



P r o f e s s i o n a l   E x p e r t i s e   D i s t i l l e d

# Oracle Business Intelligence Enterprise Edition 11g: A Hands-On Tutorial

Leverage the latest Fusion Middleware Business Intelligence offering with this action-packed implementation guide

**Haroun Khan**  
**Adrian Ward**

**Christian Screen**

**[PACKT]** enterprise   
PUBLISHING professional expertise distilled

# Oracle Business Intelligence Enterprise Edition 11g: A Hands-On Tutorial

Leverage the latest Fusion Middleware Business Intelligence offering with this action-packed implementation guide

**Haroun Khan**

**Christian Screen**

**Adrian Ward**



BIRMINGHAM - MUMBAI

# Oracle Business Intelligence Enterprise Edition 11g: A Hands-On Tutorial

Copyright © 2012 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: July 2012

Production Reference: 1090712

Published by Packt Publishing Ltd.  
Livery Place  
35 Livery Street  
Birmingham B3 2PB, UK.

ISBN 978-1-84968-566-5

[www.packtpub.com](http://www.packtpub.com)

Cover Image by Sandeep Babu ([sandyjb@gmail.com](mailto:sandyjb@gmail.com))

# Credits

**Authors**

Haroun Khan  
Christian Screen  
Adrian Ward

**Reviewers**

Daan Bakboord  
Kevin McGinley  
Ramke Ramakrishnan

**Acquisition Editor**

Stephanie Moss

**Lead Technical Editors**

Arun Nadar  
Azharuddin Sheikh

**Technical Editors**

Vrinda Amberkar  
Prasad Dalvi

**Project Coordinator**

Yashodhan Dere

**Proofreader**

Linda Morris

**Indexer**

Monica Ajmera Mehta

**Graphics**

Manu Joseph

**Production Coordinator**

Aparna Bhagat

**Cover Work**

Aparna Bhagat

# About the Authors

**Haroun Khan** is one of Europe's leading OBIEE consultants. Being a Computer Science graduate from Imperial College, London, he has been involved with OBIEE from its early days as an acquisition from nQuire by Siebel, and subsequently as part of the Oracle family. Haroun has worked as a consultant on projects worldwide for Siebel and as a Principal Consultant for Oracle over a period of 10 years. He has specialized in BI and data warehousing over a longer period including time working at MicroStrategy. Haroun now freelances in leading and designing projects in the BI and data warehousing space, combining this with entrepreneurial activities, such as his own e-commerce business JRPass.com. In his downtime, Haroun likes to spend as much time as possible climbing in the mountains, away from a computer screen and avoiding numerous requests to write a blog or get on another plane!

Haroun works through his own company Awaan and can be contacted at his company address [haroun@awaan.eu](mailto:haroun@awaan.eu).

---

It goes without saying, but I would like to thank my parents, Zainab and Ayub for their support, encouragement and for everything, really. Also I have to include my sisters Sophia and Soraya so that they can have their name in print and I can avoid admonishments! Last, but not least, thanks to my toddler niece Sharifa for providing me with boundless distractions and opportunities to procrastinate.

Thanks goes to all those people (too many to mention) who I've worked with on great projects during my time at Oracle and MicroStrategy. Thanks also go out to Packt and their editors for their work. Finally, I would also like to extend my appreciation to my co-authors for their commitment and energy. It's been tiring but fun!

---

**Christian Screen** is a Business Intelligence evangelist with over 15 years of experience in technology ranging from low-level programming, E-Commerce, Data Warehousing, Enterprise Performance Management and, of course, Business Intelligence. In his spare time, he enjoys writing technical articles, learning new technologies, developing products, writing software, spending time with his family, trying to change the world, and running his blog and podcast at ArtOfBi.com. He is an Oracle ACE, an Oracle Deputy CTO, and holds several technology and project management certifications.

---

I would like to thank my family (wife Kirsten and three super children – Riley, Jaxon, and Dylan) seems so cliché but indeed their patience and support has been essential to the process of writing my first book. I'd also like to thank all of my colleagues who have provided continual intellectual ferment for our work with Business Intelligence and Enterprise Performance Management solutions. Thanks to Greg, Jason, and Tom at Analytic Vision for giving me a chance with my first consulting gig. Thanks to Amy Mayer of BI Consulting Group and Capgemini for my second consulting gig and having an unbelievable culture that I'll take with me wherever I may go. There are a few people at Oracle to thank as well, such as Mike Hallett in the UK who handles the BI/EPM partner community for EMEA and does an amazing job at building the community and keeping even us in the USA informed, the Oracle ACE Program leaders Justin Kestelyn and Lillian Buziak (Brian Stover, of course, many thanks!) for the Oracle ACE moniker, and the Oracle BI development team Matt Bedin and Phillipe Lions for their SampleApp and other assistance. I'd like to generally thank other authors of technical books that I may have referred to over the years and the many bloggers that find it their duty to share the tidbits, prose, or other knowledge that make finding an answer to daily issues merely a Google search away – thanks for sharing. A big thanks goes to the very keen reviewers of this book – Kevin McGinley, Daan Bakboord, and Ramke Ramakrishnan. We truly appreciate the time, effort, and suggestions you've provided to make this book better. Lastly, thanks to you for reading our first book – hopefully you purchased it legally – I trust you will find it useful.

---

**Adrian Ward** started working in Siebel Analytics back in 2001 and quickly realized the potential in the technology. He formed the UK's first independent consultancy focusing purely on OBIEE (née Siebel Analytics) and Oracle BI Applications. He has led many large successful OBIEE implementations in a wide range of business sectors, from Investment Banking to Military operations. His deep technical OBIEE and BI Applications knowledge has been applied on dozens of projects throughout the globe including HR, Sales, Service, Pharma, and Custom Analytics.

He was also one of the first bloggers on Oracle BIEE and today runs the Addidici OBIEE consultancy which has operations in the UK, Europe, and South Africa.

Adrian runs one of the largest Oracle BI networking groups on LinkedIn – "Oracle Business Intelligence", and helps others to network and learn about the product and its application, including organizing networking social events in London.

In his spare time he loves sailing, skiing, enjoying life with his family, and learning new technologies.

---

Firstly, I would like to thank my wife Sarah for her enormous love and support over the years, and in particular whilst I was writing my part of this book. Thanks too, must go to my cool children, Hugh and Hatty, for their help in keeping the house quiet and delivering endless cups of tea, and also to my Mum and Dad for being the best parents you could ask for – always there when you need.

I will be eternally grateful to Narmada for her selfless support and to Chet Justice for his great sense of humor. Special thanks go to James Robinson for my first decent job, to my great clients for employing me, and the great people I have worked with over the last 20 years (including Steve Lomax, Robert Patterson, Trev Harvey, Eric Gravit, Adrian Ball, Haider Tirmizi, Luis, Piere, Andi Schloegl, Neil Ashton, Daniel and many more).

I am also indebted to Graeme Hampshire for inspiring me to get writing, and for helping to keep sailing fun, John Dunnet and the crew of The Beefeater for putting up with my captaincy, and Jon Spencer for being a great PRO.

Finally thanks to my special friends Daniel, Andy & Jenny, and Jamie & Jackie for putting up with me over the years!

---

# About the Reviewers

**Daan Bakboord** is a full Oracle BI (Applications) Consultant with extensive experience in the deployment of Oracle BI Tools (Oracle BI EE, Oracle BI Publisher, Oracle BI Applications, and so on) in general and its application in an Oracle EBS environment in particular.

Daan is employed by Ebicus – an ICT service provider in the field of Oracle/Siebel CRM and Oracle BI (Applications). Within Ebicus, Daan is responsible for the development of knowledge and the dissemination of this knowledge to the market. One of the manifestations of this is the blog (*Oracle BI By Bakboord* – <http://obibb.wordpress.com>), which is maintained by Daan. In addition, Daan is active in several (online) forums, making him a part of the major Oracle BI community.

Besides his work, Daan is a proud father of a son and a daughter. He is also active in amateur football.

**Kevin McGinley** has worked in BI/Data Warehousing in both IT and consulting since 1997. He has helped both large and small companies define and execute BI/DW roadmaps and implementations, focusing exclusively on Oracle BI since 2005. Kevin is a recognized expert on Oracle BI through speaking at many conferences, publishing articles, giving master classes, leading user group conferences, the co-host of the YouTube podcast *Real-Time BI with Kevin & Stewart*, and is the co-author of *Oracle BI Enterprise Edition Dashboard & Report Best Practices*.

---

I'd like to thank the authors for the opportunity to review their fantastic book. There are so few books on Oracle BI out in the market and it's refreshing to see a nice, straightforward, end-to-end book on OBIEE 11g available for readers new to Oracle BI.

---



**Ramke Ramakrishnan** has performed the Lead Architect and Technical Leadership roles for over 15 years on Business Intelligence, by effectively managing the project team and the business customer expectations.

He delivers hands-on capabilities in the configuration of robust Oracle database and BI architectures, Oracle's Essbase infrastructure and large scale Business Intelligence Reporting, OBI Applications and EPM implementations. He is the key contributor for Business Analytics and Enterprise Reporting by integrating various applications systems into analytics to empower business customers, executives and end users.

He deployed several Enterprise Information Integration (Ei2) architectures, the core framework for Data Warehousing, Data Marts, OLAP, Business Analytics and Enterprise Reporting.

Ramke is Oracle Implementation Certified on Business Intelligence and EPM technologies. He is an active member and designated Deputy CTO (DCTO) on the Oracle Business Intelligence Investment Partner Community (IPC). He has featured in several Oracle speaking events and Oracle Press Releases on BI topics.

Currently, he is employed as Practice Director – BI and EPM with MarketSphere. MarketSphere is a strategic advisory and technology consulting firm with a strong focus on Oracle that helps our clients to deliver integrated ERP, BI and EPM solutions to optimize business performance. For more information, visit [www.marketsphere.com](http://www.marketsphere.com).

---

I want to thank my wife Lavanya and my children Ritvik, Rasya, Rishik, and Raeya for their support and cooperation.

---

# www.PacktPub.com

## Support files, eBooks, discount offers and more

You might want to visit [www.PacktPub.com](http://www.PacktPub.com) for support files and downloads related to your book.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at [www.PacktPub.com](http://www.PacktPub.com) and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at [service@packtpub.com](mailto:service@packtpub.com) for more details.

At [www.PacktPub.com](http://www.PacktPub.com), you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<http://PacktLib.PacktPub.com>

Do you need instant solutions to your IT questions? PacktLib is Packt's online digital book library. Here, you can access, read and search across Packt's entire library of books.

## Why Subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print and bookmark content
- On demand and accessible via web browser

## Free Access for Packt account holders

If you have an account with Packt at [www.PacktPub.com](http://www.PacktPub.com), you can use this to access PacktLib today and view nine entirely free books. Simply use your login credentials for immediate access.

## Instant Updates on New Packt Books

Get notified! Find out when new books are published by following @PacktEnterprise on Twitter, or the Packt Enterprise Facebook page.



