

# **COMMERCE**

Version 11.0

Installation and Configuration Guide

Oracle ATG One Main Street Cambridge, MA 02142 USA

## Installation and Configuration Guide

Product version: 11.0 Release date: 01-10-14 Document identifier: AtgInstallGuide1402071827

Copyright © 1997, 2014 Oracle and/or its affiliates. All rights reserved.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support: Oracle customers have access to electronic support through My Oracle Support. For information, visit http:// www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

# **Table of Contents**

1. Introduction	1
Document Conventions	1
Important Terms	1
2. Installation Procedure	3
Preparing Java Application Servers	4
Oracle WebLogic Server	4
JBoss Enterprise Application Platform	5
IBM WebSphere Application Server	5
Basic Database Configuration	5
MvSOL Databases	6
Running the Platform Installer	8
Installing Other Oracle ATG Web Commerce Products	9
Configuration and Installation Manager (CIM)	. 10
Before Using CIM	. 10
Using CIM	. 11
Adding Modules to Your Application	. 11
Configuring Other Oracle Commerce Tools	. 12
Installing Demonstration Applications	14
Default User Accounts	15
Dynamo Server Admin User	15
Business Control Center Lleer Accounts	16
3 Java Application Server Configuration	. 10
General Configuration	17
Configuration	. 17
G7IP Compression	. 17
File Name Character Encoding	17
Application Server Configuration	17
	. 10 10
Mamory Allocation	10
Memory Allocation	. 19
	. 19
	. 20
JBOSS	. 20
JBoss-specific Requirements	. 20
JBoss Installation Results	22
JBOSS EAP 6	. 22
IBM WebSphere	22
WebSphere-Specific Requirements	. 22
	. 23
4. Configuring Databases and Database Access	. 25
Creating Database Tables Using SQL Scripts	. 26
Creating Database Tables for AIG Adaptive Scenario Engine	. 26
Creating Database Tables for AIG Portal	. 28
Destroying Database Tables	. 29
Destroying Database Tables for ATG Adaptive Scenario Engine	. 29
Destroying Database Tables for ATG Portal	. 31
Adding a JDBC Driver	. 32
Configuring ATG Data Sources for Data Import	32
Configuring Data Sources and Transaction Management	. 34
Configuring Data Sources for JBoss	. 34
Configuring Data Sources for WebLogic and WebSphere	. 37
Configuring Data Sources for an Oracle RAC Cluster	. 38

Setting the Transaction Timeout on JBoss	38
Setting the Transaction Timeout on WebLogic	38
Setting the Transaction Timeout on WebSphere	39
Setting the Isolation Level for Transactions in WebSphere	39
Direct SQL Deployment and Microsoft SQL Server	39
Data Source Debugging	39
Using the JDBC Browser	41
Configuring the JDBC Browser	41
Create Table Operation	41
Drop Table Operation	41
Execute Query Operation	41
Metadata Operations	42
Using Oracle ATG Web Commerce Products with an IBM DB2 Database	42
Using Oracle ATG Web Commerce Products with a Microsoft SOL Server Database	43
Moving Data from a Development Database to the Production Database	46
Transferring the Demo Data	47
Conving and Switching Databases	47
Database Conv Operations	18
Croating a DBConjor Component	-10 /10
Configuring the DBConnectionInfo	40
	49
Configuring the DBCopier	49
Setting the Native SQL Environment	50
	51
Configuring a SwitchingDataSource	51
Database Switching and Query Caching	52
Database Sorting for Localization	52
5. Platform Installation Details	53
File System Permissions	53
Performing a Maintenance Installation	53
Default Ports	53
Sun T1000 and T2000 Requirements	54
Installing the ATG Control Center on a Client Machine	54
Downloading the ACC Installer	55
Installing the ACC on a Windows Client	55
Installing the ACC on a UNIX Client	55
Removing the Oracle ATG Web Commerce Platform from Your System	55
6. Running Nucleus-Based Applications	57
Starting the SQL JMS Admin Interface	57
Starting Oracle ATG Web Commerce Web Services	58
Connecting to the Dynamo Server Admin	58
Connecting to the Business Control Center	59
Starting the ATG Control Center	59
Starting the ACC on a Server	60
Starting the ACC on a Client	62
Using the Findclass Utility	63
Stopping an Oracle ATG Web Commerce Application	63
Stopping Applications on JBoss	63
Stopping Applications on Webl ogic	64
Stopping Applications on WebSphere	64
7 Configuring Nucleus Components	65
Working with Configuration Lavers	65
Understanding Properties Files	66
Understanding Configuration Lavers	66
enselstanding configuration Eagles internet	

Accessing Configuration Layers in the ACC	67
Global Configuration Changes	. 68
Locking Configuration Layers	. 68
Finding Components in the ACC	. 68
Changing Component Properties with the ACC	. 69
Changing Component Properties Manually	. 71
Using Forward Slashes (/) and Backslashes (\)	. 72
Modifying Lists of Values	. 72
Specifying Directory Paths	73
Adding Comments to Properties Files	. 73
Using the Dynamo Component Browser	. 73
Component Browser Structure	. 73
Changing the Running Configuration	. 74
Starting Nucleus Components	. 74
Customizing the Interface	. 74
Common Configuration Changes	. 75
Modifying Environment Settings	. 75
Modifying Custom Module Resource Settings	. 76
Enabling checkFileNameCase on Windows	. 76
LogListeners	. 76
Creating Additional Oracle ATG Web Commerce Server Instances	. 77
Using the MakeDynamoServer Script	. 77
Using the Configuration Manager	. 78
Configuring a New Server Instance	. 78
Setting Up a Configuration Group	. 79
Configuration Group Properties	. 81
Storing Group Configuration Files	. 82
Downloading Group Configuration	. 84
Validating Group Configuration Properties	. 85
Session Management in Oracle ATG Web Commerce Applications	. 86
Sharing Session Information Among ATG Applications	. 87
Session Interaction Outline	. 88
Managing User Sessions	. 89
8. Configuring for Production	. 91
Enabling liveconfig Settings	. 91
Customizina liveconfia Settinas	92
Disabling Checking for Changed Properties Files	92
Disabling the Performance Monitor	92
Adjusting the nageCheckSeconds Property	93
Fine-Tuning IDK Performance with HotSpot	. 23
Configuring Benositories	. 95
Satting Cache Modes	03
Propopulating Caches on Startup	. 95
Enabling the Penesitory Cache Lock Managers	. 94
Configuring Depository Database Verification for Quicker Destarts	04
Configuring a Content Distributor System	. 94
Configuring a Content Distributor System	. 94
Conliguring Targeled E-Mail	. 95 05
Nucleus Components	. 95
Configuring web Applications	. 96
Setting Access Levels for Properties Files	. 96
Setting Logging Levels	. 9/
Limiting Initial Services for Quicker Restarts	. 98
Disabiling Document and Component Indexing	. 98

Enabling the ProtocolChange Servlet Bean	98
Setting up Clustering on JBoss	99
Configuring the HttpPort Property	99
Creating ATG Servers	99
Assembling for a JBoss Cluster	. 100
Creating and Configuring JBoss Servers	. 100
Deploying Your Application	. 100
Setting Up Clustering on WebLogic	. 100
Assembling for a WebLogic Cluster	. 101
Clustering Example	. 101
Setting up Clustering on WebSphere	. 102
Installing and Configuring WebSphere	. 103
Creating a Cluster	. 103
Creating Data Sources	. 103
Installing and Configuring Your Web Server	. 103
Installing ATG for a WebSphere Cluster	. 103
Assembling for a WebSphere Cluster	. 104
Session Management in a WebSphere Cluster	. 104
Configuring Your WebSphere Servers	105
Deploying Your Application	. 105
General Clustering Information	. 106
Specifying the drpPort Setting	. 106
Setting up localconfig and Server Configuration Files	. 106
Unique Components	. 107
Enabling Component Backup	. 107
Synchronizing Server Clocks	. 108
SecurityServlet and ParameterValidator	. 108
Configuring ParameterValidator	. 108
Creating an Overriding Parameter Validator	. 113
RedirectURLValidator	. 114
HttpOnly Attribute for Cookies	. 115
Recording Login Attempts	. 115
Login Attempt Log Information	. 116
AuthenticationMessageTrigger Components	. 117
Recording Invalid Usernames	. 117
Rotating Login Attempt Log Files	. 117
Disabling and Enabling Login Attempt Logging	. 117
Handling Multiple Servers	. 118
Logging Metric Pricing	. 118
Creating a Metric Pricing Report	. 118
Filtering Static Resource Types	. 119
Performance Diagnostics	. 121
Performance Troubleshooting Checklist	. 121
Performance Testing Strategies	. 122
Graduated Testing of Throughput	. 122
Realistic Testing Strategies	. 122
Locating Performance Bottlenecks	. 123
Monitoring System Utilization	. 123
Bottlenecks at Low CPU Utilization	. 123
Checking for Database Bottlenecks	. 123
Checking for Disk I/O Bottlenecks	. 124
Checking for Network-Limited Problems	. 124
Bottlenecks at High CPU Utilization	. 124

9.

Thread Context Switching Problems	124
System Resource Bottlenecks	125
TCP Wait Problem on Solaris	125
Server Hangs	126
Paging and Memory Allocation	126
Garbage Collection	127
Memory Leaks	127
Swap Space	128
Detecting File Descriptor Leaks	128
Using URI Hammer	128
Command Line Arguments	120
LIRI Hammer Examples	121
The script Argument	120
Pecording a Crint	122
	122
	134
URLHammer Source Files	133
10. Monitoring Site Performance	137
Performance Monitor	137
Adding PerformanceMonitor Methods to your Code	137
Performance Monitor Modes	139
Viewing Performance Monitor Data	. 140
Instrumented ATG Classes	. 141
Performance Monitor API	142
Using the Configuration Reporter	145
Configuration Reports	146
Excluding Components from the Configuration Report	146
Running the Configuration Reporter as a Standalone Utility	147
Using the VMSystem Component	149
Using a Sampler	149
Starting the Sampler	149
Sampler Information	150
Sampler Output	150
Using the Recording Servlet	150
Inserting the Recording Servlet	151
Generating Script Files	. 151
Keeping Statistics	151
Tracing Memory	151
11. Repository and Database Performance	153
Database Performance Practices	153
Repositories and Transactions	154
Repository Item Property Loading	154
Database Sorting versus Locale-Sensitive Sorting	154
Batching Database Transactions	154
Avoiding Table Scans	155
Database Caches	155
External Caching for SOL Panasitorios	150
Configuring Cohorongo for External Caching	150
Using Coherence Utilities With the ATC Server Cluster	157
Using Conference Officies With the Ard Server Cluster	159
	159
Diagnosing Database Performance Problems	159
Avoid Using Simulated Text Search Queries in Repositories	100
Drop index in OLIP Environment	100
12. Turning site Performance on JBoss	101

JBoss File Modifications	161
JSP Servlet Configuration	161
Tomcat Connector Thread Configuration	162
Tomcat Cluster Configuration	163
JBoss Logging Configuration	163
Data Source Configuration	163
Configuring run.bat/sh and run.conf	164
JBoss Application Framework Trimming	164
13. Appendix A: Data Storage and Access	167
Database Schema Best Practices	167
Data Sources	168
Repositories	169
14. Appendix B: Adjusting the FileCache Size	173
15. Appendix C: Development Tools for Eclipse	175
16. Appendix D: Using Oracle Access Management for Single Sign On	177
OAM Authentication Overview	177
Logging In to the Business Control Center	177
Logging Out of the Business Control Center and SSO	178
Session Timeouts	178
Installing the Oracle Access Manager Integration Component	179
Validating OAM Identity Assertion X.509 Certificates	179
Configuring X.509 Certificate Validation	180
Using IP Address Filtering for Verified Authentication	182
Oracle WebLogic IP Filtering	182
JBoss and Tomcat IP Filtering	183
IBM WebSphere IP Filtering	183
Recommended Deployment	184
Index	187

# **1** Introduction

This section includes information about the purpose of this guide and how to use it.

# **Document Conventions**

This guide uses the following conventions:

- <ATG11dir> represents the ATG installation directory (C:\ATG\ATG11.0, for example)
- <JBdir> represents the Red Hat JBoss home directory (C:\jboss\ jboss-eap-install-dir\jboss-as, for example)
- <WLdir> represents the Oracle WebLogic home directory (C:\Oracle\Middleware\wlserver\_10.3, for
   example)
- <WASdir> represents the IBM WebSphere home directory (C:\IBM\WebSphere\AppServer, for example)

## **Important Terms**

This section defines terms used throughout this guide.

- ATG products. Umbrella name for the software suite, particularly the platform.
- ATG installation. Collective name for the tools, files, classes, etc. used for developing and assembling JEE applications.
- ATG application. A piece of software installed independent of the platform, which can be included as a module or set of modules in a Nucleus-based application.
- ATG server. A configuration layer that is available to be added to other configuration layers by the application assembler when assembling an EAR.
- Dynamo Server Admin. Web pages used to configure and monitor the ATG installation.
- · Component. A Java object instance of a specific configuration for a JavaBean that is registered with Nucleus.
- Nucleus-based application. An assembled EAR file created out of components managed by ATG's Nucleus component manager, running on the application server.

# 2 Installation Procedure

This section provides basic instructions for installing Oracle ATG Web Commerce. Use these instructions to configure evaluation or development Oracle ATG Web Commerce applications or as an overview of the steps required to install and deploy production Oracle ATG Web Commerce applications.

Following the installation procedure in this section results in a functional server configuration but the procedure does not contain all the information you need to meet the more sophisticated requirements of live, production use. If you are installing Oracle ATG Web Commerce for production use, make sure to read the more detailed configuration information in the other chapters of this guide. This section refers to those chapters where they are needed.

To install Oracle ATG Web Commerce software:

1. Install and configure the Java Development Kit (JDK) on your host machine. You need to know the path to your JDK installation directory later in this installation procedure.

To find the required JDK version and the required versions of other supporting software, refer to the Oracle ATG Commerce Supported Environments Matrix document in the My Oracle Support knowledge base.

- 2. Install and configure the Java application server to use. See Preparing Java Application Servers (page 4).
- 3. Run the Oracle ATG Web Commerce platform installer. Follow its prompts. See Running the Platform Installer (page 8).
- 4. Run the installer programs for any other Oracle ATG Web Commerce products that you want to use. See Installing Other Oracle ATG Web Commerce Products (page 9).
- Install and configure the database application you use. Ensure that you have the latest hot fixes, fix patches or other updates needed for your database. Create database accounts for the schemas required by your Oracle ATG Web Commerce application. Start the database application. See Basic Database Configuration (page 5).
- 6. Place the Java Data Base Connectivity (JDBC) driver for your database software on your host machine.

If you have installed Oracle ATG Web Commerce on Microsoft Windows, you can find the JDBC driver for MySQL at <aTG11dir>\MySQL\mysql-connector-java-version\_number-bin.jar. You can leave the driver there and use it from that path.

- Run the Oracle ATG Web Commerce Configuration and Installation Manager (CIM) utility to configure your application, create database schemas, create data sources, and deploy them to your application server. See Configuration and Installation Manager (CIM) (page 10).
- 8. Start your Oracle ATG Web Commerce application by invoking the start functionality of your Java application server. You can use the script that CIM creates in the <aTGlldir>/home/servers/server-name directory. For example:

 $\verb|home\servers|atg\_production\_lockserver\startServerOnWeblogic.bat|$ 

**Note**: CIM does not create a start script for the IBM WebSphere application server. Use that server's starting scripts or administration application to start your Oracle ATG Web Commerce application. See the IBM WebSphere documentation for instructions.

# **Preparing Java Application Servers**

You must perform some Java application server configurations before running the Oracle ATG Web Commerce installer program. For example, you must create a base domain in your Oracle WebLogic application server.

This section provides basic information about preparing the application server that runs your Oracle ATG Web Commerce application. Prepare your application server before you install and configure Oracle ATG Web Commerce.

This section only contains the minimal information that you need to run an application for evaluation or development. See information about configuring application servers for production use in *Java Application Server Configuration* (page 17).

### **Oracle WebLogic Server**

Before you install and configure Oracle ATG Web Commerce:

- 1. Make sure that you have the username and password for the administration user.
- 2. Make sure that the administration application is running.
- 3. Configure a base domain. See information about creating a domain in the documentation for your Oracle WebLogic application server.
- 4. Make sure that you have the paths to the Oracle Middleware directory, the WebLogic application server home directory, and the Oracle WebLogic base domain directory. You need to know these paths when you run the Oracle ATG Web Commerce installer and the Configuration and Installation Manager (CIM). For example:
  - C:\Oracle\Middleware
  - C:\Oracle\Middleware\wlserver\_10.3
  - C:\Oracle\Middleware\user\_projects\domains\base\_domain

#### Bundled Oracle WebLogic Server Restricted Use

Oracle WebLogic Server Standard Edition is bundled with Oracle ATG Web Commerce. The bundled version of Oracle WebLogic is subject to restricted use.

You may use the bundled version of Oracle WebLogic Server Standard Edition to run Oracle ATG Web Commerce applications only. Oracle ATG Web Commerce and the products it contains may be modified, extended, or integrated with external systems through the use of custom Java development. If you develop custom Java applications that do not extend, modify, or integrate Oracle ATG Web Commerce applications, you must have a full-use license.

**Note**: Oracle WebLogic Server Standard Edition does not support application clustering. If you need to use clustering, you must purchase or supply Oracle WebLogic Server Enterprise Edition

#### JBoss Enterprise Application Platform

Before you install and configure Oracle ATG Web Commerce, make sure you have the path to the JBoss application server home directory. You need to provide it when you run the Oracle ATG Web Commerce installer and the Configuration and Installation Manager (CIM). For example:

```
C:\jboss-eap-install-dir\jboss-as
```

#### Set the Bind Address

If you access your Oracle ATG Web Commerce application from a different computer, set the bind address when you start the JBoss application server. By default, the JBoss application server will not serve applications to users on other computers.

For example, to access your application from any Internet Protocol (IP) address you can add -b 0.0.0.0 to the start script that is created by the CIM utility. See Configuration and Installation Manager (CIM) (page 10).

```
#!/bin/sh
/jboss-installation-directory/jboss-as//bin/run.sh \
-c server-name \
-b 0.0.0.0 $*
```

**Note**: this script is created by the CIM utility after you have installed and configured the Oracle ATG Web Commerce platform. It is not present when you are preparing the application server before the platform installation.

#### **IBM WebSphere Application Server**

Before you install and configure Oracle ATG Web Commerce:

1. Make sure you have the path to the IBM WebSphere installation directory. You need to provide it when you run the Oracle ATG Web Commerce installer and the Configuration and Installation Manager (CIM). For example:

C:\IBM\WebSphere\AppServer

- 2. Make sure you have the username and password of the IBM WebSphere administration user.
- 3. Make sure the IBM WebSphere administration application is running.

If you run IBM WebSphere as the root user of a UNIX or Linux operating system, also run CIM as the root user.

The CIM utility uses cell deployment when deploying applications to the IBM WebSphere application server. Use the IBM WebSphere Network Deployment version for this type of installation, denoted by ND.

If you have installed the IBM WebSphere Network Deployment version, when you run the ATG installer you must select the *IBM WebSphere - cluster setup* option even if you are not actually using clustering.

# **Basic Database Configuration**

Oracle ATG Web Commerce applications require one or more relational database schemas. If you are installing a server for evaluation or development use, you may need only one database schema. The Configuration

and Installation Manager (CIM) utility determines what schemas are required and configure them later in this procedure.

If you are installing Oracle ATG Web Commerce for evaluation or development use, you can use the MySQL database. This database is installed and configured automatically on Microsoft Windows. You can use the MySQL database on other operating systems but you must install and configure it yourself. See MySQL Databases (page 6).

Make sure your database software is running before you use CIM to configure your Oracle ATG Web Commerce application. To start the MySQL database on Microsoft Windows, choose ATG11.0 > Tools > Start MySQL Server from the programs section of the Windows Start menu.

If you need detailed information about databases for a production installation, see *Configuring Databases and Database Access* (page 25).

### **MySQL** Databases

Oracle ATG Web Commerce supports the MySQL database for use in evaluation and development installations. The Oracle ATG Web Commerce platform installer for Microsoft Windows automatically installs MySQL. You can also use the MySQL database on other operating systems but you must install and configure it yourself.

Use the MySQL database if you want to configure an Oracle ATG Web Commerce server quickly and you are not concerned about large-scale performance and reliability. Database configuration for production use requires more sophisticated preparation. Configuring databases for production use is covered in *Configuring Databases and Database Access* (page 25).

This section provides information about quickly configuring a MySQL database for an evaluation or development server.

- See information about using MySQL on Microsoft Windows in MySQL Installed on Microsoft Windows (page 6).
- See information about using MySQL on other operating systems in Installing MySQL on Other Operating Systems (page 7).

#### **MySQL Installed on Microsoft Windows**

The Oracle ATG Web Commerce platform installer for Microsoft Windows can also install the MySQL database. The MySQL database is preconfigured with the user accounts that you use to configure an evaluation or development application.

To install MySQL with the Oracle ATG Web Commerce platform, choose it from the list of products on the Select Products to Install screen of the installer wizard.

The Oracle ATG Web Commerce platform installer does not create the actual schemas or import required data. Use the Oracle ATG Web Commerce Configuration and Installation Manager (CIM) utility to perform these steps before using the databases to run an application. See Configuration and Installation Manager (CIM) (page 10).

To start the MySQL database that is included in a Microsoft Windows installation, choose ATG11.0 > Tools > Start MySQL Server from the programs section of the Windows Start menu.

The following table lists the MySQL database accounts and databases that the Oracle ATG Web Commerce platform installer creates for Microsoft Windows installations.

Account Name	Password	Database Name
prod	Welcome1	production_core
pub	Welcome1	publishing
switchA	Welcome1	switchinga
switchB	Welcome1	switchingb
agent	Welcome1	agent
dw	Welcome1	datawarehouse

The Oracle ATG Web Commerce installation program does not set a password for the root MySQL user. You can log in as the root user without specifying a password. Additionally, the default MySQL database does not support i18n content. See information MySQL database documentation for additional information (http://dev.mysql.com/doc/).

#### Installing MySQL on Other Operating Systems

You can use the MySQL database for evaluation and development applications on any operating system that is supported by Oracle ATG Web Commerce. However, the MySQL database software is only included with the Microsoft Windows version of the installer program. If you install Oracle ATG Web Commerce on an operating system other than Microsoft Windows, you must install and configure MySQL yourself.

To prepare a MySQL database for evaluating or developing Oracle ATG Web Commerce applications;

- 1. Install the MySQL server and MySQL client software. To find the required version, refer to the Oracle ATG Commerce Supported Environments Matrix document in the My Oracle Support knowledge base.
- 2. Open the MySQL console as the root user. If you have installed the MySQL client software, you may be able to do this by invoking the mysql executable file.

mysql -u root -p

3. Create the production core database, create a user account for the production core database, and give the user account access to the database.

```
mysql> create database production_core;
Query OK, 1 row affected (0.00 sec)
mysql> create user 'prod'@'localhost' identified by 'Welcomel';
Query OK, 0 rows affected (0.09 sec)
mysql> grant all on production_core.* to 'prod'@'localhost';
Query OK, 0 rows affected (0.01 sec)
```

4. Create any other databases that your installation requires. For example, if you use ATG Content Administration, create the publishing database. Use the MySQL commands shown in the previous step.

See a list of database types and suggested names in MySQL Installed on Microsoft Windows (page 6).

# **Running the Platform Installer**

The platform installer creates a home directory for Oracle ATG Web Commerce applications on your host machine. The default path is C:\ATG\ATG11.0. These instructions refer to this directory as <ATG11dir>.

The Oracle ATG Web Commerce platform installer is available as a self-extracting Windows executable (ATG11.0.exe) or UNIX binary file (ATG11.0.bin), which you can download from the Oracle Web site. This distribution file includes the following Oracle ATG Web Commerce products:

- Oracle ATG Web Commerce platform
- Oracle ATG Web Commerce
- ATG Content Administration
- ATG Portal

If you are installing on a Linux or UNIX operating system, **do not** run the platform installer as the root user. Installing as the root user sets file ownership incorrectly and may introduce security vulnerabilities.

Follow these steps to install the Oracle ATG Web Commerce platform:

1. Run the ATG11.0.exe or ATG11.0.bin file to start the setup program.

**Note:** If you are installing on a Linux variety that includes GCJ, in order to avoid installation errors you must specify a JVM that includes the javax.swing classes, which are not included in GCJ. Use the following command:

\$sh ./install.bin LAX\_VM path\_to\_java\_executable

For example:

\$sh ./ATG11.0.bin LAX\_VM /usr/local/j2sdkversion/bin/java

- 2. After you accept the terms of the license agreement, select the installation folder for the Oracle ATG Web Commerce software (C:\ATG\ATG11.0 or /home/ATG/ATG11.0, for example).
- 3. Select the Oracle ATG Web Commerce products you want to install.

You can choose to install the Quincy Funds and Motorprise demonstration applications in addition to the Oracle ATG Web Commerce products.

- 4. Select your application server.
- 5. If installing for **WebLogic**, enter the following configuration information:
  - The RMI port your Oracle ATG Web Commerce applications uses (defaults to 8860)
  - The listen port that WebLogic uses to listen for incoming connections (defaults to 7001)
  - The WebLogic home directory
  - The path to your WebLogic domain directory (C:\oracle\user\_projects\domains\mydomain, for example)
  - The JDK home directory (C:\j2sdkversion, for example)

If installing for JBoss, enter the following configuration information :

- The RMI port your Nucleus-based applications uses (defaults to 8860)
- The listen port that JBoss uses to listen for incoming connections (defaults to 8080)
- The JBoss home directory (C:\jboss-eap-install-dir\jboss-as, for example)
- The JDK home directory (C:\j2sdk1.6.0\_22, for example)

If installing for **WebSphere**, enter the following configuration information:

- The RMI port your Oracle ATG Web Commerce applications use (defaults to 8860)
- The port that WebSphere uses to listen for incoming connections (defaults to 9080)
- The WebSphere home directory (C:\WebSphere\AppServer, for example)
- The name of the WebSphere server (server1, for example)
- The node on which the WebSphere server is installed (Typically, the node name is the same as the host machine name.)

## Installing Other Oracle ATG Web Commerce Products

Some Oracle ATG Web Commerce products are not included in the platform installer. You must run separate installation programs in order to use these products.

Installing other Oracle ATG Web Commerce products typically adds components to the <ATG11dir> directory created by the platform installer. Make sure that you run the platform installer before individual products so that this directory is in place.

The Oracle ATG Web Commerce platform installation program installs:

- The Oracle ATG Web Commerce platform software
- · The Commerce and Merchandising products
- The Content Administration product
- · The Quincy Funds and Motorprise demonstration applications
- · Components for business intelligence

The following table lists Oracle ATG Web Commerce products that are not included in the platform installer and links to installation instructions for them.

Oracle ATG Web Commerce Product	Installation Instructions
Oracle ATG Web Commerce Service Center	Commerce Service Center Installation and Programming Guide
Oracle ATG Web Commerce Reference Store	Commerce Reference Store Installation and Configuration Guide

# **Configuration and Installation Manager (CIM)**

The Configuration and Installation Manger (CIM) reduces the complexity of configuring Oracle ATG Web Commerce applications. A series of text-based wizards guide you through configuration procedures, ensuring that necessary steps are completed, and that steps are performed in the correct order. Menus are dynamically generated based on your selections to provide choices appropriate for your installation.

The installation guides for individual products contain specific information on what CIM accomplishes for those products, but in general, CIM handles the following configuration areas:

- Oracle ATG Web Commerce product selection
- · Data source configuration
- · Database table creation and data import
- · Oracle ATG Web Commerce server instance creation and configuration
- · Application assembly and deployment

The result is a functional installation that can be used as a starting point for further configuration. CIM does not perform configuration steps while the Oracle ATG Web Commerce application is running.

**Note:** CIM does not configure a scenario or process editor server. See the *Multiple Application Integration Guide* for information on scenario editor servers.

### **Before Using CIM**

Before you use CIM to configure and deploy an application, make sure you have the following prerequisite information. See information about these basic aspects of installation in *Installation Procedure* (page 3).

• The path to your application server home directory. For example:

C:\Oracle\Middleware\wlserver\_version

 The path to the profile or similar directory for your application. Some application servers do not use profile directories. For example:

C:\Oracle\Middleware\user\_projects\domains\base\_domain

- The username and password for the administration account for your application server. If you have not
  configured security for your application server, you do not need this information.
- The account names, passwords, database/schema names, listening port, and server hostname for each database that your application requires.
- The path to the Java Data Base Connectivity (JDBC) driver for your database software.
- The password for the Dynamo Server Admin administrator account. See Connecting to the Dynamo Server Admin (page 58).
- The password for the Business Control Center administrator account. Enter this password during database imports.
- The password for the Oracle ATG Web Commerce merchandising user account. Enter this password during database imports. If you are not using Content Administration, do not configure this user account.

 If you are creating a localized installation, you should know the locale, language and region codes of the countries. This information should be added into your server's environment settings. Refer to the Modifying Environment Settings (page 75) for additional information. Also see the *Internationalizing an ATG Web Site* section of the *Platform Programming Guide*.

**Note:** To use Microsoft Windows Native Language support, you must have the correct language pack installed and have the system locale set correctly using the Region and Language > Administrative > Change System Locale control panel. Refer to your Microsoft Windows documentation on language packs and default location configuration.

### **Using CIM**

To use CIM, do the following:

- 1. Install Oracle ATG Web Commerce according to the instructions in *Installation Procedure* (page 3). Follow the procedure until you reach the step that requires CIM.
- 2. Start CIM.
  - On Microsoft Windows, choose ATG11.0 > Tools > ATG Configuration Installation Manager from the programs section of the Windows Start menu.
  - On any operating system, invoke the following script.

<ATG11Dir>\home\bin\cim.bat|sh

3. Follow the prompts to configure your installation. You need the information described in Before Using CIM (page 10). To access the online help, enter H at any point in the configuration process.

**WebSphere Note:** In order to use CIM to configure an ATG installation for WebSphere, your WebSphere installation needs to use cell deployment. Use the IBM WebSphere Network Deployment version for this type of installation, denoted by ND. Also, note that if you run WebSphere as the root user of a UNIX or Linux operating system, also run CIM as the root user.

**Note:** Ensure that you have only one instance of CIM running from the same ATG installation. When more than one instance of CIM is run, errors will occur.

### **Adding Modules to Your Application**

Use CIM to add modules to your Oracle ATG Web Commerce application. You can add modules to your application before you deploy it for the first time or after you have deployed it. If you have already deployed your application to an application server, you will need to redeploy it after adding modules.

To add a module to your Oracle ATG Web Commerce application:

- 1. Stop the application servers on which your application is installed if they are running.
- 2. Run CIM by invoking its start script:

<ATG11dir>/home/bin/cim.sh|bat

- If you are configuring your application for the first time, follow the CIM prompts until you reach the Server Instance Configuration step. If you have already configured your application with CIM, choose Server Instance Configuration from the main menu.
- 4. Choose the server to add the module to in the server instance type selection menu.

- 5. Choose Modify Calculated Module List OPTIONAL from the server instance type configuration menu.
- Choose Add A Custom Module from the module list editor options menu. Enter the name of the module you are adding.

7. If you are configuring your application for the first time, continue following the CIM prompts to deploy your application. If you have already configured your application with CIM, choose Application Assembly & Deployment from the main menu and follow the CIM prompts to redeploy it.

#### **Configuring Other Oracle Commerce Tools**

When installing Oracle ATG Web Commerce, you have the ability to implement Single Sign On (SSO) with Endeca Workbench, as described in the *Appendix D: Using Oracle Access Management for Single Sign On* (page 177) of this document or the *ATG-Endeca Integration Guide*.

When using SSO, you can configure Endeca Workbench with a link that takes a user directly to the Business Control Center :

- 1. Configure your environment, including product selection, application server, and database configuration, as described in the Using CIM (page 11) section.
- The Workbench Configuration menu is found under the Publishing Server Instance Configuration menu. Select option [W] Workbench configuration – OPTIONAL.

**Note:** If you are not using CIM to configure your environment, ensure that you include the work\_bench\_ddl.sql to your installation scripts, as it is not included in the BIZUI DDL.

- 3. Enter the Workbench configuration path, which is a location where the Workbench configuration files will be stored. For example <atGlldir>/workbench/config/.
- 4. Enter the Business Control Center host name of the Business Control Center instance that will be accessible from the Workbench server. For example, myserver.example.com.
- 5. Provide the Business Control Center server port number of the Business Control Center instance that will also be accessible from the Workbench server. For example: 8080.

This generates a /workbench directory that contains the ws-extensions.xml and ws-mainMenu.xml file, as well as a /workbench/locales directory that contains the configuration resource file.

6. These XML files must be copied or merged into the /ToolsAndFrameworks/ version/server/workspace/conf folder on the Endeca Workbench server.

If you have made no modifications to the extensions in Workbench, copy the

ws-extensions.xml file into the directory. If you have added extensions to Workbench, you must copy the bcc-home and bcc-access-control extensions into the existing ws-extensions.xml files. For example:

```
<extension id="bcc-home" defaultName="BCC" defaultDescription ="BCC"
url="http://myserver.example.com:8080/atg/bcc"
externalURL="true"/>
<extension id="bcc-access-control"
defaultName="BCC Access Control"
defaultDescription="BCC Access Control"
role="admin"
url="http://myserver.example.com:8080/ControlCenter/application/
accesscontrol"
externalURL="true"/>
```

If you have made no modifications to the menus in Workbench, copy the ws-mainMenu.xml file into the directory. If you have added menus to Workbench, you must copy the new menus into the existing ws-mainMenu.xml file. For example:

```
<menuitem id="xmgr"/>
<menuitem id="bcc-home"/>
<menunode id="user-access">
<menuitem id="user-management"/>
<menuitem id="bcc-access-control"/>
</menunode>
<menuitem id="eac-admin-console"/>
<menuitem id="eac-settings"/>
<menunode id="reports">
<menuitem id="reports.today"/>
<menuitem id="reports.daily"/>
<menuitem id="reports.weekly"/>
</menunode>
<menunode id="search">
<menuitem id="redirects"/>
<menuitem id="thesaurus"/>
<menuitem id="phrases" />
<menuitem id="relrank" />
</menunode>
<menunode id="settings">
<menuitem id="user-segments"/>
<menuitem id="preview-settings"/>
<menuitem id="report-settings"/>
<menuitem id="locks"/>
</menunode>
```

7. Merge the resource.properties file located in the /workbench/locales directory with that of the Workbench servers /locales directory. Add the Business Control Center extension, Access Control extension and User Access Menu information. For example:

```
# These resources should be localized.
atg_flags=i18n,110n
# BCC extension
tool.bcc-home.name = Business Control Center
tool.bcc-home.description = Create, preview, and deploy site content and product
catalogs.
# @i18n:begin:trans(false)
```

```
tool.bcc-home.icon = http://myserver.example.com:8080/atg/images/
BCC_home/icon_wb_bcc.png
# @il8n:end:trans
# BCC Access Control extension
tool.bcc-access-control.name = BCC Access Control
# User Access menu
menu.user-access.name = User Access
menu.user-access.description = Add and remove Oracle Endeca Workbench and Business
Control Center Users and modify their tool access and content permissions.
# @il8n:begin:trans(false)
menu.user-access.icon = /ifcr/apps/endeca/workbench-assets/images/home/icon_user.png
# @il8n:end:trans
```

**Note:** These resource configuration changes must be made for all localized files, such as the Resources\_cs.properties file.

For additional information on copying and merging these files, refer to your Endeca documentation.

# **Installing Demonstration Applications**

The Oracle ATG Web Commerce platform installation includes the Quincy Funds and Motorprise demonstration applications. You can use the Configuration and Installation Manager (CIM) to include them in an application and deploy it to your application server.

The user instructions for the demonstration applications include installation instructions. See these instructions in the documents listed in the table below.

Demonstration Application	Instructions
Quincy Funds	Quincy Funds Demo Documentation
Motorprise	Business Commerce Reference Application Guide

To install the Quincy Funds or Motorprise demonstration applications:

Install and configure Oracle ATG Web Commerce according to the instructions in *Installation Procedure* (page 3).

Choose either the Quincy Funds or Motorprise demonstration application from the list of products that you want to install.

2. While you are configuring your application in CIM, choose either Quincy Funds or Motorprise from the Include Demo Application menu during the product selection step.

-----INCLUDE DEMO APPLICATION:----enter [h]elp, [m]ain menu, [q]uit to exit Include Demo Application: [1] Quincy Funds Demo

```
[2] MotorPriseJSP
[D] Done
Select zero or one > 1
Include Demo Application:
*[1] Quincy Funds Demo
[2] MotorPriseJSP
[D] Done
Select zero or one > D
```

3. Complete the procedure in *Installation Procedure* (page 3). When your application is deployed and running, it includes the demonstration application.

The following table lists the default URLs for accessing the demonstration applications on the supported application servers.

Demonstration Application	URL and Documentation Link
Quincy Funds (Personalization)	http://hostname:port/QuincyFunds Quincy Funds Demo Documentation
Motorprise (B2B Commerce)	http://hostname:port/Motorprise Business Commerce Reference Application Guide

On WebSphere, before using a demonstration application, set the following properties in the /atg/dynamo/ servlet/pipeline/DynamoHandler.properties file:

```
fixRequestURI=true
fixServletPath=true
```

# **Default User Accounts**

This section provides information about the default user accounts for Oracle ATG Web Commerce.

#### **Dynamo Server Admin User**

Oracle ATG Web Commerce includes an administration application for server configuration. The default user name and password for this application are:

- Username: admin
- Password: your organization sets the password while using the Configuration and Installation Manager (CIM) utility to configure the server. See Configuration and Installation Manager (CIM) (page 10).

### **Business Control Center User Accounts**

Three default user profiles are provided for applications that use the Business Control Center. These profiles are designed to allow administrators to perform initial setup tasks such as creating profiles for other users. They have the following login names:

- admin -- provides access to most areas of the Business Control Center.
- merchandising -- provides access to the Oracle ATG Web Commerce Merchandising interface.
- service -- provides access to the Service Administration interface.

There are no default passwords for these accounts. Your organization sets the passwords in one of the following ways:

- Through CIM. During the post-installation setup process, CIM prompts you to set the passwords for the default profiles that your environment requires. See Configuration and Installation Manager (CIM) (page 10).
- Through the ATG Control Center (ACC). If you do not use CIM to set up your system, you must define passwords for the default profiles through the ACC.

**Important:** The default profiles are internal user profiles. You must include the DSS.InternalUsers.ACC module in your assembled application to be able to access and edit these profiles in the ACC. For more information on editing profiles through the ACC, refer to the *Personalization Guide for Business Users*.

# 3 Java Application Server Configuration

This chapter provides detailed information about configuring Java application servers for use with Oracle ATG Web Commerce.

# **General Configuration**

This section provides configuration information that applies to all Java application servers.

### **Configure SSL**

Configure your application server to use Secure Sockets Layer (SSL) connections. SSL authenticates the server to clients that connect to it and encrypts the data that they exchange. In particular, make sure that connections to internal, administration interfaces such as the Dynamo Server Admin application, the Oracle ATG Web Commerce Business Control Center, and the ATG Control Center use SSL.

SSL authentication helps to prevent session hijacking. Do not allow your Web application to transmit or receive unencrypted session identifiers. Unauthorized users can reuse session identifiers to gain access to your Web application while an authorized user is logged in. Unauthorized users cannot access session identifiers if they are encrypted by SSL.

See instructions for configuring SSL in the documentation for your Java application server. The application server controls SSL for connections to your Oracle ATG Web Commerce application.

### **GZIP** Compression

You should enable GZIP compression for static files. See your application server documentation for information.

### **File Name Character Encoding**

Make sure that your application server is configured to serve the file name characters that you use in your Web applications. If you serve files that include non-English characters in their file names, configure your application server to handle URL requests using Unicode (UTF-8) characters.

For example, if you are using the Oracle WebLogic application server, add the following element to your weblogic.xml file.

```
<charset-params>
<input-charset>
<resource-path>/*</resource-path>
<java-charset-name>UTF-8</java-charset-name>
</input-charset>
</charset-params>
```

See the documentation for your Java application server software for more information about serving files with non-English characters in their file names.

