropping a column to from an existing table. The default value is `ALTER TABLE ?1 DROP ?2`. The parameters are:
 1. the table name
 2. the column name
- **alter-column-template:** When `alter-table` is true, this SQL template specifies the syntax for dropping a column to from an existing table. The default value is `ALTER TABLE ?1 ALTER ?2 TYPE ?3`. The parameters are:
 1. the table name
 2. the column name
 3. the column type
- **alias-header-prefix:** This required element gives the prefix used in creating the alias header. An alias head-

er is prepended to a generated table alias by the EJB-QL compiler to prevent name collisions. The alias header is constructed as follows: `alias-header-prefix` + `int_counter` + `alias-header-suffix`. An example alias header would be `t0_` for an `alias-header-prefix` of `"t"` and an `alias-header-suffix` of `"_"`.

- **alias-header-suffix:** This required element gives the suffix portion of the generated alias header.
- **alias-max-length:** This required element gives the maximum allowed length for the generated alias header.
- **subquery-supported:** This required element specifies if this `type-mapping` subqueries as either `true` or `false`. Some EJB-QL operators are mapped to exists subqueries. If `subquery-supported` is `false`, the EJB-QL compiler will use a left join and `is null`.
- **true-mapping:** This required element defines *true* identity in EJB-QL queries. Examples include `TRUE`, `1`, and `(1=1)`.
- **false-mapping:** This required element defines *false* identity in EJB-QL queries. Examples include `FALSE`, `0`, and `(1=0)`.
- **function-mapping:** This optional element specifies one or more the mappings from an EJB-QL function to an SQL implementation. See Section 11.13.1 for the details.
- **mapping:** This required element specifies the mappings from a Java type to the corresponding JDBC and SQL type. See Section 11.13.2 for the details.

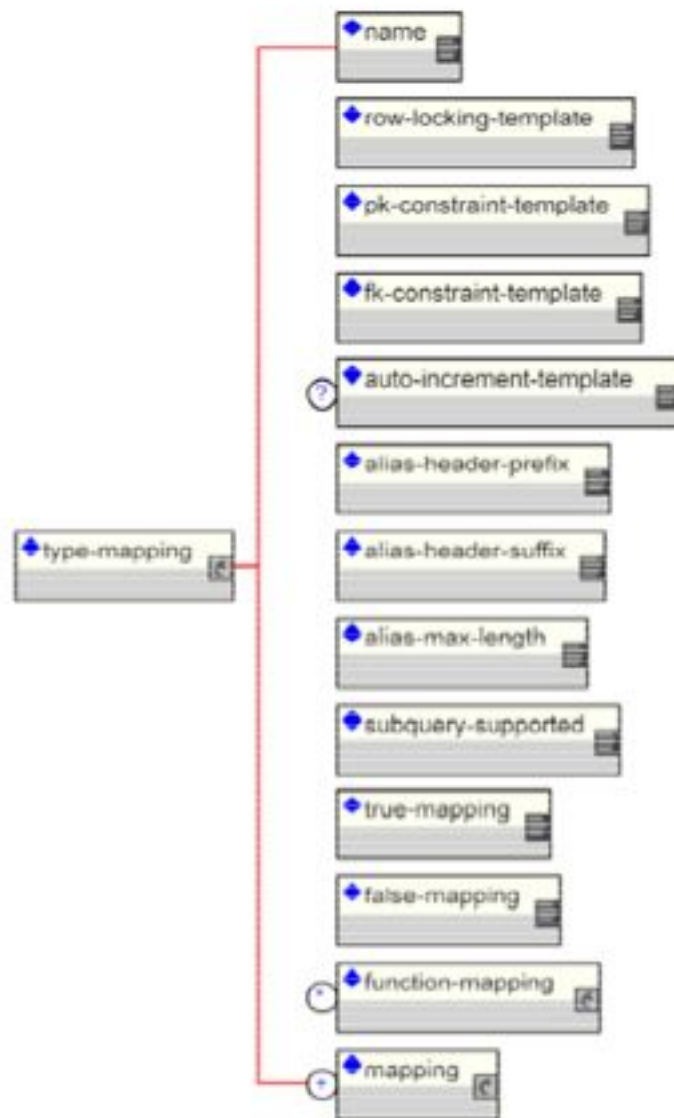


Figure 11.15. The jbosscmp-jdbc type-mapping element content model.

11.13.1. Function Mapping

- **function-name:** This required element gives the EJB-QL function name, e.g., `concat`, `substring`.
- **function-sql:** This required element gives the SQL for the function as appropriate for the underlying database. Examples for a `concat` function include: `(?1 || ?2)`, `concat(?1, ?2)`, `(?1 + ?2)`.

11.13.2. Type Mapping

A `type-mapping` is simply a set of mappings between Java class types and database types. A set of type mappings is defined by a set of mapping elements, the content model for which is shown below.




Figure 11.16. The jboss-cmp-jdbc mapping element content model.

If JBossCMP cannot find a mapping for a type, it will serialize the object and use the `java.lang.Object` mapping. The following describes the three child elements of the mapping element:

- **java-type:** This required element gives the fully qualified name of the Java class to be mapped. If the class is a primitive wrapper class such as `java.lang.Short`, the mapping also applies to the primitive type.
- **jdbc-type:** This required element gives the JDBC type that is used when setting parameters in a JDBC `PreparedStatement` or loading data from a JDBC `ResultSet`. The valid types are defined in `java.sql.Types`.
- **sql-type:** This required element gives the SQL type that is used in create table statements. Valid types are only limited by your database vendor.

An example mapping element for a short in Oracle9i is shown below.