# Table of Contents

<b>Preface</b>	<b>1</b>
<b>Chapter 1: Understanding the Oracle BI 11g Architecture</b>	<b>9</b>
Looking backward and looking forward	9
Let's look at the big picture	10
What is Oracle Fusion Middleware?	11
An application server by any other name	12
A database repository – for what?	12
Overall components	13
Java components	14
System components	15
<b>WebLogic Server</b>	<b>17</b>
A few software nuances	17
WebLogic Domain	18
WebLogic Administration Server	18
WebLogic Managed Server	19
WebLogic Node Manager	20
System tools controlled by WebLogic	21
<b>Oracle Process Management and Notification system</b>	<b>22</b>
<b>Security</b>	<b>23</b>
Backwards compatibility	24
Managing by application roles	25
Security providers	25
Identity Store	25
Credential Store	25
Policy Store	26
<b>System requirements</b>	<b>27</b>
Client Tools	27
Multiuser Development Environment	28
Certification matrix	30

<b>Scaling out Oracle BI 11g</b>	<b>30</b>
Preconfiguration run down	31
Shared storage	31
Clustering	31
Vertical expansion versus horizontal expansion	32
Oracle BI Server (system component) Cluster Controller	32
Failover and high availability	33
Enterprise deployment guide	33
<b>Directory folder structure</b>	<b>34</b>
Log files (diagnostics)	34
Configuration files	35
<b>A review – what I should now know!</b>	<b>36</b>
<b>Additional research suggestions</b>	<b>36</b>
<b>Summary</b>	<b>37</b>
<b>Chapter 2: Installing the Metadata Repository</b>	<b>39</b>
<b>Repository Creation Utility (RCU)</b>	<b>39</b>
What is the metadata store?	40
Technical metadata	40
Business metadata	40
Process metadata	41
Oracle Metadata Services (MDS)	41
Repository schemas	42
Non-repository metadata	43
Downloading the software	43
<b>Running the RCU</b>	<b>44</b>
Setting up your database	45
Oracle	45
MS SQL Server	46
Preparing for your installation	47
Linux	47
64 bit Linux	48
Standard installation steps	48
Adding repositories	60
Customizing your installation	61
Other common changes	63
Schema installations	65
Oracle CREATE USER Script	66
Silent installation	66
Example creation	68
Useful scripts	69
Usage tracking script	69
<b>A review – what I should now know!</b>	<b>71</b>
<b>Summary</b>	<b>71</b>

---

<b>Chapter 3: Installing on Windows Server 2008</b>	<b>73</b>
<b>Installation media</b>	<b>73</b>
<b>System requirements</b>	<b>75</b>
<b>Installation</b>	<b>77</b>
<b>A review – what I should now know!</b>	<b>91</b>
<b>Summary</b>	<b>92</b>
<b>Chapter 4: Installation Options</b>	<b>93</b>
<b>Oracle BI on its own server</b>	<b>93</b>
High availability and failover planning	94
<b>Simple versus Software Only Installation</b>	<b>95</b>
Software Only Install	95
Installing your own JDK	95
Simple Install	96
Silent installation	96
<b>Custom static ports</b>	<b>97</b>
Creating your own staticports.ini file	98
<b>Installing Oracle BI 11g on *Nix</b>	<b>100</b>
<b>Listening on port 80</b>	<b>100</b>
Ensuring IIS web server Role Services are installed	102
Getting the WebLogic Server Proxy IIS plugins	103
Creating and configuring an IIS Website	104
Enabling compression in IIS 7.x	107
<b>Automate starting and stopping</b>	<b>108</b>
Leveraging the WebLogic Server Windows Service installer command	109
Creating start-up and shutdown scripts	111
Creating desktop shortcuts	113
Creating boot.properties files	114
<b>Ancillary application integration awareness</b>	<b>115</b>
<b>Recommendations for further learning</b>	<b>116</b>
<b>A review – what I should now know!</b>	<b>116</b>
<b>Summary</b>	<b>117</b>
<b>Chapter 5: Understanding the Systems Management Tools</b>	<b>119</b>
<b>Let's talk about management tools</b>	<b>119</b>
<b>WebLogic Server Administration Control</b>	<b>120</b>
First access and checkpoint	121
Servers	122
Clusters	122
Machines, IP address, or DNS	123
Data sources or JDBC connections	124

Security realms	126
WebLogic Server is its own application	128
<b>Using WLST</b>	<b>129</b>
<b>Enterprise Manager Fusion Middleware Control</b>	<b>131</b>
Getting around in EM	132
BI Foundation Domain dashboard	132
Coreapplication	133
The Overview tab	134
The Availability tab	134
The Capacity Management tab	134
The Diagnostics tab	135
The Security tab	136
The Deployment tab	136
Managing Oracle BI 11g artifacts	137
Creating a new Presentation Catalog	138
Deploying an existing Presentation Catalog	140
Deploying an RPD	140
Starting/stopping system components	142
Checking the logs	143
<b>Creating the Tennis users, roles, and associations</b>	<b>144</b>
Creating users and groups in WLS	145
Assigning users to groups	146
Creating and assigning application roles	147
<b>Configuring an LDAP identity provider in WLS</b>	<b>149</b>
Setting up the identity provider	150
Establishing Fusion Middleware Control security linkage	153
Restarting the WebLogic Server and Managed Server	156
Assigning the LDAP Service Account Role Privileges	158
JMS modules for BI Publisher communication	160
Refreshing Presentation Catalog GUIDs – sync it up!	160
<b>JMX, MBeans, and Java</b>	<b>163</b>
<b>Migrating FMW Security to other environments</b>	<b>163</b>
FMW core security files	163
Project Amelia	164
Downloading Project Amelia	164
Getting the FMW Security file	165
Running the script and generating the WLST script	165
Migrating the security script and running it on the target server	166
FMW Security Import/Export utility	166
Using the Security Realm Migration utility	167
Using the migrateSecurityStore function via WLST	168
<b>Oracle BI Publisher system management</b>	<b>169</b>
<b>Monitoring system performance</b>	<b>169</b>

<b>Have a backup plan!</b>	<b>170</b>
<b>Recommendations for further learning</b>	<b>170</b>
<b>A review – what I should now know!</b>	<b>171</b>
<b>Summary</b>	<b>171</b>
<b>Chapter 6: Upgrading the RPD and Web Catalog to 11g</b>	<b>173</b>
<b>Upgrading an RPD and Web Catalog</b>	<b>173</b>
<b>Upgrade Assistant</b>	<b>175</b>
<b>Verification</b>	<b>183</b>
<b>Upgrading BI Scheduler</b>	<b>184</b>
<b>Manual migration</b>	<b>188</b>
<b>Consistency check</b>	<b>188</b>
<b>Security</b>	<b>189</b>
<b>Regression testing</b>	<b>190</b>
Unit test	190
Full regression testing	191
User Acceptance testing	191
<b>A review – what I should now know!</b>	<b>192</b>
<b>Summary</b>	<b>193</b>
<b>Chapter 7: Reporting Databases</b>	<b>195</b>
<b>Theories and models</b>	<b>196</b>
Reporting databases	197
Relational modeling	198
Dimensional modelling	198
Why is database theory important?	200
<b>Designing your database – objectives, rules, and goals</b>	<b>201</b>
Objectives	202
Rules	202
Rule 1 – complete dimensions	203
Rule 2 – build generic tables	204
Rule 3 – partition large tables	204
Rule 4 – prudent indexing	204
Rule 5 – aggregate everything	205
Rule 6 – constant analysis of usage and accuracy	205
Rule 7 – manage statistics	205
Rule 8 – understand the granularity	206
Goals	206
Goal 1 – keep it simple	207
Goal 2 – minimize type 2 slowly changing dimensions	207
Goal 3 – use data, not functions	207
Goal 4 – minimize joins	207
Goal 5 – reduce snowflaking	207
Goal 6 – make it flexible	208
Design summary	208



<b>Creating a warehouse</b>	<b>208</b>
Source system assessment	208
Warehouse design	211
Warehouse tables	211
Populate and tune	214
Monitor and maintain	214
<b>Some definitions</b>	<b>215</b>
<b>A review – what I should now know!</b>	<b>216</b>
<b>Summary</b>	<b>216</b>
<b>Chapter 8: Developing a BI Repository</b>	<b>217</b>
<b>Prerequisites</b>	<b>218</b>
<b>Repository architecture</b>	<b>218</b>
Physical layer	218
Business layer	218
Presentation layer	219
<b>Physical layer</b>	<b>219</b>
Creating an RPD and importing metadata	219
Elements of the physical layer	225
Database object	226
Connection pools	228
Physical catalog and schemas	230
Physical tables	230
Physical join	232
Consistency check	235
Table aliases and naming conventions	236
<b>Business layer</b>	<b>238</b>
Business model	238
Logical tables	239
Logical table sources	242
Logical columns	243
Logical joins	246
Dimension hierarchies	249
<b>Presentation layer</b>	<b>255</b>
Subject areas	256
Best practices in the presentation layer	258
Aliases	260
Implicit fact	260
<b>Calculated measures</b>	<b>262</b>
Logical column calculation	262
Expression Builder	264
Physical column calculation	266
Time series measures	268
Level based measure	271

---

<b>Federated and fragmented content</b>	<b>273</b>
Vertical federation – aggregation or level based	273
Horizontal federation	274
Fragmentation	274
Fragmentation example – content based	274
<b>Variables and initialization blocks</b>	<b>279</b>
<b>A review – what I should now know!</b>	<b>286</b>
<b>Additional research suggestions</b>	<b>287</b>
<b>Summary</b>	<b>287</b>
<b>Chapter 9: Features of the Presentation Catalog</b>	<b>289</b>
<b>Integrated tools</b>	<b>290</b>
Analysis	290
Dashboards	291
Published reporting	291
Actionable Intelligence	291
Performance management	292
Marketing	292
Mapping	293
Administration	293
Briefing books	293
Search	294
Help	294
Office integration	294
<b>The Home screen</b>	<b>294</b>
Common links	295
Existing object links	295
Create objects	296
Browse catalog	296
Helpful links	296
<b>Administration</b>	<b>296</b>
Groups and users	297
Privileges	298
Full privileges list	300
Session management	300
Maintenance and troubleshooting	303
Issue SQL	303
<b>The Presentation Catalog</b>	<b>304</b>
Structure of the Presentation Catalog	304
Hidden items	306
File management	306
XML files	306

Object copying	306
Multiple personal dashboards	307
Catalog deployments	307
Securing catalog objects	307
Permission inheritance	309
Practical steps to object security	309
<b>A review – what I should now know!</b>	<b>310</b>
<b>Summary</b>	<b>310</b>
<b>Chapter 10: Creating Dashboards and Analysis</b>	<b>311</b>
<b>Analysis versus Reporting</b>	<b>311</b>
<b>Creating an Analysis</b>	<b>312</b>
Analysis Editor	314
Criteria Tab	314
The Results tab	316
Filters	317
Selection Steps	320
Saving an Analysis	324
Basic table formatting	325
Bins	328
Pivot Tables	332
Graphs	337
Sectioned views	338
View prompts and section sliders	340
Conditional Formatting	342
<b>Building Dashboards</b>	<b>346</b>
Creating a dashboard	346
Dashboard builder	347
Editing a dashboard	350
Adding a page	351
Report Links	354
Dashboard Prompt	356
Presentation variable	361
Protected and "is prompted" filters	363
Repository/Session variable	364
<b>Some advanced options</b>	<b>366</b>
Column Selector	366
View Selector	368
Master-detail linking	371
Hierarchical columns	374
<b>Security</b>	<b>376</b>
Object security	377
Data security	379

<b>Design – best practices</b>	<b>385</b>
<b>Additional research suggestions</b>	<b>386</b>
<b>A review – what I should now know!</b>	<b>387</b>
<b>Summary</b>	<b>387</b>
<b>Chapter 11: Agents and the Action Framework</b>	<b>389</b>
<b>Agents</b>	<b>390</b>
<b>Actions</b>	<b>399</b>
BI navigation	399
Web navigation and passing a parameter	402
Conditionality	404
Note on invoke actions	408
<b>A review – what I should now know!</b>	<b>409</b>
<b>Additional research suggestions</b>	<b>410</b>
<b>Summary</b>	<b>410</b>
<b>Chapter 12: Developing Reports Using BI Publisher</b>	<b>411</b>
<b>Don't miss the installation integration checkpoint!</b>	<b>412</b>
<b>Where BI Publisher excels?</b>	<b>413</b>
What's all this XML talk?	413
<b>Yes, BI Publisher is now Published Reporting</b>	<b>414</b>
<b>Oracle BI Foundation versus Oracle BI Publisher</b>	<b>415</b>
<b>New features and enhancements</b>	<b>416</b>
Improved Oracle BI 11g look and feel	416
Interactive Viewer	416
Dynamic upload and source from MS Excel spreadsheets	417
Leveraging LDAP user attributes in queries	417
Sharing a Presentation Catalog with Oracle BI 11g	<b>417</b>
Data Model Editor	417
Leverage view objects – Application Development Framework	418
<b>Report design basics, terminology, and locations</b>	<b>418</b>
Report design components	418
Data model	418
Layout	419
Properties	419
Translations	419
Where to administrate BI Publisher	419
Default embedded BI Publisher configurations	420
Where to build a data model	420
Where to add a data source connection	420
What is a JNDI data source?	421

<b>Let's get publishing</b>	<b>422</b>
Administration management of BI Publisher	422
Accessing the BI Publisher Administration page	422
Verifying application roles	423
Creating the Tennis data source JDBC connection	424
Creating a File (XLS) data source	425
Verifying application role data source privileges	426
Setting up a data model	426
Creating a new Presentation Catalog folder	427
Creating a new data model	428
Creating a SQL query data set	429
Adjusting data set display names	432
Creating a parameter	433
Creating a list of values	434
Connecting the parameter to the list of values	435
Getting the sample data	436
Creating a BI Publisher report using Layout Editor	438
<b>Auditing and monitoring BI Publisher</b>	<b>441</b>
Modifying a few configuration files	442
Connecting to the Audit Framework	443
Enabling Audit Policy in the Fusion Middleware Control	
Enterprise Manager	443
Viewing the auditing log file	444
<b>BI Publisher nuances</b>	<b>445</b>
Timeout issues	445
Connecting to Oracle BI server data sources	445
BI Publisher Application Programming Interface (API)	446
BI Publisher Scheduler	446
High availability	446
<b>A review – what I should now know!</b>	<b>447</b>
<b>Additional research suggestions</b>	<b>447</b>
<b>Summary</b>	<b>448</b>
<b>Chapter 13: Customizing the Style of Dashboards</b>	<b>449</b>
<b>What's the idea?</b>	<b>449</b>
<b>Multiple skins and styles in one environment</b>	<b>450</b>
<b>A strategy for something seemingly simple</b>	<b>451</b>
Involve the Marketing department	451
Don't boil the ocean	451
Speed to implementation	452
Build a focus group	452
<b>General knowledge and tools</b>	<b>453</b>