# **Application Server Configuration**

When installing an application servers, you can install a supported version and a version that is not supported by allowed. Application server configurations are located at /atg/cim/productconfig/ appserver/SupportedApplicationServer.properties. This file allows you to add supported versions of an application server, as well as to specify a version that is allows to run. When an unsupported but allowed version of an application server runs, CIM allows the application to run, but issues a warning message in the CIM console.

The following is an example of the SupportedApplicationServer file: Note that the version numbers used in this example may not reflect the latest supported versions of the application. Refer to the Oracle Knowledgebase for the latest supported versions:

```
BLOGIC_SUPPORTED=10.3,12.1.2
WEBSPHERE_SUPPORTED=7.0.0,8.5.0
TOMCAT_SUPPORTED=6.0
JBOSS_SUPPORTED=5.1.0,5.0,5.1.2,6.0
```

JBOSS\_ALLOWED=12.1.1
WEBLOGIC\_ALLOWED=7.0.5, 8.5.3
TOMCAT\_ALLOWED=6.1
WEBSPHERE\_ALLOWED=5.1.1

# **Oracle WebLogic**

If you are using WebLogic and want to run the ACC in a dedicated VM (see Starting the ACC in a Dedicated VM (page 60) in this guide), you must add the following tag to the config.xml file for your WebLogic application server inside the <security-configuration> tag:

```
<enforce-valid-basic-auth-credentials>
  false
</enforce-valid-basic-auth-credentials>
```

See your WebLogic documentation for information on the config.xml file.

To use XA data sources with WebLogic, add the following line to your <ATG11dir>/home/servers/ servername/localconfig/GLOBAL.properties file:

localTransactionModeInitialization=false

In order to create scenarios in the ACC, you must add the <WLdir>/server/lib/wlclient.jar file to your class path. To do this, copy wlclient.jar into the /lib directory of your standalone ACC installation, then modify the bin/startClient.bat file to include wlclient.jar in the class path.

If you are planning to run SQLJMSAdmin on your WebLogic installation, you must change the sessiontimeout value in the SQLJMSAdmin webModule\WEB-INF\web.xml file from zero to a positive number. The recommended timeout value is 30.

Oracle ATG Web Commerce does not support Unicast cluster messaging mode. If you are clustering Oracle ATG Web Commerce servers, set the cluster messaging mode to multicast.

### **Memory Allocation**

To increase the performance of Oracle ATG Web Commerce applications, set the USER\_MEM\_ARGS environment variable for your Oracle WebLogic server as shown below.

USER\_MEM\_ARGS='-Xms1152m -Xmx2048m'

See information about Oracle WebLogic memory allocation and environment variables in the documentation for that product (http://www.oracle.com/technetwork/middleware/weblogic/documentation/index.html).

### **Controlling Page Recompilation on WebLogic**

When you run Oracle ATG Web Commerce applications on WebLogic, WebLogic's JSP container manages JSP compilation. If you are running WebLogic in development mode, modified pages are automatically recompiled when they are requested, ensuring that the . java files associated with the pages are up to date. To prevent performance degradation due to unnecessary page recompilation, when you run WebLogic 10 in production mode, page recompilation is automatically disabled (. jsp files should not change on a production environment, so in theory recompilation never happens; but disabling recompilation ensures that it is not be triggered by a timestamp change).

Although recent WebLogic versions automatically disable page recompilation in production mode, you may want to manually disable recompilation if you are in a testing phase, but not yet running in production mode. Unnecessary recompilation may distort performance tests and slow down your quality assurance process.

To disable page recompilation, create a weblogic.xml file (or modify an existing one) in the WEB-INF directory of each web application you want to include in your EAR file. In the weblogic.xml file, set these two parameters to -1:

pageCheckSeconds specifies the interval in seconds between stale checks for an individual JSP. When a
request for a JSP is received, if the last stale check on this page was longer ago than the number of seconds
that pageCheckSeconds is set to, a new stale check is performed, and if the page is determined to be stale,
it is recompiled. The default in development mode is 1 second. Setting this parameter to -1 disables stale
checking.

• servet-reload-check-secs specifies the interval in seconds between checks of a web application's WEB-INF/classes directory to see if any servlets have been recompiled (and therefore need to be reloaded). The default in development mode is 1 second. Setting this parameter to -1 disables checking.

The following example illustrates disabling both of these checks in the weblogic.xml file:

```
<weblogic-web-app>
  <container-descriptor>
    <servlet-reload-check-secs>-1</servlet-reload-check-secs>
    </container-descriptor>
    <jsp-descriptor>
    <page-check-seconds>-1</page-check-seconds>
    </jsp-descriptor>
</weblogic-web-app>
```

### WebLogic Installation Results

The Oracle ATG Web Commerce setup program adds a protocol.jar file to the /lib subdirectory of the WebLogic domain directory that you specified during the installation process.

Edit the setDomainEnv.sh|bat file in the <domain>/bin directory. For example, you could prepend your required values to the PRE\_CLASSPATH variable, which automatically prepends your values to the WebLogic CLASSPATH:

```
set PRE_CLASSPATH=C:\WebLogic\user_projects\domains\mydomain\protocol.jar;
%WEBLOGIC_CLASSPATH%;%POINTBASE_CLASSPATH%;%JAVA_HOME%\jre\lib\rt.jar;
%WL_HOME%\server\lib\webservices.jar;%CLASSPATH%
```

**Note**: The Configuration and Installation Manager (CIM) utility makes these configurations for you. See Configuration and Installation Manager (CIM) (page 10).

#### Working with Managed WebLogic Server Instances

When a WebLogic server is started using the WebLogic Admin console, it uses the <WebLogic\_Home>/common/ bin/commEnv script to set up environment variables, including the CLASSPATH. This causes an Unknown Protocol error. However, when a server is started from a command prompt using the startManagedWebLogic script, it calls the <WebLogic\_Domain\_Home>/

bin/setDomainEnv script, avoiding any errors. To use the startManagedWebLogic script with a server that has been started using the WebLogic Admin console, you must modify the <WebLogic\_Home>/common/ nodemanager/nodemanager.properties and set the StartScriptEnabled property to true.

## JBoss

#### **JBoss-Specific Requirements**

After installing JBoss, modify the JVM arguments. Go to <JBdir>/bin/run.conf|bat and edit the JAVA\_OPTS line. The following are recommended settings:

```
JAVA_OPTS="-server -Xms2048m -Xmx3072m -XX:MaxPermSize=768m -XX:MaxNewSize=768m -Dsun.rmi.dgc.server.gcInterval=3600000 - Dsun.rmi.client.gcInterval=3600000"
```

If you are setting up a JBoss instance that is dedicated to lock management, you can run that instance with a smaller heap size, since the lock manager does not serve pages. To do this, create a new run.bat | sh file referring to a new run.conf file.

Duplicate the run.bat | sh and run.conf files and rename the duplicates (for example, runLockMan.sh and runLockMan.conf). In the runLockMan.bat | sh file, change the following section to point to the new configuration file:

```
# Read an optional running configuration file
if [ "x$RUN_CONF" = "x" ]; then
    RUN_CONF="$DIRNAME/runLockMan.conf"
fi
if [ -r "$RUN_CONF" ]; then
    . "$RUN_CONF"
fi
```

The runLockMan.conf file should include the following settings:

```
JAVA_OPTS="-server -Xms512m -Xmx512m -XX:MaxPermSize=128m
-XX:MaxNewSize=128m
-Dsun.rmi.dgc.server.gcInterval=3600000"
```

**Note**: The Configuration and Installation Manager (CIM) utility makes these configurations for you. See Configuration and Installation Manager (CIM) (page 10).

#### **Using ACC Scenarios in JBoss**

In order to create scenarios in the ATG Control Center (ACC), you must add the following three JAR files to your class path:

```
<JBdir>/common/lib/jboss-javaee.jar
<JBdir>/common/lib/jsp-api.jar
<JBdir>/common/lib/servlet-api.jar
```

To do this, copy the files into the /lib directory of your standalone ACC installation, then modify the bin/startClient.bat file to include the three JARs in the class path.

#### **Disabling Session ID Checking in JBoss**

If you are using JBoss on UNIX and expect to run multiple Oracle ATG Web Commerce servers within a single JBoss instance (as may be the case during development or demonstrations), edit the JBoss run.conf script by adding the following line to the end of the file:

```
JAVA_OPTS="${JAVA_OPTS} -Dorg.apache.catalina.connector.Request
.SESSION_ID_CHECK=false"
```

This allows your browser to use a single jsessionid cookie for both instances, avoiding unnecessary errors.

### **JBoss Installation Results**

The Oracle ATG Web Commerce installer does not alter your JBoss installation directory.

Note: Do not deploy multiple Oracle ATG Web Commerce application EAR files to a single JBoss server.

### **JBoss EAP 6**

There are some differences between JBoss EAP 6.0.1 and previous versions that impact the way that CIM installs and configures the application server.

#### Templates

JBoss EAP 6 uses a single XML configuration file, as opposed to multiple configuration files. When using CIM to install JBoss, you must select the JBoss EAP 6 template to use:

- standalone.xml The default configuration file for web profile configuration
- standalone-full.xml Java Enterprise Edition full profile configuration
- standalone-ha.xml The default profile with clustering capabilities
- standalone-full-ha.xml The Java Enterprise Edition full profile with clustering capabilities

This file, which contains all configuration information for the server instance, is copied and renamed with the name of the application server you entered during the installation. Refer to your JBoss EAP 6 documentation for additional information.

# **IBM WebSphere**

### WebSphere-Specific Requirements

The information in the following sections applies only to those using the WebSphere Application Server.

#### **Running WebSphere on AIX**

If using WebSphere on AIX, to avoid errors when importing application data using Oracle ATG Web Commerce import scripts, you must set the following in the <arglildir>/home/localconfig/postEnvironment.sh file:

JAVA\_ARGS="\${JAVA\_ARGS} -Djava.net.preferIPv4Stack=true"

You must also set it in the WebSphere environment:

 In WebSphere Admin, go to Servers > Application servers > server\_name > Java and Process Management > Process Definition > Java Virtual Machine. 2. Under Generic JVM arguments set the following:

-Djava.net.preferIPv4Stack=true

If you do encounter this problem, you will see errors such as the following:

Error: Jan 30, 2008 12:45:01 PM javax.jmdns.JmDNScloseMulticastSocket WARNING: closeMulticastSocket() Close socketexception java.net.SocketException: The socket name is not available on this system.

#### XA Data Sources on WebSphere

To use XA data sources with WebSphere, add the following line to your <ATG11dir>/home/servers/ servername/localconfig/GLOBAL.properties file:

localTransactionModeInitialization=false

#### **Creating ACC Scenarios on WebSphere**

In order to create scenarios in the ACC, you must add the <WSdir>/AppServer/j2ee.jar file to your class path. To do this, copy j2ee.jar into the /lib directory of your standalone ACC installation, then modify the bin/startClient.bat file to include j2ee.jar in the class path.

#### Using Oracle ATG Web Commerce Multisite on WebSphere

If you are using the Oracle ATG Web Commerce multisite feature on WebSphere, in order to use virtual context roots, do the following:

- 1. In WebSphere Admin, go to Server > Server Types > WebSphere application servers > *server\_name* > Web Container settings > Web Container > Custom Properties.
- 2. Set the com.ibm.ws.webcontainer.invokefilterscompatibility property to true.

In order for session recovery to function across a multisite installation, make sure to set the following properties as indicated on each application server:

- Enable URL rewriting Enabled
- · Enable protocol switch rewriting Enabled
- HttpSessionReuse True

#### WebSphere Installation Results

If you are not using CIM to configure your installation, you must manually register the Oracle ATG Web Commerce URL providers, following this procedure:

- 1. Copy protocol.jar from the <ATG11dir>\DAS\lib directory to the \lib directory of your WebSphere installation.
- 2. Register the following URL providers in the WebSphere Admin Console (see your WebSphere documentation), using the specified settings:

name = dynamosystemresource

```
streamHandlerClassName =
atg.net.www.protocol.dynamosystemresource.Handler
protocol = dynamosystemresource
name = appmoduleresource
streamHandlerClassName = atg.net.www.protocol.appmoduleresource.Handler
protocol = appmoduleresource
```

**Note**: The Configuration and Installation Manager (CIM) utility makes these configurations for you. See Configuration and Installation Manager (CIM) (page 10).

# 4 Configuring Databases and Database Access

Before deploying your application in a production environment, configure both your application server and ATG products to use a production-quality database management system such as Oracle or Microsoft SQL Server. Your applications can then access application server resources and ATG components simultaneously. For a list of supported databases, refer to the Oracle ATG Commerce Supported Environments Matrix document in the My Oracle Support knowledge base.

The Oracle ATG Web Commerce platform includes a number of tools for creating a database, adding and removing tables, configuring data sources and connection pools, and importing and exporting data. This chapter covers the following topics:

Creating Database Tables Using SQL Scripts (page 26)

**Destroying Database Tables (page 29)** 

Adding a JDBC Driver (page 32)

Configuring ATG Data Sources for Data Import (page 32)

Configuring Data Sources and Transaction Management (page 34)

Using Oracle ATG Web Commerce Products with an IBM DB2 Database (page 42)

Using Oracle ATG Web Commerce Products with an IBM DB2 Database (page 42)

Copying and Switching Databases (page 47)

Database Sorting for Localization (page 52)

**Note**: The Configuration and Installation Manager (CIM) utility creates database schemas and data sources for you. It also imports any required data. See Configuration and Installation Manager (CIM) (page 10).

**Note:** Changing Oracle ATG Web Commerce's out-of-the-box database schemas is not recommended, although you can extend the schemas as necessary. If you do make any schema changes, you must migrate the changes manually when you upgrade to a new version of Oracle ATG Web Commerce.

**Note:** JBoss comes with its own demo database, Hypersonic (note the data source hsqldb-ds.xml in the / deploy directory). Some JBoss component require that database, so do not remove it unless you also plan to remove those components.

# **Creating Database Tables Using SQL Scripts**

The following sections explain how to create database tables for the Oracle ATG Web Commerce Adaptive Scenario Engine and Oracle ATG Web Commerce Portal.

- Creating Database Tables for ATG Adaptive Scenario Engine (page 26)
- Creating Database Tables for ATG Portal (page 28)

See the installation documentation for your other Oracle ATG Web Commerce products for information on creating database tables required for those applications.

Note: If you are using a UTF-8 Oracle database, do the following before creating tables:

- Set the database character se (called NLS\_CHARACTERSET) to AL32UTF8.
- Set the system nls\_length\_semantics to char:

alter system set nls\_length\_semantics=char;

### **Creating Database Tables for ATG Adaptive Scenario Engine**

To create the database tables for the Oracle ATG Web Commerce Adaptive Scenario Engine, run the SQL scripts provided for the DAS, DPS, and DSS modules, as described in the following sections.

- Creating the DAS Tables (page 26)
- Creating the DPS Tables (page 27)
- Creating the DSS Tables (page 27)

**Oracle ATG Web Commerce Portal Note:** The table creation scripts for Oracle ATG Web Commerce Portal also create the tables for the Oracle ATG Web Commerce Adaptive Scenario Engine; you do not need to create the DAS, DPS, and DSS tables separately. See Creating Database Tables for ATG Portal (page 28) for details.

#### **Creating the DAS Tables**

To create the database tables in the DAS module, run the das\_ddl.sql script from the following directory:

<ATG11dir>/DAS/sql/install/database-vendor

The das\_ddl.sql script is derived from the subscripts listed in the table below. If necessary, you can run these subscripts individually from the following directory:

<ATG11dir>/DAS/sql/db\_components/database-vendor

Script Name	Purpose
create_gsa_subscribers_ddl.sql	Creates tables for event-listener registrations for distributed caching mode in the GSA
create_sds.sql	Creates a table for the switching data source service

Script Name	Purpose
create_sql_jms_ddl.sql	Creates tables for the Dynamo Message System
create_staff_ddl.sql	Creates the Dynamo Staff Repository for the GSA
dms_limbo_ddl.sql	Creates tables to store delayed JMS messages
id_generator.sql	Creates a table for managing ID spaces
integration_data_ddl.sql	Creates a table for storing caching information from the integration repository
nucleus_security_ddl.sql	Creates tables for Nucleus security data

### **Creating the DPS Tables**

To create the database tables for DPS, run the dps\_ddl.sql script from the following directory:

<ATG11dir>/DPS/sql/install/database-vendor

The  $dps_ddl.sql$  script is derived from the subscripts listed in the table below. If necessary, you can run these subscripts individually from the following directory:

<ATG11dir>/DPS/sql/db\_components/database-vendor

Script Name	Purpose
logging_ddl.sql	Creates tables for the logging and reporting subsystem
logging_init.sql	Initializes the logging and reporting tables
user_ddl.sql	Creates tables for the DPS schema

#### **Creating the DSS Tables**

To create the database tables for DSS, run the dss\_ddl.sql script from the following directory:

<ATG11dir>/DSS/sql/install/database-vendor

The  $dss_ddl.sql$  script is derived from the subscripts listed in the table below. If necessary, you can run these subscripts individually from the following directory:

<ATG11dir>/DSS/sql/db\_components/database-vendor

Script Name	Purpose
das_mappers.sql	Creates tables used by sample mappers to record Oracle ATG Web Commerce startup and shutdown events
dps_mappers.sql	Creates tables used by sample mappers to record DPS events
dss_mappers.sql	Creates tables used by sample mappers to record DSS audit trail events
scenario_ddl.sql	Creates tables for the DSS Scenario Engine

### **Creating Database Tables for ATG Portal**

The install file in the <ATG11dir>/Portal/install/database-vendor directory runs a set of scripts that create the required tables for the Portal Application Framework (PAF) and baseline gears.

Note: These scripts also create the tables for the Oracle ATG Web Commerce Adaptive Scenario Engine; you do not need to run the DAS, DPS, and DSS scripts separately. Note also that the install file uses the <arglidir>/ Portal/install/minimal-data.xml file to create the minimum set of data structures necessary to run Oracle ATG Web Commerce Portal.

Use the following syntax to run the install file appropriate for your DBMS:

- install-db2 userid password database
- install-mssql userid password host database
- install-oracle userid password database

**Note**: Entering passwords as command line arguments entails some security risks. Do not enter passwords as command line arguments in situations where security is a top priority.

The table creation scripts for Oracle ATG Web Commerce Portal are located in the following directories:

```
<ATG11dir>/Portal/paf/sql/install/database-vendor
<ATG11dir>/Portal/gear_dir/sql/install/database-vendor
```

**Note:** These scripts use an Oracle ATG Web Commerce-specific JTDatasource and TransactionManager, and cannot be used with your application server's data source or transaction manager.

Script Name	Purpose
alert_ddl.sql	Creates tables for the Alerts Gear
bookmarks_ddl.sql	Creates tables for the Bookmarks Gear
calendar_ddl.sql	Creates tables for the Calendar Gear
communities_ddl.sql	Creates tables for the Communities Gear
discussion_ddl.sql	Creates tables for the Discussion Gear
Script Name	Purpose
---------------------	--------------------------------------------------------------------------------
docexch_ddl.sql	Creates tables for the Document Exchange Gear
membership_ddl.sql	Creates tables for storing membership requests
paf_mappers_ddl.sql	Creates tables used by sample mappers to record portal events
Portal_ddl.sql	Creates tables for the Portal Application Framework
poll_ddl.sql	Creates tables for the Poll Gear
profile_ddl.sql	Creates tables for storing profile data for personalized communities and pages
soapclient_ddl.sql	Creates tables for the Web Services Client Gear

## **Destroying Database Tables**

The Oracle ATG Web Commerce platform includes SQL drop scripts for destroying database tables. (If you are using Oracle ATG Web Commerce Content Administration, see the *Content Administration Programming Guide* for information on destroying the database tables for your content administration server.) Run the drop scripts in the reverse of the order used for table creation.

This section covers the following topics:

- Destroying Database Tables for ATG Adaptive Scenario Engine (page 29)
- Destroying Database Tables for ATG Portal (page 31)

### **Destroying Database Tables for ATG Adaptive Scenario Engine**

This section covers the following topics:

- Destroying the DAS Tables (page 29)
- Destroying the DPS Tables (page 30)
- Destroying the DSS Tables (page 30)

### **Destroying the DAS Tables**

To destroy all DAS tables, run the drop\_das\_ddl.sql script from the following directory:

<ATG11dir>/DAS/sql/install/database-vendor

The drop\_das\_ddl.sql script is derived from the subscripts listed in the table below. If necessary, you can run these subscripts individually from the following directory:

<ATG11dir>/DAS/sql/uninstall/database-vendor

Script Name	Purpose
drop_dms_limbo_ddl.sql	Destroys the tables used to store delayed JMS messages
drop_gsa_subscribers_ddl.sql	Destroys the tables for event-listener registrations for distributed caching mode in the GSA
drop_id_generator.sql	Destroys the table for managing ID spaces
drop_integration_data_ddl.sql	Destroys the table that stores caching information from the integration repository
drop_nucleus_security_ddl.sql	Destroys the tables for Nucleus security data
drop_sds.sql	Destroys the table for the switching data source service
drop_sql_jms_ddl.sql	Destroys the tables for the Dynamo Message System
drop_staff_ddl.sql	Destroys the Dynamo Staff Repository for the GSA

### **Destroying the DPS Tables**

To destroy all DPS tables, run the drop\_dps\_ddl.sql script from the following directory:

<ATG11dir>/DPS/sql/install/database-vendor

The drop\_dps\_ddl.sql script is derived from the subscripts listed in the table below. If necessary, you can run these subscripts individually from the following directory:

<ATG11dir>/DPS/sql/uninstall/database-vendor

Script Name	Purpose
drop_logging_ddl.sql	Destroys the tables for the logging and reporting subsystem
drop_user_ddl.sql	Destroys the tables for the DPS schema

### **Destroying the DSS Tables**

To destroy all DSS tables, run the drop\_dss\_ddl.sql script from the following directory:

<ATG11dir>/DSS/sql/install/database-vendor

The drop\_dss\_ddl.sql script is derived from the subscripts listed in the table below. If necessary, you can run these subscripts individually from the following directory:

<ATG11dir>/DSS/sql/uninstall/database-vendor

Script Name	Purpose
drop_das_mappers.sql	Destroys the DAS sample mapper tables
drop_dps_mappers.sql	Destroys the DPS sample mapper tables
drop_dss_mappers.sql	Destroys the DSS sample mapper tables
drop_scenario_ddl.sql	Destroys the DSS Scenario Engine tables

### **Destroying Database Tables for ATG Portal**

The reset file in the <ATG11dir>/Portal/install/database-vendor directory runs a set of scripts that drop the database tables for Oracle ATG Web Commerce Portal. Use the following syntax to run the reset file appropriate for your DBMS:

- reset-db2 userid password database
- reset-mssql userid password host database
- reset-oracle userid password database

**Note:** Once you run the reset file, you must run the install file again to use your database with the Portal Application Framework. See Creating Database Tables for ATG Portal (page 28) for details.

**Note**: Entering passwords as command line arguments entails some security risks. Do not enter passwords as command line arguments in situations where security is a top priority.

The drop scripts for Oracle ATG Web Commerce Portal are located in the following directories:

```
<ATG11dir>/Portal/paf/sql/uninstall/database-vendor
<ATG11dir>/Portal/gear_dir/sql/uninstall/database-vendor
```

Note that the lines in these files that drop the DSS, DPS, and DAS tables are commented out by default as a safety measure. To drop those tables, uncomment the lines before running the script.

Script Name	Purpose
drop_alert_ddl.sql	Destroys tables for the Alerts Gear
drop_bookmarks_ddl.sql	Destroys tables for the Bookmarks Gear
drop_calendar_ddl.sql	Destroys tables for the Calendar Gear
drop_communities_ddl.sql	Destroys tables for the Communities Gear
drop_discussion_ddl.sql	Destroys tables for the Discussion Gear

Script Name	Purpose
drop_docexch_ddl.sql	Destroys tables for the Document Exchange Gear
drop_membership_ddl.sql	Destroys tables for storing membership requests
drop_paf_mappers_ddl.sql	Destroys tables used by sample mappers to record portal events
drop_portal_ddl.sql	Destroys tables for the Portal Application Framework
drop_poll_ddl.sql	Destroys tables for the Poll Gear
drop_profile_ddl.sql	Destroys tables for storing portal profile data
drop_soapclient_ddl.sql	Destroys the tables for the Web Services Client Gear

## Adding a JDBC Driver

To configure the Oracle ATG Web Commerce platform to use the JDBC driver for your DBMS, first install the driver software on your system as instructed by the manufacturer. See your application server documentation for information on where the driver should be installed.

**Oracle users:** Use the Oracle JDBC driver version that matches your Oracle server version. To find supported driver versions, refer to the Oracle ATG Commerce Supported Environments Matrix document in the My Oracle Support knowledge base.

**Oracle WebLogic users:** WebLogic ships with an ojdbc6.jar located at <WLdir>/wlserver\_10.3/server/ lib/. More recent Oracle drivers may be available, in which case you should make certain that your CLASSPATH refers to the latest version, not the shipped version.

## **Configuring ATG Data Sources for Data Import**

Oracle ATG Web Commerce uses its own data sources when running data import scripts. These scripts are used for initial application configuration. The data source is based on /atg/dynamo/service/jdbc/ JTDataSource, a Nucleus service that creates new connections to a particular database.

Your running Oracle ATG Web Commerce application uses your application server's native data sources (see Configuring Data Sources and Transaction Management (page 34) in this guide).

J2EE JDBC supports the Java Transaction API (JTA) via the javax.sql.XADataSource interface. JTA allows multiple resources from different providers to participate in the same transaction. Using two-phase commits, data integrity across different services is ensured. Oracle ATG Web Commerce supplies a DataSource that sits on top of an XADataSource and returns wrapped Connections that are registered appropriately with the associated Transaction Manager. Oracle ATG Web Commerce's DataSource must get all its Connections from an XADataSource. Only a true XADataSource produces connections that behave properly in a two-phase commit controlled by JTA. XADataSources should be included in JDBC 2.0 drivers for the various database vendors.

The default DataSource connection pool, JTDataSource, uses the FakeXADataSource component. Configure the desired connection pool properties, but note that this data source should be used only to run Oracle ATG Web Commerce data import scripts.

You can set up and configure a connection pool manually by creating two files in your <code>localconfig/atg/dynamo/service/jdbc/directory:</code>

- connectionPoolName.properties
- connectionPoolNameFakeXA.properties

where *connectionPoolName* is the name of the connection pool you want to create.

The *connectionPoolName*.properties file contains properties and values similar to the following:

```
$class=atg.service.jdbc.MonitoredDataSource
min=10
max=10
blocking=true
maxFree=-1
loggingSQLWarning=false
loggingSQLDebug=false
loggingSQLInfo=false
dataSource=/atg/dynamo/service/jdbc/<connectionPoolName>FakeXA
loggingSQLError=false
```

The min property determines the number of connections that the pool starts out with. The max property determines how many connections are to be kept around in the pool. When the pool starts, it immediately creates the minimum number of connections. Whenever a service requires a connection, it takes one from the pool. If there are no connections free, then the connection pool creates a new connection, until the maximum is reached. Due to various initialization calls, Oracle ATG Web Commerce requires at least three JDBC connections on install or when started with a new database. Setting the JDBC connection pool's max property to anything less causes Oracle ATG Web Commerce to hang when starting up.

If the maximum has been reached and a service requires another connection, then the service blocks until some other service frees up a connection. If the blocking property is set to false, then instead of blocking, the connection pool fails and results in a SQL exception.

The connectionPoolNameFakeXA.properties file contains properties and values similar to the following:

```
$class=atg.service.jdbc.FakeXADataSource
server=localhost:1313
user=admin
needsSeparateUserInfo=false
URL=jdbc:mysql://localhost:3306
readOnly=false
password=admin
database=
driver=com.mysql.jdbc.Driver
```

**Note**: Entering passwords in clear text files entails some security risks. Take steps to secure files that contain clear text passwords.

These properties tell the connection pool how to make a connection. The driver parameter specifies the name of the driver that should be used. The URL property specifies the name of the database server machine, the port of the database server (optional), and the name of the database on the server (optional). The format of the URL looks like this:

jdbc:driver name[:additional server information]

By default, the connection pool's driver and URL are configured for the MySQL database, as follows:

driver=com.mysql.jdbc.Driver URL=jdbc:mysql://localhost:3306/<db-name>

The user and password properties provide the connection with login access to the database, and must be recognized by the target database.

The readOnly property determines whether the resulting connection will be used only to perform read-only operations. Some drivers may be able to improve performance if this is true. Most applications require read and write access, so this property is usually false.

Oracle ATG Web Commerce wraps the Connection object to separate out SQL warning and info messages. This lets you see the SQL statements generated by Oracle ATG Web Commerce. It also catches SQLExceptions that occur on the connection and causes the connection to be closed when it is checked by into the resource pool. In addition to the standard Application Logging log levels (loggingError, loggingWarning, loggingInfo and loggingDebug), a monitored connection lets you split off the SQL log messages with these properties:

Property	Description
loggingSQLError	logs SQL exceptions as errors
loggingSQLWarning	logs SQL warnings received by the pool
LoggingSQLInfo	logs SQL statements sent by the pool
LoggingSQLDebug	logs JDBC method calls made by the pool

By default, Oracle ATG Web Commerce turns SQL warnings off since they tend to be informational messages, not warnings. If you want to log these messages, set <code>loggingSQLWarning</code> to true.

## **Configuring Data Sources and Transaction Management**

When you deploy your sites, you should reconfigure your installation to use the data sources and transaction manager that your application server uses. Data sources for all application servers should always use the READ\_COMMITTED isolation level (on DB2, use the equivalent CURSOR STABILITY).

### **Configuring Data Sources for JBoss**

Oracle ATG Web Commerce applications running on JBoss use a JTDataSource component, which should be configured to point to a JNDI reference to a DataSource component running in JBoss.

**Note**: The Configuration and Installation Manager (CIM) utility makes these configurations for you. See Configuration and Installation Manager (CIM) (page 10).

### Where to Configure JBoss Data Sources

You should configure your data source in the localconfig, jbossconfig, or equivalent named configuration layer. See "Managing Properties Files" in the *Platform Programming Guide* for information on application-server-specific and named configuration layers.

In order to use the jbossconfig directory:

Modify the MANIFEST.MF file for the given Oracle ATG Web Commerce module to include the following
property:

ATG-JbossConfig-Path: jbossconfig

• Create a jbossconfig directory and put the properties files there.

**Note:** If JBoss configuration files are stored in the ATG-3rdPartyConfig-Path layer, you might see errors if you start up applications on other application servers, because the data sources are configured to point to JNDI names that are not set up on that application server. Data source configuration files that are specific to JBoss should be in the ATG-JBossConfig-Path rather than the ATG-3rdPartyConfig-Path of those data source configurations.

### **Configuring New JBoss Data Sources**

**Note:** The following information does not apply to JBoss EAP 6. Refer to the Configuring MSSQL Data Sources on JBoss EAP 6 (page 36) for instructions.

To configure a new data source, go to the <JBdir>\server\_name\deploy\atg-ds.xml file. Edit the following configuration settings:

JNDI name URL driver class username password transaction isolation level connection pool numbers

See your application server documentation for information on the available parameters. For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<datasources>
    <xa-datasource>
        <jndi-name>atgcore_ds</jndi-name>
        <track-connection-by-tx>false</track-connection-by-tx>
        <isSameRM-override-value>false</isSameRM-override-value>
        <min-pool-size>5</min-pool-size>
        <max-pool-size>100</max-pool-size>
        <blocking-timeout-millis>5000</blocking-timeout-millis>
        <idle-timeout-minutes>15</idle-timeout-minutes>
        <transaction-isolation>TRANSACTION_READ_COMMITTED</transaction-isolation>
```

```
<xa-datasource-class>oracle.jdbc.xa.client.OracleXADataSource</xa-</pre>
datasource-class>
    <xa-datasource-property
name="URL">jdbc:oracle:thin:@myserver:1521:oral0r2</xa-datasource-
property>
    <xa-datasource-property name="User">username</xa-datasource-property>
    <xa-datasource-property name="Password">password</xa-datasource-property>
    <!-- Uncomment the following if you are using Oracle 9i
    <xa-datasource-property name="oracle.jdbc.V8Compatible">true</xa-</pre>
datasource-property>
   -->
    <exception-sorter-class-name>
        org.jboss.resource.adapter.jdbc.vendor.OracleExceptionSorter
    </exception-sorter-class-name>
  </xa-datasource>
</datasources>
```

**Note**: Entering passwords in clear text files entails some security risks. Take steps to secure files that contain clear text passwords.

If you have changed the JNDI name, you must also change the name configured in the <ATG11dir>/home/ localconfig/atg/dynamo/service/jdbc/JTDataSource.properties file:

\$class=atg.nucleus.JNDIReference
JNDIName=JNDIDataSourceName

For example, java:/ATGOracleDS.

Note that if you are using a WatcherDataSource, this would be configured instead in a DirectJTDataSource.properties file.

### **Adding Database Class Files**

If your database driver is located anywhere other than the server's lib directory (for example, C:\jboss\jbosseap-install-dir\jboss-as\server\atg\_server\lib), you must edit <JBdir>/bin/run.sh|bat and add your database class files, such as Oracle's ojdbc6.jar, to the JBoss CLASSPATH. To do this, search for \$JBOSS\_CLASSPATH and just above it, create a line:

JBOSS\_CLASSPATH=path\_to\_ojdbc6.jar

Rebuild and redeploy your EAR file.

### Configuring MSSQL Data Sources on JBoss EAP 6

When working with JBoss data sources previous to JBoss EAP 6, you registered data sources using the atgds.xml file, as outlined in the Configuring New JBoss Data Sources (page 35) section. For JBoss EAP 6, data sources are registered in the server instance configuration file, as outlined in the JBoss EAP 6 section.

The following is an example of a JBoss data source configuration:

```
<xa-datasource-class>com.mysql.jdbc.jdbc2.optional.MysqlXADataSource
          </xa-datasource-class>
   </driver>
 </drivers>
 <xa-datasource enabled="true" jndi-name="java:/ATGProductionDS"</pre>
     pool-name="mysql" use-java-context="true">
   <xa-datasource-property name="URL">jdbc:mysql://localhost:3306/
       production</xa-datasource-property>
     <driver>mysql</driver>
     <security>
       <user-name>user</user-name>
       <password>password</password>
     </security>
     <!-- To avoid deadlocks you need set this -->
     <transaction-isolation>TRANSACTION READ COMMITTED</transaction-isolation>
 </xa-datasource>
 <xa-datasource enabled="true" jndi-name="java:/ATGPublishingDS"</pre>
     pool-name="mysql" use-java-context="true">
   <xa-datasource-property name="URL">jdbc:mysql://localhost:3306/
       publishing</xa-datasource-property>
      <driver>mysql</driver>
     <security>
       <user-name>root</user-name>
        <password>admin</password>
     </security>
     <!-- To avoid deadlocks you need set this -->
     <transaction-isolation>TRANSACTION_READ_COMMITTED</transaction-isolation>
   </xa-datasource>
 </datasources>
</subsystem>
```

### **Configuring DB2 Data Sources on JBoss EAP 6**

When you configure DB2 data sources using JBoss EAP 6, you must use the JBC4 complaint data source driver. Ensure that you have set up your configuration using the db2jcc4.jar driver. Refer to your application documentation for additional information.

### Configuring Data Sources for WebLogic and WebSphere

To configure Oracle ATG Web Commerce to use data sources for WebSphere or WebLogic, override the default configuration of each Oracle ATG Web Commerce data source, replacing it with a pointer to a WebSphere or WebLogic data source.

For example, several Oracle ATG Web Commerce repositories use as their default data source the component / atg/dynamo/service/jdbc/JTDataSource, which is of class atg.service.jdbc.MonitoredDataSource. Rather than reconfiguring the repositories individually, replace the JTDataSource with a component of class atg.nucleus.JNDIReference, so that the "data source" that the repositories now point to is just a JNDI reference to a WebSphere or WebLogic data source. To do this, you create a JTDataSource.properties file that contains these lines:

```
$class=atg.nucleus.JNDIReference
JNDIName=java:comp/env/jdbc/ATGDatasource
```

where *ATGDatasource* is the JNDI name of the WebSphere or WebLogic data source. Put this file in <ATG11dir>/home/localconfig/atg/dynamo/service/jdbc/.

**Note**: The Configuration and Installation Manager (CIM) utility makes these configurations for you. See Configuration and Installation Manager (CIM) (page 10).

### **Configuring Data Sources for an Oracle RAC Cluster**

If you use Oracle ATG Web Commerce Content Administration, you must configure data sources for the destination repositories that are used during deployment to staging and production servers. These data sources require special configuration if the following conditions are true:

- The target site database is set up as an Oracle RAC cluster with multiple nodes.
- The target site runs on WebLogic or WebSphere.

In this case, you must configure an Oracle RAC cluster so that all operations within a given transaction are directed to a single cluster instance:

- 1. Set up a database service that runs on a single instance in the production RAC cluster.
- 2. This RAC cluster instance and its database service must be referenced by the data sources of the destination repositories that Content Administration uses for deployment. To do this, configure the data sources so their JDBC URL is set as follows:

jdbc:oracle:thin:@RAC-instance:port:dbservice

For detailed information about destination repositories and how they are used for deployment, see the *Content Administration Programming Guide*.

### **Setting the Transaction Timeout on JBoss**

The default JBoss transaction timeout is 300 seconds. This may be too short for your purposes, particularly if you have a large Oracle ATG Web Commerce catalog.

To increase the transaction timeout:

- 1. Go to the <JBdir>/server/server-name/deploy/transaction-jboss-beans.xml file.
- 2. Change the <attribute name="transactionTimeout">300</attribute> to a higher number.

### **Setting the Transaction Timeout on WebLogic**

WebLogic automatically rolls back transactions that do not complete in a certain number of seconds. The default setting is 30 seconds, which may be too short for compiling certain complex pages, especially pages that embed many page fragments.

When you are developing an application, a page must be recompiled each time you change it. If your application includes complex pages (particularly if you are developing a portal with Oracle ATG Web Commerce Portal), you can avoid transaction timeouts by raising the timeout setting to 600 seconds. Before deploying the application on a production site, you should pre-compile all of the pages. You can then lower the timeout setting.

To change the setting, open the WebLogic Server Console, go to the JTA page for the domain Oracle ATG Web Commerce is installed in, and change the value in the Timeout Seconds field. Oracle ATG Web Commerce recommends setting the timeout to 120 seconds.

### Setting the Transaction Timeout on WebSphere

WebSphere automatically rolls back transactions that do not complete in a certain number of seconds. The default setting is 120 seconds, which may be too short for compiling certain complex pages, especially pages that embed many page fragments.

When you are developing an application, a page must be recompiled each time you change it. If your application includes complex pages (particularly if you are developing a portal with Oracle ATG Web Commerce Portal), you can avoid transaction timeouts by raising the timeout setting to 600 seconds. Before deploying the application on a production site, you should pre-compile all of the pages. You can then lower the timeout setting.

To change the setting, go to Servers > Application Servers > server > Transaction Service in the console.

### Setting the Isolation Level for Transactions in WebSphere

Oracle ATG Web Commerce applications require a READ\_COMMITTED isolation level for transactions. The default isolation level in WebSphere using MS SQL and DB2 is REPEATABLE\_READ. To prevent deadlocks, use the WebSphere Administration Console to set the isolation level in the atg-bootstrap.war of your Nucleusenabled application to READ\_COMMITTED. See your WebSphere documentation for instructions.

### Direct SQL Deployment and Microsoft SQL Server

If you are using the direct SQL deployment feature of Oracle ATG Web Commerce and your database software is Microsoft SQL Server, add the following elements to the data source configuration files for each Oracle ATG Web Commerce server.