Example 11.31. A sample short mapping for Oracle9i

```

<jboss-cmp-jdbc>
  <type-mappings>
    <type-mapping>
      <name>Oracle9i</name>
      <!-- ... -->
      <mapping>
        <java-type>java.lang.Short</java-type>
        <jdbc-type>NUMERIC</jdbc-type>
        <sql-type>NUMBER(5)</sql-type>
      </mapping>
    </type-mapping>
  </type-mappings>
</jboss-cmp-jdbc>
  
```

11.13.3. User Type Mappings

User type mappings allow one to map from JDBC column types to custom CMP fields types by specifying an instance of `org.jboss.ejb.plugins.cmp.jdbc.Mapper` interface, the definition of which is shown below.

Example 11.32. The `org.jboss.ejb.plugins.cmp.jdbc.Mapper` interface

```

public interface Mapper
{
    /**
     * This method is called when CMP field is stored.
     * @param fieldValue - CMP field value
     * @return column value.
     */
    Object toColumnValue(Object fieldValue);

    /**
     * This method is called when CMP field is loaded.
     * @param columnValue - loaded column value.
     * @return CMP field value.
     */
    Object toFieldValue(Object columnValue);
}

```

A prototypical use case is the mapping of an integer type to its type-safe Java enum instance. The content model of the `user-type-mappings` element consists of one or more `user-type-mapping` elements, the content model of which is shown below.




Figure 11.17. The user-type-mapping content model >

- **java-type:** the fully qualified name of the CMP field type in the mapping.
- **mapped-type:** the fully qualified name of the database type in the mapping.
- **mapper:** the fully qualified name of the `Mapper` interface implementation that handles the conversion between the `java-type` and `mapped-type`.
- **check-dirty-after-get:** This value defaults to false for primitive types and the basic `java.lang` immutable wrappers (`Integer`, `String`, etc...). For potentially mutable objects, JBoss will mark they field as potentially dirty after a `get` operation. If the dirty check on an object is too expensive, you can optimize it away by setting `check-dirty-after-get` to false.
- **state-factory:** This specifies class name of a state factory object which can perform dirty checking for this field. State factory classes must implement the `CMPFieldStateFactory` interface.



The JBoss Group and Our LGPL License

A.1. About The JBoss Group

JBoss Group LLC, is an Atlanta-based professional services company, created by Marc Fleury, founder and lead developer of the JBoss J2EE-based Open Source web application server. JBoss Group brings together core JBoss developers to provide services such as training, support and consulting, as well as management of the JBoss software and services affiliate programs. These commercial activities subsidize the development of the free core JBoss server. For additional information on the JBoss Group see the JBoss site <http://www.jboss.org/services/services.jsp>.

A.2. The GNU Lesser General Public License (LGPL)

The JBoss source code is licensed under the LGPL (see <http://www.gnu.org/copyleft/lesser.txt>). This includes all code in the `org.jboss.*` package namespace. Example A.1 gives the complete text of the LGPL license.

Example A.1. The GNU lesser general public license text

```
GNU LESSER GENERAL PUBLIC LICENSE
Version 2.1, February 1999
```

```
Copyright (C) 1991, 1999 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.
```

```
[This is the first released version of the Lesser GPL. It also counts
as the successor of the GNU Library Public License, version 2, hence
the version number 2.1.]
```

Preamble

```
The licenses for most software are designed to take away your
freedom to share and change it. By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.
```

```
This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it. You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.
```

```
When we speak of free software, we are referring to freedom of use,
not price. Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.
```

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of non-free programs enables many more people to use the whole GNUfree software. For example, permission to use the GNU C Library in operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the

users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the later must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a

table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library

creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply,

and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the library's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>

This library is free software; you can redistribute it and/or
License as published by the Free Software Foundation; either modify
it under the terms of the GNU Lesser General Public
version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public
License along with this library; if not, write to the Free Software
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school,
if any, to sign a "copyright disclaimer" for the library, if
necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the
library `Frob' (a library for tweaking knobs) written by James Random Hacker.

<signature of Ty Coon>, 1 April 1990
Ty Coon, President of Vice

That's all there is to it!

B

Book Example Installation

The book comes with the source code for the examples discussed in the book. The examples are included with the book archive. When you unzip the `JBossBook_326.zip` archive this creates an `AdminDevel` directory that contains an `examples` subdirectory. This is the `examples` directory referred to by the book.

The only customization needed before the examples may be used is to set the location of the JBoss server distribution. This may be done by editing the `examples/build.xml` file and changing the `jboss.dist` property value. This is shown in bold below:

```
<project name="JBossBook 3.2.x examples"
    default="build-all" basedir=".">

    <!-- Allow override from local properties file -->
    <property file=".ant.properties" />
    <!-- Override with your JBoss/Web server bundle dist location -->
    <property name="jboss.dist" value="/tmp/jboss-3.2.6"/>
    <property name="jboss.deploy.dir" value="${jboss.dist}/server/default/deploy"/>
```

or by creating an `.ant.properties` file in the `examples` directory that contains a definition for the `jboss.dist` property. For example:

```
jboss.dist=/usr/JBoss3.2/jboss-3.2/build/output/jboss-3.2.6
```

Part of the verification process validates that the version you are running the examples against matches what the book examples were tested against. If you have a problem running the examples first look for the output of the `validate` target such as the following:

```
validate:
  [java] ImplementationTitle: JBoss [WonderLand]
  [java] ImplementationVendor: JBoss.org
  [java] ImplementationVersion: 3.2.6RC2 (build: CVSTag=Branch_3_2 date=200409270100)
  [java] SpecificationTitle: JBoss
  [java] SpecificationVendor: JBoss (http://www.jboss.org/)
  [java] SpecificationVersion: 3.2.6
  [java] JBoss version is: 3.2.6
```