---

<b>Third-party tools</b>	<b>454</b>
Oracle JDeveloper skinning tool	454
BI Consulting Group Identity product	454
<b>Hands-on – go time!</b>	<b>455</b>
Overview	455
Let's talk about tools and assumptions	457
Locating existing styles and skins	457
<b>Developing/designing a new look and feel</b>	<b>458</b>
Getting analyticsRes deployed	459
Good artists copy, great artists steal	462
Modifying instanceconfig.xml	464
Modifying skins and styles	465
Changing the banner color	466
Changing the global header menu link's color	466
Changing the header separator bar	466
Changing the header brand name	467
Changing the Login page background	467
Restarting Presentation services	467
Restarting Presentation services from the command line	468
Restarting Presentation services from Enterprise Manager	469
Viewing the results	470
Configuring a custom message	470
Refreshing metadata files and custom messages	472
Wrap it up!	473
<b>Advanced learning topics</b>	<b>473</b>
Firebug	474
NetBeans IDE	474
Formatting CSS (prettify)	474
NQ_SESSION.SKIN and NQ_SESSION.STYLE session variables	474
AnalyticsRes Application Caveat	475
<b>A review – what I should now know!</b>	<b>475</b>
<b>Summary</b>	<b>475</b>
<b>Chapter 14: Improving the Performance</b>	<b>477</b>
<b>What is poor performance?</b>	<b>477</b>
Where can I improve the performance?	478
Hardware	479
Database	481
BI Server	482
More performance tips	484
The use of cache	484
Setting up the cache	485
Web servers on top	487
Domain setup	487

---

<b>A review – what I should now know!</b>	<b>487</b>
<b>Summary</b>	<b>488</b>
<b>Chapter 15: Using the BI Admin Change Management Utilities</b>	<b>489</b>
<b>Problems with multiple developers</b>	<b>489</b>
<b>Merges</b>	<b>490</b>
Three-way merge	490
Two-way merger	494
<b>Multi-User development</b>	<b>494</b>
Online development	495
Advantages and disadvantages	498
Multi-User development	498
Advantages and disadvantages	506
<b>A review – what I should now know!</b>	<b>507</b>
<b>Additional research suggestions</b>	<b>507</b>
<b>Summary</b>	<b>507</b>
<b>Chapter 16: Usage Tracking</b>	<b>509</b>
<b>What is usage tracking?</b>	<b>509</b>
<b>System setup</b>	<b>510</b>
Setting up the database table	510
Additional data	511
Setting up the BI Server repository	512
Updating the BI Server's configuration file	513
Analyzing the usage	516
Usage measures	516
<b>A review – what I should now know!</b>	<b>519</b>
<b>Summary</b>	<b>519</b>
<b>Chapter 17: Oracle Essbase and OLAP Integration</b>	<b>521</b>
<b>A bit about OLAP</b>	<b>521</b>
Competition	522
MOLAP, ROLAP, HOLAP, XOLAP	522
Essbase's entrenched past	523
<b>Oracle Essbase Studio</b>	<b>524</b>
<b>Oracle BI SampleApp v107+ – VM image</b>	<b>524</b>
<b>Getting started – let's get set up</b>	<b>525</b>
Prepping the VM image	525
Starting the virtual machine image	526
Starting up Essbase	526
Starting up Essbase Administration Services (EAS)	527
Prepping Essbase	528
Creating the base Essbase application and database	528
Migrating the Essbase files	531

---

Validating the Outline	531
Loading data into the cube	532
Anything needed to prep the Oracle BI Server?	533
<b>Modeling Essbase into Oracle BI</b>	<b>534</b>
A bit of Essbase to Oracle BI knowledge	534
Importing Essbase as a data source	535
A few OLAP adjustments before modeling	536
Flattening the Measure dimension	537
Getting the UDAs	538
Dimension and hierarchy types	538
Getting a quick win	540
Incremental importing of Essbase metadata	541
<b>Federation of data</b>	<b>543</b>
<b>Oracle BI/EPM roadmap</b>	<b>543</b>
Workspace integration	544
Software license combo	544
<b>A review – what I should now know!</b>	<b>544</b>
<b>Additional research suggestions</b>	<b>544</b>
<b>Summary</b>	<b>545</b>
<b>Appendix A: Programs and Definitions</b>	<b>547</b>
<b>Stress testing</b>	<b>548</b>
XML – a better approach	553
biserverxmlgen	553
biserverxmlexec	554
biserverxmlcli	555
Working example using XML	556
<b>More repository management</b>	<b>558</b>
<b>Admin tool</b>	<b>559</b>
<b>A review – what I should now know!</b>	<b>561</b>
<b>Summary</b>	<b>562</b>
<b>Appendix B: Useful Resources: Join the Oracle BI Movement</b>	<b>563</b>
<b>This book's resources</b>	<b>563</b>
OB11gBook.com forums	563
Author blogs	564
Other Oracle BI practitioner blogs	564
Oracle development team blogs	564
<b>Oracle BI user groups</b>	<b>565</b>
Oracle Development Tools Users' Group (ODTUG)	565
Oracle Applications Users' Group (OAUG)	565
Independent Oracle User Groups (IOUG)	566
International Oracle Users' Group Community (IOUC)	566

---



*Table of Contents*

---

<b>Conferences</b>	<b>566</b>
Oracle Open World (OOW)	566
COLLABORATE	567
KScope	567
RittmanMead BI Forum	567
<b>Join the movement</b>	<b>568</b>
<b>Further reading</b>	<b>568</b>
<b>Summary</b>	<b>569</b>
<b>Index</b>	<b>571</b>

---

# Preface

Oracle Business Intelligence Enterprise Edition (OBIEE) 11g is packed full of features and has a fresh approach to information presentation, system management and security. This book will introduce the reader to those features, providing a step-by-step guide to building a complete system from scratch. The aim of the book is to equip a developer or analyst with a good basic understanding of what the product contains, how to install and configure it, and how to create effective business intelligence.

## What this book covers

*Chapter 1, Understanding the Oracle BI 11g Architecture*, helps you in understanding the 11g architecture. As with any good software suite, a solid architectural foundation is required. In today's marketplace the ability for software to scale well, meet the growing needs of an enterprise, and integrate with an organization's existing Information Technology (IT) investments is expected. Having the software be transparent enough for the average IT professional to implement it is definitely a plus. Being simple enough for an end-user to use or consume the product is a must. Oracle Business Intelligence (Oracle BI) 11g fits within all of these paradigms.

*Chapter 2, Installing the Metadata Repository*, covers how to install the required database components for your Oracle BI system to use them. Before installing an OBIEE 11g system, you will need to prepare a database, not for end user reporting but for the OBIEE system itself. You will also learn how to customize the installation to change the options available, and how to use the silent installation. You cannot afford to skip this chapter if you're installing your own development system! It's crucial to understand what is possible so that you can advise the database administrators, who will be managing your production environments.

*Chapter 3, Installing on Windows Server 2008*, provides step-by-step instructions for installing Oracle BI 11g on Windows Server 2008. This installation will walk you through the Enterprise installation of Oracle BI 11g, which is one of the three possible installation options, and will give you the best of all worlds for an Oracle BI 11g platform implementation. It will allow you to work/play with all of the features seen in a production Oracle BI 11g environment.

*Chapter 4, Installation Options*, covers additional installation options. The installation conducted in *Chapter 3, Installing on Windows Server 2008*, is perfect for a sandbox or a development environment. There is one more installation option, Software Only Install, which is an advanced way to conduct the Enterprise installation option. Several advanced configuration options and many production environment considerations may be made using the Software Only Install option. Additional installation and environment configuration options are what we will cover in this chapter.

*Chapter 5, Understanding the Systems Management Tools*, goes into greater detail on the administration interfaces, explaining what these components are, what they do, and how they work together. We will dive into the navigation of these tools so that you will become more familiar with the interfaces. Finally, this chapter is crucial to the remainder of the book as it contains the security exercises for creating the users and groups that will be used to access the Tennis Repository's dashboards and reports, which you will develop in subsequent chapters.

*Chapter 6, Upgrading the RPD and Web Catalog to 11g*, looks at the upgrade process for the most fundamental parts of the system from a version 10g implementation to 11g. The upgrade process for this is extremely straightforward, as Oracle has provided an easy-to-use upgrade tool that we will step through in this chapter. If you are upgrading a current live implementation, then we must consider the wider implications of the upgrade, especially the possible effects on current functionality. Therefore, in this chapter we will also touch upon the thinking and planning that is needed prior to a full upgrade for a current live implementation.

*Chapter 7, Reporting Databases*, introduces the main concepts of a reporting database. The process of creating an efficient database is the subject of dozens of books and blogs, and therefore the details in this chapter should provide enough information to get you started in creating a database that is fit for using in an Oracle Business Intelligence system.

*Chapter 8, Developing a BI Repository*, covers the development of a simple RPD from tables in a database through to how those objects are presented to us when we move on to create an actual request. This will be carried out via the Oracle BI Administration tool, which is the primary method of accessing and modifying an RPD file. This tool provides an inviting graphical interface for developing and administering an RPD file. By the end of this chapter, you will be able to complete the major tasks associated with RPD development. We will also describe the more advanced options that are available.

*Chapter 9, Features of the Presentation Catalog*, introduces the new interface of the web catalog and the tools that are integrated into the Presentation Services. It also explores aspects of search, catalog administration, privileges, object security, and more.

*Chapter 10, Creating Dashboards and Analysis*, teaches you how to create reports and group them in dashboards. In this chapter, we will also look at the various ways of representing and formatting data that are available, along with advice on best practices gained from implementation experience. This is one of the more robust chapters in the book and provides an insightful look at dashboard and report development.

*Chapter 11, Agents and the Action Framework*, covers the new functionality that Oracle has introduced in 11g for more integration with business processes, and the actions that result from producing the analytical reports. In previous versions, we had the Delivers portion of OBIEE where you could invoke basic actions, such as the delivery of reports or dashboard alerts. 11g has drastically enhanced this capability through the **Action Framework**, through which we are now able to initiate a multitude of additional noncore actions. During this chapter, we will look at examples of the new actions that Oracle has provided in attempting to succeed in this goal.

*Chapter 12, Developing Reports Using BI Publisher*, covers some of the new features of BI Publisher 11g as well as the general functionality of BI Publisher in order to get you up-to-speed on using the tool. It is aimed at providing a crash-course that should give any reader enough hands-on exercises to get their feet wet and enough food for thought for further research.

*Chapter 13, Customizing the Style of Dashboards*, provides a step-by-step how-to guide for branding OBIEE dashboards to match your corporate look and feel. It also provides several other insights for continued development and research.

*Chapter 14, Improving the Performance*, explores some common techniques to reduce the bottlenecks that can exist in the process of delivering dashboards and reports to the users. We will look across the whole system, defining poor performance and where required, take steps to improve the performance.

*Chapter 15, Using the BI Admin Change Management Utilities*, builds on the fundamental techniques of OBIEE development learned in previous chapters explores other capabilities in the Administration tool. In larger projects we may have a group of developers accessing and modifying the same RPD. The OBIEE Administration tool provides the ability to merge multiple versions of an RPD as well as functionality for groups to manage development on a sole repository (multiuser development). In this chapter, we will go over various methods of dealing with multiuser development.

*Chapter 16, Usage Tracking*, will demonstrate how to activate the "usage tracking" feature, and create useful reports from it. One of the great features of Oracle BI is that you can use the system—Dashboards and Analysis—to monitor the system itself, which means to say that you can use an OBIEE Analysis that tells you how OBIEE is performing for your users!

*Chapter 17, Oracle Essbase and OLAP Integration*, shows how OLAP technologies integrate into Oracle BI. After exploring the options, we will then focus on Oracle Essbase as the preferred OLAP technology. This chapter will show you how to integrate Oracle Essbase as a data source in Oracle BI and define several best practices for the integration. At the end of the chapter, you should have a well-balanced sense of how Essbase integrates with Oracle BI and the added value that it can bring to an organization.