```
<new-connection-sql>SET XACT_ABORT ON</new-connection-sql>
<check-valid-connection-sql>select 1</check-valid-connection-sql>
```

See information about the direct SQL deployment feature in the Content Administration Programming Guide.

### **Data Source Debugging**

This section describes the use of the WatcherDataSource class to debug data source problems. This feature is automatically available for all application servers.

### **Using Data Source Debugging**

The default JTDataSource allows you to monitor and log data source information for debugging purposes. It does this using the WatcherDataSource class. A WatcherDataSource "wraps" another data source, allowing debugging of the wrapped data source. For example:

```
/atg/dynamo/service/jdbc/JTDataSource.properties
$class=atg.service.jdbc.WatcherDataSource
# The actual underlying DataSource.
dataSource=/atg/dynamo/service/jdbc/DirectJTDataSource
```

**Note:** Due to the potential performance impact, the features described here should be used only for debugging in a development environment. Do not use data source logging in a production environment unless absolutely necessary.

To view all logged data from the WatcherDataSource, go to /atg/dynamo/service/jdbc/JTDataSource in the Dynamo Component Browser.

### WatcherDataSource Configuration

The default WatcherDataSource configuration is:

```
showOpenConnectionsInAdmin=false
logDebugStacktrace=false
loggingDebug=false
monitored=false
loggingSQLError=true
loggingSQLWarning=false
loggingSQLInfo=false
loggingSQLDebug=false
```

This default configuration logs the following information:

- currentNumConnectionsOpen
- maxConnectionsOpen
- numGetCalls
- averageGetTime
- maxGetTime
- numCloseCalls
- averageCloseTime
- maxCloseTime
- averageOpenTime
- maxOpenTime

For additional debugging information, you can set the following properties to true:

showOpenConnectionsInAdmin—Lists currently open connections, along with the amount of time they have been held open and the thread that is holding them open. This information is useful for identifying Connection leaks. If logDebugStacktrace is also true, then stack traces are displayed as well.

**Note:** This momentarily prevents connections from being obtained or returned from the data source, and severely affects performance.

- loggingDebug—Logs debug messages on every getConnection() and close() call. These messages include interesting information such as sub-call time, number of open connections, and the calling thread. If logDebugStacktrace is also true then a stack trace is logged as well.
- logDebugStacktrace—Creates stack traces on each getConnection() call. This allows the calling code to be easily identified, which can be useful when trying to find Connection leaks, code that is holding Connections open for too long, or code that is grabbing too many Connections at a time.

Note: This is done by generating an exception, which affects performance.

• monitored—Gathers additional connection statistics and SQL logging.

## **Using the JDBC Browser**

The Dynamo Server Admin includes a JDBC Browser (http://hostname:port/dyn/admin/atg/dynamo/ admin/en/jdbcbrowser/) that enables you to examine the metadata of a database, including a listing of the tables, columns, and supported data types. The JDBC Browser also allows you to create tables, drop tables, execute queries, and examine the results of those queries.

All these operations are performed on a generic JDBC driver connection, meaning that the JDBC Browser should work with all database for which a JDBC driver exists.

### **Configuring the JDBC Browser**

The JDBC Browser obtains its JDBC connections from a JDBC connection pool service. By default, the service is set to the standard connection pool at /atg/dynamo/service/jdbc/JTDataSouce. This connection pool determines which JDBC driver and database to use.

If you want the JDBC Browser to use a different connection pool, modify the connectionPool property of / atg/dynamo/admin/jdbcbrowser/ConnectionPoolPointer so that it points to the desired connection pool service, using the following form:

/atg/dynamo/service/jdbc/your-pool-name

### **Create Table Operation**

The Create table page provides a simple way for you to define a table and create it in the database. You can fill in the names and types of up to 10 columns in the table (any columns you leave blank will not be put into the table). The column types are expressed in JDBC types, which may or may not correspond directly to your database's data types.

The Nullable, Unique, and Primary Key flags indicate properties of the column. You'll have to be careful to avoid illegal combinations; for example, most databases do not allow a primary key to be nullable.

The Additional Constraints are passed straight through to the CREATE TABLE statement. This allows you to enter additional constraints, such as foreign keys or indices.

### **Drop Table Operation**

The Drop table page drops the table you name.

### **Execute Query Operation**

The Execute query page allows you to enter an arbitrary SQL statement that is passed through the driver to the database. The results of the statement are displayed in response. If the statement generates multiple result sets and update counts, all of those result sets and update counts will be displayed.

The flag marked Show resulting column headings in long form indicates whether extra result set metadata should be shown with each column. This tends to be rather extensive and is probably not necessary for most operations.

When you submit the query, you can submit with a commit or submit with a rollback. These options are only meaningful if autoCommit is false. If autoCommit is true, then the query will always be followed by a commit. The autoCommit property is set in the connection pool service.

### **Metadata Operations**

All JDBC drivers provide metadata about the database that can be accessed through the JDBC interface. This metadata includes runtime information such as a listing of the tables and columns in the database. The metadata can also include information about the specific dialect of SQL that is understood by the database.

The JDBC Browser allows you to access all the metadata provided by the JDBC driver. Each of the metadata operations asks you to provide parameters for the requested metadata. For example, the List tables operation asks for a Catalog Name, Schema Name, and Table Name. You can leave these fields blank, in which case all the appropriate metadata is returned.

# Using Oracle ATG Web Commerce Products with an IBM DB2 Database

To use a DB2 database, you must set the parameterizedSelect and useSetBinaryStream properties of the /atg/dynamo/messaging/SqlJmsProvider component to false.

In order for some import scripts to work, you must also set the following in your <ATG11dir>/home/localconfig/GLOBAL.properties file:

```
handleRangesInMemory=true
localTransactionModeInitialization=false
```

Create at least three table spaces and buffer pools: one tablespace/buffer pool with a page size of 4KB, one with a page size of 16KB, and one with a page size of 32KB. See your DB2 documentation for more information. Oracle ATG Web Commerce recommends that you create more than one table space in each size; the numbers vary depending on your data.

The db2\_jms\_procedures\_ddl.sql file contains procedures that set the msgPollBatchSize property of SqlJmsProvider. The dms\_topic\_flag and dms\_queue\_flag procedures set a fixed batch size of 5000 (unlike Oracle or MSSQL, DB2 does not compute the batch size, but uses a fixed number).

If you find that the 5000-item configuration is not effective, you can change the setting and recompile the procedures using the following statements:

db2 connect to db2\_alias user <code>schema\_owner\_name</code> using <code>password</code> db2 <code>-td@ -v</code> <code>-ffilename</code>

### For example,

db2 -td@ -v -fdb2\_jms\_procedures\_ddl.sql > db2\_jms\_procedures\_ddl.log

**Note**: Entering passwords as command line arguments entails some security risks. Do not enter passwords as command line arguments in situations where security is a top priority.

For web-based applications such as Oracle ATG Web Commerce, the recommended isolation level READ\_COMMITTED. On Oracle and MSSQL, non-modifying transactions are allowed to read data while another transaction commits. DB2's treatment of this isolation level is different in two ways: first, it calls the isolation level CURSOR STABILITY, and second, it locks exclusively on a table that is being modified, preventing other transactions from reading data from those tables.

To modify DB2 so that it behaves the same way with CURSOR STABILITY that Oracle and MSSQL behave with READ COMMITTED, create the following registry entries for your database:

DB2\_EVALUNCOMMITTED Do not wait for uncommitted updates.

DB2\_SKIPINSERTED Do not wait for uncommitted inserts.

DB2\_SKIPDELETED Do not wait for uncommitted deletes.

## Using Oracle ATG Web Commerce Products with a Microsoft SQL Server Database

Oracle ATG Web Commerce products do not support Unicode for MS SQL Server databases. To use Microsoft SQL Server with Oracle ATG Web Commerce products, be sure the useSetUnicodeStream property of all SQL repository components is set to false (default). To ensure that no Oracle ATG Web Commerce components are configured to use useSetUnicodeStream, you can set this property in your localconfig/GLOBAL.properties file:

useSetUnicodeStream=false

If you are creating localized content, set the useSetAsciiStream property to false in your localconfig/ GLOBAL.properties file:

useSetAsciiStream=false

If you are using the Microsoft SQL Server 2005 JDBC driver, you must set sendStringParametersAsUnicode to false in your URL connection string. For example:

URL=jdbc:sqlserver://<SERVER>:<PORT>;databaseName=<DATABASE>; sendStringParametersAsUnicode=false

The sendStringParametersAsUnicode=false setting avoids Unicode character conversion and enables MS SQL Server to use indexes in queries.

In addition, to prevent deadlocks and timeout problems, you must turn on READ\_COMMITTED\_SNAPSHOT. For example:

ALTER DATABASE <database\_name> SET READ\_COMMITTED\_SNAPSHOT ON;

### Using iNet (Merlia) Drivers

If you are using iNet drivers on JBoss, bear in mind that this driver does not allow for passing information by URL; therefore, some additional information must be set in the property fields, as shown in this example:

```
<xa-datasource>
  <jndi-name>ATGProductionDS</jndi-name>
  <track-connection-by-tx/>
  <isSameRM-override-value>false</isSameRM-override-value>
  <min-pool-size>5</min-pool-size>
  <max-pool-size>100</max-pool-size>
  <blocking-timeout-millis>5000</blocking-timeout-millis>
 <idle-timeout-minutes>15</idle-timeout-minutes>
 <transaction-isolation>TRANSACTION_READ_COMMITTED</transaction-isolation>
  <xa-datasource-class>com.inet.tds.DTCDataSource</xa-datasource-class>
 <xa-datasource-property name="ServerName">server_name</xa-datasource-property>
  <xa-datasource-property name="DatabaseName">database_name</xa-datasource-</pre>
property>
  <xa-datasource-property name="User">database_username</xa-datasource-property>
  <xa-datasource-property name="Password">database_password</xa-datasource-</pre>
property>
 <xa-datasource-property name="Mode">71</xa-datasource-property>
 <!-- sql to call when connection is created -->
 <new-connection-sql>select 1</new-connection-sql>
  <!-- sql to call on an existing pooled connection when it is obtained from
<--> 100g
  <check-valid-connection-sql>select 1</check-valid-connection-sql>
  <!-- corresponding type-mapping in the standardjbosscmp-jdbc.xml -->
  <metadata>
    <type-mapping>MS SQLSERVER2000</type-mapping>
  </metadata>
  </xa-datasource>
```

**Note**: Entering passwords in clear text files entails some security risks. Take steps to secure files that contain clear text passwords.

If you are using iNet drivers with WebLogic, when you create your data source, use the following settings:

- The type should be DataDirect's MSSQL type 4 XA.
- Set the following properties:
  - url—The full connection string for your data source.
  - driver—The driver name is com.inet.tds.DTCDataSource.
  - user—User name for the database account.
  - port—Connection port used for the database.
  - mode—This should normally be set to 71, as Unicode is not supported for MS SQL.
  - serverName—The machine name of the database host.

secureLevel—Set this to 0 if you are not using SSL. If you are using SSL, see your database documentation for information.

### Using iNet Merlia Driver with JBoss EAP 6

When configuring an iNet Merlia driver for an MSSQL data source with JBoss EAP 6, you need to manually identify the correct driver for each data source. During deployment, JBoss creates two different drivers for Merlia, based on the name of the JAR file, the iNet driver class names and the version of the driver. To register your data source, you must provide the JBoss internal driver name:

- 1. Start JBoss server and deploy the driver (the Merlia.jar.dodeploy driver is deployed automatically when the server starts).
- 2. Open the JBoss Management Console at http://server-name:port and then clicking on Management Console.
- 3. Select Profile > Datasources
- 4. Select Add to add a new data source.
- 5. Enter a data source name and a JNDI name.
- 6. Select Next. JBoss displays the option to choose a driver for the data source and displays a list of installed drivers. Note that there are two different Merlia.jar drivers., which are based on the class type and version of the driver.
  - .com.inet.tds.TdsDataSource The JDBC 1 data source , a simple data source.
  - .com.inet.tds.PDataSource A pooled data source.
- 7. Identify the appropriate driver name to add to the server's configuration file.

Once you have retrieved the driver, open the server's configuration file and enter the correct driver information as outlined in the Configuring MSSQL Data Sources on JBoss EAP 6 (page 36) section. For example:

```
<xa-datasource enabled="true" jndi-name="java:/ATGProductionDS"</pre>
   pool-name="ATGProductionDS" use-java-context="true">
 <xa-datasource-property name="ServerName">localhost</xa-datasource-property>
 <xa-datasource-property name="DatabaseName">prod</xa-datasource-property>
 <xa-datasource-property name="User">sa</xa-datasource-property>
 <xa-datasource-property name="Password">password</xa-datasource-property>
 <xa-datasource-property name="Mode">71</xa-datasource-property>
 <xa-datasource-class>com.inet.tds.DTCDataSource</xa-datasource-class>
 <new-connection-sql>select 1</new-connection-sql>
 <driver>Merlia.jarcom.inet.tds.TdsDriver_8_0</driver>
 <security>
   <user-name>sa</user-name>
   <password>NB129T</password>
 </security>
 <xa-pool>
   <min-pool-size>5</min-pool-size>
   <max-pool-size>100</max-pool-size>
   <is-same-rm-override>false</is-same-rm-override>
 </xa-pool>
 <timeout>
   <blocking-timeout-millis>5000</blocking-timeout-millis>
   <idle-timeout-minutes>15</idle-timeout-minutes>
```

```
</timeout>
<validation>
<exception-sorter class-name="org.jboss.resource.adapter.jdbc.
vendor.DB2ExceptionSorter"/>
<!-- sql to call on an existing pooled connection when it is obtained from
pool -->
<check-valid-connection-sql>select 1</check-valid-connection-sql>
</validation>
<!-- To avoid deadlocks you need set this -->
<transaction-isolation>TRANSACTION_READ_COMMITTED</transaction-isolation>
</xa-datasource>
```

# Moving Data from a Development Database to the Production Database

If you want to move data from your development database to the database used by your application server, you can do this using the startSQLRepository script. This script is described in detail in the *Repository Guide*. To use this script, follow these steps:

- 1. Set the DYNAMO\_HOME environment variable to <ATG11dir>/home.
- 2. In the Dynamo Server Admin, create a new Oracle ATG Web Commerce server that uses data sources that point to the development database. This is the default configuration for a new server.
- 3. Use the startSQLRepository script to export data from the development database. Include in the startSQLRepository command the -s servername switch. For example, if the server you created in the previous step is called server1, you can export the data from all of the Oracle ATG Web Commerce repositories using this command:

```
bin/startSQLRepository -s server1 -exportRepositories all all.xml
```

4. Use the startSQLRepository script to import data from the XML file (created in the previous step) into the database used by your application server. Use the -s switch to specify a Oracle ATG Web Commerce server that is configured to use an Oracle ATG Web Commerce data source that points to that database. For example:

```
bin/startSQLRepository -s server_name -import all.xml
```

Note that the Oracle ATG Web Commerce data source must use an Oracle ATG Web Commerce-supported database driver. To see a list of supported database drivers refer to the Oracle ATG Commerce Supported Environments Matrix document in the My Oracle Support knowledge base.

**Oracle users:** Before importing the demo data, set the useSetCharacterStream property of all SQL repository components to true so that non-8859 characters are displayed correctly. You can set this property in your localconfig/GLOBAL.properties file:

useSetCharacterStream=true

**Microsoft SQL users:** In order to run the Oracle ATG Web Commerce demos with a Microsoft SQL database, you must configure the database to be case-sensitive. See your MSSQL documentation for instructions. Note that the Quincy Funds demo is not supported for MSSQL.

**Note**: Entering passwords as command line arguments entails some security risks. Do not enter passwords as command line arguments in situations where security is a top priority.

### **Transferring the Demo Data**

Use the commands in the following tables to transfer the Quincy Funds data from the development database to the your production database.

### Exporting the Demo Data from the Development Database

Use the command below (on one line, with no line breaks) to export the demo data from the development database to an XML file called all.xml.

Demo Application	Command
Quincy Funds	bin/startSQLRepository -s <i>server_name</i> -m DSSJ2EEDemo -exportRepositories all all.xml

### Importing the Demo Data to the Production Database

Use the command below (on one line, with no line breaks) to import the data contained in all.xml to the database used by your application server.

Demo Application	Command
Quincy Funds	bin/startSQLRepository -s <i>server_name-m</i> DSSJ2EEDemo -import all.xml -repository

## **Copying and Switching Databases**

In most situations, make database changes on an offline copy of the database, rather than on the database that runs your live site. Making changes on the live site can cause errors or inconsistencies or might adversely affect the performance of your live site. The Oracle ATG Web Commerce platform includes copying and switching functionality that lets you copy databases, using the database vendor's native bulk copy tools, and switch your live site between two different databases.

This section includes the following topics:

- Database Copy Operations (page 48)
- Creating a DBCopier Component (page 48)
- Configuring the DBConnectionInfo (page 49)
- Configuring the DBCopier (page 49)

- Setting the Native SQL Environment (page 50)
- Switching Databases (page 51)
- Configuring a SwitchingDataSource (page 51)
- Database Switching and Query Caching (page 52)

For information about using the database Copy and Switch features for Oracle ATG Web Commerce, see the *Commerce Programming Guide*.

### **Database Copy Operations**

The procedure for copying a database includes three basic steps:

- exporting data out of the source database to an OS file
- deleting any data in the destination database
- · importing data into the destination database from the OS file

The base class for the Oracle ATG Web Commerce database copying facility is atg.adapter.gsa.DBCopier.lt is important to note that DBCopiers use vendor-specific bulk copy and SQL utilities for speed, rather than using JDBC. This is accomplished by executing these commands in separate processes.

If the native bulk copy program operates on one table at a time, the DBCopier imports table data in the order in which the tables are specified and deletes table data in the reverse order. Thus, if there are foreign key constraints among the tables, the copy operation can still work if the tables are specified in dependency order. The various subclasses of DBCopier implement copying for different database vendors, using different vendor tools.

To use a DBCopier, follow these steps:

- 1. Create a DBCopier component. See Creating a DBCopier Component (page 48).
- 2. Configure DBConnectionInfo components for your source and destination databases. See Configuring the DBConnectionInfo (page 49).
- 3. Configure the DBCopier component as described in Configuring the DBCopier (page 49).
- 4. Set the SQL environment variables as described in Setting the Native SQL Environment (page 50).
- 5. Run the DBCopier by invoking its copy() method. For an example, see the Commerce Programming Guide.

### **Creating a DBCopier Component**

The class from which you instantiate the DBCopier depends on the database you are using. The following are subclasses of atg.adapter.gsa.DBCopier and are in package atg.adapter.gsa:

DBCopier	Vendor	Vendor Program
BcpDBCopier	Microsoft	Вср

DBCopier	Vendor	Vendor Program
DB2DBCopier	IBM	export/import
OracleDBCopier	Oracle	exp/imp

For more information about the DBCopier subclasses, see the ATG Platform API Reference.

### **Configuring the DBConnectionInfo**

The connection information about the source database (the database you are copying from) and the destination database (the database you are copying to) is maintained in a component of type atg.adapter.gsa.DBConnectionInfo.Create a DBConnectionInfo for each database and configure it with the following information:

Property	Description
Server	The name of the database server
User	A valid username to connect to the database
Password	A valid password for the username specified by the user property

**Note:** The DBConnectionInfo settings are not expressed in JDBC terms. The settings are the values of the connection parameters used by OS tools (such as bcp) when connecting to the specified database.

**Note**: Entering passwords in clear text files entails some security risks. Take steps to secure files that contain clear text passwords.

### **Configuring the DBCopier**

Set the following properties of the DBCopier:

Property	Description
Source	The DBConnectionInfo that service holds connection information for the database to copy from.
Destination	The DBConnectionInfo that service holds connection information for the database to copy into.
Tables	A comma-separated list of the names of the tables in the source database to be copied. If the native bulk copy program operates on one table at a time, the DBCopier imports table data in the order in which the tables are specified and deletes table data in the reverse order.

Property	Description
Directory	The name of a scratch directory for SQL and data files used during the copy operation. This directory must exist before the copy is launched. It is strongly recommended that no other processes or facilities use this scratch directory, especially other DBCopier instances.
CleanupDirectory	Set this to true to delete the files in the scratch directory after the copy is performed. Defaults to false.

In addition to the above properties, which are common to all DBCopier classes, each of the DBCopier subclasses has the following properties you may need to configure.

### **BcpDBCopier**

This DBCopier for MSSQL databases uses the bcp utility for copy data. Generally, you can use this copier with the default property settings, with one exception. You should set the BcpDBCopier's maxTextOrImageSize property to a value no smaller than the largest text or image column in the tables being copied. See your Microsoft documentation for details.

### **DB2DBCopier**

This DBCopier for DB2 databases uses the DB2 export and import utilities. If you are running the DB2DBCopier on UNIX or any other operating system that uses "/" as a directory separator, set the useUnixStyleDir property of the DB2DBCopier component to true. If "\" is the directory separator, set the useUnixStyleDir to false. The DB2 export utility wants to store binary objects in their own files, so make sure that the directory property points to a location in which these files can be stored temporarily. See your DB2 documentation for details.

### OracleDBCopier

This class is a DBCopier for Oracle databases. This copier uses the Oracle exp and imp utilities. You can configure OracleDBCopier to use direct path for exporting. To enable direct path for exporting, set the useDirectPathForExport property of the OracleDBCopier to true. This property is false by default.

See your Oracle documentation for more information on using direct path with the exp utility.

### **Setting the Native SQL Environment**

DBCopier components use vendor-specific bulk copy and SQL utilities for speed, rather than using JDBC. Therefore, to use a DBCopier, the native SQL environment for the database in question must be set up **before** starting your Oracle ATG Web Commerce application. This is required by the vendor tools in the database software. To use a DBCopier component, you must set up the environment in which the JVM runs as specified in the database vendor documentation. You can add this environment information to your <a href="https://www.atglidin.com">atglidin.com</a> home/localconfig/environment.sh or environment bat file. For information about the settings for your database, see the documentation from your database vendor.

For example, for Oracle you should set your environment up to look something like this:

```
ORACLE_HOME=/oracle-directory
PATH=$PATH:$ORACLE_HOME/bin
ORACLE_SID=ora8
```

### **Switching Databases**

In many database-dependent applications, you may want to make changes in an offline database and then switch over your live application so that the inactive database becomes the live database. Oracle ATG Web Commerce's switching facility is based on a class named atg.service.jdbc.SwitchingDataSource. You can use a SwitchingDataSource in place of a regular data source (such as atg.service.jdbc.MonitoredDataSource). The SwitchingDataSource can be switched between two or more underlying DataSources. All DataSource method calls are passed through to the DataSource specified by the currentDataSource property of the SwitchingDataSource. Note that each DataSource that the SwitchingDataSource points to must be of class atg.nucleus.JNDIReference, with a JNDIName property that points to an application server data source. See Configuring Data Sources and Transaction Management (page 34) for more information.

The switching database is meant to complement the DBCopier components. For example, if you are using Oracle ATG Web Commerce, you would update an inactive database, switch your live site to that database, then copy the currently-active database to the inactive database using the database vendor's native bulk copy tools.

Note: Unlike DBCopier, Oracle ATG Web Commerce's switching facility is a JDBC mechanism.

To set up and use a database switching service:

- 1. Configure DataSources that connect to your live and inactive databases.
- 2. Configure a SwitchingDataSource component, as described in Configuring a SwitchingDataSource (page 51).
- 3. Configure the Repository components that use the DataSources to point to the SwitchingDataSource. Also set the Repository components' selectiveCacheInvalication property (see Configure Selective Cache Invalidation in the Content Administration Programming Guide).

**Important:** If you have multiple independent Oracle ATG Web Commerce clusters that share a single SDSRepository, make sure each cluster uses a unique set of SwitchingDataSource names. Otherwise, the clusters interfere with each other during the switching process.

### Configuring a SwitchingDataSource

Set the following properties of the SwitchingDataSource component:

Name	Description
initialDataSourceName	The short name for the DataSource that should be used for the currentDataSource on the very first run. On subsequent runs, the initial currentDataSource is obtained from the state recorded in the SDSRepository.
dataSources	Set to a ServiceMap of DataSources. This property maps short names of DataSources to their Nucleus component path. The following example shows how you might set the dataSources property:
	dataSources=FirstDataSource=\
	/atg/dynamo/service/jdbc/FirstDataSource,\
	SecondDataSource=\
	/atg/dynamo/service/jdbc/SecondDataSource

Name	Description
repository	Set with a reference to /atg/dynamo/service/jdbc/SDSRepository.
	This refers to the switching data source repository, which keeps track of which database the switching data source points to at any time.

This sample shows the default format of the switching data source used by the product catalog in Oracle ATG Web Commerce:

```
$class=atg.service.jdbc.SwitchingDataSource
#
# A map from data source names to data sources
#
dataSources=\
    DataSourceA=/atg/commerce/jdbc/ProductCatalogDataSourceA,\
    DataSourceB=/atg/commerce/jdbc/ProductCatalogDataSourceB
#
# The name of the data source that should be used on startup
#
initialDataSourceName=DataSourceA
repository=/atg/dynamo/service/jdbc/SDSRepository
```

**Note**: The Configuration and Installation Manager (CIM) utility makes these configurations for you. See Configuration and Installation Manager (CIM) (page 10).

### **Database Switching and Query Caching**

If you are using a GSA repository and set the cacheSwitchLoadQueries property of the GSAItemDescriptor to true, the query cache is loaded for a cache switch. If false, the query cache starts out empty after a cache switch.

## **Database Sorting for Localization**

The order in which records are sorted in your Oracle ATG Web Commerce databases may control the order in which the corresponding data is presented in user interfaces. If you are configuring Oracle ATG Web Commerce for use in different language locales, consider how the sorting order of your database records affects user interfaces.

For example, if you enter non-ASCII characters in the names of content administration projects, those project names will not appear in the expected order unless you configure your database with a sorting order that corresponds to the language used.

Find information about setting the character set and multilingual sort order in the documentation for your database software.

## **5** Platform Installation Details

This section provides general information about installing the Oracle ATG Web Commerce platform.

## **File System Permissions**

Make sure the file system permissions for all Oracle ATG Web Commerce application files have the most restrictive settings possible. Do not allow users other than the user account for the application itself and the system administrator to read, write, or execute application files.

For example, if you are using a UNIX or Linux operating system, configure a dedicated user account for your Oracle ATG Web Commerce applications. Set the file permissions for the files created by that user account so that other users cannot read, write, or execute them. To do this, set the umask configuration for the user account to 077.

## **Performing a Maintenance Installation**

If you have any of the Oracle ATG Web Commerce platform products installed and would like to install additional platform products, rerun the Oracle ATG Web Commerce setup program. The maintenance installer lists the products that have not been installed yet, allowing you to select the ones you want. If you need to reinstall any of the Oracle ATG Web Commerce platform products that are currently installed on your system, you must uninstall the Oracle ATG Web Commerce platform completely (see Removing the Oracle ATG Web Commerce Platform from Your System (page 55)) and run the setup program again.

**Note:** If you have installed any Oracle ATG Web Commerce patches, you must uninstall them before running the maintenance installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install is complete, reinstall the patches. See the PatchReadme files under <a href="https://www.maintenance">wmmaintenance</a> installer. Once the maintenance install the patches. See the Patch

## **Default Ports**

This guide uses the *hostname:port* convention in URLs. The default HTTP ports for the application servers are:

• JBoss: 8080

- WebLogic: admin server 7001
- WebSphere: 9080

## Sun T1000 and T2000 Requirements

By default the Sun T1000 and T2000 systems run a server that uses port 9010. The Oracle ATG Web Commerce lock management components also use this port. If you are using lock management, you must either disable the server or change your lock manager to use a different port.

To disable the server:

- 1. Log in as root.
- 2. Enter the following command:

```
mv /etc/rc2.d/S95IIim /etc/rc2.d/K95IIim
```

3. Stop the service:

/etc/rc2.d/S95IIim stop

To change Oracle ATG Web Commerce lock manager port assignments, when you configure your lock management components, use the following settings:

1. For the ClientLockManager port assignment in <ATG11dir>/home/localconfig/atg/dynamo/service /ClientLockManager.properties:

useLockServer=true lockServerPort=39010

2. For the ServerLockManager port assignment in <ATG11dir>/home/servers/servername/ localconfig/atg/dynamo/service /ServerLockManager.properties:

port=39010

See the Locked Caching section of the Repository Guide for information on configuring lock managers.

## Installing the ATG Control Center on a Client Machine

This section explains how to install a standalone version of the ATG Control Center (ACC) on a client machine, when you do not need a full Oracle ATG Web Commerce installation. It covers the following topics:

- Downloading the ACC Installer (page 55)
- Installing the ACC on a Windows Client (page 55)
- Installing the ACC on a UNIX Client (page 55)

Note: To use the standalone version of the ACC, the client machine must have the J2SDK installed.

### **Downloading the ACC Installer**

Contact your Oracle ATG Web Commerce sales representative to obtain one of the following ACC distribution files:

- ACC11.0.exe (Windows)
- ACC11.0.bin(UNIX)

Note: You cannot use any other version of the ACC with Oracle ATG Web Commerce 11.

### Installing the ACC on a Windows Client

To install the ACC on a Windows client:

- 1. Run the ACC11.0.exe file to start the setup program.
- 2. After you accept the terms of the license agreement, select the destination folder for the ACC. The default is C:\ATG\ACC11.0. Click **Browse** to specify a different directory.
- 3. Enter a name for the ACC program folder on the Windows Start menu.
- 4. The installer displays the settings you selected. Review the setup information and click **Next** to start the installation, or **Back** to change any of the settings.

### Installing the ACC on a UNIX Client

To install the ACC on a UNIX client:

- 1. Change the permissions on the downloaded installer so you can execute it.
- 2. Run the binary:
  - ./ACC11.0.bin
- 3. Accept the license agreement.
- 4. Provide an install directory.

When finished, exit the installer.

## Removing the Oracle ATG Web Commerce Platform from Your System

Use the following methods to remove the Oracle ATG Web Commerce platform from your system.

On Windows: Use the Add/Remove Programs function in the Windows Control Panel.

**On UNIX:** Go to the <ATG11dir>/uninstall/.ASE11.0\_uninstall directory and run Uninstall\_ATG\_11.0.

## 6 Running Nucleus-Based Applications

Nucleus-based applications are assembled into EAR files that include both the application and Oracle ATG Web Commerce platform resources, and which are then deployed to your application server. The Oracle ATG Web Commerce platform installation includes the modules required to create QuincyFunds.ear, a sample J2EE application that includes the Quincy Funds demo and the Dynamo Server Admin.

Once the Oracle ATG Web Commerce installation is complete, you can assemble, deploy, and run the QuincyFunds . ear application. You can then access the Quincy Funds demo and the Dynamo Server Admin through your web browser, and connect to the application with the ACC.

This chapter covers the following topics:

Starting the SQL JMS Admin Interface (page 57)Starting Oracle ATG Web Commerce Web Services (page 58)Connecting to the Dynamo Server Admin (page 58)Starting the ATG Control Center (page 59)Stopping an Oracle ATG Web Commerce Application (page 63)

## **Starting the SQL JMS Admin Interface**

The Oracle ATG Web Commerce platform includes a browser-based administration interface for its SQL JMS message system. This interface makes it easy to view, add, and delete SQL JMS clients, queues, and topics. To use the SQL JMS Admin interface, include the SQLJMSAdmin module in your application.

To access the interface, point your browser to the following URL:

http://hostname:port/sqlJmsAdmin

To learn more about the SQL JMS system, see the Platform Programming Guide.

## **Starting Oracle ATG Web Commerce Web Services**

The Oracle ATG Web Commerce platform includes a number of preconfigured web services that provide remote access to Oracle ATG Web Commerce repositories and various personalization and commerce features. (For detailed information about these services, see the *Repository Guide*, *Personalization Programming Guide*, and *Commerce Programming Guide*.) These services are packaged in three separate applications:

<ATG11dir>/DAS/WebServices/repositoryWebServices.ear
<ATG11dir>/DPS/WebServices/userprofilingWebServices.ear
<ATG11dir>/DCS/WebServices/commerceWebServices.ear

You can include any of these web services in an assembled EAR file by including the module that contains the desired services. For example, to include the Commerce services, specify the DCS.WebServices module when you invoke the runAssembler command (see the Assembling Applications section of the Platform Programming Guide for information on using runAssembler).

Note: The Oracle ATG Web Commerce platform also provides REST Web services. See Web Services Guide.

## **Connecting to the Dynamo Server Admin**

The Dynamo Server Admin gives you quick access to the following features:

### **Configuration Manager**

Modify configuration for Oracle ATG Web Commerce server instances.

### **Component Browser**

Browse the Nucleus component hierarchy.

### **ATG Control Center Administration**

Start up the ACC.

### **Password Management** Change administrator passwords.

### **JDBC Browser**

Browse a database through a JDBC connection, examine database metadata, create and drop tables, and execute database queries.

### Performance Monitor

View performance statistics on Oracle ATG Web Commerce applications.

### Web Service Administration

Create and manage web services.

### **Batch Compiler**

Precompile JHTML pages to prevent any delay the first time they load.

### **Configuration Reporter**

Display reports about Oracle ATG Web Commerce component properties and environment.

### **Personalization Administration**

Find, edit, and create user profiles. If you have access to the Business Control Center, that should be used instead.

**Configuration Reporter** Display reports about Dynamo component properties and environment

**Personalization Administration** Find, edit, and create user profiles

**Commerce Administration** 

Administer product inventory, order fulfillment, and the product catalog

**Sitemap Administration** Generate Sitemap and Sitemap index information

**ATG Platform Documentation** Read the Dynamo documentation set

**Running ATG Products** View the products currently running

You can access the Dynamo Server Admin at http://hostname:port/dyn/admin. The Dynamo Server Admin is password protected. The username is "admin." Your organization sets the password when you use the Configuration and Installation Manager (CIM) utility to configure an Oracle ATG Web Commerce server. See Configuration and Installation Manager (CIM) (page 10)

For information about including the Dynamo Server Admin when you assemble an EAR file, see the *Including the Dynamo Server Admin* section of the *Developing and Assembling Nucleus-Based Applications* chapter of the *Platform Programming Guide*.

## **Connecting to the Business Control Center**

If your application includes the BIZUI module, you can use the Oracle ATG Web Commerce Business Control Center to create, preview, approve, deploy, and revise site content, as well as to access other Oracle ATG Web Commerce applications. To access the Oracle ATG Web Commerce Business Control Center, point your browser to the following URL:

http://hostname:port/atg/bcc

To learn more about the Oracle ATG Web Commerce Business Control Center, see the *Content Administration Guide for Business Users*.

## **Starting the ATG Control Center**

You can start the ACC in several ways, depending on whether you're starting it locally in relation to your Nucleus-based application, or on a separate client.

**Note:** If you are using a UNIX variant, the shell from which you start the ACC must support X11 forwarding. Depending on your client, you may need to install X11 packages, or use Xming or equivalent tools.

**Note:** Due to a Java bug, if you are running Java 6, you cannot run the ACC in the same virtual machine as an IBM WebSphere application server. You can run the ACC in a dedicated VM, or install the following IBM iFix:

http://ww-01.ibm.com/support/docview.wss?uid=swg24027328

To connect to a Nucleus-based application from a client machine, you must use the client version of the ACC (see Installing the ATG Control Center on a Client Machine (page 54)). Note that for a Nucleus-based application to accept connections from the ACC, all of the following must be true:

- The application includes the DAS-UI module.
- The rmiEnabled property of the /atg/dynamo/Configuration component is set to true.
- The adminPort property of the /atg/dynamo/Configuration component is set to the listen port of your application server (for example, the JBoss default is 8080).

These settings are all part of the default configuration created by the Oracle ATG Web Commerce installer, so you generally do not need to configure them.

In addition, to enable the client version of the ACC to connect to an application, the application must include the DafEar.Admin module. This module is not included by default, so you must explicitly specify it when you assemble the application. See *Including the Dynamo Server Admin* in the *Platform Programming Guide* for more information.

**JBoss Note:** In order to connect to your running Oracle ATG Web Commerce application from any remote location (that is, not using localhost), you must start your JBoss server using the -b option. For example, on Windows use the following command:

run.bat -b 0.0.0.0

See your JBoss documentation for information on this and other settings.

### Starting the ACC on a Server

If you're starting the ACC on the machine that's running your application server, you can run the ACC either in a dedicated VM or in the same VM as the application server.

**Note:** Starting the ACC in a dedicated VM requires more memory than starting the ACC in the same VM as the application server. Running the ACC and the application server simultaneously on a production server is not recommended, as it could affect performance.

### Starting the ACC in a Dedicated VM

To start the ACC in a dedicated VM:

### On Windows:

On the Start menu, click the **ATG Control Center** icon in the Tools folder of the ATG 11.0 program group.

### On UNIX:

Go to <ATG11dir>/home/bin and type the command startACC.

You can also start the ACC in a dedicated VM through the Dynamo Server Admin:

1. Open the Dynamo Server Admin (http://hostname:port/dyn/admin, by default), and click the ATG Control Center Administration link.

The Start ACC page appears, indicating the server VM on which the ACC will be started and the machine on which the ACC will be displayed.

2. Click the Start ACC in Separate VM button.

When the ACC starts up, it displays the Connect to Server screen. Enter a valid user name, password, and the RMI port number. By default, the initial settings are:

- User Name: admin
- Password: your organization sets the password while using the Configuration and Installation Manager (CIM) utility to configure the server. See Configuration and Installation Manager (CIM) (page 10).
- Locale: English (United States)
- Port: 8860

Note that the host name appears as localhost. This value is not editable. To start up the ACC on a remote client machine, see Starting the ACC on a Client (page 62).

### Starting the ACC in the Same VM as the Application Server

To start the ACC in the same VM as your application server, use the Dynamo Server Admin:

1. Open the Dynamo Server Admin (http://hostname:port/dyn/admin, by default) and click the ATG Control Center Administration link.

The Start ACC page appears, indicating the server VM on which the ACC will be started and the machine on which the ACC will be displayed.

2. Click the Start ACC in Server VM button.

When the ACC starts up, it displays the Connect to Server screen. Enter a valid user name and password. By default, the initial settings are:

- User Name: admin
- Password: your organization sets the password while using the Configuration and Installation Manager (CIM) utility to configure the server. See Configuration and Installation Manager (CIM) (page 10).

Note that you cannot specify the host name, locale, or RMI port. The ACC automatically uses the values set in the Nucleus-based application.

### **Exporting RMI Objects**

If the ACC displays an error message while trying to connect to the server, you may need to modify the arguments passed to the Java Virtual Machine by configuring Java Remote Method Invocation (RMI) to export RMI objects on a particular IP address. This can happen under either of the following conditions:

- · The server or the client is running on a machine with multiple host addresses; or
- Oracle ATG Web Commerce is running on a machine that has a primary IP address other than localhost, but the IP address is not functional because the machine is offline.

If Oracle ATG Web Commerce is running on a multi-homed server, you can enable RMI to export objects to a particular address by including the following switch in the JAVA\_ARGS environment variable:

-Djava.rmi.server.hostname=IP\_Address

For the IP address, specify the IP address or name of the host that the client uses to connect to the server. Alternatively, you can specify the name of the server instead:

-Djava.rmi.server.hostname=hostname

If Oracle ATG Web Commerce is running on a machine whose IP address is not functional because the machine is offline, use the following switch:

-Djava.rmi.server.hostname=localhost

### Troubleshooting

If you encounter any errors while using the ACC, check the <ATG11dir>/home/data/acc.log file for information.

### Starting the ACC on a Client

To start the ACC on a client machine and connect to an Oracle ATG Web Commerce application running on a remote application server:

### On Windows:

Click the **Start ATG Control Center** icon in the ATG Control Center 11.0 program group on the Start menu.

### On UNIX:

Go to the ACC 11.0 installation directory and run bin/startClient.

When the ACC starts up, it displays the Connect to Server screen. Enter a valid user name and password, and the RMI port number. By default, the initial settings are:

- User Name: admin
- Password: your organization sets the password while using the Configuration and Installation Manager (CIM) utility to configure the server. See Configuration and Installation Manager (CIM) (page 10).
- Locale: English (United States)
- Port: 8860

In addition, you must specify the name of the host machine on which the Nucleus-based application is running. This is the name used to identify the machine on a network.

### Logging in to a Different Nucleus-Based Application

When the client ACC connects to a Nucleus-based application, it compiles information about the modules in that application (see the *Working with Application Modules* chapter in the *Platform Programming Guide* for

information). To disconnect the ACC from one application and connect to an application that includes a different combination of modules, close down the ACC and restart it to ensure that the ACC compiles all the necessary information.

### Troubleshooting

If you encounter any errors while using the client ACC, check the /data/acc.log file in the ACC installation for information.

## **Using the Findclass Utility**

Oracle ATG Web Commerce provides a utility that finds the .class or JAR file from which a Java class has been loaded. It also prints the CLASS\_VERSION information if found.

To use the findclass utility, open the following page on your Oracle ATG Web Commerce server.

http://server:port/dyn/dyn/findclass.jhtml

Enter the name of the class in the Class Name field and click Find Class. The utility accepts class names in several formats. For example:

- atg.droplet.Cache
- /atg/droplet/Cache.java
- atg.droplet.Cache.class
- atg/droplet/Cache.class

You can append &debug=true to the URL of the findclass utility to print debugging information.

## **Stopping an Oracle ATG Web Commerce Application**

How you stop an Oracle ATG Web Commerce application depends on your application server.

### **Stopping Applications on JBoss**

To stop an application, you can remove it from the deploy directory or shut down the application server. To shut down the server, go to <JBdir> and enter the following command:

#### Windows:

bin\shutdown -s hostname

### UNIX:

bin/shutdown.sh -s hostname

On Windows, you can also use CTRL+C to shut down the JBoss server.

### **Stopping Applications on WebLogic**

You can stop an Oracle ATG Web Commerce application through the WebLogic Server Console or by invoking the server shutdown scripts provided with the application server. You do not need to shut down the application server to stop the application.

### **Stopping Applications on WebSphere**

You can stop an Oracle ATG Web Commerce application through the WebSphere administrative console or by invoking the server shutdown scripts provided with the application server. You do not need to shut down the application server to stop the application.
## 7 Configuring Nucleus Components

This chapter explains how to configure Nucleus components in your Oracle ATG Web Commerce installation. Components represent a particular configuration for a class. Many different components can be based on a single class, each representing a different set of properties for that class. When the class is instantiated from the component, it uses the component properties to configure it.

You can configure components in the following ways:

- Using the ACC
- Manually editing properties files
- Using the Dynamo Configuration Manager in the Dynamo Server Admin (changes are limited)
- Using the Component Browser in the Dynamo Server Admin (live components only, changes do not persist beyond restart)

This chapter covers the following topics:

Working with Configuration Layers (page 65)
Finding Components in the ACC (page 68)
Changing Component Properties with the ACC (page 69)
Changing Component Properties Manually (page 71)
Using the Dynamo Component Browser (page 73)
Common Configuration Changes (page 75)
Creating Additional Oracle ATG Web Commerce Server Instances (page 77)
Setting Up a Configuration Group (page 79)
Session Management in Oracle ATG Web Commerce Applications (page 86)

Most of the information in this chapter applies only for applications running in development mode (see the *Developing and Assembling Nucleus-Based Applications* chapter of the *Platform Programming Guide* for the differences between development and standalone modes).

## **Working with Configuration Layers**

Before changing the configuration of Nucleus-based applications, you should be familiar with the concept of *configuration layers*. This section covers the following topics:

- Understanding Properties Files (page 66)
- Understanding Configuration Layers (page 66)
- Accessing Configuration Layers in the ACC (page 67)
- Global Configuration Changes (page 68)
- Locking Configuration Layers (page 68)

#### **Understanding Properties Files**

Oracle ATG Web Commerce application modules use properties files to configure Nucleus components. The base properties files are normally stored in the config subdirectory of the module, either as individual plain text files or as part of a JAR file (see Modifying Custom Module Resource Settings (page 76) to configure alternative configuration paths). For example, much of the default configuration is determined by properties files stored in <ATGI1dir>/DAS/config/config.jar.

**Note:** Do not modify the properties files in these JAR files to change configuration settings, or your changes will be overwritten when you install a new Oracle ATG Web Commerce platform distribution.

To see the properties file in your Oracle ATG Web Commerce installation, do the following:

- 1. Start the ACC.
- 2. Select Pages and Components > Components by Path from the navigation menu.
- 3. Open the /atg/dynamo/Configuration component. When the ACC Component Editor opens, click the Configuration tab.

Note that there are several Configuration.properties files. You can view the contents of these properties files by double-clicking the file names.

#### **Understanding Configuration Layers**

Oracle ATG Web Commerce platform configuration layers allow you to make configuration changes and preserve them locally, without modifying the base configuration. Layers contain properties files, and can be stacked in a variety of ways to create different configurations for different purposes. The configuration stack is determined from the MANIFEST.MF files for the Oracle ATG Web Commerce application modules included in the application.

Nucleus locates configuration properties by examining the properties files in the directories and JAR files specified by the configuration path or paths (a module can have any number of configuration paths). The paths for all modules used in your application are aggregated and ordered based on the module dependencies. The result is a combination of the property values found in each of the files or directories in the configuration paths. If the same property value is defined in more than one properties file, values found later in the configuration path (as determined by the module dependencies) override the values found earlier. The localconfig directory usually appears last in the configuration path, so that any properties defined there override default system settings.

For example, suppose you change the port number for Oracle ATG Web Commerce's internal RMI server, by setting the rmiPort property of the /atg/dynamo/Configuration component, and save the new value in the localconfig directory. The next time you start the application, Nucleus takes the value of the rmiPort property from localconfig, because it is the last directory in your configuration path.

Any changes you make to localconfig are preserved when you install a new Oracle ATG Web Commerce version.

For more information on modules, configuration layers, and properties files, see the *Nucleus: Organizing JavaBean Components* and the *Working with Application Modules* chapters of the *Platform Programming Guide*.

#### **Accessing Configuration Layers in the ACC**

When you modify a component's properties in the ACC, the updated properties file is stored in one of the following locations:

- The ACC's default configuration directory, initially set to <ATG11dir>/home/localconfig
- A server-specific directory if the component already has a configuration in that layer. For example, if you run an application that does not use the default Oracle ATG Web Commerce server, and you modify a component using the ACC, the updated properties file is stored in the localconfig directory for the Oracle ATG Web Commerce server used by that application.

**Note:** The ACC shows only the configuration layers used by the application to which you are currently connected.

#### **Resetting the Default Configuration Layer**

Unless you specify otherwise, the ACC editor saves all updates to a component's configuration in the default configuration layer. Components that you create, duplicate, or paste are also placed there.

The installation initially sets the default configuration layer to <aTGlldir>/home/localconfig. You might want to change the default configuration directory if you have multiple servers running different applications. For example, you might have one server running a customer service application and another running an online store.

You can set any unlocked configuration layer as the default.

You can change the default configuration layer on the server, so it affects all server clients and persists across all editing sessions; or only on the local client. If you change the default layer locally, the setting remains in effect until you shut down the host.

1. Navigate to the configuration layer that is currently set as the default, and open its CONFIG.properties file, or create one if it does not yet exist.

For example, the Oracle ATG Web Commerce installation initially sets the default configuration layer to <aTGlldir>/home/localconfig/. Therefore, open this file:

<ATG11dir>/home/localconfig/CONFIG.properties

- 2. Set defaultForUpdates to false.
- 3. Navigate to the desired configuration directory and open its CONFIG.properties file.

For example, to set <ATG11dir>/home/servers/myNewServer as the default configuration directory, open this file:

<ATG11dir>/home/servers/myNewServer/CONFIG.properties

4. Set defaultForUpdates to true.

To temporarily reset the default configuration layer within the ACC:

- 1. In the ACC, select Set Update Layer from the Tools menu.
- 2. When the **Set a Default Configuration Layer** dialog opens, select the configuration layer that you want to open by default.

#### **Changing a Component in a Non-Default Configuration Layer**

To change a component in a non-default configuration layer:

- 1. Select the component to edit.
- 2. Choose File > Open Component in Layer.

The dialog box **Select a Configuration Layer** opens, listing the name and path of each configuration layer. Check marks identify the layers currently in use.

3. Select the layer to open and click **OK**. The component opens in a separate **Component Editor** window.

#### **Global Configuration Changes**

Global configuration settings are configured in the GLOBAL.properties (located in config/config.jar) file. The settings in this file control logging and log listeners and apply to all components in the config tree except those that set these properties explicitly themselves. To change these values, you must edit this file manually (see Changing Component Properties Manually (page 71) later in this chapter), or override them by adding your own GLOBAL.properties file in another configuration layer.

#### **Locking Configuration Layers**

Locked configuration layers such as Dynamo Base are marked with a padlock icon. Properties in a locked layer cannot be edited. To lock a configuration layer, modify the CONFIG.properties file for that layer as follows:

- 1. Open the CONFIG. properties file for the layer to lock.
- 2. Add the following line to CONFIG.properties:

readOnly=true

## **Finding Components in the ACC**

When changing Oracle ATG Web Commerce component configuration, you can use the ACC to search for components by name, class or interface.

To search for a component:

 Choose File > Find Component in the main ACC window. The Find Component dialog box opens, as shown below.

Find: Component	<u>_   ×</u>
Name & Location	Find
	Stop
Search for:	
Look in: Browse	Cancel
<u> M</u> atch case	Help
Include subfolders	

- 2. Click the radio button that indicates the way you want to search: Search by component name or Search by class or interface name.
- 3. Type the component name in the **Search for** field. You can search for partial names by using the asterisk (\*) or question mark (?) wildcard symbols. If you want your search to be case-sensitive, check the **Match case** box.
- Type the location you want to search in the Look in field or click the Browse button to select a directory from the component hierarchy. To search all folders within this directory, make sure the Include subfolders box is checked.
- 5. Click the Find button. The search results appear at the bottom of the Find : Component dialog box.

## **Changing Component Properties with the ACC**

The ACC provides a simple way to change many configuration settings. This section uses an example in which you change the port number of Oracle ATG Web Commerce's internal RMI server.

To change the port number:

- 1. Start the ACC.
- 2. Select **Pages and Components** > **Components by Path** from the navigation menu.
- 3. Open the /atg/dynamo/Configuration component.
- 4. When the ACC Component Editor opens, click the **Properties** tab. Scroll down to the rmiPort property:

Component Editor - /atg/dynamo/Configuration			
<u>File Component Window</u>	Help		
Configuration		E 🗊 🗗 😭	?
Class: atg.service.dynamo.DAFConfiguration Config Layer: Local Pathname: /atg/dynamo/Configuration Module: Dynamo Category: Configuration			
Properties Methods Events Basics Configuration			
Property Name	Configured Value	Live Value	
nucleus		in /	
rmiEnabled	true	true	
rmiPort	8860 @	8860	
rmiSystemPathPrefix	/dyn/admin/systemreso	/dyn/admin/systemreso	
root		💼 I	
running		true	
serviceConfiguration			
serviceInfo		DAFConfiguration	
servletPaths		/null-request,/exittrackin	-
rmiPort Class: int Description: The port used	I by the RMI server		

**Note:** Certain expert-level properties are visible only if you select the **Show expert-level information** check box in the **Preferences** > **Tools** > Edit Preferences dialog box.

If the component has been started (indicated by a red dot ), the Properties tab displays two columns of property values: the *Configured Value* and the *Live Value*, described in the table below. You can edit the value of any non-shaded property by clicking in its value cell and entering a new value.

Configured Value	Live Value
The value specified by the component's properties file	The current value, which may be different from the configured value
Changes to the value appear in the ACC immediately, but the changed values are not used to configure the component until you restart the Oracle ATG Web Commerce platform	Changes to the value take place immediately, but are not retained if you stop the component

**Note:** If you are configuring a live component and change properties that are referred to by another component, the references are not updated until you restart the application; they are not updated when you stop or restart the component. For example, Component A has a status property, the value of which is linked to the status property of Component B, changes to the value of the Component B status property are not reflected

in Component A. Stopping or restarting a referenced component leaves the application in an unstable state, and is not recommended.

Editing options depend on the type of property:

- String values provide a text field for editing. You can type values directly into this field or click the ... button to open a pop-up editing window.
- The int, long, float, and double values provide a number field for editing.
- Boolean values provide a pull-down list with true/false options.
- · Enumerated values provide a pull-down list of options.
- Array, hash table, and component values have a ... button that opens a corresponding pop-up editing window.
- All property types have a @ button that lets you set the property value by linking to another component or component property.

In the case of our example, the port number for the RMI server is set by the rmiPort property (of type int). To change the port number, click in the value cell and type the new port number.

After you make changes, choose **File** > **Save** in the Component Editor window. If the component is live, a dialog box appears, asking if you want to copy your configuration changes to the live state. If you copy the changes, restart the Oracle ATG Web Commerce application to ensure that the changes take effect.

## **Changing Component Properties Manually**

As an alternative to using the ACC or Configuration Manager, you can always edit properties files manually. A few configuration properties can only be configured manually, and are not accessible through the ACC or Configuration Manager.

Note, however, that when configuring properties manually, no errors are generated if you specify a property name incorrectly. The component may generate an error if it cannot find the value; in this case, check your properties file for typos.

To manually edit a properties file, do the following:

1. Create a new properties file in <ATG11dir>/home/localconfig with the same name and path structure as the original file. For example, the defaultFrom property in the /atg/dynamo/service/SMTPEmail component specifies the e-mail address from which messages will be sent via SMTP. To modify defaultFrom, create a new file called SMTPEmail.properties in the path <ATG11dir>/home/localconfig/atg/ dynamo/service.

**Note:** Step 1 is not necessary for the Configuration.properties file because a file of this name is created in the <aTG11dir>/home/localconfig/atg/dynamo directory during the installation process.

2. Add the desired property to the new file. For example, to change the setting for defaultFrom, such as to test@example.com, add the following line to the SMTPEmail.properties file in <ATG11dir>/home/localconfig/atg/dynamo/service:

```
defaultFrom=test@example.com
```

For example, to change the port number of Oracle ATG Web Commerce's RMI server to 8862 manually, open your <ATG11dir>/home/localconfig/atg/dynamo/Configuration.properties file and add (or modify) the following line:

```
rmiPort=8862
```

When specifying values for a property, you can add a manual line break using the backslash (\) line continuation character:

```
myList=valueOne,\
    valueTwo,\
    valueThree
```

This can help with readability when configuring lists of values.

Save the Configuration.properties file and restart the application. Because you made the change in the localconfig directory, the new port number overrides the original value (still stored in the config/atg/ dynamo/Configuration.properties file) and will be preserved when you install a new Oracle ATG Web Commerce platform distribution.

#### Using Forward Slashes (/) and Backslashes (\)

When specifying values for file properties, Nucleus translates the forward slash (/) to the file separator for your platform (for example, Windows uses a backslash (\) as a file separator).

The backslash (\) is the escape character for properties files, so if you edit a properties file by hand, you must use two consecutive backslashes (\\) to specify a value that contains a backslash. For example:

```
documentRoot=\\WebServer6.1\\docs
```

The ACC Component Editor handles the escape character automatically; if you change properties using the ACC, use single backslashes.

#### **Modifying Lists of Values**

When adding to a list of values for a property in a properties file, use the += appending operator. This operator is commonly used in localconfig/atg/dynamo/Initial.properties to specify the components to create at startup time. For example:

initialServices+=/StartComps/services/comp1

The += operator specifies that you want to append /StartComps/services/compl to the value of initial services set elsewhere in the configuration path, rather than replace the value.

Similarly, you can use the -= operator to remove an item from a value list. This allows you to avoid re-declaring a list when you only want to remove one member. Note that in order for values to be removed, they must match

exactly; if you specify 2.0 for removal, 2.00 is not removed. If the item to be removed is not found, no errors are generated.

#### **Specifying Directory Paths**

When you specify a directory path as a value in a component, you can do so either relative to the <arglidir>/ home directory, or relative to your Oracle ATG Web Commerce server's directory.

#### Adding Comments to Properties Files

To add comments to a properties file that you've edited manually, you must add the comment in the \$description field. If you preface the comment with a pound sign (#), the comment will be deleted if you subsequently modify the properties file using the ACC.

### Using the Dynamo Component Browser

The Dynamo Component Browser, an element of the Dynamo Server Admin, is a window into Oracle ATG Web Commerce's Nucleus framework. From the Component Browser, you can view and modify components in a running Nucleus-based application.

To open the Component Browser, connect to the Administration UI using this URL:

http://hostname:port/dyn/admin

Enter the username and password. The username is "admin." Your organization set the password while using the Configuration and Installation Manager (CIM) utility to configure the server. See Configuration and Installation Manager (CIM) (page 10). When the Administration UI opens, click the **Component Browser** link.

The following topics are covered in this section:

- Component Browser Structure (page 73)
- Changing the Running Configuration (page 74)
- Starting Nucleus Components (page 74)
- Customizing the Interface (page 74)

#### **Component Browser Structure**

The Dynamo Component Browser is set up so that you can view and edit component properties. The Component Browser main page shows a list of components (called services in the Admin UI) currently running in Nucleus, such as Initial. When you click **Initial**, you see a page that shows the hierarchical location and class reference of that service:

```
Service/Initial/
Class atg.nucleus.InitialService
```

The forward slash character (/) separates the elements of the name into a hierarchy you can click through.

Note: Clicking the first forward slash (/) character brings you back to the main Nucleus service page.

Below the beginning information, you see tables with Properties, Event Sets, and Methods. The current service's property names and values are listed. Continuing on the Initial service page, if you click loggingDebug in the Properties table, you see a page that shows the properties of loggingDebug; you can edit these properties on this page. For example, to enable debugging statements to be logged, go to New Value and select **true**. Then click the **Change Value** button. To see the changes listed back on the Initial service page, click your browser's **Reload** button to refresh the view of the Properties table.

**Note:** Avoid changing system property values unless you know what they do. Changes set here remain in effect while this Oracle ATG Web Commerce instance is running.

#### **Changing the Running Configuration**

You can change the configuration of a running Nucleus-based application from the Dynamo Component Browser. For example, on the Initial service page, click loggingDebug in the Properties table to see the properties of loggingDebug. To enable logging for debugging statements, go to New Value and select **true**, then click the **Change Value** button. To see the changes listed back on the Initial service page, click your browser's **Reload** button to refresh the view of the Properties table.

**Note:** Avoid changing system property values unless you know what they do. Values changed in the Dynamo Component Browser are not written to the properties files; when you stop and restart the application, configuration properties revert to those in the configuration properties file. To make permanent changes to configuration, make the change in development mode using the ACC, then redeploy the application.

#### **Starting Nucleus Components**

In addition to browsing for running components to change their configuration, you can use the Component Browser to start a Nucleus component that is not currently running. To start a stopped component, enter the full Nucleus path of the component in your browser. For example, you can start the OrderRepositoryPipelineDriver by going to this URL:

http://hostname:portnumber/dyn/admin/nucleus/atg/reporting/datawarehouse/loaders/OrderRepositoryPipelineDriver

#### **Customizing the Interface**

By default, the Dynamo Component Browser displays a component by listing its contained children and the values of the component's properties. You might want to customize a component's administrative interface, for example to show more information about a service. To do this, override the methods in the default administrative servlet, atg.nucleus.ServiceAdminServlet. The Scheduler service, for example, extends the standard administration servlet to show information about all the tasks the scheduler is running. To see a list of these tasks, go to the following URL:

http://hostname:port/dyn/admin/nucleus/atg/dynamo/service/Scheduler

To customize an administrative interface, create a subclass of atg.nucleus.ServiceAdminServlet.

## **Common Configuration Changes**

This section outlines several common configuration changes.

#### **Modifying Environment Settings**

Oracle ATG Web Commerce's startup behavior is affected by its CLASSPATH, the Java arguments passed to the Java Virtual Machine, and any custom environment variables you define. You can modify the startup behavior of these parameters as follows:

- CLASSPATH: Oracle ATG Web Commerce's CLASSPATH includes a <ATG11dir>/home/locallib directory, which you can use for any Java class files you create (classes should be in exploded form). Any classes stored in this directory are picked up by Oracle ATG Web Commerce automatically. For more information, see the Nucleus: Organizing JavaBean Components chapter of the Platform Programming Guide.
- Java Arguments: You can set or add to the arguments passed to the Java Virtual Machine by setting the environment variable JAVA\_ARGS.

To customize the CLASSPATH and JAVA\_ARGS settings, as well as define custom environment variables, see your application server documentation.

The following table lists some common values for JAVA\_ARGS:

Java Argument	Description
-Djava.rmi.server.hostname= IP_Address	Configures Java Remote Method Invocation (RMI) to export RMI objects on a particular IP address; for more information, see Starting the ATG Control Center (page 59) in the <i>Running Nucleus-Based Applications</i> chapter
-Djava.compiler=NONE	Turns off the just-in-time compiler so that stack traces include full line number information
-Xms <i>size</i>	Minimum size of memory heap for Java Virtual Machine on startup
-Xmxsize	Maximum size of memory heap for Java Virtual Machine
-Xnoclassgc	Prevents garbage collection of classes
-verbose[:class gc jni]	Enables verbose output about each class loaded, garbage collection, or Java Native Interface (JNI) messages
-Duser.country=country -Duser.language=language -Duser.region=region	This allows you to set the Java Virtual Machine locale, which is used by the Locale.getDefault method to set the region. This setting allows you to configure your installation for localization by specifying the region and language used by the JVM. The following example sets the JVM locale to Russian: -Duser.country=RU -Duser.language=ru -Duser.region=RU

For more information about arguments you can use with the java command, enter the command java -help.

**Note:** When setting CLASSPATH be careful to append or prepend your values onto the original value of the environment variable rather than replace it, or you will omit directories that Oracle ATG Web Commerce needs to start properly.

#### **Modifying Custom Module Resource Settings**

If you create a custom module (see the *Platform Programming Guide*), you can use the module's MANIFEST.MF file to specify paths to the module's resources, as follows:

• ATG-Class-Path: Specify a space-delimited set of paths to module resources that contain classes required by the module. For example:

ATG-Class-Path: lib/resources lib/classes.jar

Oracle ATG Web Commerce adds the ATG-Class-Path value to the CLASSPATH as each module is processed.

• ATG-Config-Path: Specify a space-delimited set of paths to module resources that provide Nucleus configuration files needed by the module's server application components. For example:

ATG-Config-Path: config/config.jar config/oca-ldap.jar

Oracle ATG Web Commerce adds the ATG-Config-Path value to the configuration path.

**Note:** The path names in a module's ATG-Class-Path and ATG-Config-Path settings are relative to the module's root, not to the <ATG11dir> install directory.

In the MANIFEST. MF file, the ATG-Required attribute specifies which modules the custom module requires to start up. ATG-Required ensures that a given module's manifest is processed *after* it processes all the modules that the module depends on. For example, if you want to place the config directory for your custom module after the DPS config directories in the configuration path, configure the attributes as follows:

ATG-Config-Path: config/ ATG-Required: DPS

#### Enabling checkFileNameCase on Windows

In order to prevent Nucleus from creating new components unnecessarily during development, you can configure Oracle ATG Web Commerce to check the case of file names by setting the checkFileNameCase property of the Nucleus component to true. This prevents Nucleus from creating new components if, for example, you create a component named Person and then mistakenly refer to it as person.

The checkFileNameCase property has no effect on UNIX platforms. It imposes a small performance cost on Windows. Therefore, once your application is no longer in active development and you are not creating new components often, you should set the checkFileNameCase property back to false (the default).

The recommended deployment configuration (false) is set in the <code>liveconfig</code> configuration layer. To learn more about liveconfig settings, see Enabling liveconfig Settings (page 91) in the Configuring for Production chapter.

#### LogListeners

Oracle ATG Web Commerce's global configuration settings are configured in the GLOBAL.properties file (located in config/config.jar). The settings in this file control logging and log listeners and apply to all

Oracle ATG Web Commerce components in the config tree except those that set these properties explicitly themselves. If you want to edit this file, you must edit it manually.

The components listed in the logListeners property receive messages from components that send log events. By default, two log listeners are set: ScreenLog and LogQueue. ScreenLog writes messages to the console, while LogQueue puts messages into the log files.

```
logListeners=\
    atg/dynamo/service/logging/LogQueue,\
    atg/dynamo/service/logging/ScreenLog
```

On JBoss, the ScreenLog component is an instance of the CommonsLoggingLogListener class, which logs via the Apache commons logging APIs.

Normally, commons logging uses the class name of the class doing the logging. Oracle ATG Web Commerce has changed this slightly to provide the component's Nucleus path, prefixed with nucleusNamespace and separated by periods. The prefix prevents collisions with actual class names, and makes it clear that the logging component is a Nucleus component.

For example, the /atg/dynamo/Configuration component would have a commons logging class name of nucleusNamespace.atg.dynamo.Configuration. By default, you see the short name of the component in the JBoss log (Configuration, in this example). To see the entire Nucleus path, set the useFullPaths property to true in the Global.properties file. The logging system then prints out atg/dynamo/Configuration as the short class name.

To disable global logging to the console, set the loggingEnabled property of the ScreenLog component to false.

## Creating Additional Oracle ATG Web Commerce Server Instances

**ATG server** is the term for a specific collection of configuration information, which can then be included with your Nucleus-based application when you assemble the EAR file. It can include information such as machine names and ports, system paths, and connection pools.

The Oracle ATG Web Commerce platform installation comes configured with a default server instance in the <ATG11dir>/home/servers/original directory. You can create additional, individually configurable Oracle ATG Web Commerce servers by running the <ATG11dir>/home/bin/makeDynamoServer script, or through the Configuration Manager in the Dynamo Server Admin when the default server is running. If you are using CIM to configure your installation, CIM creates Oracle ATG Web Commerce servers for you (see Configuration and Installation Manager (CIM) (page 10)).

#### Using the MakeDynamoServer Script

Run the makeDynamoServer script with the following syntax:

makeDynamoServer.bat new\_server\_name rmi\_port\_number drp\_port\_number

This script creates a new <ATGIldir>/home/servers/new\_server\_name directory with the following subdirectories and properties files:

|--- data
|--- j2ee
|--- runtime
|--- localconfig (includes CONFIG.properties)
|--- atg
|--- dynamo (includes Configuration.properties)
|--- logs
|--- archives
|--- pagebuild
|--- sessionswap

It sets the name property in the localconfig/CONFIG.properties file. For example:

name=Server myServer

It also sets the rmiPort, rmiEnabled, and drpPort properties in the localconfig/atg/dynamo/ Configuration.properties file. For example:

```
rmiEnabled=true
rmiPort=9001
drpPort=9002
```

The DRP port value uniquely identifies the instance; the port itself is not used for communication.

#### **Using the Configuration Manager**

To open the Configuration Manager, connect to the Dynamo Server Admin using this URL:

http://hostname:port/dyn/admin

Enter your username and password; the initial username is "admin." Your organization sets the password while using the Configuration and Installation Manager (CIM) utility to configure the server. See Configuration and Installation Manager (CIM) (page 10).

When the Dynamo Server Admin opens, click the **Configuration Manager** link to see your configuration options.

To add a new server, click **Add**, **Delete**, **or Reset Servers**. Unless you explicitly set its properties, the new server inherits the properties of the original default server.

#### **Configuring a New Server Instance**

The Configuration Manager's server list shows the Oracle ATG Web Commerce servers registered with the Configuration Manager. Any changes you make to the *default configuration* affect all Oracle ATG Web Commerce servers that are using the default configuration for that setting.

To configure an individual server, click the server's name in the list. To configure a cluster, see the *Configuring for Production* (page 91) chapter.

The Changing Component Properties with the ACC (page 69) section includes an example of how to change the port number of Oracle ATG Web Commerce's internal RMI server. To make that same change using the Configuration Manager, do the following:

1. Click the name of the server you want to configure (for example, Default Configuration).

The Server page opens, listing the configuration properties that you can modify in various categories.

- 2. In the Configure Internal Servers section, click the RMI Server link.
- 3. When the Configure RMI Service page opens, type the new port number in the RMI service port field.
- 4. Click Apply Changes.

The change is written to a properties file in your Oracle ATG Web Commerce installation, but does not affect the currently running Nucleus-based application. For a development-mode application, restart the application for the change to take effect. For a standalone application, reassemble and redeploy the EAR.

## **Setting Up a Configuration Group**

A configuration group provides a mechanism for ensuring consistent configuration among Oracle ATG Web Commerce server instances. At startup, instances that are members of a configuration group download group configuration properties from the group's master server. At runtime, group members can periodically download updates that pertain to their group.

**Note:** Like other configuration changes, group configuration changes generally take effect only on instance startup; they have no effect on a running Nucleus component.

In order to join a group, an Oracle ATG Web Commerce instance must define itself as a group client or server by setting a ConfigurationClient or a ConfigurationServer component:

 A ConfigurationClient component obtains its group configuration settings from an Oracle ATG Web Commerce server instance that is designated as the group master. Each /atg/dynamo/service/ groupconfig/ConfigurationClient component is an instance of this class:

atg.service.configuration.group.ConfigurationClient

• A ConfigurationServer component maintains group configuration settings and ensures that those settings are uniform among all group members. Each /atg/dynamo/service/groupconfig/ ConfigurationServer component is an instance of this class:

atg.service.configuration.group.ConfigurationServer

One ConfigurationServer is designated as the default group master. Changes to group settings must be set on the master ConfigurationServer; it then distributes those changes to other ConfigurationServers and ConfigurationClients in the group.

A group can have one or more ConfigurationServers. If the primary master fails, another ConfigurationServer assumes the role of group master until the primary master resumes operation. The order of succession is established by the primary master and distributed to other ConfigurationServers.

**Note:** An Oracle ATG Web Commerce instance that serves as a configuration server can also act as a configuration client, and typically does so.

#### **Requirements**

To use group configuration, the following requirements apply to each Oracle ATG Web Commerce server instance in the group:

- The instance must be assembled with the DafEar.Admin module.Its /atg/dynamo/ Configuration.adminPort property must be set to the port where the HTTP server is listening and can service the Dynamo Server Admin (http://host:port/dyn/admin/).
- For each Oracle ATG Web Commerce server instance, set its group properties in the Configuration.properties file, as described later in this section.

#### **Group Identifiers and Node Types**

A configuration group is identified by its group name, where each ConfigurationClient and ConfigurationServer in the group is configured with the same Configuration.groupName property. Settings that are specific to a group are known only to the member Oracle ATG Web Commerce instances. An Oracle ATG Web Commerce instance can belong to only one group at a time.

**Note:** A configuration group can overlap multiple Oracle ATG Web Commerce server clusters — for example, publishing and production clusters.

Within a group, Oracle ATG Web Commerce server types are differentiated through their node types. For example, a configuration group might contain these servers:

- Commerce production
- Commerce staging
- Asset management

In order to differentiate settings among server types, each server's configuration client sets its Configuration.nodeType property to a value that corresponds to its server type. Given the previous server types, you might set their respective Configuration.nodeType properties as follows:

- commerce-production
- commerce-staging
- publishing

#### **Dynamo Server Admin Administrator Authentication**

Enter the Dynamo Server Admin administrator account and password in the properties of the HttpConfigurationServerAccess component.

Set the following property in the <ATG11dir>/home/servers/server-name/atg/dynamo/service/ groupconfig/HttpConfigurationServerAccess.properties file.

adminPassword=myadminusername

#### WebLogic Application Server Administrator Authentication

If the servers in your configuration group run on the Oracle WebLogic application server, enter the WebLogic server administrator account and password in the properties of the HttpConfigurationServerAccess component.

Set the following two properties in the <ATG11dir>/home/servers/server-name/atg/dynamo/service/ groupconfig/HttpConfigurationServerAccess.properties file.

weblogicAdminUsername=myadminusername weblogicAdminPassword=myadminpassword

If your WebLogic server does not require authentication, set the sendWeblogicAuth property to false.

sendWebLogicAuth=false

#### **Configuration Group Properties**

In order to configure an Oracle ATG Web Commerce server instance to participate in a configuration group, the following properties must be set before startup in Configuration.properties, in one of the following locations:

<ATG11dir>/home/localconfig/atg/dynamo/service/groupconfig/

```
<ATGlldir>/home/servers/serverName/localconfig/atg/dynamo/
service/groupconfig/
```

Property	Туре	Description
groupName	String	A unique string that defines the group. All group members must set this property to the same value.
defaultMasterServer	boolean	Set to true for one Oracle ATG Web Commerce instance in the group, designates this configuration server to serve as the primary master. If set to true, the property serverEnabled must also be set to true. All other configuration servers in the group should set this property to false.
serverEnabled	boolean	Set to $true$ for all primary and backup configuration server instance in the group.
clientEnabled	boolean	Set to true on every configuration client, enables an Oracle ATG Web Commerce instance to participate in a configuration group
clientNodeType	String	Required for all enabled configuration client instances, associates a configuration client with settings that are specific to that node type. All configuration clients of the same node type must set this property to the same value.

For example, you might configure a master configuration server as follows:

```
groupName=myUniqueGroupName
defaultMasterServer=true
serverEnabled=true
clientEnabled=true
clientNodeType=generic
```

#### **Optional Properties**

You can also set the following properties in Configuration.properties:

Property	Туре	Description
autoDiscoveryEnabled	boolean	Specifies whether auto-discovery is enabled.
httpPort	int	The HTTP port for this Oracle ATG Web Commerce instance.
httpsPort	int	The HTTPS port for this Oracle ATG Web Commerce instance.
startingServerUrls	URL[]	The array of starting servers. This property is required if your configuration group spans network subnets, or Auto-Discovery (page 85) is disabled.

#### **Configuration Server and Configuration Client Properties**

You can set all required group configuration properties in Configuration.properties, as described earlier. If desired, you can fine-tune the behavior of configuration servers and configuration clients by setting their properties directly. For example, after detecting the failure of the master configuration server, by default a backup configuration server immediately assumes the master role or looks for another backup configuration server by setting a latency period for a given configuration server by setting its WaitBeforeBecomingServerTimeout property.

#### **Viewing Group Properties**

At runtime, you can use the Dynamo Administration Component Browser to view the properties of all configuration client and configuration server components, in this Nucleus directory:

/atg/dynamo/service/groupconfig/

After the master configuration server collects all configuration properties from configuration clients in its group, you can review configuration errors by pointing the Component Browser at the master configuration server component. You can also review the settings that are currently in effect for the group.

#### **Storing Group Configuration Files**

Configuration files for a configuration group are stored in one of these directories:

- <ATG11dir>/home/groupconfig
- <ATGlldir>/home/server/server/ame/groupconfig (for named Oracle ATG Web Commerce servers)

The groupconfig directory contains server and client subdirectories, which are used by the local ConfigurationServer and ConfigurationClient, respectively. The client directory obtains its content from the master ConfigurationServer and should not be edited. You should only update the server directory content on the master ConfigurationServer.

#### **Node-Type Configuration**

Configuration for a given node type is stored in the subdirectory of the same name. For example, if a configuration group defines two node types, production and staging, the master ConfigurationServer stores settings for them in two subdirectories as follows:

../groupconfig/server/nodetype/production

../groupconfig/server/nodetype/staging

For example, a production setting for /atg/dynamo/service/jdbc/FakeXADataSource is stored on the master ConfigurationServer in this directory:

\$ATG\_HOME/groupconfig/server/nodetype/production/atg/dynamo/service/jdbc/ FakeXADataSource.properties

This file is propagated to production clients as follows:

\$ATG\_HOME/groupconfig/client/nodetype/production/atg/dynamo/service/jdbc/ FakeXADataSource.properties

At startup, a client's nodetype directory is added to its configuration path and is read before its instance and localconfig directories.

#### Instance Configuration

Configuration settings that are specific to an Oracle ATG Web Commerce server instance can be stored on this path:

../groupconfig/client/instance/host-name+server-name

The master configuration server stores configuration settings for each Oracle ATG Web Commerce server instance in this subdirectory:

../groupconfig/server/instance/host-name+server-name

For example, a master configuration server maintains instance settings for server production1 on host saturn in this directory:

../groupconfig/server/instance/saturn+production1

At startup, a server's instance directory is added to its configuration path and is read immediately before any localconfig property settings.

**Note:** In order to avoid ambiguity among instances on a given host, each instance subdirectory has its own instance.id.properties file. The properties in that file uniquely identify the given instance, so the master configuration server can differentiate among multiple Oracle ATG Web Commerce instances on the same host, if necessary.

#### **Downloading Group Configuration**

At startup, each configuration client and backup configuration server downloads the full groupconfig directory structure from the master configuration server. Local replication of all subdirectories enables a configuration client to start up in the absence of a configuration server, and to start as any of the defined node types.

For example, the layout of a groupconfig directory might look like this:

```
groupconfig
   server
    instance
        saturn+production1
        saturn+lockmgr
        jupiter+publishing1
        jupiter+publishing2
          . . .
    nodetype
       production
       publishing
         . . .
   client
    instance
        saturn+production1
        saturn+lockmgr
        jupiter+publishing1
        jupiter+publishing2
          . . .
    nodetype
       production
       publishing
         . . .
```

The master configuration server can be configured to create a group configuration JAR file as needed after startup, which other group members can download. To do so, set the master's autoCreateConfigJars property to true. You can also create a JAR file on the configuration server manually, by invoking one of these methods from the Dynamo Component Browser:

- createGroupConfigJar()
- createGroupConfigJarIfNeeded()

Each configuration client periodically checks the master configuration server for updates to the group configuration, according to the value set on its ConfigurationClient.schedule property. by default, every 60 seconds. You can also manually download updates to a client at any time by invoking its method downloadConfigUpdate().

#### **Finding a Group Configuration**

Each configuration client caches information about known configuration servers. When required, it checks for configuration updates as follows:

- 1. Reads through its cached list of known configuration servers, starting with the last-known master configuration server.
- 2. If no previously known configuration server can be found from the cached list, uses Zeroconf—via the component /atg/dynamo/service/jmdns/ ClusterBroadcaster —to find a member of its configuration group. It uses that member's published information to find the current master configuration server and downloads its configuration.
- 3. If starting up and all attempts to auto-discover a configuration server fail, starts up with the previously downloaded group configuration.
- 4. If starting up for the first time and no previous group configuration is available, logs an error message and starts up without it.

#### **Auto-Discovery**

As installed, the group configuration system uses Zeroconf to advertise the existence of Oracle ATG Web Commerce server instances in the configuration group. Configuration clients and configuration servers notify Zeroconf of their existence, which also compiles and maintains a list of the group's configuration servers and the order of master succession. Zeroconf maintains the following information about each group member:

Published Data	Description
hostName	Host name, included in broadcast messages
port	HTTP admin port, included as the port in broadcast messages
httpsPort	HTTPS admin port, if any
serverDirectory	Oracle ATG Web Commerce server directory of this instance, truncated to 255 characters (as required by the DNS Service Discovery specification)
atgVersion	Oracle ATG Web Commerce version string of this instance. For example: "11.0".
groupName	Name of the configuration group
clientNodeType	Node type of this configuration client
isGroupServer	boolean, specifies whether this instance is a configuration server
isGroupClient	boolean, specifies whether this instance is a configuration client
isMaster	boolean, specifies whether this instance is defined as the master configuration server
commandLineModules	The list of modules specified on the command line, truncated to 255 characters (as required by the DNS Service Discovery specification)

#### **Validating Group Configuration Properties**

A configuration group can be used to validate Nucleus properties across various configuration clients, through one or more configuration validators that run on the master configuration server. The master configuration

server collects live property values from running configuration clients for validation. Validation errors and warnings are written out via Nucleus logging; these are also accessible on the master configuration server via the Dynamo Component Browser.

#### **Installed Validators**

All validators are registered with the component atg/dynamo/service/groupconfig/validation/
ValidatorRegistry. The Oracle ATG Web Commerce distribution provides three validator components, in this Nucleus directory:

/atg/dynamo/service/groupconfig/validation/

Validator	Description
UniquePortNameValidator	Verifies that a configured port name is unique to a given host machine. This validator is useful for checking settings such as the drpPort, and can be used by Oracle ATG Web Commerce services to compose a unique ID for an Oracle ATG Web Commerce instance.
LiveConfigValidator	Verifies whether all configuration clients of a given node type have the same liveconfig setting.
RepositoryValidator	Checks all known repositories to determine whether the following settings are consistent:
	- Client lock manager settings for repositories that use locked or distributedHybrid caching
	- Subscriber repository and event sender settings for repositories that use distributed caching

All registered validators are scheduled to run according to the value set on the property ConfigurationServer.schedule — by default, every 30 seconds. You can also manually execute all registered validators by invoking the runValidators() method on the master configuration server. Errors are logged every five minutes on the master configuration server.

# Session Management in Oracle ATG Web Commerce Applications

This section discusses topics relating to session management in Oracle ATG Web Commerce applications running on third-party application server.

The J2EE specification defines that each web application has its own session object and any attributes added to the session are only accessible from within that web application. The application server is entirely responsible for managing session life cycles; it generates a unique session ID, creates the session, invalidates it, fails it over, etc. An "ATG session" refers to session-scoped components. Also, keep in mind that Nucleus

components have a tree structure, and can include multiple scopes, with each scope being rooted at a particular component. The root for session-scoped components is /atg/dynamo/servlet/sessiontracking/ GenericSessionManager/sessionid/ where sessionid is generated by the application server.

#### Sharing Session Information Among ATG Applications

You can run multiple ATG applications in the form of WAR files within a single EAR. In this case, you should share session-scoped Nucleus components so that your application always has access to the same instance of session scoped components. By default, J2EE servers hand out different session objects in each web application visited, even if all requests came from the same browser. Sharing sessions across ATG applications ensures that you can build a J2EE application consisting of multiple WAR files in a single EAR, and each WAR has access to the same session-scoped components. Note that you should never run more than a single ATG EAR per application server instance.

When multiple web applications exist in the ATG EAR file, one of them must be designated as the parent application. Being the parent means that that application's session ID is used as the basis for creating the ATG session scope root.

By default, ATG makes the <ATG11dir>\DafEar\base\j2ee-components\atg\_bootstrap.war file the parent web application. The parent context path is /dyn. No additional configuration is required to use this, but your web applications should define the atg.session.parentContextName and atg.dafear.bootstrapContextName parameters in their web.xml to point to the parent web-application as shown:

```
<context-param>
<param-name>atg.session.parentContextName</param-name>
<param-value>/dyn</param-value>
</context-param>
<param-name>atg.dafear.bootstrapContextName</param-name>
<param-value>/dyn</param-value>
<description>The name of the DAF bootstrap WAR context.</description>
</context-param>
```

The context path the context-param points to must be for a WAR file with the SessionNameContextServlet defined in its web.xml:

```
<servlet>
    <servlet-name>SessionNameContextServlet</servlet-name>
    <servlet-class>atg.nucleus.servlet.SessionNameContextServlet
    </servlet-class>
</servlet>
```

Note that there can be only one parent web application specified per EAR file. Therefore, if you change the parent application, be sure to set the context-param to the same values in all web.xml files within your EAR file:

```
<context-param>
<param-name>atg.session.parentContextName</param-name>
<param-value>/portal</param-value>
</context-param>
```

**Note:** This information applies only to session-scoped Nucleus components, and does not affect HTTP sessions obtained using atg.servlet.ServletUtil.getDynamoRequest(request).getSession(), which retain a context particular to the current web application.

#### **Session Interaction Outline**

This section describes the request process and how a Nucleus session name context is associated with that request.

- 1. When a request comes in without a session ID in the cookie or in the URL, the application server creates a new session for the requested web application.
- 2. The ATG PageFilter determines if the session has been failed over and needs to be restored, or is a new session.
- 3. If the request is for the parent web application, a session name context is created with the current session ID and added to the Nucleus component /atg/dynamo/servlet/sessiontracking/ GenericSessionManager.

View that component in the Nucleus Component browser to see a list of current ATG session name contexts and the web applications that share those name contexts.

If the request is for a child web application, the parent application's session ID is resolved in one of two ways, depending on the application server.

Some application servers maintain a single session ID between web applications for the same client (browser), in which case the session name context ID is the current web application's session ID. This behavior is controlled by the /atg/dynamo/

servlet/sessiontracking/GenericSessionManager.singleSessionIdPerUser property, which is
set to one of the following default values in the DafEar sub-module configuration layer:

- WebLogic false
- JBoss-true
- WebSphere-true

Note: Do not change these values from their defaults.

When the singleSessionIdPerUser value is true, the application server uses the same session ID for all web applications, so lookup is not required. Note that the application server hands out the same session id, but **not** the same HttpSession object.

When singleSessionIdPerUser is false, a lookup determines the session name context ID. This is done by the atg.nucleus.servlet.

SessionNameContextServlet servlet (included in atg\_bootstrap.war), using a
RequestDispatcher.include() call. The SessionNameContextServlet does two things:

- Sets the parent session ID as a request attribute that can then be used by the child web application to bind to the correct session context.
- For application servers that do not allow request attributes to be shared between web applications, it also sets a cookie named ATG\_SESSION\_ID with the session ID. This behavior is controlled by the /atg/ dynamo/servlet/sessiontracking/GenericSessionManager.useSessionTrackingCookie property, which is pre-configured with the correct value for each application server.
- 4. The atg.parent.session.id session attribute is set to the parent session ID to avoid repeating the lookup.

- 5. A new session-scoped context of type atg.servlet.SessionNameContext now exists under the GenericSessionManager. Because the ATG Nucleus components live outside the application server's session, an atg.servlet.SessionBindingReporter object is added to each web application session as an attribute. According to the J2EE spec, this object must be notified by the application server when the session is started (its valueBound method invoked) or invalidated (its valueUnbound method invoked).
- 6. The SessionBindingReporter increments a counter in the SessionNameContext it belongs to. This counter keeps track of the number of child web application session references to the Nucleus session scope. As each child is requested during the course of the browser session, this number increases.
- 7. When the application server expires a session, either because of a user request (session.invalidate invoked) or due to a session timeout, it unbinds all the session attributes and invokes the atg.servlet.SessionBindingReporter.valueUnbound() method.
- 8. The valueUnbound decrements the SessionNameContext counter.
- 9. When the counter reaches 0, all the child and parent web application sessions have been expired and it is safe for the ATG Nucleus session scope to be removed.

**Note:** Because the only link to the underlying session is through the SessionBindingReporter attribute, session management is a common cause for memory leaks. One such leak occurs on IBM WebSphere in a clustered environment, where the session invalidation can occur in a different JVM instance than where the session originated. See the Session Management in a WebSphere Cluster (page 104) section.

#### **Managing User Sessions**

You can manage user sessions from the Dynamo Component Browser for debugging or administrative purposes. To access the Session Manager, click through the hierarchy:

/atg/dynamo/servlet/sessiontracking/

Click **GenericSessionManager** to view sessions. Choose the selection criteria, then click the **View** button. Click an individual session to see its properties.

## 8 Configuring for Production

The default configuration settings in your ATG installation are designed for evaluation and development. When your ATG application moves from development to live deployment, you should change some of these configuration settings as described in this chapter for better performance.

This chapter covers the following topics:

Enabling liveconfig Settings (page 91)Fine-Tuning JDK Performance with HotSpot (page 93)Configuring Repositories (page 93)Configuring Targeted E-Mail (page 95)Setting Access Levels for Properties Files (page 96)Setting Logging Levels (page 97)Limiting Initial Services for Quicker Restarts (page 98)Disabling Document and Component Indexing (page 98)Setting up Clustering on JBoss (page 99)Setting Up Clustering on WebSphere (page 102)General Clustering Information (page 106)

## **Enabling liveconfig Settings**

The settings in the ATG base configuration layer are optimized for application development, but are not appropriate for a production environment. When you're ready to deploy your Nucleus-based application in a production environment, enable the settings in the <code>liveconfig</code> configuration layer. This layer overrides many of the default configuration settings with values that are more appropriate for a deployed site. For example, the <code>liveconfig</code> configuration layer improves performance by reducing error checking and detection of modified properties files.

To enable liveconfig, you can use the -liveconfig argument for runAssembler or add the following line to the WEB-INF/ATG-INF/dynamo.env file in the atg\_bootstrap.war module of your EAR file:

atg.dynamo.liveconfig=on

**JBoss Note:** If you are using ATG Portals with JBoss, and you use the <code>-liveconfig</code> flag when you create your EAR file, you must also have a lock manager configured and running in order to create or edit communities. See the *Locked Caching* section of the *Repository Guide* for information on lock management.

To disable liveconfig in an application in which it is currently enabled, either reassemble the application without the -liveconfig flag, or remove or set the liveconfig value to off in WEB-INF/ATG-INF/ dynamo.env file in the atg\_bootstrap.war module.

#### **Customizing liveconfig Settings**

You can add your own configuration files or directories to the <code>liveconfig</code> configuration layer in your ATG installation. It is best to put any such custom settings in a separate directory from the <ATG11dir>/ATG\_module/liveconfig directories, in order to keep track of which liveconfig settings are ATG settings and which are your own custom settings. For instance, if you have a custom application module named MyModule, you could create a MyModule/liveconfig directory in your module, and include in that directory any configuration settings that you want to take effect when the <code>liveconfig</code> configuration layer is enabled.

To add an entry to the liveconfig configuration layer, include it in your module's manifest in an entry like this:

ATG-LiveConfig-Path: liveconfig

#### **Disabling Checking for Changed Properties Files**

Some disk access and memory allocation overhead can be eliminated by setting the configurationCheckMilliseconds property of the Nucleus component (with a Nucleus component path of /) to -1. This property controls whether or how often ATG rereads .properties files or .java files the next time that an instance of a given component is created. The default is 1000. This feature is useful during development, but we recommend disabling it once a site goes live for better performance. The value -1 disables the reloading of properties and .java files altogether.

**Note:** If you subsequently make changes to .properties files or .java files on your live site (which you generally should not do), you need to restart your application server before changes are picked up. If you change property settings using the ACC, you may need to restart to fully register changes that may affect interdependent components.

The recommended configuration is enabled in the <code>liveconfig</code> configuration layer.

#### **Disabling the Performance Monitor**

The Performance Monitor (/atg) can be used to gather statistics about the performance of specific operations in ATG components. However, this information gathering can itself have a negative effect on performance. Therefore, for deployment, disable the Performance Monitor by setting its mode property to 0:

mode=0

The Performance Monitor is disabled in the liveconfig configuration layer.

For more information about the Performance Monitor, see the Monitoring Site Performance (page 137) chapter.

#### Adjusting the pageCheckSeconds Property

ATG's Page Processor compiles JHTML pages into .java files (JSP compilation is handled by your application server). The page processor, located at /atg/dynamo/servlet/pagecompile/PageProcessor, checks for new JHTML pages that need to be compiled. You can improve performance by increasing the Page Processor's pageCheckSeconds property. The page compile servlet uses this property value to determine whether to check for new JHTML pages that need to be recompiled. If a request occurs within this time interval (measured in seconds) for the same page, ATG will not check the date on the file. This improves performance in serving pages.

A value of 0 causes Oracle ATG Web Commerce to check for new pages on each request. The default value is 1. The liveconfig value is 60.

## **Fine-Tuning JDK Performance with HotSpot**

Oracle's Java HotSpot technology is available in a Client Virtual Machine (VM) and a Server VM. The default Client implementation can be considerably slower than the Server implementation, so ATG recommends using the Server JVM.

## **Configuring Repositories**

On a production site, it is critical that your ATG repositories be configured properly to ensure that data storage and retrieval are performed accurately and efficiently. Poorly configured repositories can result in performance bottlenecks and errors. In particular, repository caches must be tuned properly to ensure that data is retrieved quickly and is up to date. This section discusses repository settings to configure on your sites.

#### **Setting Cache Modes**

The ATG SQL repository offers a choice of cache modes. When you have only a single ATG instance installed, you can use the simple cache mode with no problems, since there is no chance of two servers using inconsistent copies of a repository item due to caching. However, when you deploy multiple ATG instances, you need to choose an appropriate cache mode for each item descriptor used by your application. See the *Repository Guide* for more information.

In particular, if your sites are running more than one ATG server, it is highly recommended that you use locked mode caching for the individualScenario item descriptor. See the *Personalization Programming Guide* for more information.

Remember that if you use locked mode caching, you must also enable lock manager components. See Enabling the Repository Cache Lock Managers (page 94) in this chapter.

#### **Prepopulating Caches on Startup**

ATG performance typically improves after an application has been running a while, because more requests can be satisfied from caches. Under some circumstances, it may make sense to prepopulate your caches, so that you get the benefit of the caches immediately. Note, however, that this benefit may come at the cost of slower startup times.

You can prepopulate caches in a SQL Repository by using <query-items> tags in a repository definition file. For more information, see the *Repository Guide*.

#### **Enabling the Repository Cache Lock Managers**

If you are using a SQL Repository with locked mode caching, you must enable the ClientLockManager (/ atg/dynamo/service/ClientLockManager) on each ATG server and enable the ServerLockManager (/ atg/dynamo/service/ServerLockManager) on one or more ATG servers. You may want to dedicate an ATG server only to lock management. Note that elements of the ATG Adaptive Scenario Engine are configured in the liveconfig layer to use locked mode caching by default.

Add the ServerLockManager component to the initialServices property of the /atg/dynamo/Initial component, and make sure that the ClientLockManager points to the correct host. The ClientLockManager should be configured like this:

Property	Description
lockServerAddress	The hostname of the machine running the ServerLockManager.
lockServerPort	The port configured in the ServerLockManager component (9010 by default).
useLockServer	True

The ClientLockManager is enabled by the liveconfig configuration layer. For more information about cache lock managers, see the *Repository Guide*.

#### **Configuring Repository Database Verification for Quicker Restarts**

By default, each SQL Repository component verifies each of the tables in its database on startup with a simple SQL query. These verification queries can slow the ATG startup routine. There are several approaches you can take to modify the SQL Repository startup procedures that can result in dramatically faster start times. In particular, you may wish to set the updateSchemaInfoCache property to true in your atg.adapter.gsa.GSARepository components, such as /atg/dynamo/service/jdbc/ ProfileAdapterRepository. For details, see the SQL Types and Repository Data Types section in the SQL Repository Item Properties chapter of the Repository Guide.

#### **Configuring a Content Distributor System**

ATG includes a content distributor system that allows you to cache content from repositories to an HTTP server. Using this system can significantly speed up request handling on a site. By default, only ATG Commerce uses this system, but it can be used by any ATG application. The content distributor system is described in the *Platform Programming Guide*. If you are using an HTTP server such as Apache, no additional configuration of the content distributor system is required. If you are using your application server as your HTTP server, however, you need to configure the system to prepend the context path of the  $atg_bootstrap.war$  application (by default, /dyn) to the URL of any file it sends to the server.

The class atg.distributor.DistributorSender has a property named documentRootContextPath that you can set to specify the string to prepend. For example, for the distributor system used by ATG Commerce, set this property in the component /atg/commerce/catalog/ContentDistributor, either through the ACC or by adding the following line to the properties file of that component:

documentRootContextPath^=/atg/dynamo/Configuration.defaultDynamoPrefix

## **Configuring Targeted E-Mail**

When running on your application server, ATG's targeted e-mail system makes loopback requests back to the server to render the e-mail template for each e-mail recipient. ATG makes one loopback request to create an HTTP session, and uses that session's ID when making subsequent loopback requests to render the template.

#### **Nucleus Components**

The components /atg/userprofiling/email/TemplateEmailSender and /atg/scenario/ IndividualEmailSender (both of class atg.userprofiling.TemplateEmailSender) have several properties used for configuring loopback requests. The following table lists these properties and their defaults when running ATG on your application server:

Property and Default	Purpose
<pre>siteHttpServerName^=/atg/dynamo/ Configuration.siteHttpServerName</pre>	Server name for loopback requests.
<pre>siteHttpServerPort^=/atg/dynamo/ Configuration.siteHttpServerPort</pre>	Port number for loopback requests.
applicationPrefix <sup>*</sup> =/atg/dynamo/ Configuration.dynamoEarContextRoot	The context path of the application.
initSessionURL=/init-session	The URL pattern used by InitSessionServlet (see below).
sessionManager=/atg/dynamo/servlet/ sessiontracking/GenericSessionManager	Used to find the session from the session ID.
loopbackRequestsEnabled=true	Determines whether loopback requests are performed. Can be set to false if you are using this TemplateEmailSender only with DSP templates (see below).

Property and Default	Purpose
contextPathPrefix <sup>+</sup> =/atg/dynamo/ Configuration.defaultDynamoPrefix	String to prepend to template URLs. Default is /dyn/dyn/.

If you are using JHTML templates exclusively, you can disable loopback requests by setting the loopbackRequestsEnable property to false. In addition, you should set the contextPathPrefix property
to null, and set the setupLoopbackTemplateEmailRequests property of the /atg/dynamo/servlet/
pipeline/DynamoServlet component to false.

#### **Configuring Web Applications**

To enable targeted e-mail in an application, the application must run an instance of atg.nucleus.servlet.InitSessionServlet.For example, the web.xml file for the atg\_bootstrap.war application includes the following lines:

```
<servlet>
  <servlet-name>InitSessionServlet</servlet-name>
  <servlet-class>atg.nucleus.servlet.InitSessionServlet</servlet-class>
</servlet>
  <servlet-mapping>
      <servlet-name>InitSessionServlet</servlet-name>
      <url-pattern>/init-session</url-pattern>
</servlet-mapping>
```

This servlet handles the requests to /dyn/init-session, and, as its name implies, initializes a session.

## **Setting Access Levels for Properties Files**

ATG components are configured with plain text properties files. You should set access levels on your properties files so they cannot be altered or viewed by unauthorized users. Only site administrators should have **read** and **write** permission. ATG must be invoked from an account with these permissions as well. The properties files that contain sensitive information typically reside in each server's localconfig directory. The most important properties files to protect include:

Component	Description
/atg/dynamo/Configuration.properties	Basic configuration for ATG
/atg/dynamo/security/BasicSSLConfiguration.properties	Default configuration for any service that uses SSL
/atg/dynamo/service/jdbc/FakeXADataSource.properties	Distributed transaction data source

Component	Description
/atg/dynamo/service/jdbc/JTDataSource.properties	JTA participating and pooling data source
<b>Note:</b> Multiple versions of this component may exist in your installation; all of them may contain information that should be protected.	
/atg/dynamo/service/POP3Service.properties	Checks the POP server for bounced e-mail

The most important ATG Commerce properties files to protect include:

Component	Description
atg/commerce/jdbc/ProductCatalogFakeXADataSourceA.properties	A distributed transaction DataSource
atg/commerce/jdbc/ProductCatalogFakeXADataSourceB.properties	A distributed transaction DataSource

These ATG Commerce properties files are located in a .jar file at <ATG11dir>/DCS/config/ config.jar.For more information on ProductCatalogFakeXADataSourceA.properties and ProductCatalogFakeXADataSourceB.properties, refer to the Commerce Programming Guide.

## **Setting Logging Levels**

By default, ATG sends all log events to two log listener components: /atg/dynamo/service/logging/ LogQueue (which directs output to log files) and /atg/dynamo/service/logging/ScreenLog (which directs output to the console screen). Logging to the screen can cause performance problems on a production site. Note that JBoss and Tomcat extend ExternalLogSystemLogListener and do not use the PrintStreamLogger, which is used by ScreenLog to redirect output to the console screen.

It is recommended that you redirect console logging to a file to prevent loss of information. When creating the file, it is helpful to script the console log file name to contain the current time stamp, such as console\_2013-03-03\_16-39-22.log. Refer to your Web application documentation for information on redirecting console log information to a file.

If you want to disable logging entirely, or specify different logging levels, you can do that in the GLOBAL.properties file. For example:

```
loggingError=true
loggingWarning=true
loggingInfo=true
```

The loggingDebug log generates large numbers of messages, which can impair performance, so loggingDebug should be set to false on a live site. You still have the option of overriding the global settings for a specific component. For example, if loggingDebug is set to false in the GLOBAL.properties file, you can still enable it for an individual component by setting that component's loggingDebug property to true.

See the Logging and Data Collection chapter of the Platform Programming Guide for more information.

## **Limiting Initial Services for Quicker Restarts**

When you restart an ATG application, it starts up the services specified by the initialServices property of the /atg/dynamo/Initial component. You may add services to this list while you are developing your application. These services may in turn start up other components. Starting up a service at the same time as ATG ensures that the service is created and ready when it is first called upon. However, if too many services are configured to start up at the same time, then the ATG startup routine can become time-consuming and server restarts may be slow, which might make it more difficult to recover and restart if a server runs into problems. If server startups seem to be taking too long, consider whether some services can be started up on some other schedule than immediately on ATG startup.

If you set the loggingInfo property of the Nucleus component (with a Nucleus path of /) to true, and then start up ATG, the resulting info messages display an indented list of when each service starts up. From this list, you can determine which component causes which other components to be started.

## **Disabling Document and Component Indexing**

The ACC creates and maintains indexes of documents and components. For sites with large numbers of documents or components, indexing can take time and CPU resources. Once your sites are deployed and relatively stable, you may want to limit or eliminate the indexing of documents or components.

The document and component indexes are maintained incrementally once built, and are rebuilt completely once a day at 1 a.m. by default. An index is rebuilt at startup only if it does not exist at all.

You can selectively exclude portions of the document tree from indexing by adding absolute pathname prefixes to the excludeDirectories property of the /atg/devtools/DocumentIndex component. The same is true for component indexing, but the component is /atg/devtools/ComponentIndex instead. To improve performance on a live site, you can turn off all document and component indexing by setting the enabled property of the DocumentIndex and ComponentIndex components to false.

## **Enabling the ProtocolChange Servlet Bean**

ATG includes a servlet bean named /atg/dynamo/droplet/ProtocolChange. The ProtocolChange servlet bean lets pages switch between secure and non-secure HTTP servers. The ProtocolChange servlet bean takes

a URL as input and renders a URL that uses either the HTTP protocol or the HTTPS protocol, depending on the output parameter specified. The default configuration is:

```
secureHost^=/atg/dynamo/Configuration.siteHttpServerName
nonSecureHost^=/atg/dynamo/Configuration.siteHttpServerName
securePort=443
nonSecurePort^=/atg/dynamo/Configuration.siteHttpServerPort
secureProtocol=https
nonSecureProtocol=http
enable=false
```

When the enable property is false, the servlet bean renders the URL without changing the protocol. To enable this servlet bean to change the protocol, set the enable property to true. Also, ensure that the secureHost and securePort properties are set to values appropriate for your sites.

### Setting up Clustering on JBoss

A cluster is a set of JBoss servers working together, serving pages at the same port. From the user's point of view, all of the servers function as a single server; it does not matter which server handles a given request. JBoss documentation refers to a cluster as a partition.

Virtually all production sites use clustering. Clustering provides much better performance and reliability than running on a single server. For example, if one server in a cluster goes down, the user may not be aware of it, because the other servers in the cluster can take over the sessions it was handling.

Setting up clustering of JBoss servers running ATG applications involves the steps described in the following sections.

#### **Configuring the HttpPort Property**

When running ATG server instances in a JBoss cluster, you must configure the httpPort property in the /atg/ dynamo/Configuration.properties component to match the port set in the siteHttpPort property. If this is not done, the ATG email sender fails. For example:

```
siteHttpPort=8080
httpPort=8080
```

#### **Creating ATG Servers**

The first step is to create your ATG servers, using the Configuration and Installation Manager (CIM) or the makeDynamoServer script (see Creating Additional Oracle ATG Web Commerce Server Instances (page 77)).

A typical production environment includes: a server lock manager, process editor server, workflow process manager, etc., for services that require a dedicated server, plus several servers that handle page requests. The servers you need to create depend on which ATG applications you are using, and on your unique site requirements.

#### **Assembling for a JBoss Cluster**

When you assemble your application, the application assembler includes all of the ATG servers you have configured (see "Assembling Applications" in the *Platform Programming Guide* for information on application assembly). This means that you can build your application once for each JBoss partition, deploy it on each partition, and enable the appropriate ATG server on each instance simply by changing the value of the atg.dynamo.server.name system property when you start up JBoss:

```
bin\run or bin/run.sh -c server_name -Datg.dynamo.server.name=
ATG_server
```

To assemble and configure your ATG application to run on a JBoss partition, when you invoke the application assembler, use the -liveconfig, -standalone, -distributable, and -pack flags in runAssembler as in the example:

```
bin/runAssembler -liveconfig -standalone -distributable -pack
output_file_name.ear -m module-list DafEar.Admin
```

The -pack flag is optional. The -distributable flag is required to enable JBoss session failover. Do not use the -server flag to specify an ATG server configuration. If you are using a named configuration layer, specify that as well (see "Managing Properties Files" in the *Platform Programming Guide* for information on named configuration layers).

#### **Creating and Configuring JBoss Servers**

Create and configure your JBoss servers; see the JBoss documentation for configuration information.

- 1. Use the <JBdir>/server/all server as a template for creating JBoss instances, since it is set up for clustering.
- 2. Make configuration changes, such as removing unneeded JBoss services (see JBoss Application Framework Trimming (page 164) in this guide).
- 3. Copy the /all server for each corresponding ATG server.

#### **Deploying Your Application**

See the JBoss documentation for information about deploying to JBoss clusters.

## **Setting Up Clustering on WebLogic**

A cluster is a set of WebLogic servers working together, serving pages at the same port. From the user's point of view, all of the servers function as a single server; it does not matter which server handles a given request.

Virtually all production sites use clustering. Clustering provides much better performance and reliability than running on a single server. For example, if one server in a cluster goes down, the user may not be aware of it, because the other servers in the cluster can take over the sessions it was handling.
Setting up clustering of WebLogic servers running ATG applications involves the following steps:

- 1. Create a group of WebLogic servers for serving pages, and assign them to a cluster.
- 2. Create additional WebLogic servers for the ATG lock manager, process editor server, workflow process manager and any other services that require a dedicated server. Assign these servers to a different cluster from the page servers.
- 3. For each WebLogic server, create a corresponding ATG server configuration.
- 4. Assemble your ATG application, and deploy it on each WebLogic server in both clusters. Configure the application on each WebLogic server to use the ATG server configuration that corresponds to that server.

See the Oracle WebLogic documentation for information about creating WebLogic servers and clusters. For information about creating ATG server configurations, see Creating Additional Oracle ATG Web Commerce Server Instances (page 77).

### Assembling for a WebLogic Cluster

When you assemble your application, the application assembler includes all of the ATG servers you have configured. This means that you can build your application once, deploy it on each WebLogic server, and enable the appropriate ATG server on each instance simply by changing the value of the atg.dynamo.server.name system property when you start up WebLogic.

Follow these steps to assemble and configure your ATG application to run on a WebLogic cluster:

1. When you invoke the application assembler, use the -standalone flag to assemble the application in standalone mode, so it is not dependent on your ATG installation.

Note: You cannot use <code>-pack</code> and <code>-standalone</code> in combination on WebLogic.

In addition, use the -liveconfig flag to enable the liveconfig configuration layer. Do not use the - server flag to specify an ATG server configuration.

- 2. Deploy the application on each WebLogic server.
- 3. On each WebLogic server, enable the corresponding ATG server configuration by creating the atg.dynamo.server.name property for the JVM the server is running on and setting the property to the name of the ATG server. For example:

startManagedWebLogic.bat myWebLogicServer Datg.dynamo.server.name=myserver

## **Clustering Example**

Suppose you want to set up a site consisting of an Administration Server, three servers that serve pages, one server that runs the ATG lock manager, and one server that runs the process editor server. Here's an example of how you might do this:

- 1. Start up WebLogic Server using the startWebLogic script. This starts up the WebLogic Administration Server (default name myserver, default port 7001).
- 2. In the WebLogic Console, create a server named pageServer. Assign it port number 7700. Assign an IP address used by no other server in the domain.
- 3. Create a cluster named pageCluster. Put pageServer1, pageServer2, and pageServer3 into this cluster.

- 4. Create servers named procedit and lockmgr. Assign each server the port number 7800. Assign each server a unique IP address.
- 5. Create a cluster named serviceCluster. Put procedit and lockmgr into this cluster.
- 6. Assign the two clusters different multicast addresses.
- 7. Using either the Configuration and Installation Manager (CIM) or the makeDynamoServer script, create ATG servers named pageServer1, pageServer2, pageServer3, procedit, and lockmgr. (You do not need to give the ATG servers the same names as the WebLogic servers, but it is a good idea do so.)
- 8. Configure the ATG lockmgr server to run the ATG ServerLockManager. (See Enabling the Repository Cache Lock Managers (page 94) for more information.)
- 9. Configure the ATG Scenario Manager to run the process editor server on the ATG procedit server.
- 10.Set up ATG session backup, as discussed in Enabling Component Backup (page 107).
- 11.Assemble your application, deploy it on each server in both clusters, and configure each instance to use the ATG server corresponding to the WebLogic server the instance is running on. (This process is discussed in Assembling for a WebLogic Cluster (page 101).)
- 12.Un-deploy any applications that are deployed on the Administration Server.
- 13.Configure your HTTP server to serve pages from each server in pageCluster (but **not** any of the other servers).
- 14.Shut down the Administration Server and then restart it. This ensures that all of the changes you made take effect.
- 15.Start up the managed servers you created, using the startManagedWebLogic script. The syntax of this script is:

startManagedWebLogic WebLogicServer adminURL
-Datg.dynamo.server.name=myserver

where *WebLogicServer* is the name of the WebLogic server, and *adminURL* is the URL of the WebLogic Administration Server. Let's assume that the hostname for the Administration Server is myMachine. To start up the WebLogic pageServer1, the command would be:

```
startManagedWebLogic pageServer1 http://myMachine:7001/
-Datg.dynamo.server.name=pageserver1
```

**Note**: Oracle ATG Web Commerce does not support Unicast mode. Make sure that you set the cluster messaging mode to multicast.

# Setting up Clustering on WebSphere

A cluster is a set of WebSphere servers working together, serving pages at the same port. From the user's point of view, all of the servers function as a single server; it does not matter which server handles a given request.

Virtually all production sites use clustering. Clustering provides much better performance and reliability than running on a single server. For example, if one server in a cluster goes down, the user may not be aware of it, because the other servers in the cluster can take over the sessions it was handling.

# Installing and Configuring WebSphere

The first step in setting up a clustered deployment is to install and configure the WebSphere cluster. See the IBM WebSphere documentation for information.

- 1. Install WebSphere Network Deployment.
- 2. Run the Profile Creation Wizard to create a Deployment Manager profile. While you are installing, take note of the following information for use during your ATG installation:
  - Deployment manager profile name
  - Deployment manager cell name
  - Administration console port
  - SOAP port

# **Creating a Cluster**

Use the WebSphere Administration Console to create your clusters. The recommended topology for running ATG products on a WebSphere cluster has the following characteristics:

- · Includes one Deployment Manager profile and at least one custom profile
- · Separates page serving instances and non-page-serving instances into different clusters
- Includes the web servers in the Deployment Manager cell for management

## **Creating Data Sources**

Create your data sources in the WebSphere Administration console; see the documentation for WebSphere and your database solution for information.

**Note:** The JNDI lookup for your data source must be consistent with the JNDI name configured in your Nucleusbased application. Make sure you set the data source's scope correctly.

# **Installing and Configuring Your Web Server**

Install your web server and configure it for use with WebSphere; see your web server documentation for information.

Take note of the path used for installing web server configuration files. It is extremely important to copy the web server configuration files to your WebSphere servers, so that your application can be targeted to the appropriate web servers and clusters (refer to your WebSphere documentation for more information).

## Installing ATG for a WebSphere Cluster

Follow these steps to install the ATG platform to run in a clustered environment:

1. Run the installer program for the Oracle ATG Web Commerce platform. This is an executable file with the file name extension \*.exe (Windows) or \*.bin (UNIX).

- 2. After you accept the terms of the license agreement, select the installation folder for the ATG software (C: \ATG\ATG11.0 or /home/ATG/ATG11.0, for example).
- 3. Select the ATG products you want to install.
- 4. Select IBM WebSphere Cluster as your application server.
- 5. Enter the WebSphere home directory (C:\WebSphere\AppServer, for example).
- 6. Select the Deployment Manager profile and cell.
- 7. Deploy and install your application (see the WebSphere documentation for information).

# **Assembling for a WebSphere Cluster**

When you invoke the application assembler, use the following flags:

 $-{\tt standalone},$  to assemble the application in standalone mode, so it is not dependent on your ATG installation

-liveconfig, to enable the liveconfig configuration layer

 $_{-\texttt{pack}}$  , because the WebSphere application installation wizard does not recognize an exploded EAR file (see Note)

If you are using a named configuration layer, specify that as well (see "Managing Properties Files" in the *Platform Programming Guide* for information on named configuration layers).

**Note:** It is possible to deploy an exploded EAR through the WAS admin. To do so, in the WAS Deployment Wizard, click the radio button for **Server path** instead of **Local path**, then type in the full path of the EAR directory and submit the form. Note that in order for the WAS deployment wizard to recognize the Server path you provide, the directory must exist on a file system accessible to the server that is serving the WAS admin pages.

Do **not** use the -server flag to specify an ATG server configuration.

See the Assembling Applications section of the Developing and Assembling Nucleus-Based Applications chapter in the Platform Programming Guide for more information on assembly.

# **Session Management in a WebSphere Cluster**

When a session is persisted in a database, WebSphere does not correctly invoke the valueUnbound() method when that session expires, resulting in memory leaks when running ATG applications. The /atg/dynamo/ servlet/sessiontracking/SessionInvalidationService component handles this problem by checking the current set of child sessions known to ATG and comparing the last accessed time to the session's configured timeout, as specified by the application server. If the child session has timed out, it is removed from the list of sessions. When all children of a parent session have been removed, all session-scoped components for that session are cleaned up. The SessionInvalidationService runs on a configurable schedule, with a default of every 5 minutes.

Note that this component does not invalidate the session or interfere in any way with the application server's own cleanup work; it touches only ATG-created items. For even more safety, you can set the additionalTimeoutMinutes property, in which case the service waits the specified additional number of minutes above the application server's configured session timeout before performing the cleanup.

If debugging is turned on, the SessionInvalidationService component indicates when it performs a check, and the last accessed time for each child session. The component is defined in the DafEar.WebSphere module, which is run automatically on WebSphere.

# **Configuring Your WebSphere Servers**

When you assemble your application, the application assembler includes all of the ATG servers you have configured. This means that you can build your application once, deploy it on each WebSphere application server, and enable the appropriate ATG server on each WebSphere instance simply by changing the value of the atg.dynamo.server.name system property when you start up WebSphere.

Do the following for each server.

- 1. In the WebSphere Administration console, go to Server > Application Servers.
- 2. Click the link for the server you want to configure.
- 3. On the right hand side, go to Server Infrastructure > Java and process management > Process Definition > Additional Properties > Java Virtual Machine.
- 4. Enter an initial and maximum heap size; the recommended value is at least 512/512.
- 5. Return to the Java Virtual Machine page.
- 6. Go to **Custom Properties** > New.
- 7. Create a new system property named atg.dynamo.server.name. The value should be the ATG server instance you want to associate with this WebSphere server.
- 8. If applicable, return to the **Java Virtual Machine** page to enable server mode. In the Generic JVM Arguments field, enter -server.
- 9. Save all changes to the master repository; make sure sync node is enabled.

# **Deploying Your Application**

Use the WebSphere Administration console to deploy your EAR file to a cluster. Each Nucleus-based application needs to be installed as follows:

- If you are deploying your web application to a page-serving cluster, that application should also be deployed to a web server instance.
- If you are deploying the application to a cluster that does not serve pages, but runs the application, do **not** deploy the application to a web server instance.
- The web server should only route application requests to instances on the web serving instances node, but non-web serving instances will also run the application.

To deploy an application:

1. Using the Administrative Console for the Deployment Manager, install the application.

If your web application includes a resource reference for your data source, in the WebSphere application installation wizard make sure the reference binding and JNDI name match and are consistent with the name configured in the JTDatasource component (excluding the java:/comp/env prefix).

2. Regenerate the web server plug-in. In the WebSphere Administration Console, go to **Servers** > **Web servers**. Select the entry corresponding to your web server.

On IHS with remote web server management enabled, click **Propagate Plug-In**. The plug-in is propagated automatically.

For all other web servers, click **Update Plug-In**, locate the plug-in on the deployment manager's file system and transfer it to the web server host; overwrite the existing plug-in.

# **General Clustering Information**

The information in this section applies to all application servers.

# Specifying the drpPort Setting

For each ATG server you create, you **must** edit the Configuration.properties file in the <aTG11dir>/home/ servers/servername/localconfig/atg/dynamo directory. Set the adminPort property to the listen port of the corresponding application server, and give the drpPort property a unique value. For example, for the ATG procedit server, you might use these settings:

adminPort=7800 drpPort=8851

Note that DRP ports are not enabled when you run ATG applications, but the port numbers are still needed to identify scenario server instances. Therefore, you must specify a unique value for the drpPort property for the server.

## Setting up localconfig and Server Configuration Files

Set up your localconfig and server configuration files under <ATG11dir>/home/servers and <ATG11dir>/ home/localconfig to configure the default and server specific behaviors of your Nucleus-based application. These files are included in your EAR when it is generated.

- 1. Using either the Configuration and Installation Manager (CIM) or the makeDynamoServer script, create one ATG server configuration for each application server.
- 2. Configure the ATG lock manager server to run the ATG ServerLockManager. (See Enabling the Repository Cache Lock Managers (page 94) earlier in this chapter for more information.)
- 3. Configure the ATG Scenario Manager to run the workflow and process editor servers. There should be exactly one instance of your ATG application running each of these components, and all other instances should be aware of them. (See the *Personalization Programming Guide* for more information.)

**Note:** The JNDI lookup for your data source must be prefixed with java:comp/env/. For example, if your data source has a JNDI name ATGDB, the JNDIName property of the JTDataSource should be java:/comp/env/ ATGDB. Set your transaction manager to use your application server's implementation. See the Creating Data Sources (page 103) section for additional JNDI naming constraints.

## **Unique Components**

The ATG product suite contains several components that must be unique within an ATG server cluster. If you enable and start up more than one instance of these components, errors can result. These unique components are:

- Fulfillment Module (page 107) used by ATG Commerce
- Process Editor Server (page 107) used by the Scenario Manager
- Workflow Process Manager (page 107)

### **Fulfillment Module**

Only one instance of the ATG Commerce Fulfillment module should run on the system. Only one ATG server instance should be started with the ATG Commerce Fulfillment module. To learn more about the Fulfillment module, see the *Commerce Programming Guide*.

### **Process Editor Server**

A cluster of ATG servers should only contain one process editor server. Make sure you have one process editor server configured and that all other ATG instances are aware of it. See the *Personalization Programming Guide* information about setting up scenario servers.

Because running the global scenario server places an additional burden on your ATG server, this instance should not serve any pages.

### Workflow Process Manager

A cluster of ATG servers should always contain exactly one workflow process manager. Make sure only one workflow process manager is configured and that all other ATG instances are aware of it. See the *Personalization Programming Guide* information about setting up a workflow process manager.

## **Enabling Component Backup**

The ATG platform implements a session backup facility that allows you to specify a set of session-scoped or window-scoped Nucleus components and properties that should be backed up after every request. This session backup mechanism saves these components and properties, and restores them when the application server migrates a session to another server.

To enable ATG's backup, set the backingUpSessions property to true in the /atg/dynamo/ Configuration.properties file in the localconfig layer.

backingUpSessions=true

By default, the user's profile and shopping cart (if one exists) are backed up. To back up additional sessionscoped components, set the sessionBackupServerPropertyList property in the /atg/dynamo/ Configuration.properties file to a comma-separated list of Nucleus component properties.

Keep in mind when backing up additional information that the more you back up, the more data the app server must save, which could affect performance.

Each component or property specified in sessionBackupServerPropertyList must implement java.io.Serializable (or Externalizable). If a component is listed without any properties, the entire component is backed up.

# **Synchronizing Server Clocks**

Make sure that all server clocks in a cluster are synchronized. Unsynchronized clocks within the cluster can lead to unexpected results.

# SecurityServlet and ParameterValidator

Oracle ATG Web Commerce includes a component, /atg/dynamo/servlet/dafpipeline/ SecurityServlet, that monitors query parameters and stops processes if they appear suspicious. The SecurityServlet component uses the /atg/dynamo/servlet/security/ParameterValidator component to check query parameters.

The SecurityServlet component is enabled by default. You can disable it by removing /atg/dynamo/ servlet/dafpipeline/SecurityServlet from the insertableServlets property of the /atg/dynamo/ servlet/dafpipeline/DynamoHandler/ component.

# **Configuring ParameterValidator**

Use the configuration properties described in the following table to control the way the ParameterValidator component checks query parameters.

Property	Explanation
alwaysAddDefaults	Controls whether the ParameterValidator component adds sets of default string values to the illegalTagNames and illegalAttributeNames properties.
	Set the value of the property to $true$ to include the default values. Set it to false to omit the default values.
	See Default Values for illegalTagNames (page 109) and Default Values for illegalAttributeNames (page 110).
illegalTagNames	A list of HTML elements that are disallowed in URL query parameters.
	See Default Values for illegalTagNames (page 109).
onlyDisallowIllegalTagNames	Controls whether strings that appear to be HTML elements are allowed in URL query parameters.
	Set the value of this property to true to allow strings that appear to be HTML elements as long as they are not included in the illegalTagNames property.
	Set the value of this property to false to disallow any string that appears to be an HTML element.

Property	Explanation
illegalAttributeNames	A list of HTML attributes that are disallowed in URL query parameters.
	See Default Values for illegalAttributeNames (page 110).
onlyDisallowIllegalAttributeNames	Controls whether strings that appear to be HTML attributes will be allowed in URL query parameters.
	Set the value of this property to true to allow strings that appear to be HTML attributes as long as they are not included in the illegalAttributeNames property.
	Set the value of this property to false to disallow any string that appears to be an HTML attribute.
illegalStrings	A list of specific strings that are disallowed in URL query parameters.
illegalRegexes	A list of Java regular expressions that describe strings that are disallowed in URL query parameters.
overridingValidators	A list of custom components that function as URL query validators for specific conditions.
	Set the value of this property to a comma separated list of the nucleus paths for each of the custom query validator components. For example:
	overridingValidators=/mycompany/validators/ MyValidator,/mycompany/validators/ MyOtherValidator
	See Creating an Overriding Parameter Validator (page 113).

# Default Values for illegalTagNames

The ParameterValidator component includes a set of default string values that it will include in the list of illegal tag names if you set the alwaysAddDefaults property to true. These default string values are shown in the following list.

- img
- script
- iframe
- frame
- applet
- embed
- object

### • meta

### **Default Values for illegalAttributeNames**

The ParameterValidator component includes a set of default string values that it will include in the list of illegal attribute names if you set the alwaysAddDefaults property to true. These default string values are shown in the following list.

- onabort
- onactivate
- onafterprint
- onafterupdate
- onafterupdate
- onbeforeactivate
- onbeforecopy
- onbeforecut
- onbeforedeactivate
- onbeforeeditfocus
- onbeforepaste
- onbeforeprint
- onbeforeunload
- onbeforeupdate
- onbeforeupdate
- onblur
- onbounce
- oncellchange
- oncellchange
- onchange
- onclick
- oncontextmenu
- oncontrolselect
- oncopy
- oncut
- ondataavailable
- ondatasetchanged

- ondatasetcomplete
- ondblclick
- ondeactivate
- ondrag
- ondragend
- ondragenter
- ondragleave
- ondragover
- ondragstart
- ondrop
- onerror
- onerrorupdate
- onfilterchange
- onfinish
- onfocus
- onfocusin
- onfocusout
- onhashchange
- onhelp
- oninput
- onkeydown
- onkeypress
- onkeyup
- onload
- onlosecapture
- onmessage
- onmousedown
- onmouseenter
- onmouseleave
- onmousemove
- onmouseout

- onmouseover
- onmouseup
- onmousewheel
- onmove
- onmoveend
- onmovestart
- onoffline
- ononline
- onpaste
- onpropertychange
- onreadystatechange
- onreset
- onreset
- onresize
- onresizeend
- onresizestart
- onrowenter
- onrowexit
- onrowsdelete
- onrowsinserted
- onscroll
- onsearch
- onselect
- onselectionchange
- onselectstart
- onstart
- onstop
- onsubmit
- onunload
- onunload
- src

### Removing Default Values for illegalTagNames and illegalAttributeNames

To remove an individual, default string value from the illegalTagNames or illegalAttributeNames property, add the string value to either illegalTagNames or illegalAttributeNames with a minus sign character before it.

For example, to remove the value onscroll from the illegalAttributeNames property, add the following to the configuration for that property.

illegalAttributeNames+=-onscroll

**Note:** To remove all of the default values from both the illegalTagNames and illegalAttributeNames properties, set the alwaysAddDefaults property to false.

## **Creating an Overriding Parameter Validator**

Configure overriding parameter validator components if you need to enforce different rules for query parameter validation in specific situations. For example, you can configure validation rules that apply to a specific part of your web application and not to others.

The overridingValidators property of the /atg/dynamo/servlet/security/ParameterValidator component holds a list of the nucleus paths of overriding parameter validator components. When the ParameterValidator validates query parameters, it checks each component named in the overridingValidators property. It calls the canValidateRequest method of each component. If the component returns the value true, ParameterValidator delegates responsibility for checking the query parameters to that component. Once one of the overriding validator components returns true to indicate that it can validate the request, ParameterValidator will stop checking the remaining overriding validator components. If none of the overriding validator components returns true, it will validate the query parameters itself.

To create an overriding parameter validator:

1. Create a custom class that implements the atg.servlet.security.param.RequestParameterValidator interface. See Custom Validator

Class (page 113).

- 2. Create a nucleus component based on the custom class.
- 3. Add the nucleus path of the component to the overridingValidators property of the ParameterValidator component.

### **Custom Validator Class**

Base overriding parameter validator components on a custom Java class that implements the atg.servlet.security.param.RequestParameterValidator interface. The custom class must implement the following two methods:

- canValidateRequest This method takes an atg.servlet.DynamoHttpServletRequest object as its
  argument and returns a boolean value to indicate whether the validator component will validate the request.
- areParamValuesSuspicious This method takes the String name of a query parameter and the parameter values in an array of String objects. It returns a boolean value to indicate whether the query parameters are dangerous and justify stopping the process.

The following example class shows the required methods.