*Appendix A, Programs and Definitions*, describes the main OBIEE command utilities and provides some reusable examples. One of the powerful features of OBIEE 11g administration is its ability to be controlled by User Interfaces (web browsers and Admin tools) as well as by command-line utilities. Many of the manual tasks that you undertake each day can be scripted and therefore automated. In this chapter you will learn about those automations.

*Appendix B, Useful Resources: Join the Oracle BI Movement*, lists some of the best books, events, groups, blogs for further reading and further practice on OBIEE 11g. Over the last decade, the number of resources focusing on Oracle Business Intelligence has skyrocketed. Oracle's documentation of the software has become increasingly more useful and user-friendly. Blogs all over the globe have popped-up in large numbers (though some better maintained and better written than others). Use this chapter to find where to learn more about OBIEE.

## What you need for this book

The book assumes that you have no prior knowledge of Oracle Business Intelligence 11g. A general IT understanding will assist but is not a pre-requisite. To follow along with the most technical material in this book, you will need a workstation or server running Microsoft Windows 7, Microsoft Windows Server 2003, or Microsoft Windows Server 2008 along with the ability to download the Oracle BI and Oracle Database installation files over the internet. If you currently have a working Oracle BI 11g system to which you have access you will also be equipped for the material in this book.

To get the most out of this book, it is also advisable that you have a basic grounding in traditional data warehouse design and business intelligence concepts. For example, you are familiar with facts, dimensions, and star schemas.

## Who this book is for

This book is for IT professionals, business analysts, project managers, and/or newcomers to Business Intelligence who wish to learn from self-paced professional guidance and actual implementation experience. Ultimately this book is for anyone who needs a solid grounding in the subject of Oracle Business Intelligence.

## Conventions

In this book, you will find a number of styles of text that distinguish between different kinds of information. Here are some examples of these styles, and an explanation of their meaning.

Code words in text are shown as follows: "The following is a list of configuration files based on the central Oracle BI 11g instance path of <FMW\_HOME>\instance\instance1\config\."


A block of code is set as follows:


```
ALTER database [dbname_here] SET SINGLE_USER WITH ROLLBACK IMMEDIATE;  
ALTER database [dbname_here] SET READ_COMMITTED_SNAPSHOT ON;  
ALTER database [dbname_here] SET MULTI_USER;
```

Any command-line input or output is written as follows:

```
cd C:\Windows\System32\inetsrv
```

**New terms** and **important words** are shown in bold. Words that you see on the screen, in menus or dialog boxes for example, appear in the text like this: "From the Oracle Business Intelligence (11.1.1.x) Downloads page, scroll down to the **Required Additional Software** section and click on the link for your operating system."

 Warnings or important notes appear in a box like this.

 Tips and tricks appear like this.

## Reader feedback

Feedback from our readers is always welcome. Let us know what you think about this book – what you liked or may have disliked. Reader feedback is important for us to develop titles that you really get the most out of.

To send us general feedback, simply send an e-mail to [feedback@packtpub.com](mailto:feedback@packtpub.com), and mention the book title through the subject of your message.

If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, see our author guide on [www.packtpub.com/authors](http://www.packtpub.com/authors).

## Customer support

Now that you are the proud owner of a Packt book, we have a number of things to help you to get the most from your purchase.

## Downloading the example code

You can download the example code files for all Packt books you have purchased from your account at <http://www.packtpub.com>. If you purchased this book elsewhere, you can visit <http://www.packtpub.com/support> and register to have the files e-mailed directly to you.

## Errata

Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you find a mistake in one of our books – maybe a mistake in the text or the code – we would be grateful if you would report this to us. By doing so, you can save other readers from frustration and help us improve subsequent versions of this book. If you find any errata, please report them by visiting <http://www.packtpub.com/support>, selecting your book, clicking on the **errata submission form** link, and entering the details of your errata. Once your errata are verified, your submission will be accepted and the errata will be uploaded to our website, or added to any list of existing errata, under the Errata section of that title.

## Piracy

Piracy of copyright material on the Internet is an ongoing problem across all media. At Packt, we take the protection of our copyright and licenses very seriously. If you come across any illegal copies of our works, in any form, on the Internet, please provide us with the location address or website name immediately so that we can pursue a remedy.

Please contact us at [copyright@packtpub.com](mailto:copyright@packtpub.com) with a link to the suspected pirated material.

We appreciate your help in protecting our authors, and our ability to bring you valuable content.

## Questions

You can contact us at [questions@packtpub.com](mailto:questions@packtpub.com) if you are having a problem with any aspect of the book, and we will do our best to address it.





# 1

## Understanding the Oracle BI 11g Architecture

As with good software suite, a solid architectural foundation is required. In today's marketplace the ability for software to scale well, meet the growing needs of an enterprise, and integrate with an organization's existing **Information Technology (IT)** investments is expected. Having the software be transparent enough for the average IT professional to implement it is definitely a plus. Being simple enough for an end-user to use or consume the product is a must. **Oracle Business Intelligence (Oracle BI)** 11g fits within all of these paradigms.

This chapter provides an overview of the Oracle BI architecture and its place in the Oracle **Fusion Middleware (FMW)** stack.

### Looking backward and looking forward

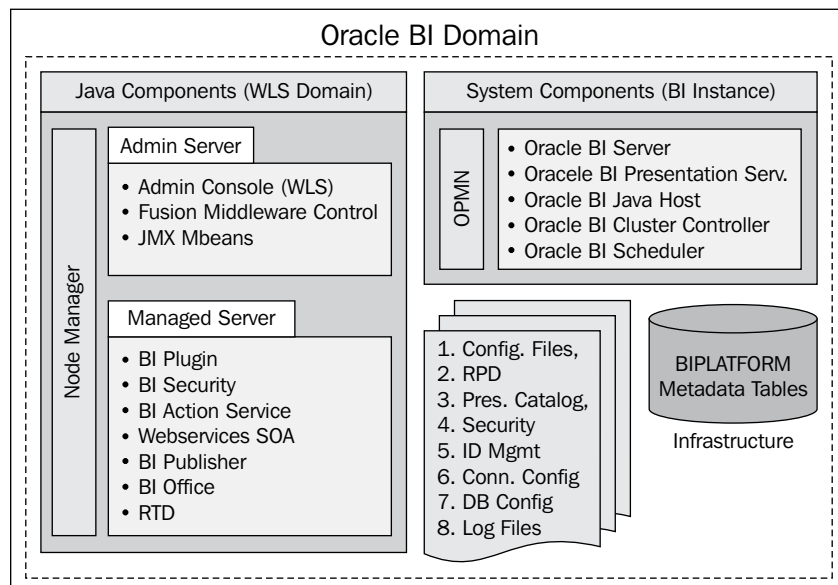
Oracle BI has a history forged by acquisition and brilliant advances in both technology and market share (refer to the document called **Magic Quadrants for Business Intelligence Platforms** on Gartner's site at <http://www.gartner.com/id=1531017>). Without going into much historical detail about the early beginnings of the software suite, as a reader, you are either new to Oracle BI or have experienced the tool in one of its former versions. Those versions can be Siebel Analytics or Oracle BI 10g.

The old adage of knowing where you've been to know where you are going doesn't much apply to understanding most of Oracle BI 11g. From architecture of the **Graphical User Interface (GUI)** the majority of the tool has been rewritten. Although, the principal of the tool remains the same; to provide an end-to-end enterprise analytics solution; the core components of Oracle BI have changed substantially enough for it to feel like a brand new tool to those familiar with its previous versions. With that being said, let's talk about Oracle BI 11g's architecture.

## Let's look at the big picture

The following illustration shows some of Oracle BI 11g's infrastructure components from a core architecture perspective:


- **Oracle BI Domain:** It is the core architecture of Oracle BI 11g.
- **WebLogic Server:** It is the chosen application server for Oracle BI 11g.
- **Java components:** These are the components which have been written in Java for Oracle BI 11g. They are deployed to the application server and WebLogic Server.
- **System components:** These are the components which have been written mainly in C++ for Oracle BI 11g. They are managed by the Oracle Process Management and Notification Server.
- **Oracle BI relational repository:** It is a set of database schemas (BIPLATFORM and MDS) that store metadata related to a specific Oracle BI 11g instance.
- **Oracle BI filesystem:** It is the instructional set of physical files and directories containing configuration, logs, and metadata concerning the Oracle BI 11g instance:



Once Oracle BI 11g is installed and configured, the architecture seen in the preceding illustration will exist. The components, pointed out in the preceding illustration, will be the areas where most of the backend, or day-to-day platform maintenance, and troubleshooting takes place. Some IT resources may already have insight into maintaining some of these components if they have experience in working with other Oracle products in the FMW stacks such as WebLogic, WebCenter, SOA, and so on. However, to most Oracle BI aficionados, this environment will be new.

## What is Oracle Fusion Middleware?

Oracle Fusion Middleware is taking on the enterprise challenge of bringing together the Oracle database and Oracle applications stacks. It is the middle-tier between them. **Fusion Middleware** is Oracle's go-forward foundation for the fusion of products between the database and application stack. Oracle has acquired many companies over the last decade for their technology or market share. This has taken them to a position of having excellent software. Oracle didn't communicate effectively before that. To achieve interoperability, a common layer had to be formed to fuse together the existing technologies, create efficiencies, and provide consistent delivery of software applications. The following image illustrates the main categories of products, making up the current Oracle product stack:



The Fusion Middleware product category contains Oracle **Fusion Middleware (FMW)**, which forms the core of Oracle's **Application Integration Architecture (AIA)**. It is the foundation for Oracle's fusion applications and software suites, such as Oracle BI 11g, Oracle Hyperion EPM, and so on.

## An application server by any other name

In previous versions of Oracle BI, the default application server delivered with the product suite was **Oracle Container for Java (OC4J)**. This Java application server was actually a slimmed-down version of the better-known **Oracle Application Server (OAS)**. However, with the release of Oracle BI 11g—as with the mass majority of the Oracle Fusion Middleware applications—**Oracle WebLogic Server (WLS)** became the core application server. Previously, it was known as BEA WebLogic.

The WLS is a robust and scalable Java application server and it has been ranked as one of the top application servers in the market. Oracle has made a strong investment in WebLogic Server's atlas-like position as the foundation to which its Fusion Middleware stack is raised. With the current Oracle BI 11.1.1.6 release, the consideration to leverage IBM WebSphere as an alternate application server is under review. However, no guidance has been provided by Oracle on this topic. For prior releases of Oracle BI 11g, no other Java application server has been certified. So, an application server by any other name just won't do.

## A database repository – for what?

Oracle BI 11g is a system that has evolved—and continues to evolve—based on expansive user requests, a market that dictates stronger integration points, and more powerful BI tools. As such, Oracle BI 11g now incorporates, or better yet, requires a relational database repository to hold metadata concerning the installation, report scheduling, usage tracking, auditing, and other aspects of the environment.

Actually, the Oracle BI 11g installation process cannot begin until these repositories are created by the **Repository Creation Utility (RCU)** and accessible on a database server. *Chapter 2, Installing the Metadata Repository*, goes into greater detail about this crucial repository structure, better known by two database schemas—**Metadata Services (MDS)** and **BIPLATFORM**. The installation and configuration of these two repositories are required primarily for integration of Oracle BI 11g with the Oracle Fusion Middleware stack.



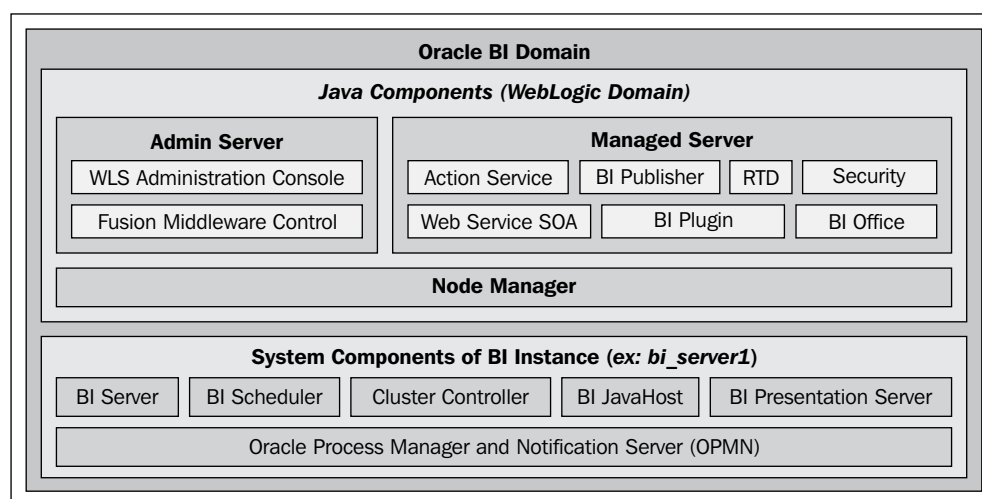
One interesting fact to note, however, is that the Oracle metadata repository (RPD) is still file based.

The **RPD** is the metadata storage mechanism in which Oracle BI developers model and map physical data sources to logical business representations in order for the resulting analytics to be easily consumed by the end users.

## Overall components

There are a lot of components that comprise Oracle BI 11g. Oracle BI 11g can categorize these components by classifying them based on the programming languages in which they were developed. The programming languages are mainly either Java or C++. In comparison with previous versions of Oracle BI, where it seemed to be a somewhat compact environment, Oracle BI 11g is much more integrated into the larger platform of Fusion Middleware, which adds both value and complexity. Some of the similar legacy components' names persist. However, it should not be taken lightly because most of the legacy components have been renamed, removed, or placed under new management processes.

The following diagram provides a high-level overview of the main components that comprise the Oracle BI 11g architecture. The components are clearly segmented by the processes in which they are managed, each of which ultimately comprises the Oracle BI Domain:



The term **Oracle BI Domain**, as noted previously and shown in the illustration, is used as a way to group all Oracle BI 11g components within the Fusion Middleware architecture. This should not be confused with the WebLogic Application Server domain which is given the default name, `bifoundation_domain`, while following the default Oracle BI 11g installation options. The latter is a WebLogic Java Application Server term. The former is a Fusion Middleware term. Since Fusion Middleware is so closely related with the WebLogic Application Server, it is good to keep it in mind from a technical perspective.



When learning about the overall component composition, a very important detail to keep in mind is the matter in which the components are managed. Since WebLogic Server is a Java application server, it manages all of the Oracle BI components developed in the Java programming language. Another management system, the **Oracle Process Management and Notification (OPMN)** system, handles the other components, which are referred to as the System Components.

## Java components

In relation to Oracle BI 11g, the Java components are those which have been developed in the Java programming language. Those components are as follows:

- **Action service:** Primarily used by the Action Framework it executes actions on behalf of Presentation Services and Oracle BI Scheduler. Actions may be invocations of third party web services, or invocations of user supplied Java code executed as **Enterprise JavaBeans (EJB)**.
- **Administrative components: Java Management Extensions (JMX)** and **Managed Beans (MBean)** allow dynamic **Application Programming Interface (API)** functionality for managing, configuring, and administering Oracle BI 11g.
- **Web Service SOA:** This provide a web service interface to the contents of the Oracle BI Presentation Catalog. The tree of objects in Oracle BI Presentation Catalog is exposed as a tree of web services, defined by a **Web Services Inspection Language (WSIL)** tree with **Web Service Definition Language (WSDL)** leaves.
- **Oracle BI Office:** It provides integration between Microsoft Office and Oracle BI 11g.
- **Oracle Real-time Decisions (RTD):** RTD provides a decision making rules engine that enables real-time business intelligence predictions and outcome analysis.
- **Oracle BI Presentation Service plugin:** Presentation Services run as a process, not as a web server, and does not communicate using any web server plugin API. The Oracle BI Presentation Services plugin forwards HTTP requests to Presentation Services. The HTTP requests are the requests from the browser-based user interface, or SOAP requests. This is ultimately just a servlet.
- **Oracle BI Publisher:** It is the enterprise reporting solution used for authoring and delivering highly formatted documents.

- **Security Services:** It provides standards-based authentication and population services. It enables the Oracle BI server to integrate with the Fusion Middleware security platform which includes the Credential Store Framework and the Identity Store.

## System components

In relation to Oracle BI 11g, the system components are those which are developed in a non-Java programming language. Most of the system components have been developed in the C++ programming language as mentioned in the previous section. Here is a list of those components:

- **Oracle BI Server:** This is a C++ process that performs the data manipulation and aggregates data from data sources. You can configure multiple Oracle BI Server processes, which share the load. No session replication takes place between the Oracle BI Server processes.

The Oracle BI Server does not maintain a user session state. For high availability deployments, query results are cached in the global cache.

- **Oracle BI Presentation Server:** This is a C++ process that generates the user interface pages and renders results sets on behalf of the Oracle BI Scheduler. You can configure multiple Presentation Services, which share the load. No session replication takes place between the Presentation Services.

Presentation Services are almost stateless. The only significant state is the client authentication. If Oracle BI is configured to use a single sign-on for authentication purposes, users do not have to reauthenticate after a failover. For all other authentication schemas, when failover occurs, clients will have to reauthenticate. The client will see an interruption of service and will be redirected to a login page.



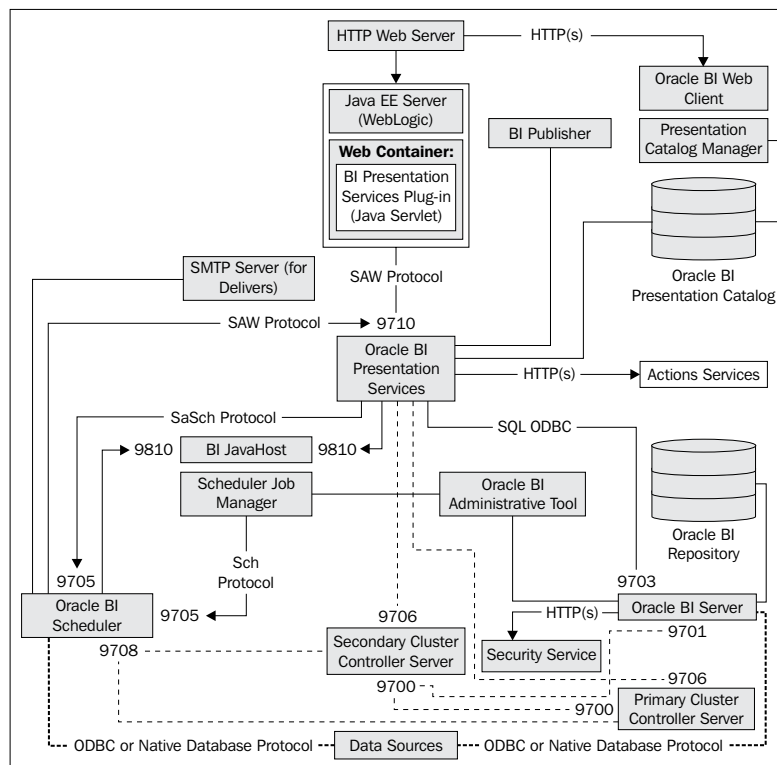
Oracle BI has some very basic capabilities, which can store the state or current session activity for a user logged in to the system. This is referred to as **Session Management**. Since the BI Presentation Server maintains the authentication state, the users do not have to log in on each subsequent dashboard that they visit. If a system is stateless, the application cannot easily remember the information about the user or actions performed previously.

- **Oracle BI Scheduler:** It is a C++ process that runs the jobs according to a configurable frequency. Jobs can be created by agents in Oracle BI Presentation Catalog, or jobs can be created by the Job Manager. The Oracle BI Scheduler differs from the Quartz Scheduler that the Oracle BI Publisher leverages.



- **Oracle BI JavaHost:** It is a Java process that includes resource-intensive graphs and PDF rendering. It also allows Oracle BI Presentation Services to support Oracle BI Publisher and Java tasks within the Oracle BI Scheduler. You can configure multiple JavaHost processes, which share the load. No session replication takes place between the JavaHost processes. JavaHost is a stateless process.
- **Oracle BI Server Cluster Controller:** It is a C++ process, which manages the population of Oracle BI Servers and Oracle BI Schedulers. It also distributes the requests to the Oracle BI Server and ensures that requests are evenly load balanced across scaled-out Oracle BI Servers in the domain.

In general, it is important to understand how all of the components interact within the Oracle BI environment. Understanding some general concepts, such as "which port numbers are defined to communicate within the default Oracle BI architecture", "how the Oracle BI Administration Tool communicates with the Oracle BI database repository" will be quite helpful in your journey for becoming an Oracle BI professional. The following illustration shows each of the components comprising the core Oracle BI architecture, the communication ports, and the communication direction:



## WebLogic Server

Let's talk a little more in detail about the enterprise application server that is at the core of Oracle Fusion Middleware—WebLogic. The Oracle WebLogic Server is a scalable, enterprise-ready **Java Platform Enterprise Edition (Java EE)** application server. Its infrastructure supports the deployment of many types of distributed applications. It is also an ideal foundation for building service-oriented applications.

The WebLogic Server is a robust application in itself. In previous versions of Oracle BI, the Oracle BI administrator and other developers took less effort to modify, configure, or otherwise maintain the Java application server. In Oracle BI 11g, the WebLogic Server is more crucial to the overall implementation, not only for installation but also throughout the Oracle BI 11g life-cycle. Learning the management components of the WebLogic Server, which ultimately controls the Oracle BI components, is critical for the success of an implementation. These management areas within the WebLogic Server are referred to as the **WebLogic Administration Server**, **WebLogic Manager Server(Servers)**, and **WebLogic Node Manager**.

## A few software nuances

Before we move on to the description of each of the management areas within WebLogic, it is important to understand that the WebLogic Server software, which is used for the installation of the Oracle BI product suite, carries a limited license. Although the software itself is the full enterprise version—logically containing 100 percent of the product's functionality—the license that ships with Oracle BI 11g is not a full enterprise license for the WebLogic Server. This book hardly deals with software licensing, but here are a few nuances that one should keep in mind as they go about an Oracle BI 11g implementation:

- The WebLogic Server license, which is provided with Oracle BI 11g, does not grant horizontal scale-out. An enterprise WebLogic Server license needs to be obtained for this advanced functionality.
- The WebLogic Server does not provide a separate HTTP server with the installation. The Oracle BI Enterprise Deployment Guide ([http://docs.oracle.com/cd/E21764\\_01/doc.1111/e15722/toc.htm](http://docs.oracle.com/cd/E21764_01/doc.1111/e15722/toc.htm)) discusses the separation of the application tier from the web/HTTP tier and suggests **Oracle HTTP Server (OHS)**. OHS is part of the Oracle FMW web tier and must be downloaded separately (<http://www.oracle.com/technetwork/java/webtier/downloads/index.html>). The other web/HTTP servers that can be used are Apache and Windows IIS, which we will discuss in the *Chapter 4, Installation Options*.

These items are simply a few nuances of the product suite in relation to Oracle BI 11g. Most software products have a very short list of nuances like the preceding one. However, once you understand the nuances, it will be easier to ensure that you have a more successful implementation. It also allows your team to be prepared for implementations. Be sure to consult your Oracle sales representative to assist you with the licensing concerns.



In order to learn more about the installation features, configuration options, administration, and maintenance of the WebLogic Server, we recommend that you reference the documentation of the WebLogic Server itself and not just the material on how it relates to Oracle BI 11g. The core of the WebLogic Server doesn't change just because Oracle BI 11g integrates into it. Understanding this approach should provide you with more efficient results.

## WebLogic Domain

The highest unit of management for controlling the WebLogic Server installation is called a **domain**. A domain is a logically related group of WebLogic Server resources that you manage as a unit. A domain always includes, and is centrally managed by, one Administration Server. Additional WebLogic Server instances which are controlled by the Administration Server for the domain are called **Managed Servers**. The configuration for all the servers in the domain is stored in the configuration repository, the `config.xml` file, which resides on the machine hosting the Administration Server. The `config.xml` files, by default, are stored in the path `<FMW_HOME>\user_projects\domains\bifoundation_domain\` where `<FMW_HOME>` is the path on the server to which you have installed Oracle BI 11g.

Upon installing and configuring Oracle BI 11g, the domain named `bifoundation_domain` is established within the WebLogic Server. This domain is the recommended name for each Oracle BI 11g implementation and should not be modified.

## WebLogic Administration Server

The WebLogic Server is an enterprise software suite that manages a myriad of application server components mainly focused on Java technology. It is also comprised of many ancillary components that enable the software to scale well, and also make it a good choice for distributed environments and high availability.

Clearly, it is good enough to be at the core of Oracle Fusion Middleware. One of the most crucial components of the WebLogic Server is the WebLogic Administration Server. When installing the WebLogic Server software, the WebLogic Administration Server is automatically installed with it. It is the Administration Server that not only controls all subsequent WebLogic Server instances called Managed Servers, but also controls aspects such as security, Persistence Stores, and other application server-related configurations.