```
package com.mycompany.validators;
import atg.servlet.DynamoHttpServletRequest;
import atg.servlet.security.param.RequestParameterValidator;
public class MyOverridingValidator implements RequestParameterValidator {
    @Override
    public boolean canValidateRequest(DynamoHttpServletRequest pRequest) {
        /* Implement tests to determine whether this component should
            validate a particular request. */
    }
    @Override
    public boolean areParamValuesSuspicious(String pParamName, String[] pValues) {
        /* Implement tests to determine whether the parameter values
            are suspicious. */
     }
}
```

# RedirectURLValidator

Oracle ATG Web Commerce includes a component /atg/dynamo/servlet/pipeline/ RedirectURLValidator, that will prevent URL redirection to hostnames other than the hostname used in the current HTTP request. You can configure a list of hostnames that are allowed by RedirectURLValidator even if they do not match the hostname in the client's original request.

The RedirectURLValidator component includes the properties described in the following table.

Property	Description
allowLocalHost	If this boolean property is set to true, RedirectURLValidator will allow redirection to localhost and synonyms for it. Synonyms for localhost are defined by the /atg/dynamo/service/LocalHostConfiguration component.
	Even when this property is set to true, the LocalHostConfiguration component may not successfully discover all aliases for the localhost. If this happens, add aliases for your localhost to the allowedHostNames property of the RedirectURLValidator component.
allowAllSiteURLs	If this boolean property is set to true, RedirectURLValidator will allow redirection to any hostname and port that match the URLs of your multisite Web sites.
allowedHostNames	RedirectURLValidator will allow redirection to any hostname that you include in this String array property.

Property	Description
allowedHostRegexes	RedirectURLValidator will allow redirection to any hostname that matches one of the regular expressions that you include in this String array property.
enabled	This property is set to true by default. To turn RedirectURLValidator off, set this property to false.

If the RedirectURLValidator component prevents an attempt to redirect a URL, it will generate a server log message similar to the following.

\*\*\*\* Warning Tue Sep 13 15:52:22 EDT 2011 1315943542589 /atg/dynamo/servlet/
pipeline/RedirectURLValidator
Not allowing redirect of the URL "http://bad.com:7103/somedir/somepage.jsp". Adjust
settings of this component
(such as the "allowedHostNames", "allowLocalHost", and "allowAllSiteURLs" properties)
to allow.
Will not warn again for URLs of host "bad.com" for 5 minutes.

# **HttpOnly Attribute for Cookies**

By default, Oracle ATG Web Commerce sets the HttpOnly attribute when it adds ATG cookies to Web application clients. The HttpOnly attribute restricts use of ATG cookies to HTTP or HTTPS requests and prevents access by JavaScript. Note that this attribute does not affect the jsessionid cookie, which is controlled by the application server.

You can control this behavior by setting the createHttpOnlyCookie property of the /atg/dynamo/servlet/ ServletUtil component. If the value of the boolean createHttpOnlyCookie property is true (the default), Oracle ATG Web Commerce will set the HttpOnly attribute when adding cookies. If the value is false, it will not.

# **Recording Login Attempts**

Oracle ATG Web Commerce includes a login auditing feature that records each attempt to log into its administration applications. The login auditing feature writes information about each attempt in a file on your Web application server's file system. Use the audit log file to help detect unauthorized attempts to access your administration applications.

The login auditing feature records attempts to log into the following administration applications:

- Dynamo Server Admin
- Business Control Center

• ATG Control Center (ACC)

The login auditing feature saves audit log files in the Oracle ATG Web Commerce installation or ATG-Data directory. Each administration application records login attempts in the directory for its server. The file path of the audit log is:

<ATG11dir>/home/servers/server-name/logs/authentication\_audit.log

Secure access to the audit log files on the server file system. Make sure that the log files are not deleted or altered by unauthorized users. If the log file is deleted, Oracle ATG Web Commerce will not create a new file until you restart the server.

## **Login Attempt Log Information**

Oracle ATG Web Commerce records information about each attempt to log into its administration applications in a new log file line. The following example shows lines in the log file.

DynAdmin 11/27/2012 15:59:54 -0500 123.4.5.6 HOSTNAME:8850 "admin" SUCCESS ACC 11/27/2012 16:29:07 -0500 123.4.5.6 HOSTNAME:8850 "admin" FAILURE ACC 11/27/2012 16:29:22 -0500 123.4.5.6 HOSTNAME:8850 "UnknownUser" FAILURE DynAdmin 11/28/2012 13:37:38 -0500 123.4.5.6 HOSTNAME:8850 "admin" SUCCESS

Each line in the log file includes a set of information fields that are separated by tab characters. The following table describes the fields in each line of the log file.

Field	Description
Application name	The administration application that the user attempted to log into.
Date and time stamp	The date, time, and time zone of the attempt.
Originating IP address	The Internet Protocol (IP) address from which the user attempted to log in.
Server name	The host name and access port number of the server.
User name	The identifier of the administration application user. You can configure your Web application to record invalid user identifiers for an administration application. By default, this field contains "UnknownUser" if the user identifier is invalid. See Recording Invalid Usernames (page 117).
Result	Indicates whether the user logged in successfully. Contains SUCCESS or FAILURE.

**Note**: In some situations, the standalone version of the ACC may record multiple log file entries for a single attempt to log in.

# AuthenticationMessageTrigger Components

Oracle ATG Web Commerce includes an AuthenticationMessageTrigger component for each of its administration applications. These AuthenticationMessageTrigger components control the way that the administration application records information about attempts to log in.

The AuthenticationMessageTrigger components for the administration applications are:

- Dynamo Server Admin: /atg/dynamo/servlet/adminpipeline/AuthenticationMessageTrigger
- Business Control Center: /atg/userprofiling/ InternalProfileAuthenticationMessageTrigger
- ACC:/atg/devtools/DevAuthenticationMessageTrigger

## **Recording Invalid Usernames**

You can configure your Web application to record invalid user identifiers for an administration application. By default, this field contains "UnknownUser" if the user identifier is invalid.

**Note:** Authorized users may accidentally enter valid passwords in the user name field. If you choose to record invalid user identifiers, the audit log file may contain those unencrypted passwords. Make sure you secure access to the audit log file if it may contain valid passwords.

To disable recording of invalid usernames, set the disableUserName property of the AuthenticationMessageTrigger component for the administration application to true. See AuthenticationMessageTrigger Components (page 117).

If you do not record invalid usernames, you can choose the string that is used in place of them in the audit log file. Set the string in the defaultUserName property of the AuthenticationMessageTrigger component.

## **Rotating Login Attempt Log Files**

Oracle ATG Web Commerce periodically creates a new audit log file and saves the previous file for archiving. It will save a specific number of archived log files and delete older files to save storage space. Rotating log files consists of switching the log file, saving the previous file, and deleting the oldest file. The /atg/dynamo/ service/logging/authentication/AuthenticationFileLogger Nucleus component performs log file rotation for the login audit feature.

By default, the AuthenticationFileLogger component will rotate the audit log file on the first Sunday of each month at midnight.

You can configure the schedule used for log file rotation and the number of archived files that will remain on the server file system. By default, Oracle ATG Web Commerce servers retain 10 archive files. Configure the schedule property of the /atg/dynamo/service/logging/authentication/AuthenticationFileLogger component to adjust the schedule. Configure the maximumArchiveCount property to adjust the number of archived files. See information about RotatingFileLogger components in the *Platform Programming GuidePlatform Programming Guide*.

## **Disabling and Enabling Login Attempt Logging**

Set the auditLogin property of the AuthenticationMessageTrigger component for each administration application to disable or enable login attempt logging. See AuthenticationMessageTrigger Components (page 117).

- Set auditLogin to false to disable login attempt logging.
- Set auditLogin to true to enable login attempt logging.

### **Handling Multiple Servers**

Each administration application for each server instance writes login attempt information to its own authentication audit log file. If your Oracle ATG Web Application consists of multiple production servers, you can collect and concatenate these separate files if your organization requires a single, comprehensive file.

# **Logging Metric Pricing**

Oracle ATG Web Commerce licenses are issued for a specific amount of usage per time period. Usage is defined as pipeline requests. Processes that run ATG applications, such as such as the ACC or TemplateEmail, are filtered out of data collection and are not counted in metric totals. For additional information on filtering, refer to the Filtering Static Resource Types (page 119)section.

The data tracking collects the number of times a pipeline process executes. The data is stored in the /atg/tracking/UsageTrackingRepository. The repository item usageMetric contains the following properties:

- ServerIdentifier The server from which the data is obtained
- usageDate The date associated with the data
- usageHourOfDay The hour of the day associated with the data
- usageValue The pricing metric value

# **Creating a Metric Pricing Report**

You create metric price reports using the Usage Tracking Reporter in the Dynamo Server Admin. It allows you to generate a report starting on the day that you access the service.



Select the number of days for which the report will return data. The report is generated on the same page in a table:

Oracle ATG Application Usage
Date Range: 08/27/2013 - 09/02/2013
Number of unique hosts during reporting time period: 100
Date Request Count
08/27/2013: 5
08/28/2013: 15
08/29/2013: 10
08/30/2013: 15
08/31/2013: 10
09/01/2013: 15
09/02/2013: 30
Peak Day in range of 08/27/2013 - 09/02/2013: 09/02/2013

# **Filtering Static Resource Types**

By default, the following static resource types are filtered, as they do not have full Commerce capabilities:

File Extension	Description
qmd	Bitmap graphic
CSS	Cascading Style Sheet
gif	GIF
htm	HTML
html	HTML
ico	lcon
jpeg	JPEG
bā	JPEG
js	JavaScript
png	Portable Network Graphic

Although not recommended, if you have configured your Web application server to send additional non-Commerce requests through the Commerce pipeline, you can add to the above list using the staticResourceExtensions configuration property. These requests will be passed through a static pipeline

chain that does not log metric pricing. Note that such requests will not have access to most Commerce capabilities.

# **9** Performance Diagnostics

This chapter includes a checklist that can help you identify performance problems in a systematic way. It also describes tools you can use to look for problem areas and discusses how to analyze the results you get from these tools. This chapter includes the following sections:

Performance Troubleshooting Checklist (page 121)
Performance Testing Strategies (page 122)
Locating Performance Bottlenecks (page 123)
Server Hangs (page 126)
Paging and Memory Allocation (page 126)
Detecting File Descriptor Leaks (page 128)
Using URLHammer (page 128)

# **Performance Troubleshooting Checklist**

As your application nears its launch date, you should test the sites as extensively as possible, using tests that simulate the expected site load as realistically as possible.

If you run into performance problems, you can best identify and correct the source of the problem by taking a systematic approach. The following checklist can help you identify the most common sources of performance problems:

- Have you properly configured memory for your Java Virtual Machines? Have you set your -xms and -xmx arguments the same? Do all ATG heap sizes fall within the limits of physical memory?
- Has one or more servers stopped responding? There could be a number of causes, including a Java deadlock. See Server Hangs (page 126).
- Are you seeing many IDExceptions with the message "Too many open files"? You may have a file descriptor leak. See Detecting File Descriptor Leaks (page 128).
- At maximum throughput, look at the CPU utilization, database CPU utilization, I/O activity, and paging activity. See Monitoring System Utilization (page 123).
- If CPU utilization is low, then you may have an I/O or database bottleneck. See Checking for Disk I/O Bottlenecks (page 124), Checking for Network-Limited Problems (page 124), and *Repository and Database Performance* (page 153).

- If CPU utilization is high, then the bottleneck is most likely in the application code. Use a performance profiling tool to try to locate bottlenecks in the code. Review your code to make sure it uses good Java programming practices.
- If paging is occurring, adjust the memory allocated to your Java Virtual Machines. See the Swap Space (page 128) topic in the Paging and Memory Allocation (page 126) section.
- Look at the I/O and CPU utilization of the database. If utilization is high, database activity is probably slowing down the application. See the *Repository and Database Performance* (page 153) chapter.
- · Are you receiving page compilation errors? You may not have enough swap space for page compilation.

If your sites develop performance problems, you need to test several paths through your sites to determine the source or sources of the problems. To generate meaningful test results, you need to test sites with loads that achieve maximum throughput.

# **Performance Testing Strategies**

Since your server may be handling requests for different URLs at the same time, there is no way to get throughput statistics on a page-by-page basis. Instead, you may want to run tests with different sequences of URLs to determine how much throughput varies based on what the user is doing on your sites. Some bottlenecks may occur only in certain page sequences. Your ATG installation includes a test utility named URLHammer that you can use to create and run test scripts. See Using URLHammer (page 128).

# **Graduated Testing of Throughput**

When you test the performance of your sites, you will get the clearest results if you start with very simple tests. Once you know that individual pages or sequences are performing adequately, you can work toward tests that exercise the full range of functionality on your sites. For example, you might structure your throughput tests as follows:

- a minimal, "hello world" page (tests pipeline/request logging)
- home page (tests a single real page)
- login process
- · the 10 most frequently requested pages
- every page

## **Realistic Testing Strategies**

When you load test your sites, be sure to use realistic tests.

- Do not rely on throughput data from a test script in which 100 separate clients make an identical request at the same moment.
- You will want to test cases where request threads do and do not accept cookies. However, be aware that if you run a performance test in which *every* request does not accept cookies, your results will reflect a high performance cost from the need to create a session object for each request. You can easily exhaust memory by creating too many sessions.

# **Locating Performance Bottlenecks**

Once you have brought your ATG system up to maximum throughput, you can look at the components of the system to determine which components are limiting factors in performance.

# **Monitoring System Utilization**

Use a program like top (on Solaris), the Windows Performance Monitor, or a more sophisticated tool to keep track of information like:

- CPU utilization
- paging activity
- disk I/O utilization
- network I/O utilization

A well-performing site will have high CPU utilization when the site is achieving its maximum throughput and will not be doing any paging. A site with high I/O activity and low CPU utilization has some I/O bottleneck.

# **Bottlenecks at Low CPU Utilization**

If your sites have low CPU utilization when achieving maximum throughput, the bottleneck is likely either:

- database limited (if database output is maxed out); see Checking for Database Bottlenecks (page 123)
- disk I/O limited (if I/O output is maxed out); see Checking for Disk I/O Bottlenecks (page 124)
- network I/O limited (if I/O output is maxed out); see Checking for Network-Limited Problems (page 124)
- database or I/O activity in a synchronized method (if database or I/O output is not maxed out); see System Resource Bottlenecks (page 125)

If your site is in this situation, CPU profiling tools are not that useful. Thread dumps taken while the system is under load can give you better information. If you take a few of these, you can get a quick idea of which parts of your application are the slowest. That may help you direct your efforts to the right part of your application. You should be able to tell, for example, whether threads are waiting for a response from the database, a write to the client, or a read from a local file system. If many threads are waiting for the same resource, this is an indication of a potential bottleneck on that resource. Here is some information on what to do about resource bottlenecks for various resources:

# **Checking for Database Bottlenecks**

If your site has low CPU utilization at maximum throughput, check whether the database is limiting performance.

- Get a JVM thread dump and examine it to see if there are many threads waiting for a response from the database.
- Check the CPU utilization and disk I/O utilization of your database server.

• Check the network bandwidth between the ATG server and the database server.

For more information about improving database performance with ATG, see the *Repository and Database Performance* (page 153) chapter.

# **Checking for Disk I/O Bottlenecks**

Make sure that your JVM really is waiting for file I/O, not paging activity. Check for paging with your operating system's monitoring tools.

If the source of slow performance is file I/O, it will show up in JVM thread dumps. The cause could be either some application-specific code that you have, or else the file I/O that ATG does itself.

# **Checking for Network-Limited Problems**

One way to identify network-limited performance problems is by getting your JVM to dump out stack traces while your system is under load. You can tell if your system is network limited because your thread dump will show lots of threads waiting in socket reads or writes.

Some ways to address network-limited problems include:

- Reduce the size of your HTML files by limiting comments and white space or redesigning the content of especially large pages.
- Increase the number of request handling threads. This will not improve the latency experienced by a user who requests a large file, but it will improve total throughput.
- Get a faster network connection.
- · Locate and correct network bottlenecks.

## **Bottlenecks at High CPU Utilization**

If your site CPU utilization is close to 100%, you can use a Java profiler tool like JProfiler or JProbe Profiler to help determine slow points of your code.

In some instances, profilers cannot handle large sites running under load. If so, another way to identify deadlocks and bottlenecks is to get your JVM to dump out stack traces while your system is under load. If you examine 5 or 10 of these stack traces, you can start to see a pattern and find places in your site that are consuming CPU resources or causing deadlocks.

An alternative to stack dumps is the HPROF utility provided with the JDK. See Oracle's Java documentation for information on this utility.

## **Thread Context Switching Problems**

Check how many simultaneous requests are typically being handled when you have a large number of clients trying to access your application. Thread dumps can be useful to see where these threads are waiting. If there are too many threads waiting, your site's performance may be impaired by thread context switching. You might see throughput decrease as load increases if your server were spending too much time context-switching between requests. Check the percentage of System CPU time consumed by your JVM. If this is more than 10%

to 20%, this is potentially a problem. Thread context switching also depends in part on how your JVM schedules threads with different priorities.

You can also reduce overhead from thread context switching by making sure you have at least one CPU for each process involved in handling the majority of requests: one CPU for your HTTP server, one for ATG, one for the database server.

You might see throughput go down as load increases in cases where all of your request handler threads were busy waiting for some resource at the same time. For example, you might have one page on your site that makes a very long-running database query. If you increase the number of clients well beyond 40, you might see all 40 threads waiting for the response to this query. At this point, your throughput will go down because your CPU is idle. You should either speed up the slow requests (perhaps by adding caching of these queries) or increase the number of request threads to increase the parallelism. Of course, at some point, the database may become the bottleneck of your site (which is likely before you have 40 simultaneous queries running).

Context switching can also occur when you have a network protocol which synchronizes too often (such as sending a request and waiting for a response).

Typically, these context switches can be overcome by increasing the parallelism in your site. If there are just too many of these synchronization points, though, this will not work. For example, if you have 40 synchronous RPC calls for each HTTP request, you'd need to context switch processes 80 times for each request if you handled one request at a time. If you handled 2 requests at a time, you'd cut the number of context switches in half. This is in addition to the number of handlers that you'd need to hide any I/O or database activity so the number can add up fast.

## System Resource Bottlenecks

If your site has not maxed out either CPU utilization, database server utilization, or I/O subsystem, the problem may result from synchronized access to one of your system's resources (such as disk, network, database, etc.). This situation occurs when you access this resource from within a synchronized method in Java. All other requests wait for this monitor lock while you do the I/O, thus wasting both CPU and I/O resources. The only ways around this problem are to recode the Java (the right solution) or add more ATG instances (the wrong solution).

The easiest way to find these problems is to test your site when it is serving pages under load and get a JVM thread dump. By examining the thread dump, you may see one thread waiting for a response from the OS (database or I/O) and a set of other threads waiting on a monitor lock that this other thread has.

### **Lower Thread Priorities**

If you have a rarely used feature that uses a lot of CPU resources, you can lower the priority of the thread that handles requests for that feature. Use the setPriority() method of java.lang.Thread to temporarily lower the thread priority. This will result in higher latency for users of that expensive feature, but prevents that feature from hurting performance of other users.

## **TCP Wait Problem on Solaris**

In some testing situations involving a very large number of requests from a single client on the Solaris platform, you may see a dramatic and periodic decline in throughput. You may be able to correct this by modifying the tcp\_close\_wait\_interval setting in the /dev/tcp module. You can do this in two different ways:

- Start ndd, access the /dev/tcp module, and change the value of tcp\_close\_wait\_interval to 60000 (60 seconds).
- Edit the /etc/init.d/inetinit file and include the following line:

# **Server Hangs**

If one or more servers on your site stops responding unaccountably after running under load for a certain period of time, there are a few possible causes:

- HTTP servers not sending requests to your application.
- A Java deadlock.
- Some resource that your application depends on is itself hung (such as the database or some service with which the application communicates via sockets). For example, if a single client opens up hundreds of connections to request pages and then stops reading the response data, this could lock up a server without any real failure of any ATG components.
- You may also have consumed all of the memory in your JVM. If this happens, you'll usually see OutOfMemory errors in your console right before the server hangs. This may appear as a hang because the server will do a garbage collection to reclaim a few bytes, run a few lines of code, then walk through the heap again trying to find another few bytes to reclaim.
- · An infinite loop in some code.

Here are some steps you can take to attempt to identify the cause of the server hang.

- Check the CPU utilization of the machine and particularly the Java process running your ATG application. If CPU utilization is 100%, it is either an OutOfMemory problem or a CPU burning thread.
- Check the server logs to see if any errors right before the hang indicate why the server has failed. You might see a "server not responding" message or an OutOfMemory error.
- · Get a thread dump from your Java VM. A thread dump can help you recognize all of these problems.

If all threads are waiting in system calls such as socket read/write, then they are waiting for a resource to respond (for instance, the database or the network). You should look to this resource for answers. If the resource is a database, try using a third party database tool to make a query. It is possible that the tables used by your ATG application are locked by some other operation so they will wait until that operation has completed.

# **Paging and Memory Allocation**

If you see any paging activity, increase system memory or decrease the size of the JVMs. Be aware that decreasing heap sizes may increase the overhead of garbage collection. Each time a full garbage collection is performed, all of the memory needs to be scanned for garbage. Garbage collections occur more frequently with smaller heaps, which could waste CPU time.

You can check the size of your JVM heaps or cause garbage collection with the ATG VMSystem component at:

http://hostname:port/dyn/admin/nucleus/VMSystem/

# **Garbage Collection**

Set your JVM to monitor how much time is spent doing garbage collection. You can do this by adding the -verbose:gc parameter to the JAVA\_ARGS passed to the JVM when ATG starts up. The -verbose:gc option causes the JVM to output a message each time garbage collection is performed, including:

- how much memory was reclaimed
- the amount of free memory
- the total heap size
- how much time the garbage collection operation took

If you see your garbage collections happening too often or occupying a significant percentage of your CPU, you should either increase the Java heap size arguments or look for places in your application that are allocating memory unnecessarily.

If the garbage collection takes a very long time to complete, you may have configured your heap size to be too large for the amount of memory your system has. If your system is spending more than 20% of its CPU time on garbage collection, you have a significant performance problem that must be corrected. Use your OS monitoring tools to see if you are paging and check the process size of all of the processes running on your system. Compare this with the physical memory of your machine.

If your heap is very large, your garbage collections may occur very infrequently but may take a long time (30 seconds or more) to complete. This is a structural limitation of Java that is difficult to work around. When a full garbage collection occurs, it typically acquires the heap lock, which prevents all requests from being served during this time interval. You can potentially reduce the time of these garbage collections by forcing them to occur more frequently.

If your garbage collections take a significant percentage of your overall CPU time, you may have places in your code that allocate memory inefficiently. It is a good idea to reuse objects where possible, rather than creating them over and over again. You can use a memory profiler to determine where and how much memory is allocated by which places of the code. Only the allocation side of the garbage shows up in the stack traces and profiling. You have to factor in the time spent reclaiming garbage as well. You can also use the Performance Monitor to trace the memory allocation of various operations in your system. See Performance Monitor (page 137) in the *Monitoring Site Performance* (page 137) chapter.

### Memory Leaks

When the Java VM runs low on memory, you should see two behaviors:

- · very slow performance, as garbage collections occur more frequently and absorb a greater share of CPU time
- occasional OutOfMemory errors.

To confirm the presence of a memory leak, add -verbose:gc to your JAVA\_ARGS and monitor the number of sessions on your site (see the Garbage Collection (page 127) section for details). If you see free memory decrease over time as your site has a constant number of sessions, you may have a memory leak. Before deciding that you have a memory leak, make sure you have given the system enough time to fill all caches and reach a stable state after startup.

Memory leaks in Java are caused by data structures that hold onto objects that are no longer needed. This is often due to a Collection (such as a Vector or Hashtable) that is not coded correctly. For example, if you store objects in a Hashtable using a session ID as a key, but you do not remove these objects when the session expires, this Hashtable will grow without bounds.

You can use memory profilers to help find these errors. Another way to detect when a Hashtable or Vector is growing without bounds is to use a modified version of the standard Hashtable and Vector that is instrumented to print a message each time the 10000th, 20000th, etc. element is added. Of course, if you use a different Collection class, this will not find that problem.

One frequent cause of Java memory leaks is the use of an addXXXListener() method without a corresponding removeXXXListener() method. Review your code to make sure you have not made this mistake.

# **Swap Space**

In order for ATG to fork a javac compiler to compile a JHTML page, it requires two times the current process size in swap space for a short period of time until it executes the new process. If you receive an error message like this:

```
/atg/dynamo/servlet/pagecompile/PageCompileServlet
atg.servlet.pagecompile.PageCompileResources->
pageCompileServletErrorCompiling :
Error compiling page: /path of page> :
Unable to execute the command '/page compile command>'
Make sure that you have the 'bin' directory for your JDK in your PATH
variable before starting ATG and that you have enough swap space.
```

then you probably do not have enough swap space for page compilation. Increase your swap space.

# **Detecting File Descriptor Leaks**

It is important to ensure that files that are opened always get closed. Failing to close files can result in file descriptor leaks. You can detect a file descriptor leak in two different ways:

- You may notice a lot of IOExceptions with the message "Too many open files."
- During load testing, you periodically run a profiling script, such as lsof (on UNIX), and you notice that the list of file descriptors grows continually.

File descriptor leaks can also lead to a variety of failures on attempts to open properties files, sockets, etc. If your error log contains a lot of chaotic-looking error messages, the presence of a file descriptor leak is one thing to check.

# **Using URLHammer**

The URLHammer program is a Java utility. URLHammer makes repeated page requests, allowing you to simulate the effects of load on your ATG application. The utility detects and reports HTTP errors, but performs no

validation of the HTTP response itself. URLHammer supports HTTP cookies. You can use it to submit forms by playing back scripts (see Using the Recording Servlet (page 150)). URLHammer is run from the DOS or UNIX command line. It runs in a separate JVM from ATG. For the best results, we recommend running URLHammer on a separate machine from the server you are testing.

To run the URLHammer program:

- 1. Set your CLASSPATH to include the directory <atgl1dir>/Das/lib/classes.jar.
- 2. Run the following command:

java atg.core.net.URLHammer [arguments]

For example:

java atg.core.net.URLHammer http://examplehost:8840/ 5 10 -cookies

This creates five different threads, each of which represents a separate session that requests the specified URL 10 times (50 requests total).

You can configure URLHammer using several command line arguments that are described below; you can also use the -usage argument to get the current list of arguments. The -cookies argument makes URLHammer parse the Set-cookie headers and return cookies that it receives in all subsequent requests. If you do not use the -cookies argument, then ATG creates new sessions for each request. Each thread has its own set of cookies. Thus, the above example creates 5 sessions and executes 10 requests in each.

### **Command Line Arguments**

URLHammer takes a number of command line arguments so that you can implement your tests in the manner that best fits your site. Use the following syntax:

```
java atg.core.net.URLHammer URL | script_pathname threads iterations
[optional arguments]
```

The following URLHammer arguments are required:

<b>Required Arguments</b>	Description
URL or <i>script_pathname</i>	The URL to use in each request. The URL must begin with http://. (Note: https is not supported.) If you use the -script argument, then instead of a URL, specify the pathname of the script to execute. See The - script Argument (page 132).
threads	Number of independent thread connections to create to the HTTP server. Use a value from 1 to 20. All threads run concurrently.
iterations	Number of requests to issue on each thread. If the <code>-script</code> argument is used, this represents instead the number of times each thread executes the entire script.

The following URLHammer arguments are optional:

Optional Arguments	Description
-addCookie name=value	Enables you to set a cookie. For example:
	-addCookie FOO=Zippy
-addHeader name=value	Enables you to define a header. You can define multiple headers; for example:
	-addHeader LOGIN=Zappa -addHeader PASS=nan00k
-cookies	Returns Set-cookie headers sent by the server. Note: <pre>path= and</pre> expires= are not processed by URLHammer.
-htmlStats <i>HTML file</i>	Output statistics to the specified HTML file. This argument gives detailed statistics about the amount of time consumed by each individual URL you requested. It also gives summary statistics about the number of errors encountered. By default, URLHammer outputs these statistics to the console.
-maxRequests	Limits the number of redirects that can be generated.
-nopause	Use only with the -script argument. Ignores pause information in script files by default. When using the Recording Servlet, the time between the server's receipt of one page request and the server's receipt of the next request is recorded in the script file (in milliseconds). Each URLHammer thread sleeps for this number of milliseconds before requesting the URL. If you use the -nopause argument, URLHammer instead requests each subsequent URL as soon as the previous output is received.
-password	Use only with the $-user$ argument. Supplies a user password if needed to log in to any pages.
	<b>Note</b> : Entering passwords as command line arguments entails some security risks. Do not enter passwords as command line arguments in situations where security is a top priority.
-pause	Use only with the -script argument. Pause for the specified time between each request (number of milliseconds). For example, the following argument causes URLHammer to pause 1 second between each request: -pause 1000 If you use a negative value, then URLHammer pauses for a random amount of time, not to exceed the absolute value of the value you use. For example, the following argument causes URLHammer to pause a random amount between 0 and 550 milliseconds between each request:
	-pause -550

Optional Arguments	Description
-randomStop	Simulates the browser's stop button by randomly closing the connection to the server for 20% of the requests.
-recordAll	Outputs statistics for each request. Use this argument with the -htmlStats argument and the HTML file will contain the statistics broken down for each request, as well as in summary form. It also keeps track of which requests had errors and prints (error) next to the time for that request.
-runningStats	Prints information periodically for each thread. This allows you to get an idea of how long runs are proceeding.
-script	Instead of making a request to a single URL, each thread instead executes a script of user browser actions. See The -script Argument (page 132).
-server name:port-number	Name of the server and the port number to use if you are using the - script argument. If you do not specify a server, localhost:80 is used as the default.
-stop <n></n>	Simulates the browser's stop button by closing the connection to the server for $\langle n \rangle$ % of the requests. This argument is useful to make sure that your site is robust with respect to aborted requests.
-substitute	Use only with the <code>-script</code> argument. Performs keyword substitution in your script file. This facility allows you to generate more flexible form processing scripts. You can place keywords <u>RANDOM</u> , <u>COUNTER</u> , and <u>TIME</u> into your script file's URLs and POST data sections. (Note that these keywords are preceded and followed by two underscore characters.)
	Before each request, URLHammer substitutes these keywords with a random string, a continually incremented counter, or the current time in milliseconds. You can use this argument, for example, to generate unique login IDs when load testing login forms.
-user	Use only with the -password argument. Supplies a username if needed to log in to any pages.
-verbose	Dumps the complete output of the request (including request headers). This argument is very valuable when testing a new script or the first time you execute a command, so that you can inspect the output generated.

# **URLHammer Examples**

The following examples use UNIX syntax. Adjust the syntax accordingly for Windows. We also presume that your CLASSPATH includes <code>\$DYNAMO\_HOME/lib/classes.jar</code>.

### **Checking Availability of ATG**

Suppose you want to see whether your ATG application is responding. A single request on a single thread, using a very simple page, would be sufficient for this test:

java atg.core.net.URLHammer http://hostname:8080/index.jsp 1 1

If your application is responding, you should see output like the following (the times will vary):

```
Time = 521 ms (1.91 requests/s; average latency = 521 ms)
0 errors out of 1 request
```

The time output reports the total elapsed time in milliseconds, the number of requests per second, and the average time per request, in milliseconds.

### **Generating a Typical Load**

Using multiple concurrent threads, each making repeated requests, will generate a sustained load on the ATG server:

java atg.core.net.URLHammer http://hostname:8080/test.jsp 10 25

In this example, 10 threads are used, each making 25 requests, for a total of 250 requests, each of which uses its own session.

### **Playing Back a Script**

The previous examples generate a number of simultaneous requests for the same page. For a more realistic usage scenario, you can use URLHammer to run a script of more complex user behavior. A script file can be as simple as a list of relative URIs (one per line). See Recording a Script (page 133) for a simple way to construct a script, and Editing a Script (page 134) for details on the syntax and semantics. The following command plays back the script myscript.txt one time, using one thread, making requests from the default ATG server port:

java atg.core.net.URLHammer myscript.txt 1 1 -script -server examplehost:8080

## **The -script Argument**

The -script argument treats the URL argument as the name of a script file on the local system. This script file can contain any of the following:

- URLs
- URLs with POST data
- URLs with POST data and session ID arguments

You can write your own script files, or you can use ATG's Recording Servlet, which records script files that replay a previously recorded set of user actions. See Using the Recording Servlet (page 150) in the *Monitoring Site Performance* (page 137) chapter.

Script files are line-oriented ASCII text files. Each line can be in one of the following formats:

```
#include another_script_file
```

```
URL
```

URL time\_in\_milliseconds

```
URL time_in_milliseconds #_lines_of_post_data
post_data
post_data
post_data
...
```

```
URL time_in_millis #_lines_of_post_data session_id
post_data
post_data
post_data
...
```

If a line specifies a number of lines of POST data, URLHammer reads that number of lines and passes them as URLencoded POST data to the specified URL. Typically, lines of this form are generated by the Recording Servlet.

Note that the URLs in a script file need not contain the http://hostname:port prefix, since the full URL can be constructed using the host and port number specified by the -server command line argument. This allows you to reuse the same script to test different servers.

## **Recording a Script**

You can use the ATG Recording Servlet facility as an aid in constructing a test script. This is particularly helpful in tests of form submission (such as requests with the POST method) because the script must supply the data for the form. Follow these steps to record a test script:

- 1. Open the /atg/dynamo/servlet/pipeline/RecordingServlet component in the ACC.
- 2. If the RecordingServlet component is not running, start it by clicking the Start button.
- 3. Change the live value of the recording property to true.
- 4. Perform the actions you wish to record (for example, page requests and submitting forms).
- 5. Change the live value of the recording property to false.
- 6. Copy the <ATG11dir>/home/logs/record.log file to another filename to save its contents.

You can also use the Recording Servlet with the Dynamo Server Admin:

1. Browse the Recording Servlet in the Dynamo Server Admin:

http://hostname:port/dyn/admin/nucleus/atg/dynamo/servlet/pipeline/ RecordingServlet

- 2. Click the name of the recording property.
- 3. Set the value to true and click the Change Value button.
- 4. Perform the actions you wish to record (for example, page requests and submitting forms).
- 5. Return to the Recording Servlet page in the Dynamo Server Admin.
- 6. Click the name of the recording property.
- 7. Set the recording value to false and click the Change Value button.

8. Copy the <ATG11dir>/home/logs/record.log file to another filename to save its contents.

See also Using the Recording Servlet (page 150) in the Monitoring Site Performance (page 137) chapter.

# **Editing a Script**

A request in a script file is specified using this syntax:

Relative\_URI [ Delay\_ms [ POST\_lines [ Session\_ID ] ] ]

### where:

- Relative\_URI is the relative URI of the file to request, with optional parameters
- Delay\_ms is the number of milliseconds to pause
- POST\_lines specifies the number of following lines to use as POST data
- Session\_ID designates an ATG session ID

The URIs in a recorded script must be relative to the document root. Note also that when the -cookies option is used, all of the session IDs in a script are replaced by the current session ID for the given thread; each thread will have a new unique session created for it.

### **Comments in Scripts**

A line that begins with the # character is considered a comment and will be ignored (with the exception of lines that begin with #include; see next section). You can add comments to your scripts to document the purpose, author, usage, etc.

### **Including Scripts within Scripts**

A line that begins with the #include keyword includes a specified script within the current script. For example:

#include subfile.txt

adds the contents of the script subfile.txt to the current script at that position. This is especially useful for simplifying a long script into a hierarchy of easy-to-understand parts.

# **URLHammer Source Files**

ATG includes the source for URLHammer, together with source for implementation classes, in:

<ATG11dir>/DAS/src/Java/atg/core/net/

You may want to modify or extend URLHammer for your own testing purposes. However, ATG does not guarantee backward compatibility in future releases of URLHammer. If you make modifications to the code, you should change the class and package names to avoid potential conflicts with future versions we may release.
# **10 Monitoring Site Performance**

ATG includes a variety of diagnostic and administrative tools to help you keep your site up and running smoothly. This chapter covers the following topics:

Performance Monitor (page 137) Using the Configuration Reporter (page 145) Using the VMSystem Component (page 149) Using a Sampler (page 149) Using the Recording Servlet (page 150)

### **Performance Monitor**

ATG's Performance Monitor component provides a tool you can use to monitor the performance of regions of your code. To use the Performance Monitor:

- Instrument your Java code with static methods that enable the Performance Monitor to gather information about performance (see Adding PerformanceMonitor Methods to your Code (page 137)).
- View the Performance Monitor page in the Dynamo Server Admin to inspect information gathered (see Viewing Performance Monitor Data (page 140)).

The Performance Monitor can run in different modes. In normal (default) mode it causes negligible overhead, but allows you to globally turn on one or more monitoring options which give more diagnostic information. These monitoring options would typically be used during load testing but are not suitable for running on a live site under heavy load. See Performance Monitor Modes (page 139).

### Adding PerformanceMonitor Methods to your Code

To enable the Performance Monitor to monitor a section of your Java code:

- 1. Import the atg.service.perfmonitor.\* package.
- 2. Declare an OpName parameter to label the section of the code. This parameter is displayed in the Performance Monitor page under the **Operation** heading.
- 3. (Optional) Declare a parameter name if you want to gather data on individual executions of an operation.

- 4. Call the startOperation method at the beginning of the operation whose performance you want to be able to measure.
- 5. Call the endOperation method at the end of the operation whose performance you want to be able to measure.
- 6. Optionally, call the cancelOperation method if an exception occurs. This causes the results of the current execution to be ignored.

For details about the Performance Monitor's startOperation, endOperation, and cancelOperation methods, see Methods for Storing Performance Data (page 143).

For example:

```
String opName = "render jsp";
String parameter = "foo.jsp";
boolean exception = false;
PerformanceMonitor.startOperation(opName, parameter);
try {
    ... code to actually render foo.jsp
} catch (Exception e) {
    PerformanceMonitor.cancelOperation(opName, parameter);
    exception = true;
} finally {
    if (! exception)
        PerformanceMonitor.endOperation(opName, parameter);
    }
```

These methods can be nested with different or the same opNames. For example:

```
private final String RENDER_JSP = "Render JSP page";
private final String EXECUTE_SQL = "Execute SQL Query";
private String mPageName = "page.jsp";
private String mSQLQuery = "select * from table";
PerformanceMonitor.startOperation(RENDER_JSP, mPageName);
... source code to start render
PerformanceMonitor.startOperation(EXECUTE_SQL, mSQLQuery);
... source code to read from table 1 in database
PerformanceMonitor.startOperation(EXECUTE_SQL);
... source code to read from database
PerformanceMonitor.endOperation(EXECUTE_SQL);
... more source code to read from table 1 in database
PerformanceMonitor.endOperation(EXECUTE_SQL);
... more source code to read from table 1 in database
PerformanceMonitor.endOperation(EXECUTE_SQL, mSQLQuery);
... more source code to finish render
PerformanceMonitor.endOperation(RENDER_JSP, mPageName);
```

Note that the calls to startOperation are nested within other calls to startOperation. You must place the endOperation and cancelOperation calls in the code in opposite order that the startOperation calls were placed. If this requirement is not followed, then the endOperation or cancelOperation call throws a PerfStackMismatchException. This exception tells you that the calls to endOperation are not being matched up. Either they were not called in the correct order or the arguments were not exactly the same as those that were passed into the methods.

To ensure that endOperation is always called, wrap the Performance Monitor methods in a try ... finally block, as in this example:

```
boolean exception = false;
try {
    PerformanceMonitor.startOperation(OP_NAME);
    performOperation (pParameter);
} catch (Exception e) {
    PerformanceMonitor.cancelOperation(OP_NAME);
    exception = true;
} finally {
    try {
        if (!exception)
            PerformanceMonitor.endOperation(OP_NAME);
        } catch (PerfStackMismatchException e) {
            System.out.println(e);
        }
}
```

### **Performance Monitor Modes**

The Performance Monitor code can run in one of four modes:

- DISABLED. When the Performance Monitor is disabled, its diagnostic methods immediately return without doing any additional work.
- NORMAL. In this mode, the Performance Monitor keeps track only of the current stack of operations. This mode is useful in identifying the location in the code of hung or active threads.
- **TIME**. In this mode, in addition to the current operation stack, the Performance Monitor maintains dictionaries for each operation. These dictionaries store the number of times each operation has been performed, and the minimum, maximum and average time to process that operation.
- TIME mode is not meant to be used on a live system for an extended period of time. This mode is for gathering data on the amount of time spent in various parts of the code.
- MEMORY. In this mode, the Performance Monitor maintains the information specified for NORMAL and TIME
  mode. In addition, the Performance Monitor maintains dictionaries that store the number of times each
  operation has been performed, and the minimum, maximum and average amount of memory required to
  process that operation. These statistics are estimates and do not take into account asynchronous processing
  activity that may be occurring. Do not rely on data from only one or two samples, since the Performance
  Monitor may generate anomalous data that can be ignored.

MEMORY mode causes all requests to the server to be serialized and could possibly cause deadlock. This mode is provided for diagnostics during development only and is not suitable for use on a live system.

### **Setting the Mode**

Set the Performance Monitor's operating mode at the Performance Monitor Configuration page of the Dynamo Server Admin:

```
http://hostname:port/dyn/admin/atg/dynamo/admin/en/
performance-monitor-config.jhtml
```

Click the radio button for the mode you want, and then click the Change Mode button.

You can also set the Performance Monitor's operating mode by setting the mode property of the component at / atg/dynamo/service/PerformanceMonitor. The value of the mode property is an int corresponding to the mode:

mode	int value
disabled	0 (default)
normal	1
time	2
memory	3

### **Viewing Performance Monitor Data**

You can view the information collected by the Performance Monitor on the Performance Monitor's page of the Dynamo Server Admin at:

http://hostname:port/dyn/admin/atg/dynamo/admin/en/
performance-monitor.jhtml

This page displays any information recorded by the Performance Monitor. Under the **Threads** heading, the Performance Monitor page displays the operation stack of the current thread.

If you have configured the Performance Monitor to run in TIME mode, then the Performance Monitor page displays under the **Performance Data** heading a Time Performance Data table with a list of operations that have been recorded (such as Invoke Servlet, Compile Page, Service Request, etc.) along with the number of times the operation was executed and the minimum, maximum, average, and total time for each.

For example, the Time Performance Data table might look like this:

Operation	Number of Executions	Average Execution Time (msec	Minimum Execution Time (msec	Maximum Execution Time (msec	Total Execution Time (msec
Handle HTTP Request	1	223	223	223	223
Invoke Servlet	4	8	0	19	35
Invoke Form Handler	1	108	108	108	108
Compile Page	1	3	3	3	3
Service Request	1	123	123	123	123

The name of each operation is a link to another administration page that provides the detailed parameterized information, if any (for example, for each URL, the number of times requested, the minimum, maximum, and average times).

If you have configured the Performance Monitor to run in MEMORY mode, then the Performance Monitor page displays under the **Performance Data** heading Time and Memory Performance Data tables that includes all the TIME mode information described above, and in addition displays the minimum, maximum, average, and total *memory* used by each operation.

### Instrumented ATG Classes

Several common ATG operations have already been instrumented with Performance Monitor startOperation and endOperation methods. By default, this includes all scheduled jobs handled by the Dynamo Scheduler. These operations appear grouped together under the line **Scheduled Jobs** in the Performance Monitor page. Clicking on this link lets you drill down and see the statistics for each job separately. If you do not want performance monitoring of scheduled jobs, you can set the Scheduler's performanceMonitorEnabled property to false to disable this behavior.

In addition, ATG's instrumented methods include:

Class Name	Method	Operation Name
atg.targeting.TargetingArray	getTargetArray()	Perform Targeting
atg.servlet.pipeline. HeadPipelineServlet	service()	Service Request
atg.servlet.pagecompile.SubServlet	<pre>serviceByName()</pre>	Invoke Servlet
atg.servlet.pagecompile. PageSubServlet	<pre>serviceServlet()</pre>	Invoke Servlet
atg.servlet.pagecompile. PageCompileServlet	service()	Render Page
atg.service.resourcepool. MonitoredStatement	<pre>executeQuery() executeUpdate()</pre>	Execute Query Execute Update
atg.service.resourcepool. MonitoredPreparedStatement	<pre>executeQuery() executeUpdate()</pre>	Execute Query Execute Update
atg.server.http.HttpConnection	handleRequest()	Handle HTTP Request
atg.nucleus.NucleusNameResolver	createFromName()	Create Component
atg.droplet.DropletEventServlet	sendEvents()	Invoke Form Handler
atg.service.pipeline.PipelineManager	runProcess()	Run Pipeline Chain
atg.service.pipeline.PipelineLink	runProcess()	Run Pipeline Processor
atg.adapter.gsa.GSARepository	createNewItem()	GSA createItem

Class Name	Method	Operation Name
atg.adapter.gsa.GSAItemDescriptor	getPersistentItem()	GSA Uncached getItem

### **Performance Monitor API**

The main class for the Performance Monitor is atg.service.perfmonitor.PerformanceMonitor.This class contains all the static methods for interacting with the Performance Monitor. In addition, it stores the data structures that contain the performance data. The Performance Monitor's methods have the following functions:

- Methods for Controlling the Performance Monitor (page 143)
- Methods for Storing Performance Data (page 143)
- Methods for Accessing Stack Data (page 144)
- Methods for Accessing Performance Data (page 144)
- Exception Summary (page 145)

The PerformanceMonitor component contains two primary data structures. One stores the runtime stack data for all registered threads. The other stores the performance data for operations and parameterized operations on those registered threads.

#### **Runtime Stack Data Structure**

This structure is a Hashtable where the key is a registered thread and the element is a java.util.Stack of atg.service.perfmonitor.PerformanceStackData objects. This data is what is recorded and tracked in NORMAL mode. When a stack becomes empty, then all the performance operations have completed in that thread. This data structure is used in all modes except for DISABLED.

#### **Performance Data Structure**

This data structure stores all the time and memory performance related data for operations and parameterized operations. It is only used when the mode for the Performance Monitor is set to TIME or MEMORY. The structure is a Hashtable where the key is an operation name and the element is a PerformanceHashtable. The PerformanceHashtable is a subclass of Hashtable. In addition to providing the services of a Hashtable, it also stores the totals for all the parameterized operations contained in the Hashtable in an atg.service.perfmonitor.PerformanceData object. The Hashtable in the superclass of this object contains the parameterized operation name in the key and a PerformanceData object as the element.

There are also two data structures for holding pools of PerformanceStackData and PerformanceData objects. These exist to avoid allocation and improve performance. When startOperation is called, a new PerformanceStackData object is retrieved from the pool, populated and pushed on the stack. When endOperation is called, the top element in the stack is compared for mismatch and then popped off the stack, assuming there was no mismatch. At this time, the corresponding PerformanceData object for the operation in the PerformanceStackData object which is stored in the performance data structure is updated with number of times executed and total execution time (min and max will also be updated if the most current execution requires it). In addition, the global PerformanceData object for the operation is updated. If endOperation is updated. If endOperation is updated. If the PerformanceData object for the operation or parameterized data does not exist, then a new PerformanceHashtable will be created and PerformanceData object will be retrieved from the pool and inserted.

#### **Methods for Controlling the Performance Monitor**

You can control the Performance Monitor programmatically using the methods listed in this section. Most often, however, you will configure the Performance Monitor using the Performance Monitor Configuration page in the Dynamo Server Admin (http://hostname:port/dyn/admin/atg/dynamo/admin/en/performance-monitor-config.jhtml) or through the ACC.

public int getMode();

Returns the mode that the Performance Monitor is running in. The return value is an int that refers to one of DISABLED, NORMAL, TIME, OR MEMORY.

public void setMode(int pMode);

Allows a user to dynamically set the mode of the Performance Monitor. The mode is normally set in the Performance Monitor's properties file, but can be changed during runtime using the ACC.

public void resetPerformanceData();

Resets all the performance data back to 0. This means that the TIME mode and MEMORY mode minimum, maximum, and total statistics will be reset to 0 for all operations and parameterized operations.

### **Methods for Storing Performance Data**

The startOperation and endOperation methods designate the start and end of an operation. These methods need to bracket the code that performs the designated function.

public static final void PerformanceMonitor.startOperation
(String pOpName);

public static final void PerformanceMonitor.startOperation(String pOpName, String
pParameter);

The startOperation method tells Performance Monitor that a new operation is starting. The pOpName parameter is the name of the operation. This parameter should be short and as descriptive as possible. The next parameter, pParameter, is optional data that gives the Performance Monitor more detailed information on exactly what object it is performing the given operation on. The parameterized version of this method records data for the operation on the given parameter and the global operation. The non-parameterized version of this method records performance data to the operational level only.

public static final void PerformanceMonitor.endOperation(String pOpName)
 throws PerfStackMismatchException;

public static final void PerformanceMonitor.endOperation();

The endOperation method tells Performance Monitor that a previously started operation has come to completion. The pOpName parameter must be exactly the same as the pOpName parameter that was passed into the corresponding startOperation method. The pParameter is optional data which gives the

Performance Monitor more detailed information on the object it completed the operation on. The call to endOperation must have exactly the same parameters that the call to startOperation did. Otherwise, a PerfStackMismatchException (an extension of RuntimeException) is thrown.

You can also call endOperation without any arguments to mark the end of the most recent operation for which monitoring has started, but not yet ended. In this case, there is no need to supply it with the same arguments that were passed at the start of the operation. Accordingly, it will never throw an exception.

The cancelOperation method cancels an operation and discards any performance statistics.

public static final void PerformanceMonitor.cancelOperation(String pOpName) throws PerfStackMismatchException;

public static final void PerformanceMonitor.cancelOperation(String pOpName, String pParameter) throws PerfStackMismatchException;

The cancelOperation method tells Performance Monitor that a previously started operation should be cancelled. Canceling an operation means that statistics from this operation execution are discarded. The pOpName parameter must be exactly the same as the pOpName parameter that was passed into the corresponding startOperation method. The pParameter is optional data which gives the Performance Monitor more detailed information on the object on which it completed the operation. The call to cancelOperation must have exactly the same parameters that the call to startOperation did. Otherwise, a PerfStackMismatchException is thrown.

The isEnabled method indicates whether the Performance Monitor is enabled or not.

public static final boolean PerformanceMonitor.isEnabled();

Returns a boolean that specifies whether the Performance Monitor is enabled or not.

#### **Methods for Accessing Stack Data**

The stack data contains the runtime location of all the threads currently registered in the Performance Monitor. This data is stored in objects of type PerformanceStackData. The PerformanceStackData object is contained in a java.util.Stack object. The PerformanceStackData object alone is not useful; it becomes useful when it is placed inside the context of a java.util.Stack. The PerformanceStackData has the following methods you can use:

public String getOperation();
Returns the operation name within the PerformanceStackData object.

public String getParameter();
Returns the parameter operation name within the PerformanceStackData object.

```
public long getStartTime();
```

Returns the start time of the operation as the number of milliseconds since Jan 1, 1970. This method is used internally by the Performance Monitor and is not very useful outside of it, but it is provided.

#### **Methods for Accessing Performance Data**

The performance data is stored in read-only properties in objects of type PerformanceData. This object is a JavaBean that contains the following data:

Property	Description
minimumExecutionTime	Minimum execution time
maximumExecutionTime	Maximum execution time
totalNumberOfExecutions	Number of times operation has been executed
averageExecutionTime	Total execution time
minimumMemoryRequired	Minimum memory required
maximumMemoryRequired	Maximum memory required
totalMemoryRequired	Total memory required

The PerformanceData object has get methods that correspond to each of these properties. The average execution time and memory required can be derived from number of times and total execution time or memory required.

### **Exception Summary**

```
PerfStackMismatchException
```

Thrown when endOperation is called with out of order arguments or different arguments than what was expected.

## **Using the Configuration Reporter**

The ATG product suite has vast possibilities for configuration and customization. These possibilities are multiplied when you consider the different platforms, HTTP servers, and database software you might use in your site. These myriad possible combinations can make it difficult to describe your Nucleus-based web application's overall configuration in a concise way. The Configuration Reporter compiles a description of your ATG configuration, so that useful troubleshooting information is gathered in a single place.

The Configuration Reporter can generate reports in several different forms that you can use to help identify configuration problems. These reports also make it possible to e-mail configuration information to ATG support.

You can access the Configuration Reporter from the link on the Dynamo Server Admin home page, or navigate to it directly at:

http://hostname:port/dyn/admin/atg/dynamo/admin/en/conf-reporter.jhtml

The heart of the Dynamo Configuration Reporter is the service located at /atg/dynamo/service/ ConfigurationReporter. The Configuration Reporter service works by browsing the hierarchy of components, starting at the root, gathering information, and outputting it in various formats.

### **Configuration Reports**

The Configuration Reporter can generate the following four reports:

- HTML Component Browser Report A report on the components in the component hierarchy in the form of HTML files. This report is more or less like printing out the entire Dynamo Administration Component Browser.
- Bean Representation Report A list of each ATG component, with each of its properties and property values, in the form of a serialized file.
- Property Representation Report Like the Bean Representation Report, but it includes only those components and properties whose values have been set through properties files (including properties set in the ACC).
- · CONFIGPATH Report A text file that lists the configuration path of the ATG server.

### **Excluding Components from the Configuration Report**

By selecting a custom report, rather than a basic report, you can configure the Configuration Reporter to exclude selected components:

1. Set the restrictedComponents property of the /atg/dynamo/service/ConfigurationReporter service. This property is a comma-separated list of Nucleus component paths of components and directories that should be excluded from configuration reports.

If a Nucleus component path included in the restrictedComponents property is a folder, neither it nor any of its children will be included in custom configuration reports.

Make a file that lists the components to include. Go to the Output Dynamo Component Hierarchy to File page at:

http://hostname:port/dyn/admin/atg/dynamo/admin/en/ config-reporter-output-hierarchy-titled.jhtml

- 3. In the Output File field, enter the pathname of a file to receive the list of components to include.
- 4. Click the **Create Dynamo Component File** button. The Configuration Reporter will generate the component list and output it to the file you specified in step 3.
- 5. Select Custom Report from the report page for the type of report you want to generate.
- 6. In the **Component file** field, enter the pathname of the file you created in step 4.
- 7. In the Serialization output file field, enter the pathname of a file to receive the serialized report file.
- 8. Click the Create Serialization Output File button.

The Bean Representation Report and Property Representation Report generate information in the form of serialized files. After you create a serialized report, you can output a more readable version of the information, using the XML Representation Report options:

- 1. Check the **Output all property values** box if you want to view the property names and values, and not just the list of components.
- 2. In the Serialization output file field, enter the name of the serialized report file you created.
- 3. In the XML output file field, enter the pathname of the file for the XML output.
- 4. Click the Create XML File button.

### **Running the Configuration Reporter as a Standalone Utility**

You can run the Configuration Reporter as a standalone utility. This allows you to generate configuration reports even if ATG is not running. Before you run the Configuration Reporter as a standalone utility, you need to create two files:

- A file that contains a list of the components to include in the report. See Creating the Component File (page 147).
- A file that contains the configuration path. See Creating the Configuration Path File (page 147).

In addition, you should also set certain properties in /atg/dynamo/service/ConfigurationReporter. See Configuring the Configuration Reporter (page 148).

#### **Creating the Component File**

You can create the component file by running the report on the Output Dynamo Component Hierarchy to File page at:

http://hostname:port/dyn/admin/atg/dynamo/admin/en/ config-reporter-output-hierarchy-titled.jhtml

Add the name of the file thus created to the componentFileName property of /atg/dynamo/service/ ConfigurationReporter.

As an alternative, you can create a component file by hand. The component file format is as follows:

```
<component>/Initial</component>
```

<component></atg/dynamo/service/Scheduler</component>

A component file is not expected to be well-formed XML. Anything other than what is between the component start and end tags is ignored. Anything between the tags is treated as a component name. Folders can be included between component tags; the Configuration Reporter includes all components in such a folder. Add the name of the component file to the componentFileName property of /atg/dynamo/service/ ConfigurationReporter.

### **Creating the Configuration Path File**

You can create the configuration path file by running the CONFIGPATH report on the Output Configuration Path to File page at:

```
http://hostname:port/dyn/admin/atg/dynamo/admin/en/
config-reporter-conf-path-titled.jhtml
```

Add the name of the file thus created to the dynamoConfigurationPathFileName property of /atg/ dynamo/service/ConfigurationReporter.

As an alternative, you can create a configuration path file by hand. The configuration path file format is as follows:

```
<configuration_path_item>c:\ATG\ATGl1.0\DAS\config
</configuration_path_item>
```

```
<configuration_path_item>c:\ATG\ATG11.0\home\localconfig
</configuration_path_item>
<configuration_path_item>c:\ATG\ATG11.0\MyModule\config
</configuration_path_item>
```

A configuration path file is not expected to be well-formed XML. Anything other than what is between the <configuration\_path\_item> start and end tags is ignored. Anything between the tags is treated as an element of the Dynamo CONFIGPATH. Elements of the CONFIGPATH should be listed in the configuration path file in the order that they appear in the Dynamo CONFIGPATH. Add the name of the configuration path file to the dynamoConfigurationPathFileName property of /atg/dynamo/service/ConfigurationReporter.

### **Configuring the Configuration Reporter**

As described in the previous sections, you need to set the componentFileName and dynamoConfigurationPathFileName properties of /atg/dynamo/service/ConfigurationReporter. In addition, set the serializedPropertiesFileName property to the pathname of the file you want to output.

You can set these properties using the ACC, or by adding a properties file like this at <ATG11dir>/home/localconfig/atg/service/ConfigurationReporter.properties:

```
$class=atg.service.configurationreporter.ConfigurationReader
componentFileName=
dynamoConfigurationPathFileName=
serializedPropertiesFileName=
```

### **Running the Configuration Reader**

To run the Configuration Reporter as a standalone utility, use the following command:

```
java atg.service.configurationreporter.ConfigurationReader
-saveProperties config_directory
```

The *config\_directory* argument is the directory that holds your ConfigurationReporter.properties file. A typical value would be localconfig.

This command generates a serialized output file. When you run this utility, the Configuration Reader reads the following input properties from properties file /atg/dynamo/service/ ConfigurationReporter.properties.

Property	Description
dynamoConfigurationPathFileName	The name of a file that contains the Dynamo CONFIGPATH.
componentFileName	The name of the component file to read the list of Dynamo components from.
serializedPropertiesFileName	The name of the serialized file to output.

After you run the Configuration Reader utility with the -saveProperties argument, you can run it in this form to output an XML representation of the properties report:

-outputRepresentationToXML SourceFile OutputFileName OutPutPropertyValues=true|false

The *SourceFile* argument is the name of the output file (serializedPropertiesFileName) and the *OutputFileName* argument is the name of the file where the Configuration Reader should output the XML representation of the serialized output file. Use the OutPutPropertyValues=true flag to output the property values as well as the component names; use the OutPutPropertyValues=false flag to omit the property values.

### Using the VMSystem Component

The ATG component located at /VMSystem provides a way for you to access the Java memory manager. You can monitor the status of the Virtual Machine and call methods on it. An interface to the VMSystem component is included in the Dynamo Server Admin at:

http://hostname:port/dyn/admin/nucleus/VMSystem/

From this page, you can conduct the following VM Operations:

- Perform garbage collection
- Run finalizations
- · Show memory information
- · List system properties
- List thread groups
- List threads
- Stop the VM

### **Using a Sampler**

When testing your site, it is useful to automatically sample performance to understand throughput as a function of load. ATG includes a Sampler component at /atg/dynamo/service/Sampler.

### **Starting the Sampler**

You can start the Sampler component by opening it in the ACC and clicking the **Start** button.

You can also start the Sampler component from the Dynamo Server Admin by requesting this URL:

http://hostname:port/dyn/admin/nucleus/atg/dynamo/service/Sampler

The first time you request this page, ATG instantiates the Sampler component, which begins recording statistics.

You can configure ATG to start the Sampler whenever ATG starts by adding the Sampler to the initialServices property of the /atg/dynamo/service/Initial component:

initialServices+=Sampler

#### **Sampler Information**

The Sampler outputs information to the file <ATG11dir>/home/logs/samples.log. For each system variable that it samples, it records the following information in the log file:

- the current value
- the difference between the current value and the value recorded the last minute
- the rate of change of the value

You can adjust values recorded by the Sampler, but the default set is comprehensive in monitoring ATG request handling performance. The Sampler's output includes the following:

Value	Description
handledRequestCount	Total number of requests handled by this ATG server.
averageRequestHandlingTime	Average time spent handling requests since the sampler was started.

### **Sampler Output**

If you collect enough real data of your site under varying loads, your Sampler output gives you the answers to the following important questions:

- What is the peak throughput of your site in pages per minute for each ATG server?
- Does the peak throughput of your site go down as load increases beyond a certain threshold?
- How many sessions can each server handle while maintaining a comfortable latency (such as, latency < 1 second)?

### **Using the Recording Servlet**

The Recording Servlet is a servlet that you place in your request handling pipeline that records the amount of time spent handling each URL on your site. It performs two distinct functions:

- Records script files used in conjunction with URLHammer. You can record scripts of actual user activity on your site, then use URLHammer to execute a script repeatedly, simulating actual system load. For information on URLHammer, see Using URLHammer (page 128) in the *Performance Diagnostics* (page 121) chapter.
- Records performance information for a single user, including the minimum, maximum, and average time spent handling each URL on your site during the recording interval.

### **Inserting the Recording Servlet**

The Recording Servlet must be enabled before you can use it. You can enable it in one of three ways:

- Open the Recording Servlet in the ACC Component Editor at /atg/dynamo/servlet/pipeline/ RecordingServlet and set the recording property to true.
- Request the following URL in your administration interface:

http://hostname:port/dyn/admin/nucleus/atg/dynamo/servlet/pipeline/Record ingServlet

Set the recording property to true.

• Add the Recording Servlet to the initialServices property of the /atg/dynamo/servlet/Initial component, so that the Recording Servlet is added to the servlet pipeline automatically each time your server is started:

initialServices+=pipeline/RecordingServlet

### **Generating Script Files**

To generate a script file from the Recording Servlet, use the Component Browser to modify the value of the recording property. Set this to true to start recording or false to stop recording.

Then, use your web browser to make a series of requests from your site, in the pattern of user behavior that you want to record. Each of your requests becomes part of the script.

The script is saved to the file specified by the Recording Servlet's recordFile property. By default, the script is saved to <ATG11dir>/home/logs/record.log. Each time you start recording, the old script file is overwritten. So be sure to copy the script before you enable recording for a second time.

### **Keeping Statistics**

The Recording Servlet is also used to maintain per-URL performance statistics. To turn on this feature, set the keepingStatistics property to true. While this property is on, the minimum, maximum, and average times used to serve each requested page will be maintained and displayed in the component browser's page for the Recording Servlet component.

#### **Tracing Memory**

You can use the Recording Servlet to get an approximate reading on the amount of memory each request consumes. Set the Recording Servlet's tracingMemory property to true to turn on this feature. The Recording Servlet records memory information only for those URLs that run through the server one at a time; it is not appropriate for use on a live site.

# 11 Repository and Database Performance

Most ATG applications require database access, which represents another area where performance bottlenecks can occur. To effectively tune a large production database, your team should include an experienced database administrator.

This chapter includes the following sections:

Database Performance Practices (page 153)

Repositories and Transactions (page 154)

Repository Item Property Loading (page 154)

Database Sorting versus Locale-Sensitive Sorting (page 154)

Batching Database Transactions (page 154)

Avoiding Table Scans (page 155)

Database Caches (page 156)

Diagnosing Database Performance Problems (page 159)

### **Database Performance Practices**

Follow these practices in designing and developing your site to avoid database performance problems:

- Use Repository caching features to optimize database access. See SQL Repository Caching in the Repository Guide.
- Use queues to batch database transactions, rather than performing each transaction individually. See Batching Database Transactions (page 154).
- Avoid using database queries that might result in table scans of large tables. See Avoiding Table Scans (page 155).
- Run your database server on a separate machine from your application servers, or at least allocate a separate CPU.

### **Repositories and Transactions**

By default, if you do not have a JTA transaction in place, each SQL Repository operation that affects the state of a repository item creates and commits a transaction around the operation. This is generally not the most efficient way to handle repository item updates. It is generally most efficient to ensure that all of the method calls in creating or updating a repository item are performed in a single transaction. ATG offers several different techniques for transaction demarcation that you can use to group repository method calls into a single transaction. You can use transaction demarcation in a Java Server Page using the Transaction servlet bean. You can demarcate a transaction programmatically. These are described in detail in the *Transaction Management* chapter of the *Platform Programming Guide*. You can also use ATG's Repository Form Handler and TransactionalFormHandler classes to improve the transactional behavior and performance of repository operations. Refer to the *Platform Programming Guide* and *Repository Guide* for more information.

### **Repository Item Property Loading**

By default, whenever the SQL Repository calls getItem, it loads from the database (or the cache) not just the repository ID of the item, but all repository item properties that are stored in the primary database table for that item's item descriptor. For some applications, this may result in too much database activity. For other applications, you may want to load repository item properties that appear on other tables. You can adjust how the SQL Repository loads repository item properties by grouping properties, using the group attribute in property tags in the repository definition file. All properties with the same group attribute are loaded whenever one property of the group is loaded. For more information, see the *Repository Guide*.

### **Database Sorting versus Locale-Sensitive Sorting**

SQL Repository components include a localeSensitiveSorting property that controls how query results are sorted. If this property is set to true, query results are sorted using locale-sensitive String comparison (via java.text.Collator). Since most databases cannot handle sorting with multiple locales, setting this option to true also means that the repository will perform all sorting in memory. If localeSensitiveSorting is set to false (the default), database sorting (via ORDER BY) is used where applicable and Strings are compared using String.compareTo(). If database sorting is adequate for your purposes, leaving this property set to false will result in better performance. For more information, see the *Repository Guide*.

### **Batching Database Transactions**

If you have large volumes of data to insert or update, you should wherever possible perform those operations in batched transactions. It is more expensive to start a new transaction for every change than it is to attempt to make many changes in a single database transaction. For example, a request handler might log every single hit to a log table. Suppose that it takes 50 milliseconds to write a row in a log table. If that is the case, then the request handler cannot serve requests any faster than 20 per second, even if the rest of the request handling mechanism is blazingly fast. But writing an entry in a log table is not a critical part of the request handling operation, and thus should not be such a limiting factor.

The solution to this problem is to introduce a queue between the request handler and the database facility. When the request handler wants to make a database entry, it places the log entry on the queue, then continues handling the rest of the request. A separate component reads sets of log entries and writes the whole set in a single database transaction. This arrangement decouples the request handlers from the loggers, thereby eliminating the bottleneck introduced by the database.

### **Avoiding Table Scans**

A table scan is the reading of every row in a table and is caused by queries that do not properly use indexes. Table scans on large tables take an excessive amount of time and cause performance problems.

Make sure that, for any queries against large tables, at least one WHERE clause condition:

- · refers to an indexed column and
- · is reasonably selective

You should be concerned primarily with queries against large tables. If you have a table with a few hundred rows, table scans are not a problem and are sometimes faster than indexed access.

During initialization, systems like ATG may front-load caches to avoid unnecessary database operations later. You may see queries with large results during this time, but that is okay. Within reason, lengthy database operations at startup are acceptable. However, if you see frequent, large, or slow queries issuing from ATG during the course of normal operation, then you have a design problem that must be addressed to achieve acceptable performance.

For example, suppose your database has a large table that holds products such as this:

```
CREATE table product

( sku char(6) not null,

type char(1) not null,

name varchar(50) not null,

description varchar(200) null )
```

and has these indexes:

```
CREATE unique index il on product(sku)
CREATE index i2 on product(name)
CREATE index i3 on product(type)
```

The following query is fine:

```
SELECT *
FROM product
WHERE sku = 'a12345'
```

That query will not cause performance problems because the WHERE clause refers to a very specific condition on a column with an index.

Here is an example of a query that is likely to cause problems:

```
SELECT *
FROM product
WHERE description LIKE '%shoes%'
```

This query causes a table scan, since the indexes cannot help the database to optimize the query. Queries like this on a large table will result in an unacceptable performance drag and therefore should not be allowed in a production system.

Here are some more queries that are likely to cause performance problems. The following query is inadvisable because, although it refers to the indexed sku column, it is not very selective and could return millions of rows:

```
SELECT *
FROM product
WHERE sku > 'abc'
```

The following query is bad because, although it is relatively selective, it will cause a table scan on most DBMSs. A LIKE query with a leading wildcard typically cannot be optimized:

```
SELECT *
FROM product
WHERE name LIKE '%stereo'
```

### **Database Caches**

If you are using the SQL Repository, see how multiple requests of the same behavior affect cache usage. The first time your application references database information, the request causes a SQL database operation, but subsequent requests will use the cache. Try to optimize cache usage. Consider the best caching mode to use for each of the item descriptors in your SQL repositories. See the *Repository Guide* for more information.

When you are testing the system, make sure you think about real-world usage of your data. If your system could potentially have tens of thousands or millions of rows of data, make sure you test that scenario. If you test only against small sets of data, some performance bottlenecks will be masked, because the database can cache the entire dataset into memory.

### **External Caching for SQL Repositories**

Use external caching to improve performance by storing SQL repository data in the memory of a separate distributed cache application. This section explains how to configure the Oracle ATG Web Commerce platform to connect to a distributed cache application. See information about configuring the way your application uses external caching in the *Repository Guide*.

The Oracle ATG Web Commerce platform includes an adapter for the Oracle Coherence data grid application. Install and configure Oracle Coherence before using the external caching feature. See information about Oracle Coherence at http://www.oracle.com/technetwork/middleware/coherence/documentation/index.html.

### **Configuring Coherence for External Caching**

Oracle Coherence operates in a clustered arrangement. Configure individual Oracle Coherence instances to locate each other by joining the same cluster. Then that cluster will automatically coordinate the distribution of cached data among its members. Configuring Oracle Coherence clusters is explained in detail in the documentation for that product (http://www.oracle.com/technetwork/middleware/coherence/documentation/index.html).

You can add members to the Oracle Coherence cluster as needed to improve your application performance. Add Oracle Coherence cluster members either as part of any additional Oracle ATG Web Commerce server that is configured to use external caching or as a standalone instance of Oracle Coherence.

When an Oracle ATG Web Commerce server starts up, it creates an Oracle Coherence cluster member in its own Java virtual machine. To do so, it invokes code in the coherence.jar file that is provided in the Oracle Coherence distribution. Standalone instances of Oracle Coherence interoperate with Oracle ATG Web Commerce servers and must have access to Java classes provided in the Oracle ATG Web Commerce distribution in the file <a href="https://www.atgslice.com"></a> ATG11dir>/DAS/lib/atg-coherence-classes.jar.

#### **Connecting Oracle ATG Web Commerce to an Oracle Coherence Cluster**

To connect an Oracle ATG Web Commerce server to an Oracle Coherence cluster:

1. Extract the files tangosol-coherence.xml and coherence-cache-config.xml from the top level of the coherence.jar file that is provided with the Oracle Coherence distribution.

Place the files in a directory that is accessible by your Oracle ATG Web Commerce application server. Add the directory to the application server's class path environment variable. Do not add the files themselves.

Customize the XML files with the Oracle Coherence cluster connection information and your caching configuration. See the Oracle Coherence documentation for information about these customizations (http://www.oracle.com/technetwork/middleware/coherence/documentation/index.html).

See the sample configuration file sample-coherence-cache-config.xml in <ATG11dir>/DAS/lib/ classes.jar.

- 2. Add [coherence-install-directory]/lib/coherence.jar to the CLASSPATH of your Oracle ATG Web Commerce application server. Add this file path after the directory that holds the Oracle Coherence configuration files from the previous step.
- 3. Start the Oracle ATG Web Commerce application server and confirm that it has created an Oracle Coherence cluster member and that the cluster member joined the cluster that you expected. Example output from an application server log file is shown below. Note that messages from Oracle Coherence are labeled as errors by some application servers. These error messages may be routine and benign.

```
2011-08-05 16:03:23.118/428.352 Oracle Coherence GE 3.6.0.4 <Info>
(thread=[ACTIVE] ExecuteThread: '0' for queue: 'weblogic.kernel.Default (self-
tuning)':ipaddr=10.64.200.102;path=/dyn/admin/nucleus//atg/commerce/catalog/
ProductCatalog-ver/;sessionid=6NkZT8MFrhvGw4QLCS8DQkw7yJp1lG8mfVCC2q90QmTKDJ119vpZ!
762240401!1312574469587, member=n/a): Started cluster Name=ecCluster
Group{Address=123.4.5.6, Port=36000, TTL=4}
MasterMemberSet
```

```
(
ThisMember=Member(Id=1, Timestamp=2011-08-05 16:03:19.554,
Address=12.34.567.890:48088, MachineId=57228,
Location=machine:myserver,process:19539, Role=WeblogicServer)
OldestMember=Member(Id=1, Timestamp=2011-08-05
16:03:19.554, Address=12.34.567.890:48088, MachineId=57228,
Location=machine:myserver,process:19539, Role=WeblogicServer)
ActualMemberSet=MemberSet(Size=1, BitSetCount=2
Member(Id=1, Timestamp=2011-08-05 16:03:19.554, Address=12.34.567.890:48088,
MachineId=57228, Location=machine:myserver,process:19539, Role=WeblogicServer)
)
RecycleMillis=1200000
RecycleSet=MemberSet(Size=0, BitSetCount=0
)
)
TcpRing{Connections=[]}
IpMonitor{AddressListSize=0}
```

#### Adding an Oracle Coherence Cluster Member to Your Cluster

To add an Oracle Coherence cluster member to your cluster:

1. Extract the files tangosol-coherence.xml and coherence-cache-config.xml from the top level of the coherence.jar file that is provided with the Oracle Coherence distribution.

Place the files in a directory that is accessible from your Oracle Coherence installation directory.

Customize the XML files with the Oracle Coherence cluster connection information and your caching configuration. See the Oracle Coherence documentation for information about these customizations (http://www.oracle.com/technetwork/middleware/coherence/documentation/index.html).

**Note**: You can skip this step if you already have these files configured for the cluster member that is created by your Oracle ATG Web Commerce server.

See the sample configuration file sample-coherence-cache-config.xml in <ATG11dir>/DAS/lib/ classes.jar.

2. Edit the file bin/cache-server.sh (or the corresponding \*.cmd file) in your Oracle Coherence installation directory.

Add the directory that holds the XML configuration files to the class path used by the script's JAVAEXEC command. Do not add the files themselves. The class path is specified by the -cp argument. Make sure you add the directory before \$COHERENCE\_HOME/lib/coherence.jar.

Add the directory that holds <code>atg-coherence-classes.jar</code> to the class path used by the script's <code>JAVAEXEC</code> command. This file is available in your Oracle ATG Web Commerce installation directory at <a href="https://directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.directory.di

- 3. Change to the root directory of your Oracle Coherence installation directory and invoke the file bin/cacheserver.sh (or the corresponding \*.cmd file).
- 4. Confirm that the cluster member joined the cluster you expected. See the example terminal output that shows cluster information below.