The WebLogic Server gets installed on the operating system and ultimately runs as a service on that machine. The WebLogic Server can be managed in several ways. The two main methods are via the **Graphical User Interface (GUI)** web application called the WebLogic Administration Console or via the command line using the **WebLogic Scripting Tool (WLST)**. You can access the WebLogic Administration Console from any machine using a web-based client (that is, web browser) that can communicate with the WebLogic Administration Server through the network and/or firewall.

The WebLogic Administration Server and the WebLogic Server are basically synonymous. If the WebLogic Server is not running, the WebLogic Administration Console will be unavailable as well. If the WebLogic AdminServer is not running, no administrative tasks can be made to the system, although concessions are made for a High Availability configuration.


## WebLogic Managed Server

Web applications, **Enterprise Java Beans (EJB)**, and other resources are deployed on to one or more WebLogic Managed Servers in a WebLogic Domain. A WebLogic Managed Server is an instance of a WebLogic Server in a WebLogic Server domain. Each WebLogic Server domain has at least one instance which acts as the WebLogic Administration Server which we have just discussed. Only one Administration Server per domain must exist, but one or more Managed Servers may exist in the WebLogic Server domain. Having one or more managed servers, allow for deployed JEE applications to be logically delineated. They also provide a means to independently configure application server port numbers and they provide a barrier for runtime issues such as a server crash. You can deploy applications, EJBs, and other resources on the WebLogic Managed Servers and use the WebLogic Administration Server only for configuration and management purposes.

In a production deployment, Oracle BI 11g is deployed into its own Managed Server. The Oracle BI 11g installer comes with three installation types – simple, enterprise, and software. The latter two installation types configure two WebLogic Server instances, the Administration Server and another Managed Server called `bi_server1`. Oracle BI 11g is deployed into the Managed Server called `bi_server1` and is configured by default to resolve to port 9704. The simple installation type configures only the administration server, deploys Oracle BI 11g into it, and resolves to port 7001. For the simple installation type, only one WebLogic Server instance exists.

 The simple installation type is not recommended for anything more than sandbox, test environment development, or demonstrations.

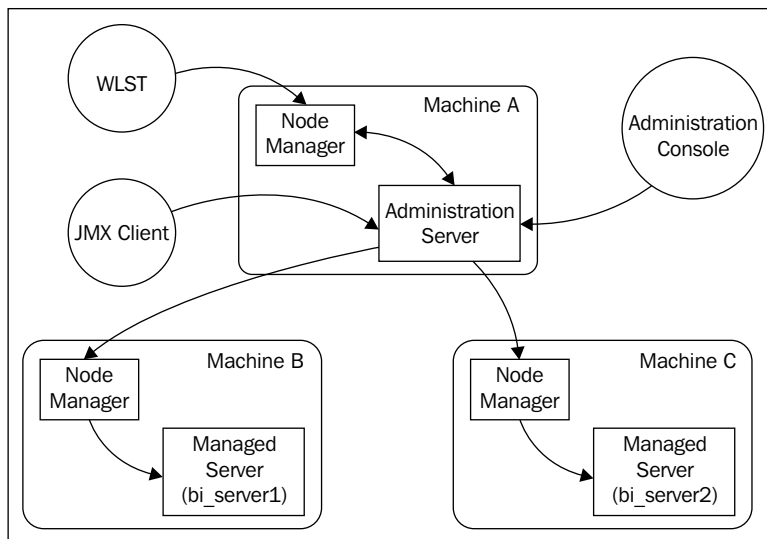
When administering the WebLogic Server via the Administration Console, the WebLogic Administration Server instance appears in the same list of servers that also includes the Managed Servers. The WebLogic Administration Server should be used only for configuration and management of the WebLogic Server and should not contain any additionally deployed applications, EJBs, and so on.

 One thing to note is that the Enterprise Manager Fusion Control is actually a JEE application deployed to the Administration Server instance, which is why its web client is accessible under port 7001. It is not necessarily a native application deployment to the core WebLogic Server, but gets deployed and configured during the Oracle BI 11g configuration. In the deployment's page within the Administration Console, you will find a deployment named `em`.

## WebLogic Node Manager


The general idea behind the Node Manager is that it takes on somewhat of a middle-man role. That is to say, the Node Manager provides a communication tunnel between the WebLogic Administration Server and any WebLogic Managed Servers configured within the WebLogic Domain. When the WebLogic Server environment is contained on a single physical server, it may be difficult to recognize the need for a Node Manager. However, its real power comes into play when Oracle BI 11g is scaled out horizontally on one or more physical servers. Each scaled-out deployment of WebLogic Server will contain a Node Manager. If the Node Manager is not running on the server on which the Managed Server is deployed, then the core Administration Server will not be able to issue start or stop commands to that server. As such, if the Node Manager is down, communication with the overall cluster will be affected.

The following diagram shows how machines A, B, and C are physically separated – each of them contains a Node Manager. You can see that the Administration Server communicates with the Node Managers and not directly to the Managed Servers:



## System tools controlled by WebLogic

We have briefly discussed the WebLogic Administration Console which controls the administrative configuration of the WebLogic Server Domain. This includes the components managed within it such as security, deployed applications, and so on. The other management tool which provides control on the deployed Oracle BI application ancillary deployments, libraries, and several other configurations is called the **Enterprise Manager Fusion Middleware Control**.





This seems to be a long name for a single web-based tool. As such, the name is often shortened to **Fusion Control** or **Enterprise Manager**. Reference to either abbreviated title in the context of Oracle BI should ensure fellow Oracle BI teammates understand what you mean.

To discuss the vast amount of individual configuration points contained within the WebLogic Administration Console and Fusion Control, you can warrant an entire chapter devoted to this subject. In fact, a subsequent chapter, *Chapter 5, Understanding the Systems Management Tools* is dedicated to it.

## Oracle Process Management and Notification system

The **Oracle Process Management and Notification (OPMN)** system is not a new concept or tool within the Oracle product line. It has acted as a management service in many other Oracle products for years and is now a fitting utility for the cross-platform deployment architecture that Oracle BI 11g allows. The OPMN controls the Oracle BI 11g system components. Those are the components primarily developed in the C++ programming language. The OPMN not only allows each of the five system components to be started and stopped by calling a single command, it will also monitor those system component processes at runtime. It can even attempt to restart a component if it detects a failure. To start and stop the system components, you can use the Oracle Enterprise Manager Fusion Middleware Control or a command-line interface. The following screenshots represent the status retrieval of the System Components using the Fusion Middleware Control GUI and command-line interface, respectively:



The command-line executable is deployed with the Oracle BI 11g installation. It resides in the default filesystem directory locations – /u01/FMW/Instances/instance1/bin/opmnctl or C:\oracle\fmw\instances\instance1\opmnctl.bat.

The Enterprise Manager Fusion Middleware Control, as per the default Oracle BI 11g installation, is accessible via the URL – `http://<bi_server_name>:7001/em`.

Either approach used for stopping, starting, or checking the status of the system components is valid. Either approach will work properly, however it is recommended – whenever possible – to leverage the Fusion Control interface for these actions in order to achieve the most consistent results.

## Security

It would be difficult to discuss the overall architecture of Oracle BI 11g without at least giving some mention to how the basics of security, authentication, and authorization are applied. By default, installing the Oracle WebLogic Server provides a default **Lightweight Directory Access Protocol (LDAP)** server referred to as the **WebLogic Server Embedded LDAP server**.

This is a standards-compliant LDAP system which acts as the default authentication method for out-of-the-box Oracle BI 11g. Integration of secondary LDAP providers, such as **Oracle Internet Directory (OID)** or **Microsoft Active Directory (MSAD)** is crucial in order to leverage most organizations' identity management systems.

The combination of multiple authentication providers is possible, it is in fact common. For example, a configuration may wish to have users that exist in both the Embedded LDAP server and MSAD to authenticate and have access to Oracle BI 11g.

Oracle BI 11g security incorporates the Fusion Middleware Security model – **Oracle Platform Security Services (OPSS)**. This is a positive influence over managing all aspects of Oracle BI 11g as it provides a very granular level of authorization and a large number of authentication and authorization integration mechanisms. OPSS also introduces, to Oracle BI 11g, the concept of managing privileges by application role instead of directly by user or group. It abides by open standards to integrate with security mechanisms that are growing in popularity. Those security mechanisms are **Security Assertion Markup Language (SAML)** 2.0 and so on. Other well known single sign-on mechanisms such as SiteMinder and Oracle Access Manager have already preconfigured integration points within the Oracle BI 11g Fusion Control.



Chapter 5, *Understanding the Systems Management Tools* will go into an exercise for creating new users, groups, and assigning application roles, but for now, here are a few key concepts to know about security:

- Oracle BI 11g security is managed completely in a different way from the previous versions, although the Oracle BI 10g security model is still allowed for backwards compatibility.
- An Oracle BI 11g best practice is to manage security by Application Roles.
- Understanding the differences between the Identity Store, Credential Store, and Policy Store is critical for advanced security configuration and maintenance.

The following sections discuss these few key concepts at a high level. Understanding these concepts is not critical at this moment for you to continue with the remainder of the book, however, once you complete the book and are ready to engage in more advanced discovery, you will need to research and understand these items to be more versed in managing Oracle BI 11g security.

## Backwards compatibility

Let's first disclaim that the best practice for Oracle BI 11g security is to use the default Oracle Fusion Middleware security model. However, the legacy approach to manage security via the Oracle BI metadata repository (RPD) is still allowed. This backwards compatibility for security allows environments running on previous versions of Oracle BI (for example, Oracle BI 10g) to leverage the investments in complicated metadata repository architectures or unique identity solutions while still taking advantage of Oracle BI 11g's new functions and features.

This backwards compatibility can potentially provide a false sense of comfort. However, since the official Oracle BI roadmap is only to support the Fusion Middleware security model, backwards compatibility merely bridges a gap in migrating to Oracle BI 11g where rearchitecture of security would otherwise delay, or prevent, some organizations from taking advantage of the new offering. There are certain requirements to leverage this backwards compatible security architecture. The Oracle BI 11g product documentation discusses the right way to incorporate this technique in a section called **Alternative Security Administration Options** ([http://docs.oracle.com/cd/E21764\\_01/bi.1111/e10543/legacy.htm#CHDCEFBC](http://docs.oracle.com/cd/E21764_01/bi.1111/e10543/legacy.htm#CHDCEFBC)).

## Managing by application roles

In previous releases of Oracle BI security, groups and the relationship of a managed Identity Store's (that is LDAP or custom relational table) users with an Oracle BI group was managed within the RPD. A group was the highest level of organization for specific sets of users. This goes for both the metadata repository and the Web Catalog. This legacy approach was limited to a single software solution, Oracle BI.

In Oracle BI 11g, the default security model is the Oracle Fusion Middleware security model which has a much broader vision and scope. General Information Technology security administration best practice is to set permissions or privileges to a specific point of access on a group and not individual users. The same idea applies here except there is another enterprise level of user/group aggregation called an **Application Role**. Application roles can contain other application roles, groups, or individual users. Access privileges to a certain object such as a folder, web page, or column should always be assigned to an application role. Application roles for Oracle BI 11g can be managed in the Oracle Enterprise Manager Fusion Middleware Control interface.

## Security providers

Fusion Middleware security can seem complex at first, but knowing the correct terminology and understanding how the most important components communicate with each other it becomes easier. The application at large is extremely important as it relates to security management. Oracle BI 11g uses three main repositories for accessing authentication and authorization information, all of which are explained in the next sections.

### Identity Store

This is the authentication provider. A simple mnemonic here is that this store tells Oracle BI how to identify any users attempting to access the system. An example of creating an Identity Store would be to configure an LDAP system such as the Oracle Internet Directory or Microsoft Active Directory to reference users within an organization.

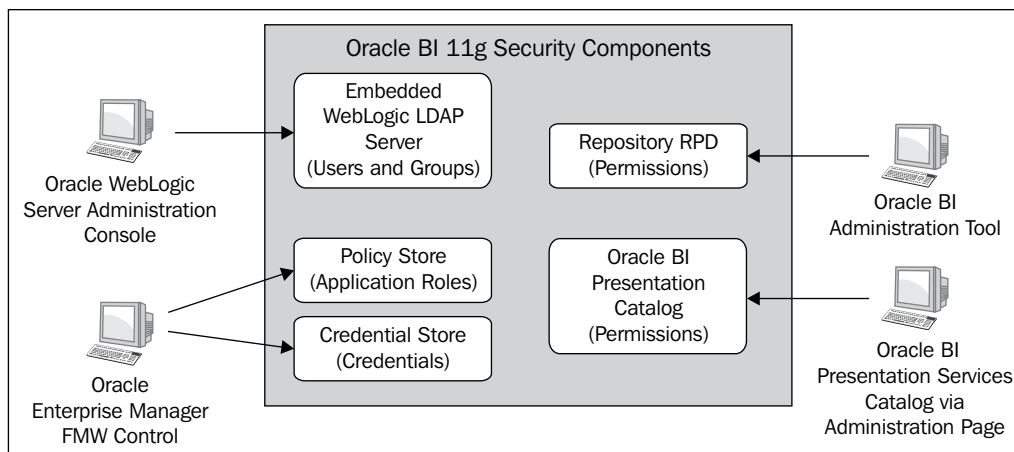
### Credential Store

The credential store is ultimately for advanced Oracle configurations. You may touch upon this while establishing an enterprise Oracle BI 11g deployment, but not much thereafter unless integrating with the Oracle BI 11g Action Framework, or something equally complex. Ultimately, the Credential Store does exactly what its name implies – it stores credentials.

Specifically, it is used to store credentials of other applications, which the core application (that is, Oracle BI) may access later without having to re-enter said credentials. An example of this would be integrating Oracle BI 11g with the **Oracle Enterprise Management (EPM)** suite. In this example, let's pretend that there is an internal requirement at Company XYZ for users to access an Oracle BI dashboard. Upon viewing said dashboard if a report with discrepancies is viewed, the user requires the ability to click on a link which opens an Oracle EPM Financial Report containing more details about the concern. If all users accessing the Oracle BI dashboard do not have credentials to access the Oracle EPM environment directly, how could they open and view the report without being prompted for credentials? The answer would be that the credential store would be configured with the credentials of a central user having access to the Oracle EPM environment. This central user's credentials (encrypted, of course) are passed along with the dashboard viewer's request and hey presto, access!

## Policy Store

The policy store is unique to Fusion Middleware security and leverages a security standard referred to as **eXtensible Access Control Markup Language (XACML)**, which ultimately provides granular access and privilege control for an enterprise application. This is one of the reasons why managing by application roles becomes so important. It is the individual application role to which assigned policies are defining access for the information within Oracle BI. Stated another way, the application privileges, such as the ability to administer the Oracle BI 11g RPD, are assigned to a particular application role, and these associations are defined in the policy store. The following illustration shows from where each area of security management is controlled:



These three types of security providers within Oracle Fusion Middleware are integral to the Oracle BI 11g architecture. A chapter or more could be written on each provider but that is outside the scope of this book. Further recommended research on this topic would be to look at Oracle Fusion Middleware Security, OPSS, and the **Application Development Framework (ADF)**.

## System requirements

The first thing to recognize with infrastructure requirements prior to deploying Oracle BI 11g is that its memory and processor requirements have been increased since previous versions. The Java application server, WebLogic Server, gets installed with the full version of its software (though under a limited/restricted license as already discussed). A multitude of additional Java libraries and applications are also deployed. Ultimately, the authentication **Application Development Framework (ADF)** used to develop much of the platform accounts for a larger overall footprint. Be prepared for a recommended minimum 8 GB **Read Access Memory (RAM)** requirement for an enterprise deployment and a 4 GB minimum requirement for a developer workstation deployment. Other system requirement information can be found within the Oracle documentation at [http://docs.oracle.com/cd/E23943\\_01/doc.1111/e15722/overview.htm#CJAHADHD](http://docs.oracle.com/cd/E23943_01/doc.1111/e15722/overview.htm#CJAHADHD).

## Client Tools

Since release 11.1.1.5, Oracle BI 11g has a separate Client Tools installation that requires Microsoft Windows XP or a more recent version of the Windows OS. The Oracle BI 11g Client Tools provides the majority of client to server management capabilities required for normal day-to-day maintenance of the Oracle BI repository and related artifacts. The Client Tools installation is usually reserved for Oracle BI developers who design and maintain the Oracle BI metadata repository, better known as RPD, which stems from its binary file extension (.rpd). Compared to previous versions of the product, there are two tools now, which have been removed from the Client Tools installation:

- Oracle BI ODBC Manager
- Oracle BI Catalog Manager (available with version 11.1.1.6)

The Oracle BI 11g Client Tools installation provides each workstation with the Administration Tool, Job Manager, and all command-line Application Programming Interface (API) executables.

## **Multiuser Development Environment**

One of the key features of Oracle BI development is the ability for multiple metadata developers to develop simultaneously. Although the use of the term "simultaneously" can vary amongst the technical communities, the use of concurrent development within the Oracle BI suite requires Oracle BI's **Multiuser Development (MUD)** environment configuration.

The configuration itself is fairly straightforward and ultimately relies on the Oracle BI administrator's ability to divide metadata modeling responsibilities into projects. Projects – which are usually defined and delineated by logical fact table definitions – can be assigned to one or more metadata developers.

In previous versions of Oracle BI, a metadata developer could install the entire Oracle BI product suite on an up-to-date laptop or commodity desktop workstation and successfully develop, test, and deploy an Oracle BI metadata model. The system requirements of Oracle BI 11g prevent developers from installing the full Oracle BI 11g server suite on a legacy workstation as the minimum memory requirement is 4 GB. Most of the 32-bit workstations operate with 3 GB or less requirement.

If an organization currently leverages the Oracle BI MUD environment, or plans to leverage with the current release, this raises several questions:

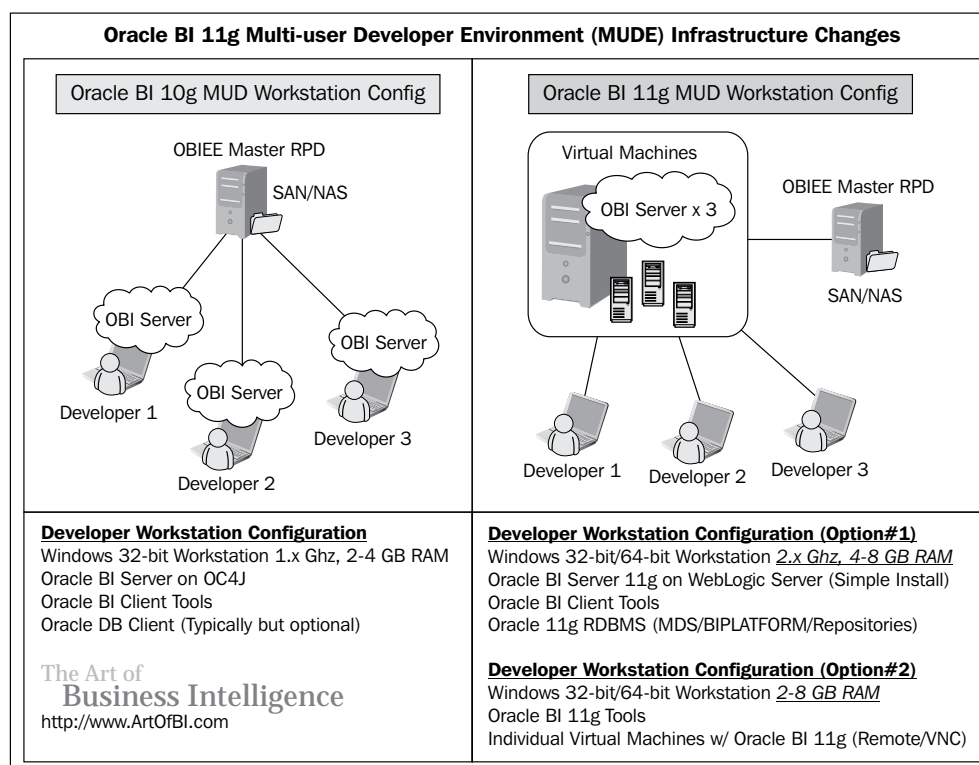
- How do we get our developers to the best environment suitable for developing our metadata?
- Do we need to procure new hardware?

Most of the developers' desktop workstations or laptops run 32-bit Microsoft Windows XP or Windows Vista. Microsoft Windows is a requirement for the Oracle BI Client Tools. However, the Oracle BI Client Tools does not include the server component of the Oracle BI 11g environment. It only allows for connecting from the developer's workstation to the Oracle BI server instance. In a multiuser development environment, this poses a serious problem as only one metadata repository can exist on any one Oracle BI server instance at any given time. If two developers are working from their respective workstations at the same time and wish to see their latest modifications published in a **rapid application development (RAD)** cycle, this type of iterative effort fails as one developer's published changes will overwrite the others in real time.

To resolve the above issue there are two recommended solutions. The first is an obvious localized solution. This solution merely upgrades the Oracle BI developers' workstations or laptops to comply with the minimum requirements for installing the full Oracle BI environment on said machines. This upgrade should be both memory (RAM) and processor (MHz) centric. A dual-core processor and a 4 GB plus RAM are recommended.

However, in order for a Windows operating system to acknowledge and use more than 4 GB memory, a 64-bit operating system kernel is required. Without an upgraded workstation, from which to work, Oracle BI 11g metadata developers will sit at a disadvantage for general iterative metadata development and especially be disenfranchised if interfacing within a multiuser development environment.

The second solution is one that takes advantage of **virtual machines (VM)**. The virtual machines have become a staple within most IT departments as they are versatile and allow speedy proposition of server environments. For this scenario, it is recommended to create a virtual machine template of an Oracle BI 11g environment, from which individual virtual machine images for each metadata developer on the Oracle BI development team can be duplicated or stood up. This solution effectively provides each metadata developer with their own Oracle BI development environment server which contains the fully deployed Oracle BI server environment. Then, developers have the ability to develop and test iteratively by connecting to their assigned virtual server without fear that their efforts will conflict with another developer's. The following diagram illustrates how an Oracle BI 11g MUD environment can leverage either upgraded developer workstation hardware or VM images to facilitate development:



This book does not cover the installation, configuration, or best practices for developing in a MUD environment. However, the Oracle BI development team deserves a lot of credit for documenting these processes in unprecedented detail. The Oracle BI 11g MUD environment documentation provides a case study ([http://docs.oracle.com/cd/E21764\\_01/bi.1111/e10540/case\\_study.htm#CHDGIBBD](http://docs.oracle.com/cd/E21764_01/bi.1111/e10540/case_study.htm#CHDGIBBD)), which conveys best practices for managing a complex Oracle BI development lifecycle. When you are ready to deploy a MUD environment, it is highly recommended to peruse this documentation first.

The information in this section seeks to convey these best practices in facilitating a developer's workstation when using a MUD environment, but the mentioned resources will delve further into this.

## **Certification matrix**

Now that the Oracle BI 11g is part of the larger enterprise Fusion Middleware stack, the Oracle BI tool suite complies largely with the overall Fusion Middleware infrastructure. This common foundation allows for a centralized model to communicate with operating systems, web servers, and other ancillary components which are compliant. The certification matrix for Oracle BI 11g can be found on the Oracle website at <http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>.

The certification matrix document is usually provided in a Microsoft Excel format and should be referenced before you begin any project or deployment of Oracle BI 11g. This will ensure that infrastructure components such as the selected operating system, web server, web browsers, LDAP server, and so on will actually work when integrated with the product suite.

## **Scaling out Oracle BI 11g**

There are several reasons why an organization may wish to expand their Oracle BI 11g footprint. This can range anywhere from requiring a highly available environment to achieving high levels of concurrent usage over time. The number of total end users, the number of total concurrent end users, the volume of queries, the size of the underlying data warehouse, and cross-network latency are even more factors to consider.

Scaling out an environment has the potential to solve performance issues and stabilize the environment. When scoping out the infrastructure for an Oracle BI 11g deployment, there are several crucial decisions to be made. These decisions can be greatly assisted by preparing and properly using Oracle's recommended guides for clustering and deploying Oracle BI 11g on an enterprise scale.

## Preconfiguration run down

Configuring the Oracle BI 11g product suite, specifically while involving scaling out or setting up **high availability (HA)**, needs preparation. Proactively taking steps to understand what it takes to correctly establish or preconfigure the infrastructure required to support any level of fault tolerance and high availability is critical. Even if the decision to scale out from the initial Oracle BI 11g deployment has not been made and if the potential exists, proper planning is recommended. Proper planning for HA can be achieved by following the Oracle Enterprise Deployment Guide for Oracle BI at [http://docs.oracle.com/cd/E23943\\_01/doc.1111/e15722/toc.htm](http://docs.oracle.com/cd/E23943_01/doc.1111/e15722/toc.htm).

## Shared storage

We would remiss the most important concepts of scaling out Oracle BI 11g if we do not highlight it, specifically for high availability — shared storage. The idea of shared storage, is that in a fault tolerance environment there are binary files and other configuration metadata that needs to be shared across the nodes. If these common elements are not shared, then if one node is to fail there is a potential loss of data.