```
2011-08-05 14:10:08.886/4.737 Oracle Coherence GE 3.7.0.0 <Info> (thread=main, member=n/a): Started cluster Name=cluster :0x96AB
```

```
Group{Address=123.4.5.6, Port=37000, TTL=4}
MasterMemberSet
(
ThisMember=Member(Id=2, Timestamp=2011-08-05
14:10:07.679, Address=12.34.456.89:8090, MachineId=57105,
Location=site:oh.my.domain.com,machine:myserver,process:14292, Role=CoherenceServer
```

### Using Coherence Utilities With the ATG Server Cluster

Your Oracle Coherence distribution includes utilities such as the <code>query.sh</code> script that opens a Coherence query language terminal. In order to use utilities such as this one, you must configure it to join your Coherence cluster and give it access to the Oracle ATG Web Commerce classes provided in <code>atg-coherence-classes.jar</code>.

Configure Oracle Coherence utility scripts such as query. sh in the same way that you configured cacheserver.sh to create additional cluster members. See Adding an Oracle Coherence Cluster Member to Your Cluster (page 158).

### Using Portable Object Format (POF) Serialization

Oracle ATG Web Commerce includes classes and a configuration file to enable the Portable Object Format (POF) feature of Oracle Coherence. POF serialization is faster and uses less memory than standard Java serialization.

The configuration provided with Oracle ATG Web Commerce is intended as a starting point. Make further customizations based on the way you intend Oracle Coherence to use POF serialization. See information about configuring Oracle Coherence to use POF in the documentation for that product (http://www.oracle.com/technetwork/middleware/coherence/documentation/index.html).

To use POF serialization:

 Extract the file atg/adapter/gsa/externalcache/config/atg-pof-config.xml from the file <ATG11dir>/DAS/lib/atg-coherence-classes.jar.

Place the file in a directory that is accessible by your Oracle ATG Web Commerce application server. Add the directory to the application server's class path environment variable. Do not add the file itself.

- Customize the atg-pof-config.xml file according to the way you intend to use POF serialization. See the Oracle Coherence documentation for information about these customizations (http://www.oracle.com/ technetwork/middleware/coherence/documentation/index.html).
- 3. Configure Oracle Coherence to enable POF serialization and to use the atg-pof-config.xml configuration file. See the Oracle Coherence documentation for information about these customizations (http://
  www.oracle.com/technetwork/middleware/coherence/documentation/index.html).

### **Diagnosing Database Performance Problems**

Make use of performance analysis tools offered by your database and application server vendor. These tools typically enable you to measure transactions per second and memory, cache, and disk utilization. Check the CPU utilization and I/O utilization of your database server. If they are near maximum levels, this is a strong indication that the database is limiting the performance of your site.

To understand database performance, you must know your data and the operations you are performing on it. The first step is to get a copy of the DDL for all the tables in your database and get a good estimate of how many rows are in each table. Most major database systems have tools that can tell you this quickly. In a pinch, you can issue the following query for each table:

```
SELECT count(*) FROM <table-name>
```

This query might take some time for large tables, so it is best to use the vendor-supplied tools or commands. In addition to this information, you'll need a list of the indexes on each table.

### **Avoid Using Simulated Text Search Queries in Repositories**

As a convenience feature, a SQL Repository can simulate full text searches using the SQL LIKE operator. If full text searching is not available for your database, you can substitute pattern matching queries for text search queries by setting the following property in the GSARepository component:

simulateTextSearchQueries=true

The SQL Repository will then convert text search queries into CONTAINS pattern match queries, which are implemented using the SQL LIKE operator.

Simulated text search queries are useful for demos and standalone development when you want to put in place the createTextSearchQuery() API calls without having to set up a text search engine. However, simulated text queries are extremely inefficient and are not supported for production systems. A simulated text search query using LIKE will typically cause a table scan, so you should not use simulated queries in production.

### **Drop Index in OLTP Environment**

The order\_lastmod\_idex index in the dcspp\_order table can cause performance degradation in OLTP environments. If your environment is experiencing performance issues, you can drop this index.

Note that the lastModifiedDate property, on which this index is based, is used in queries accessed by the atg.commerce.order.OrderQueries class, as well as in the CommercProfileTools. loadShippingCarts methods. For information on the OrderQueries class, refer to the Business Layer Classes section of the Commerce Programming Guide. For information on the CommerceProfileTool, refer to the Profile Tools and Property Manager Extension section of the Commerce Programming Guide.

Additionally, this index is used when performing manual index requests of orders when creating full text order searches in the Dynamo Server Admin. For information on full text searches and orders, refer to the *Configuring Order and Profile Search* section of the *Commerce Service Center Installation and Programming Guide*.

# **12 Tuning Site Performance on JBoss**

This chapter describes configuration steps you can perform which might improve performance of your ATG software running on JBoss. Note that these are suggestions only; JBoss configuration is a complex topic, and no recommendations can be applied globally. Work with your JBoss representative to fine-tune your application's performance.

Tuning suggestions are divided into two sections:

JBoss File Modifications (page 161)

JBoss Application Framework Trimming (page 164)

### **JBoss File Modifications**

This section describes changes you can make to JBoss configuration files to improve application performance.

### JSP Servlet Configuration

This section concerns changes you can make to your <JBdir>/server/configdir/deploy/ jbossweb.deployer/conf/web.xml file.

Add the following to the web.xml file under the JSP servlet (search for <servlet-name>jsp</servlet-name>) and make changes in that context.

```
<init-param>
        <param-name>trimSpaces</param-name>
        <param-value>false</param-value>
</init-param>
        <param-name>genStrAsCharArray</param-name>
        <param-value>true</param-value>
</init-param>
        <param-name>classDebugInfo</param-name>
        <param-value>false</param-value>
</init-param>
        <param-value>false</param-value>
</init-param>
```

### **Tomcat Connector Thread Configuration**

Thread pools used by Tomcat are configured on a per connector basis. The changes in this section are applied to the <JBdir>/server/configdir/deploy/jbossweb.deployer/server.xml file.

The default configuration is shown in this sample:

Thread pools can be monitored using the Tomcat monitor at http://hostname:http\_port. The Tomcat status link is under the JBoss Management heading, for example:

```
Tomcat status (full) (XML)
```

#### **Reducing the HTTP Connector Thread Pool**

This connector is only used when you connect to Tomcat directly from your web browser. In this example, the thread pool for the HTTP connector was reduced from 250 to 20.

The maxThreads setting should reflect the expected maximum number of users that can simultaneously use the system. This number should also drive the maximum number of database connections in the data source \*-ds.xml file.

Full documentation for the HTTP Connector configuration can be found at http://tomcat.apache.org.

#### Increasing the AJP Connector Thread Pool

This is the primary means of contacting the server for a user (via Apache and  $mod_jk$ ). In this example, the thread pool for the AJP connector is increased:

```
<!-- A AJP 1.3 Connector on port 8009 -->

<Connector port="8009" address="${jboss.bind.address}"

maxThreads="250" strategy="lf" minSpareThreads="50"

emptySessionPath="true" enableLookups="false" redirectPort="8443"
```

```
bufferSize="10240" maxHttpHeaderSize="8192" tcpNoDelay="true"
protocol="AJP/1.3"/>
```

Full documentation for the AJP Connector and a complete dictionary of the AJP connector configuration can be found at http://tomcat.apache.org.

### **Tomcat Cluster Configuration**

The <JBdir>/server/configdir/deploy/jboss-web-cluster.sar/META-INF/jboss-service.xml file contains the session replication settings. Consider the following options to improve performance:

- Use the replication strategy REPL\_ASYNC.
- Under the UDP protocol stack ensure that the mcast\_addr is the same on all cluster members.
- Under the UDP protocol stack ensure that the mcast\_port is the same on all cluster members.
- When running under Windows 2003, ensure that the loopback attribute of the UDP protocol stack is set to true. For Linux this should be set to false. See the comment about this in the file.

### JBoss Logging Configuration

JBoss uses Log4j wrapped in an MBean as a logging service. This means that an independent logging library does not need to be bundled with the application.

All logging configuration is done in the <JBdir>/server/configdir/conf/jboss-log4j.xml file. For more information on Log4j, see http://logging.apache.org/log4j/docs/manual.html.

You can adjust class specific logging in the category elements toward the end of the log4j configuration file. Each category can have a priority assigned to it. For example:

```
<category name="org.jboss">
    <priority value="DEBUG" />
    <appender-ref ref="FILE"/>
</category>
```

### **Data Source Configuration**

In any -ds.xml files used by ATG, edit the <min-pool-size> and <max-pool-size> settings to reflect the expected maximum number of simultaneous connections.

Note: Your file may have a different name or location, depending on your configuration.

```
<min-pool-size>50</min-pool-size><max-pool-size>75</max-pool-size>
```

Data source connections can be monitored using the JMX-Console at:

```
http://hostname:port/jmx-console
```

Look for ATG and ManagedConnectionFactory. The MBean monitor page shows how many connections exist and how many are being used.

### Configuring run.bat/sh and run.conf

You may want to add the following JVM tuning parameters to the JAVA\_OPTS in the bin/run.conf (UNIX) or run.bat (Windows) file:

• -Dtomcat.util.buf.StringCache.byte.enabled=true

Enables the byte array to String conversion caching.

• -Dtomcat.util.buf.StringCache.char.enabled=true

Enables the char array to String conversion caching.

-Dtomcat.util.buf.StringCache.cacheSize=2000

The maximum number of String objects that will be cached, according to their usage statistics.

The effectiveness of the StringCache can be checked using the JMX-Console. Look for StringCache under Catalina in the JMX-Console page.

For the JVM command-line, the following settings can be used:

- · Memory set at just over 1G for each server
- MaxPermSize adjusted to 256m

## **JBoss Application Framework Trimming**

Removing non-required services can reduce the memory footprint as well as simplifying configuration for your application. To remove JBoss services, consider deleting the services listed below from the deploy (or deploy-hasingleton) directory.

Warning: The jboss-service.xml found in the *configdir*/conf directory should never be deleted or moved.

Consider whether you might be able to remove the following services:

- JBoss Mail (mail-ra.rar, mail-service.xml)
- HA-JMS (in the deploy-hasingleton directory of the all configuration)
- HA-JNDI (in all/deploy/cluster-service.xml, search for HAJNDI)
- UUID Key Generator (used only for CMP, uuid-keygenerator.sar)
- Monitoring (monitor JMX changes, in monitoring-service.xml)
- Scheduling (schedule tasks to execute, in schedule-manager-service.xml and schedulerservice.xml)

• EJB3 related services; see <JBdir>/server/configdir/conf/jboss-service.xml

# 13 Appendix A: Data Storage and Access

This appendix describes the recommended configuration for storing and accessing ATG data. It covers the following:

Database Schema Best Practices (page 167)

Data Sources (page 168)

Repositories (page 169)

### **Database Schema Best Practices**

Your DBA has ultimate control over the arrangement of ATG database schemas. However, ATG recommends as a best practice that your installation include the databases described in the following list. ATG documentation for data sources and other components uses this division as the frame of reference.

- Production Schema—Data to be accessed or affected by external users, such as product catalogs and customer profiles, and the loader tables for the data warehouse.
- Management Schema—Data required for ATG administrative applications to run, including versioned repositories and internal users.
- Agent Schema—Data to be accessed by internal users of the customer service applications, such as ATG Knowledge solutions and profile data for internal users.
- Warehouse Schema—All of the data warehouse data. This schema should be created in a database optimized for data warehousing, and on a high-performance machine. see the Business Intelligence Data Warehouse Guide.

ATG documentation may also refer to a "local" schema. This schema contains the platform tables created by the das\_ddl.sql script (see Creating the DAS Tables (page 26) in this guide).

See the Multiple Application Integration Guide for additional information on system architecture.

## **Data Sources**

The following table lists the data sources components available for use by ATG applications. The data sources you configure will depend on which applications you are using.

Data Source Component Name	Module Defined In	Configured In
/atg/dynamo/service/jdbc/JTDataSource	DAS	config
This data source is always configured to point to the		
core schema for the server instance on which it is		
running. For instance, on a Content Administration		
server, it points to the management schema; on a		
production server, it points to the production schema.		
/atg/dynamo/service/jdbc/	DAS	config
JTDataSource_production		
This data source points to a production schema, but		
runs on a non-production server instance, such as		
asset management or agent.		
/atg/dynamo/service/jdbc/	DafEar.base	configlayers/
JTDataSource_staging		stagingandprod
This data source points to a staging schema, but		
runs on a non-staging server instance, such as asset		
management, production, or agent.		
/atg/reporting/datawarehouse/loaders/	ARF.base	config
JTDataSource		
/atg/reporting/datawarehouse/	ARF.DW.base	config
JTDataSource		
/atg/commerce/jdbc/ProductCatalog	DCS	config
SwitchingDataSource		
Used for switching. See Configuring a		
SwitchingDataSource (page 51) in this guide.		
/atg/commerce/jdbc/ProductCatalog	DCS	config
DataSourceA		
Used for switching. See Configuring a		
SwitchingDataSource (page 51) in this guide.		
/atg/commerce/jdbc/ProductCatalog	DCS	config
DataSourceB		
Used for switching. See Configuring a		
SwitchingDataSource (page 51) in this guide.		
		1

Data Source Component Name	Module Defined In	Configured In
/atg/search/service/SearchJTDataSource	DAF.Search.Base	config
/atg/dynamo/service/jdbc/ JTDataSource_agent	DAS	config
/atg/dynamo/service/jdbc/ JTDataSource_management	DAS	config
/atg/dynamo/service/jdbc/ eServerJTDataSource	Service.migration	config
/atg/dynamo/service/jdbc/ SelfServiceReportingJTDataSource	Service.SelfService DataWarehouse	config
/atg/campaign/communication/reporting/ JTDataSource	ACO.communication.DW	config

## Repositories

The following table lists the repositories used by ATG applications, which data source they use by default, and server-dependent conditions for use:

Repository Component Name	Data Source Component Configuration Varies Depending on Server
/atg/commerce/catalog/ProductCatalog	/atg/dynamo/service/ JTDataSource_production
/atg/commerce/atg/	/atg/dynamo/service/
ClaimableRepository	JTDataSource_production
/atg/commerce/gifts/GiftLists	/atg/dynamo/service/ JTDataSource_production
/atg/commerce/inventory/	/atg/dynamo/service/
InventoryRepository	JTDataSource_production
/atg/commerce/jdbc/	/atg/dynamo/service/
ProductCatalogDataSourceA	JTDataSource_production
/atg/commerce/order/OrderRepository	/atg/dynamo/service/ JTDataSource_production
/atg/commerce/pricing/priceLists/	/atg/dynamo/service/
PriceLists	JTDataSource_production

Repository Component Name	Data Source Component Configuration Varies Depending on Server
/atg/scenario/ScenarioClusterManager	/atg/dynamo/service/
	o i Davabour oo_produceren
/atg/userprofiling/ ProfileAdapterRepository	/atg/dynamo/service/ JTDataSource_production
/atg/userprofiling/ PersonalizationRepository	/atg/dynamo/service/ JTDataSource_production
/atg/search/service/ SearchJTDataSource	Generic Reference to /atg/dynamo/service/ JTDataSource_production
/atg/search/repository/ IncrementalItemQueueRepository	/atg/dynamo/service/ JTDataSource_production
/atg/content/media/MediaRepository	/atg/dynamo/service/JTDataSource
/atg/dynamo/messaging/SqlJmsProvider	/atg/dynamo/service/JTDataSource
/atg/dynamo/service/ClusterName	/atg/dynamo/service/JTDataSource
/atg/dynamo/service/IdGenerator	/atg/dynamo/service/JTDataSource
/atg/dynamo/service/ ObfuscatedIdGenerator	/atg/dynamo/service/JTDataSource
/atg/dynamo/service/jdbc/SDSRepository	/atg/dynamo/service/ JTDataSource_production
/atg/dynamo/service/jdbc/ SQLRepository	/atg/dynamo/service/JTDataSource
/atg/integrations/repository/ IntegrationsRepository	/atg/dynamo/service/JTDataSource
/atg/webservice/security/ NucleusSecurityRepository	/atg/dynamo/service/JTDataSource
/atg/epub/PublishingRepository	/atg/dynamo/service/ JTDataSource_management
/atg/epub/process/ProcessDataRepository	/atg/dynamo/service/ JTDataSource_management
(atg/opub/progogg/	/ata/dumamo/gorvice/
/acg/epub/process/	TTDataSource management
ver prolimaliager vehopt for Å	
/atg/epub/process/PortalRepository	/atg/dynamo/service/ JTDataSource_production
/atg/userprofiling/	/atg/dynamo/service/
InternalProfileRepository	JTDataSource_management

Repository Component Name	Data Source Component Configuration Varies Depending on Server
/atg/dynamo/service/jdbc/	/atg/dynamo/service/
BCRJTDataSource	JTDataSource_management
/atg/reporting/datawarehouse/	/atg/reporting/datawarehouse/
LogicalOrganizationReportRepository	JTDataSource
/atg/reporting/datawarehouse/	/atg/reporting/datawarehouse/
RmClsRoutingReportRepository	JTDataSource
/atg/reporting/datawarehouse/	/atg/reporting/datawarehouse/
RMReportRepository	JTDataSource
/atg/userprofiling/	/atg/dynamo/service/
InternalProfileRepository	JTDataSource_agent
/atg/commerce/custsvc/CsrRepository	/atg/dynamo/service/ JTDataSource_production
/atg/svc/ServiceRepository	/atg/dynamo/service/ JTDataSource_production
/atg/svc/ui/framework/	/atg/dynamo/service/
ServiceFrameworkRepository	JTDataSource_production
/atg/svc/option/UserOptionRepository	/atg/dynamo/service/ JTDataSource_production
/atg/svc/option/OptionRepository	/atg/dynamo/service/ JTDataSource_production
# 14 Appendix B: Adjusting the FileCache Size

ATG's servlet pipeline includes servlets that are used for JHTML pages, and which use a FileCache component to store files that ATG has read from disk, so that subsequent accesses for those files can be delivered directly from memory instead of being read from disk. Using the FileCache component improves performance by reducing disk accesses. For maximum performance, you want the FileCache to be large enough to hold all the files that ATG serves frequently. Set the totalSize property of this component at:

/atg/dynamo/servlet/pipeline/FileCache

to an appropriate value, measured in bytes, such as the following:

```
# size in bytes (2 million bytes)
totalSize=2000000
```

One approach in sizing the FileCache is to batch compile the entire document root and set the file cache to the resulting size. Make sure, however, that you account for the size of your FileCache when you set the size of your JVM. You can preload the FileCache by creating a script that accesses every page on your site and running the script on startup.

You can view statistics on how the file cache is used, as well as the contents of the FileCache in the Dynamo Administration page at *hostname:port/dyn/admin/nucleus/atg/dynamo/servlet/pipeline/* FileCache.

# 15 Appendix C: Development Tools for Eclipse

Oracle ATG Web Commerce offers a set of development tools for the open source Eclipse platform (http:// www.eclipse.org). These tools are a beta product and are not a supported component of Oracle ATG Web Commerce.

The Eclipse development tools are included in the Oracle ATG Web Commerce installation directory in the following file.

<ATG11dir>/Eclipse/ATGUpdateSite.jar

To install the Oracle ATG Web Commerce development tools for Eclipse:

- 1. Install the Eclipse IDE for Java EE Developers. If you have already installed the IDE, uninstall any previous versions of the Oracle ATG Web Development Tools.
- 2. In the Eclipse IDE, choose Help > Install New Software.
- 3. In the Install dialog, click the Add button next to the Work with text box.
- 4. In the Add Repository dialog, click the Archive button.
- 5. Navigate to the <ATG11dir>/Eclipse/ATGUpdateSite.jar file in the Repository Archive window. Click the Open button.
- 6. Click OK in the Add Repository dialog.
- 7. Select the checkbox next to Oracle ATG Web Development Tools in the Install dialog. Click Next, accept the license agreement, and click Finish to complete the installation. Restart Eclipse when you are prompted to.
- After Eclipse restarts, choose Window > Preferences. Enter ATG in the search text field and choose ATG Preferences.
- 9. Browse to select your Oracle ATG Web Commerce installation directory and click Apply. Restart Eclipse when you are prompted to.

To learn more about using the development tools, see the Oracle ATG Web Commerce Development Tools documentation under Help > Help Contents in the Eclipse IDE.

# 16 Appendix D: Using Oracle Access Management for Single Sign On

Oracle Access Management (OAM), which is a component of the Oracle Identity and Access Management (OIM) suite, is an enterprise-level security application that allows you to configure a number of security functions, including Single Sign On (SSO). OAM enables SSO through a common engine used across multiple protocols and applications.

You can use OAM SSO to authenticate internal users of Oracle ATG Web Commerce Business Control Center and Oracle Endeca Workbench using a single authentication step. The following section describes how to configure Oracle ATG Web Commerce to use OAM SSO to access the ATG Business Control Center. Refer to the Oracle Endeca Workbench documentation for information on installing and configuring these applications for OAM SSO.

For detailed information on Oracle Access Management and its tools, refer to the Oracle Identity and Access Management documentation web page.

## **OAM Authentication Overview**

OAM recognizes user populations and LDAP directory stores as *identity domains*. Each identity domain maps to a configured LDAP User Identity Store that is registered with OAM. This allows the LDAP directory to act as a master record for any user entered into OAM.

**Note:** Oracle ATG Web Commerce supports only OAM LDAP authentication. For detailed information on creating and maintaining LDAP repositories in Oracle ATG Web Commerce, refer to the *LDAP Repositories* section of the *Repository Guide*. For information on creating LDAP profile repositories, refer to the *Setting Up an LDAP Profile Repository* section in the *Personalization Programming Guide*.

#### Logging In to the Business Control Center

When you have configured Oracle ATG Web Commerce to work with OAM, requests to log into the Business Control Center are filtered by OAM.

- If OAM receives a request that is not recognized as an existing valid SSO session, the user is directed to the centralized login form configured through OAM. Once there, the user must provide their log in credentials to access the Business Control Center. Refer to your OAM documentation on creating a log in screen.
- If OAM does recognize an existing valid SSO session, the HTTP request is sent through the DAF servlet
  pipeline to the Business Control Center without directing the user to the OAM login form.

Once a user has been authenticated, OAM inserts a header name into the HTTP request. By default, OAM sets the header as <code>OAM\_REMOTE\_USER</code>. This header name can be configured to another value in OAM, or when you are configuring OAM integration with CIM. The header value is stored in the <code>userIdHttpHeaderName</code> property of the <code>/atg/dynamo/servlet/dafpipeline/</code>

OamRemoteUserServlet.

As the HTTP request is processed by the DAF servlet pipeline, the /atg/userprofiling/ ProfileRequestServlet component reviews the HTTP request to obtain the user profile based on the user ID provided in the header. The user profile is then loaded and made active when initiating the Business Control Center.

Note that when using OAM SSO authentication, the standard Oracle ATG Web Commerce login and user authentication process is disabled.

#### **Logging in Transient Profiles**

If the HTTP request does not have an ATG session cookie, and a new session has been created, the profile is considered to be transient. The /atg/dynamo/servlet/dafpipeline/OamRemoteUserServlet sets the value of the user ID from the HTTP request header into the remoteUser property of the DynamoHttpServletRequest. The servlet pipeline continues when the ProfileRequestServlet invokes the /atg/userprofiling/ProfileRequest to extract the user ID from the remoteUser property and uses it to load the associated user profile.

For detailed information on the ProfileRequestServlet, refer to the Platform Programming Guide and the Personalization Programming Guide.

For information on the DAF servlet pipeline, refer to the *Request Handling with Servlet Pipelines* section in the *Platform Programming Guide*.

#### Logging Out of the Business Control Center and SSO

When you configure OAM to work with ATG Oracle Commerce, you can configure an optional WebGate Logout URL. A WebGate is a Web server plug-in for OAM that intercepts HTTP requests and forwards them to the Access Server for authentication and authorization. For additional information on WebGates, refer to the *Installing WebGates for Oracle Access Manager* documentation.

When a user logs out of the Business Control Center, their Business Control Center session is terminated, and the logout request is directed to the WebGate Logout URL. When the OAM WebGate receives an HTTP request that contains this URL, it triggers a logout handler, which removes the SSO authentication cookie and ends the SSO session. Refer to the *Configuring Centralized Logout* OAM documentation for additional information.

When you provide this field in OAM, you ensure that both the Business Control Center and the SSO sessions are terminated when the user logs out of the Business Control Center.

#### **Session Timeouts**

If an OAM SSO session times out, the user is redirected to the OAM login form to re-authenticate. When passing the request to the Business Control Center, OAM inserts a new header into the HTTP request that contains the user ID of the authenticated user. Similar to the log in process, this user ID must match the Business Control Center user ID. Once the authenticated user ID is passed to the Business Control Center, the Business Control Center uses this header value to load the user's profile and make the user active.

If the Business Control Center session times out, yet the OAM SSO is still active, the Business Control Center user will not be required to re-authenticate, but a new application session is created. Note that this may result in session information being lost.

Some parts of the Business Control Center, such as parts of the Asset Manager framework, are built in Flex. Other parts of the Business Control Center use JavaServer Page (JSP) technology. If the Business Control Center user is working in a part of the Business Control Center that was written in Flex, the user will be presented with a dialog box warning that their session is about to expire.

# Installing the Oracle Access Manager Integration Component

To enable OAM SSO integration, use the Configuration and Installation Manager (CIM).

- 1. Run CIM, as outlined in the Using CIM (page 11) section.
- 2. Select the OAM Authentication option from the menu, which configures CIM to install and configure the Oracle ATG Web Commerce integration components for OAM.
- 3. Enter the following information in CIM:
  - The name of the OAM Remote User HTTP header This identifies the name of the header used to hold the OAM-authenticated user's userId. The value you provide CIM is used by the userIdHttpHeaderName property of the /atg/dynamo/servlet/dafpipeline/OamRemoteUserServlet file. By default, OAM uses the OAM\_REMOTE\_USER header. Note that the value of this property cannot contain white space.
  - The OAM Web Server host name This value is used to identify the allowedHostNames property in the / atg/dynamo/servlet/pipeline /RedirectURLValidator file, as well as the webgateHost property in the /atg/userprofiling/oam/ Configuration file.
  - The port number of the OAM Web Server This value is used to provide the webgatePort property of the /atg/userprofiling/oam/
     Configuration file.
  - The WebGate Logout URL This optional value is used to identify the logout URL to which the user will be directed upon successfully logging out of the Business Control Center. This value is stored in the logoutSuccessURL property of the /atg/userprofiling/oam/Configuration file. This ensures that both the ATG and OAM SSO user sessions are terminated when the user logs out of the Business Control Center.

Before using the Oracle ATG Web Commerce OAM integration, ensure that you have configured and installed OAM. Refer to your Oracle OAM documentation for instructions.

## Validating OAM Identity Assertion X.509 Certificates

When OAM authenticates a new user login, it adds an IdentityAssertion header into the authenticated HTTP request. This IdentityAssertion is a SAML-compliant structure that contains an X.509 certificate.

X.509 is a data structure that sends a public key to a receiving party. Certificates are issued by certificate authorities (CA), which verify an entity's identity and grants a certificate, signing it with the CA's private key. The CA publishes its own certificate, which includes its public key.

Each network entity has a list of the certificates of the CAs it trusts. Before communicating, this list is used to verify that the signature of other certificates comes from a trusted CA. For detailed information on X.509 certificates, refer to the RFC 3280: Internet X.509 Public Key Infrastructure Certificate and CRL Profile document on the IETF Web site.

You can configure ATG to perform certificate validation using the /atg/userprofiling/oam/ OamCertificateVerifier component. When this component is running, ATG extracts the certificate from the IdentityAssertion in the header and uses the trusted key to verify the authenticity of the HTTP request.

#### **Configuring X.509 Certificate Validation**

Before you can configure the OAMCertificateVerifier, you must have the OAM Trusted Key for the embedded certificate that OAM Keystore uses. Once you identify the OAM Trusted Key, you must export the certificate to a local keystore that the OAMCertificateVerifier can access. Refer to the OAM documentation for detailed information on obtaining the OAM Trusted Key and exporting a certificate.

To export the OAM Identity Assertion X509 Certificate, perform the following. These steps assume that you are using Oracle WebLogic and may differ depending on your environment. Refer to your OAM documentation for information on performing these steps for other applications:

- Retrieve the Oracle Keystore Password using the WLST connect() command and then the listCred(map="OAM\_STORE", key="jks") command. Make note of the password and location of the keystore.
- 2. Export the certificate using the following syntax:

```
keytool -exportcert -keystore $domain_home/config/fmwconfig/
default-keystore.jks -storetype JKS -alias keyname -file $cert_file
```

Provide the location of the keystore noted in Step 1, as well as the keystore type and the name of the file that will hold the certificate. The following example extracts the certificate from the base\_domain/config/fmwconfig/directory and provides an alias of assertion-key, identifies the key store as oamkeystore and the certificate file as assertion.cer:

```
keytool -exportcert -v -alias assertion-key -storetype JKS -keystore .oamkeystore -file assertion.cer
```

- Once the certificate has been exported to the file you identified in Step 2, copy the certificate file to the ATG server.
- 4. Import the certificate into the <code>oamkeystore.ks</code> on the ATG server by running the <code>importcert</code> command. The <code>importcert</code> command uses the following syntax:

```
keytool -importcert -keystore $domain_home/config/fmwconfig/
default-keystore.jks -storetype JKS -alias $trusted_cert_alias
-file $trusted_cert_file
```

#### For example:

```
keytool -importcert -v -keystore C:\all\oamkeystore.ks
-storetype JKS -alias assertion-key -file C:\all\assertion.cer
```

#### This produces output similar to the following:

```
Enter keystore password:
Re-enter new password:
```

```
Owner: CN=OAM User Assertion Issuer CA Root
Issuer: CN=OAM User Assertion Issuer CA Root
Serial number: 66
Valid from: Tue May 21 05:54:34 CST 2013 until: Fri May 19 05:54:34 CST 2023
Certificate fingerprints:
MD5: 03:A3:6D:C7:AC:36:D7:30:01:6B:34:52:97:B0:DD:6B
SHA1: 37:A7:CB:F0:3E:BF:99:D9:93:51:0D:B3:9C:AA:9C:43:0A:0C:30:79
Signature algorithm name: MD5withRSA
Version: 1
Trust this certificate? [no]: y
Certificate was added to keystore
[Storing c:\all\oamkeystore.ks]
```

When prompted, enter the keystore password noted in Step 1. When asked to trust the certificate, answer yes to add the certificate to the keystore.

5. Configure the OamCertificateVerifier.properties file in your /localconfig directory to include the keystore locations:

```
$class=atg.userprofiling.oam.security.CertificateVerifier
keystoreLocation=C:/all/oamkeystore.ks
keystoreType=JKS
certificateHeaderName=oam_identity_assertion
keyStorePassword=oampass
```

6. Configure the OamRemoteUserServlet.properties file in your /localconfig directory to add the OamCertificateVerifier:

```
# add the OamCertificateVerifier to the trustVerifiers property
trustVerifiers=\
/atg/userprofiling/oam/OamCertificateVerifier
```

Once you have configured the X.509 certificate validation, you can provide additional validation by configuring the system to recognize OAM authentication cookies.

#### **Configuring OAM Authentication Cookies**

By default, OAM maintains a session cookie that uses the OAMAuthnCookie<host:port><random number> format. For detailed information on OAM authentication cookies and their formats, refer to your OAM documentation.

The ATG platform can be configured to check for the presence of OAM authentication cookies by performing the following steps:

1. Configure the /atg/userprofiling/oam/OamCookieVerifier.properties file in your /localconfig directory:

```
$class=atg.userprofiling.oam.security.OAMCookieVerifier
# For 11g webgates the cookieNamePrefix property value is as below.
# For 10g webgates use ObSSOCookie.
cookieNamePrefix=OAMAuthnCookie_
webgateHost^=Configuration.webgateHost
webgatePort^=Configuration.webgatePort
```

2. Configure the /atg/dynamo/servlet/dafpipeline/
 OamRemoteUserServlet.properties file to use the OAMCookieVerifier:

```
# add the OAMCookieVerifier to the trustVerifiers property
trustVerifiers=\
/atg/userprofiling/oam/OAMCookieVerifier
```

Note that OAM adds this cookie into all HTTP requests. This is different than the Identity Assertion header, which is only added to requests that result in SSO user authentication and the creation of new OAM SSO sessions.

Note: It may be necessary to add or upgrade the user-defined parameter filterOAMAuthnCookie in the WebGate 11g configuration with a value of false. Refer to the *Administrator's Guide for Oracle Access Manager* for details on the user-defined parameter filterOAMAuthnCookie. It may be necessary to restart this server for the change to take effect.

## **Using IP Address Filtering for Verified Authentication**

In some cases, you may not want to use OAM Identity Assertion outlined in the Validating OAM Identity Assertion X.509 Certificates section. Or your Web server may not support this OAM feature. To ensure verified authentication requests you can configure IP address filtering for your specific Web server.

#### **Oracle WebLogic IP Filtering**

WebLogic provides a default connection filter that can perform IP address filtering. The weblogic.security.net.ConnectionFilterImpl filter accepts all incoming network connections and provides static methods that allow the server to obtain the current network connection filter. This filter can be configured using the WebLogic Administration Console, or by editing the domain config.xml file.

#### To Configure IP Filtering Using the WebLogic Administration Console

- 1. Open the WebLogic Administration Console and navigate to the domain.
- 2. Select Security > Filter from the tab menu.
- In the Connection Filter property, enter the connection filter class weblogic.security.net.ConnectionFilterImpl.
- 4. In the Connection Filter Rules property, enter the connection filter rule. When entering a rule, enter them in following format:

targetAddress LocalAddress localPort action protocols

5. Restart your servers to initiate the configuration changes.

#### To Configure IP Filtering By Editing the config.xml File

- Open the WebLogic\_home/user\_projects/domain/domain\_name/config/ config.xml file.
- 2. Edit the file to add the following filter configuration:

<security-configuration>

```
...
<connection-filter>weblogic.security.net.ConnectionFilterImpl
</connection-filter>
<connection-filter-rule>192.168.0.0 127.0.0.1 7001 allow http https t3
t3s</connection-filter-rule>
<connection-filter-rule>192.168.0.0 127.0.0.1 7003 deny
</connection-filter-rule>
</security-configuration>
```

3. Save the config.xml file and restart your server.

The following is an example of a network connection filter rule:

For further examples, and additional information, refer to your WebLogic documentation for the ConnectionFilterImpl filter.

#### JBoss and Tomcat IP Filtering

JBoss and Tomcat's servlet container provide a built-in filter that can be used to configure IP Address Filtering. The org.apache.catalina.filters.RemoteAddrFilter can be configured in the web.xml file of any application to which you want to restrict access. The filter configuration allows you to specify initial parameters that allow or deny a set of IP addresses or a range of IP addresses represented by regex patterns.

For example:

```
<filter>
<filter>
<filter-name>Remote Address Filter</filter-name>
<filter-class>org.apache.catalina.filters.RemoteAddrFilter</filter-class>
<init-param>
<param-name>deny</param-name>
<param-value>127\.\d+\.\d+\:\l0:0:0:0:0:0:0:0:0:1</param-value>
</init-param>
</filter>
<filter-mapping>
<filter-name>Remote Address Filter</filter-name>
<url-pattern>/*</url-pattern>
</filter-mapping>
```

For detailed information on the RemoteAddrFilter, refer to the RemoteAddrFilter API documentation on the http://tomcat.apache.org Web site.

#### IBM WebSphere IP Filtering

WebSphere TCP transport chains allow you to reject or allow connections to an application from a list or range of IP addresses. To set up an inbound transport chain, perform the following steps:

- 1. Open the WebSphere Administration Server Console.
- 2. Select Server > Server Types > WebSphere Application Server > server\_name > Ports.
- 3. Click View Associated Transports to see the ports that are associated with the TCP transport channel of the application to which you want to restrict access.
- 4. Select a transport chain.
- 5. Select TCP Inbound Channel (TCP 2).
- 6. The configuration page for the TCP transport chain will be displayed. Lists or ranges of IP addresses can be added under the Address Exclude List and Address Include List to exclude or include IP addresses that can access the application.

For detailed information on transport chains, refer to your IBM WebSphere documentation.

## **Recommended Deployment**

The following diagram displays the recommended deployment for OAM SSO integration.

For detailed information on OAM deployment methodologies, refer to the Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management.



# Index

## A

ACC (see ATG Control Center (ACC)) access levels, 96 accounts, 96 **Business Control Center, 16** default user, 15 MySQL, 6 ATG Control Center (ACC), 54 and component properties, 69 defining passwords, 16 downloading, 55 editing configuration layers, 67 installing, 54 on UNIX, 55 on Windows, 55 scenarios with JBoss, 21 with WebLogic, 19 with WebSphere, 23 searching with, 68 starting, 59 on a client, 62 on a dedicated VM, 18, 60 troubleshooting, 62, 63

### С

caches, 155 FileCache size, 173 loading, 52 locked mode, 94 memory leaks and, 127 modes, 93 prepopulating, 94 repository, 93, 94 server configuration, 84 caching, 156 external, 156 JVM parameters, 164 CIM (see Configuration and Installation Manager (CIM)) CLASSPATH, 75

and URLHammer, 131 customizing, 75 JBoss, 36 WebLogic, 20, 32 Configuration and Installation Manager (CIM), 10, 25 adding modules to, 11 and JBoss versions, 22 and WebSphere, 4 installing demonstrations, 14 prerequisites, 4, 10 configuration group, 79, 80, 82 configuring, 81 validating, 85 configuration layers, 65, 66 locking, 68 resetting default, 67 cookies, 115, 122, 129, 129, 134, 181

## D

data sources, 168 connection pool, 33 debugging, 39 import scripts, 32 JBoss, 35 logging, 39 Oracle, 38 switching, 51 WebLogic, 37 WebSphere, 37 XA, 19, 23, 32 databases configuration, 5, 25 copying, 48 creating tables, 26 DB2, 42 defining tables, 41 destroying tables, 29 moving to production, 46 MSSQL, 43 MySQL, 6 Oracle, 26 supported, 25 switching, 51 tables DAS, 26 DPS, 27 DSS, 27 troubleshooting, 159 DBCopier, 48 configuring, 49 creating, 48 DB2, 50 MSSQL, 50

native SQL environment, 50 Oracle, 50 debugging, 39 logging, 74 drivers connection pool, 33 DB2, 37 iNet on JBoss, 44 on WebLogic, 44 JDBC, 32, 41, 42 MySQL, 3 WebLogic, 32 location of database, 36

### Ε

e-mail, 71, 95 targeted, 96 Eclipse, 175 escape character, 72

### F

FileCache, 173 full text searches, 160

## G

GZIP compression, 17

### Η

HotSpot, 93 HttpOnly, 115 httpPort, 99 HTTPS protocol, 99

#### I

indexing, 98, 155 iNet Merlia driver, 45 installing, 3 JBoss, 20 isolation levels, 34 recommended, 43 WebSphere, 39

#### J

JAVA\_ARGS, 75 customizing, 75 JBoss, 161 and DB2, 37 and MSSQL, 36 bind address, 5 CLASSPATH, 36 clustering, 99

configuration, 8 data sources, 34 demo database, 25 iNet, 44, 45 installing, 20 IP filtering, 183 lock management, 21 Log4j, 163 logging, 77 remove services, 164 sessionID checking, 21 transaction timeout, 38 version differences, 22 JDBC Browser, 41 connections, 33 drivers, 32, 32, 41, 42 MSSQL, 43 with Oracle, 32 with WebLogic, 32 URL on Oracle RAC, 38 JVM, 75 and DBCopier, 50 garbage collection, 127 JBoss, 20 locale, 75 memory errors, 126 on Linux, 8 on WebSphere, 23, 89 paging size, 126 stack traces, 124 thread dumps, 124 tuning, 164

### L

liveconfig, 91 lock management, 99 JBoss, 21, 92 repositories, 93, 94 Sun T1000, 54 logging, 74, 76, 97, 117, 118, 118, 163 disable, 77 disabling, 117

#### Μ

messaging mode, 19, 102 metric pricing reports, 118 modules, 16, 26, 76 customizing, 76 Motorprise, 14 MS SQL Server, 43 and Unicode, 43 JDBC drivers, 43 MSSQL, 39 MySQL, 6, 7 on Microsoft Windows, 6

## 0

OLTP, 160 Oracle Access Management, 12, 177

## Ρ

Performance Monitor, 92, 142 performance sampling, 149 Portal Application Framework (PAF), 28 properties, 58, 97 properties file, 66 properties files, 97 ProtocolChange servlet bean, 98

## Q

Quincy Funds, 14 tranfer data, 47

### R

Remote Method Invocation (RMI), 61, 75 repositories, 93, , 169

## S

scenarios, 19, 21, 23, 26 schemas, 167 Secure Sockets Layer (SSL), 17 session backup, 107 Single Sign On, 12 Solaris, 125 startSQLRepository, 46 SupportedApplicationServer, 18

## Т

transaction timeout, 38, 38, 39

### U

Unicode, 43 uninstalling, 55 URLHammer, 128

#### W

WebLogic, 183 -standalone flag, 101 509 certificate, 180 ACC in dedicated VM, 18 CLASSPATH, 20 clustering, 100 configuration, 8

configuration group, 80 configuring, 4 file name encoding, 18 iNet, 44 IP filtering, 182 JDBC driver, 32 memory allocation, 19 page recompilation, 19 protocol.jar, 20 Server Standard Edition, 4 SQLJMSAdmin, 19 startManagedWebLogic, 20 transaction timeout, 38 XA data sources, 19 WebSphere and CIM, 11 clustering, 102 configuration, 9 configuring, 5 multisite, 23 on AIX, 22 protocol.jar, 23 transaction isolation level, 39 transaction timeout, 39 XA data source, 23

## X

X.509 certificates, 179