Most importantly, in a highly available Oracle BI 11g environment, there can be only one WebLogic Administration Server running for that environment at any instance in time. An HA configuration makes one Administration Server active while the other is made passive. If the appropriate preconfiguration steps for shared storage (as well as other items in the high availability guide) are not properly completed, one should not expect accurate results from his/her environment.

## Clustering

A major benefit of Oracle BI 11g's ability to leverage the WebLogic Server as the Java application server tier, is that for every default installation, Oracle BI gets established in a clustered architecture. There is no additional configuration necessary to set this architecture in motion. Clearly, installing Oracle BI 11g on a single server only provides a single server with which to interface, however, upon doing so Oracle BI is installed into a single node clustered application server environment. Then, additional clustered nodes of Oracle BI 11g can be configured to establish and expand the server either horizontally or vertically.



## **Vertical expansion versus horizontal expansion**

With respect to the enterprise architecture and infrastructure of the Oracle BI 11g environment, a clustered environment can really expand in one of the two ways – horizontally and vertically. A horizontal expansion is the typical expansion type when clustering. It is represented by installing and configuring the application on a separate physical server with reference to the main server application. A vertical expansion is usually represented by expanding the application on the same physical server under which the main server application resides.

There are benefits to both scaling options. A vertical scale out can provide an advantage of a single machine's processor or memory power and save the cost of a separate physical machine. Horizontal scale out provides the advantages of failover, multiple machines, and redundancy. The decision to scale out in one way or the other is usually predicated on cost of additional physical servers, server limitations, and peripherals such as memory, processors, or an increase in usage activity by the end users. Some considerations, which may be used to assess which approach is the best for your specific implementation, might be as follows:

- Load balancing capabilities and need for an Active-Active versus Active-Passive architecture (suggests horizontal scaling)
- Need for a failover or high availability (suggests horizontal scaling)
- Cost for processor and memory enhancements versus cost of new servers (suggests vertical scaling)
- Anticipated increase in concurrent user queries (horizontal or vertical scaling)
- Realized decrease in performance due to an increase in user activity (horizontal or vertical scaling)

## **Oracle BI Server (system component) Cluster Controller**

When discussing scaling out of the Oracle BI server cluster, it is a common mistake to confuse the WebLogic Server application clustering with the Oracle BI Server Cluster Controller. In an attempt to clarify, it helps to remember that the Oracle BI Server System Component is the service or server engine controlled by the OPMN. There is also an Oracle BI Managed Server which is controlled by the WebLogic Server. Currently, Oracle BI 11g can only have a single metadata repository (RPD) reference associated with an Oracle BI server deployment instance at any single point in time. Because of this, the Oracle BI server engine leverages a failover concept to ensure some level of high availability that exists when the environment is scaled out.

In an Oracle BI scaled out and clustered environment, a secondary node, which has an instance of Oracle BI 11g installed, contains a secondary Oracle BI Server engine. From the main Oracle BI Managed Server containing the primary Oracle BI server instance, the secondary Oracle BI server instance is established as the failover server engine, using the Oracle BI Server Cluster Controller. This configuration takes place in the Enterprise Manager Fusion Control console. Based on this configuration, the scaled-out Oracle BI Server engine acts in an Active-Passive mode. That is, when the main Oracle BI 11g server engine instance fails, the secondary or passive Oracle BI server engine then becomes active to route requests and field queries.

## Failover and high availability

With high availability, it is often very hard to achieve multiples of nine (that is, 99.999 percent) for uptime of any server or application server environment. High availability is the type of architecture associated with an environment when attempting to maintain a high level of application availability and minimize downtime.

**Failover** is the process that takes place when a server node in a cluster fails and application traffic, otherwise intended for the down server, flows to the other active clustered server nodes. Failover also requires some level of load balancing and the concept can vary depending on the desired architecture within an organization, but the general concept should be roughly the same in most topologies.

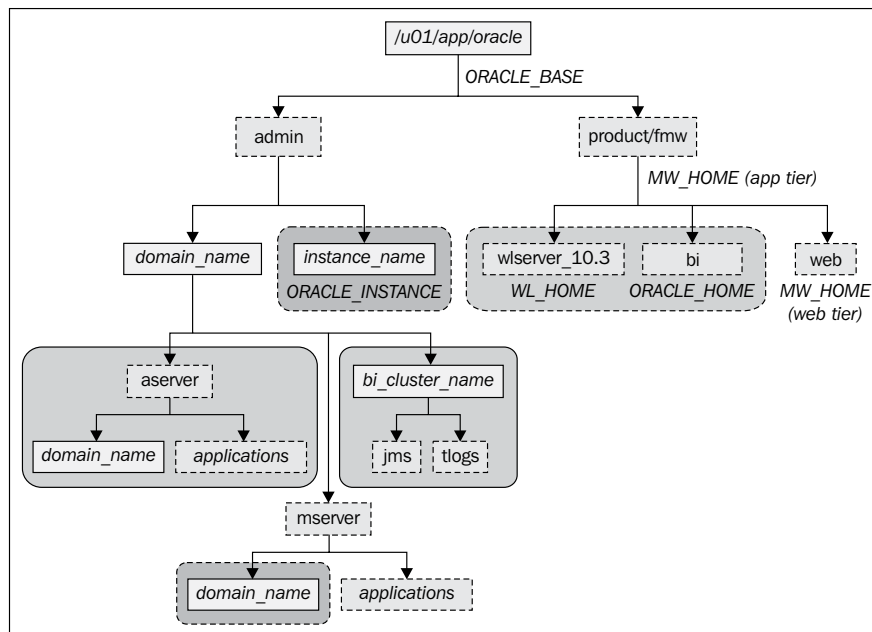
As part of an enterprise deployment strategy, taking failover and high availability into consideration is usually part of the architecture planning process. A step-by-step configuration for HA or a failover environment is an advanced infrastructure topic and is beyond the scope of this book. However, it is important to note that because Oracle BI 11g is part of the Fusion Middleware stack, it has the ability to capitalize on all fault tolerance features offered by that common architecture.

## Enterprise deployment guide

In an effort to relay best practices and strategic deployment of a large-scale enterprise, Oracle BI deployment, Oracle lends a big helpful hand and provides a topology referred to as the **Enterprise Deployment Guide (EDG)**. This guide should not be taken lightly. When deciding on major factors of a full-scale enterprise wide Oracle BI deployment, this topology is the one that should be referenced first. Use the EDG to plan for required resource skills, procurement of hardware, and as a gauge to estimate the effort involved in achieving the architecture that your requirements demand. The topology includes pertinent information regarding load balancing, virtual IP addresses, separation of HTTP servers from application servers, and other fully vetted infrastructure recommendations. It is especially recommended to view this guide before embarking on any deployment involving extranet access to your Oracle BI implementation or a large user base of an internal deployment.

## Directory folder structure

As you get started with installing, configuring, and deploying Oracle BI 11g in the subsequent chapters, you will see several references to files inside the Fusion Middleware folder structure and the Oracle BI home folder structure. Again, taking advantage of a common architecture, Oracle BI 11g leverages the Fusion Middleware foundation to organize the filesystem structure. This consistency is a benefit for all implementers and administrators that will eventually maintain the platform. It is recommended that, as you progress in your learning of Oracle BI 11g, you should note the folders which contain files pertinent for modifying the environment, or assisting with troubleshooting efforts. The following diagram illustrates the standard logical deployment structure for Oracle BI 11g:



## Log files (diagnostics)

As you begin developing and deploying your Oracle BI 11g solutions, you will eventually run into some issues down the road. Face them, you are implementing a technology! Issues, no matter how minor, are bound to arise. No one plans for them. But here is a list of log files' locations where it is best to begin troubleshooting in Oracle BI 11g:

```
<OBI_LOGS> = /u01/fmw/instances/instance1/diagnostics/logs
```

Log Name (.log)	Location
NQQuery	<OBI_LOGS>/OracleBIServerComponent/coreapplication_obis1
NQServer	<OBI_LOGS>/OracleBIServerComponent/coreapplication_obis1
NQSAdminTool	<OBI_LOGS>/OracleBIServerComponent/coreapplication_obis1
NQSUDMLExec	<OBI_LOGS>/OracleBIServerComponent/coreapplication_obis1
RPD migration utility	<OBI_LOGS>/OracleBIServerComponent/coreapplication_obis1
SAW log	<OBI_LOGS>/OracleBIPresentationServicesComponent/ coreapplication_obips1
JavaHost	<OBI_LOGS>/OracleBIJavaHostComponent/coreapplication_ objh1
NQScheduler	<OBI_LOGS>/OracleBISchedulerComponent/coreapplication_ obisch1
NQCluster	<OBI_LOGS>/OracleBIClusterControllerComponent/ coreapplication_obiccs1
ODBC	<OBI_LOGS>/OracleBIODBCComponent/coreapplication_obips1

All logs regarding the OPMN system can be found at the following centralized location:

```
<OBI_LOGS>/OPMN/opmn/
```

## Configuration files

Those who are familiar with the previous versions of Oracle BI might be surprised that several legacy named physical configuration files still reside in the Oracle BI 11g architecture. These files can still be manually manipulated to configure the Oracle BI 11g environment. However, much of the basic configuration is handled via the Oracle BI Enterprise Manager which will be discussed in more detail in *Chapter 5, Understanding the Systems Management Tools*. If there is a need to locate these configuration files, they can be found at the following location:

```
<FMW_HOME>\instance\instance1\config\component_type\bi_component_name
```

The following is a list of configuration files based on the central Oracle BI 11g instance path of <FMW\_HOME>\instance\instance1\config\:

- \OPMN\opmn\opmn.xml
- \OracleBIClusterControllerComponent\coreapplication\_obiccs1\ccslogconfig.xml

- \OracleBIJavaHostComponent\coreapplication\_obijhl\logging\_config.xml
- \OracleBIPresentationServicesComponent\coreapplication\_obips1\instanceconfig.xml
- \OracleBISchedulerComponent\coreapplication\_obisch1\instanceconfig.xml
- \OracleBIServerComponent\coreapplication\_obis1\logconfig.xml

## **A review – what I should now know!**

For a self review and recap of this chapter, here are a few questions based on important topics covered in this chapter. There is no answer key. These questions are for your own reflection on the chapter material.

1. What is Oracle Fusion Middleware?
2. What are some of the reasons to scale out an Oracle BI deployment?
3. List each of the Oracle BI system components. How many are there?
4. What is the name of the guide that should be referenced while planning an enterprise deployment of Oracle BI 11g?
5. What is the purpose of the Identity Store?

## **Additional research suggestions**

The following list would help you to continue your learning:

- Oracle BI 11g Enterprise Deployment Guide: [http://download.oracle.com/docs/cd/E21764\\_01/doc.1111/e15722/toc.htm](http://download.oracle.com/docs/cd/E21764_01/doc.1111/e15722/toc.htm)
- Multiuser Development Environment: [http://download.oracle.com/docs/cd/E21764\\_01/bi.1111/e10540/lifecycle.htm#CIAFAFGE](http://download.oracle.com/docs/cd/E21764_01/bi.1111/e10540/lifecycle.htm#CIAFAFGE)
- A case study: [http://download.oracle.com/docs/cd/E21764\\_01/bi.1111/e10540/case\\_study.htm#CHDGIBBD](http://download.oracle.com/docs/cd/E21764_01/bi.1111/e10540/case_study.htm#CHDGIBBD)

## Summary

This chapter provided a high-level overview of the Oracle BI 11g architecture. Oracle BI 11g is an overhaul to prior versions of the tool. This chapter discussed how Oracle Fusion Middleware has made Oracle BI 11g an extremely robust tool incorporating open standards and best practices where possible. An effort was made to discuss topics, which are high on the priority list for most organizations, such as security and scaling out the architecture. Recommendations were provided around these topics where possible. Although there is a lot more detail that could be written into each subject section of this chapter, the chapter conveyed enough information to capture the bulk of the Oracle BI 11g architecture, so that you can speak intelligently about the subject after reading this chapter.

The architecture described in this chapter discussed the core foundation of Oracle BI 11g and did not discuss Oracle Real-time Decisions (RTD), the Oracle decision making engine, nor will the remainder of this book discuss it (an entire book could be devoted merely to this subject; so, we chose to leave it out of this).

The next several chapters take you on a journey through setting up Oracle BI 11g, prepping data sources, modeling a custom analytics solution, and getting to know Oracle BI intimately. Keep the momentum going and just know that by the end of this book, you'll be well on your way to becoming an Oracle BI aficionado.