Deployment Guide

Broadcast Agent 5.5

BusinessObjects 5.1

WebIntelligence 2.7

Windows/UNIX

Business Objects ™
Version 5.1

Deployment Guide

No part of the computer software or this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from Business Objects. The information in this document is subject to change without notice. If you find any problems with this documentation, please report them to Business Objects in writing at documentation@businessobjects.com. Business Objects does not warrant that this document is error free.

Copyright © Business Objects 2002

All rights reserved. Printed in France.

Trademarks:

The Business Objects logo, BusinessMiner, BusinessQuery, and WebIntelligence are registered trademarks of Business Objects SA in the US and/or other countries.

The Business Objects tagline, Broadcast Agent, BusinessObjects, Personal Trainer, Rapid Deployment Templates, and Set Analyzer are trademarks of Business Objects SA in the US and/or other countries.

Microsoft, Windows, Windows NT, Access, Microsoft VBA and other names of Microsoft products referenced herein are trademarks or registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation. All other names of Oracle products referenced herein are trademarks or registered trademarks of Oracle Corporation.

AIX is a registered trademark of International Business Machines (IBM) Corporation in the United States and other countries. All other names of IBM products referenced herein are trademarks or registered trademarks of International Business Machines Corporation.

Sun, Sun Microsystems, the Sun l.ogo, Solstice and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

HP-UX is a registered trademark of Hewlett-Packard Company in the United States and other countries. All other names of Hewlett-Packard products referenced herein are trademarks or registered trademarks of Hewlett-Packard Company.

UNIX is a registered trademark of The Open Group.

All other product and company names mentioned herein are the trademarks of their respective owners.

This software and documentation is commercial computer software under Federal Acquisition regulations, and is provided only under the Restricted Rights of the Federal Acquisition Regulations applicable to commercial computer software provided at private expense. The use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in subdivision (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at 252.227-7013.

U.S. Patent Numbers 5,555,403 and 6,247,008

Part Number: 373-50-510-01

Edition: 9

Contents

Finding the Information You Need **ix**

About this Guide x
Business Objects References xi
Your Comments Are Welcome xxii

Chapter 1 Overview **23**

What We Mean by "Deployment" 24
A Standard SQL Reporting Configuration 25
A Secure Business Objects Configuration 26
The Business Objects Product Line 28

Chapter 2 What's New in the Latest Release? **45**

Overview 48
Zero Admin BusinessObjects 48
Enhanced Audit and Tracing Facilities 56
CORBA Externalization 56
Scheduling BusinessObjects Documents from InfoView 57
Multiple WIGenerators on a Single Server 57
Administration Tools on Windows Platforms 57
Managing Default Options for New Users 59
Document Agnosticism 60
Saving BusinessObjects Documents as PDF Files 61
Adding Hyperlinks to BusinessObjects Documents 61
BusinessObjects OLAP Connect 62
The PDF File Cache 63
My InfoView 64
Choosing the Platform for Document Processing 65
Scheduling WebIntelligence Documents With Prompts 65

Improved Date and Time Handling 66
Broadcast Agent in UNIX-Only Clusters 66
Importing/Exporting Universes from SUPERVISOR 67
Printing Repository User and Resource Data to Table Format 67
Purging Inbox Documents 67
New WebIntelligence SDK Object Models 68
Broadcast Agent Publisher 68
What's New in Documentation 68
Multihomed Configuration Support 72

Chapter 3 Planning Your Deployment 73

Overview 74
Pre-Deployment Checklist 74
Planning Your Pilot Project 80
Conclusion 83

Chapter 4 Deployment Environments 85

Deployment Architectures 87
Client-Server or Distributed Architecture? 87
Distributed Architecture Overview 91
Distributed Environments and Clusters 116
Intranets and Extranets 122

Chapter 5 First Steps: Installing and Setting Up the System 127

Introduction 129
Fresh Installation or Update? 129
System Requirements 130
Business Objects Products and Windows 2000 130
Installing Business Objects Software 136
Installing and Setting Up Enterprise Server Products 139
Installing Desktop Products (Master Installation) 143
Creating Your Database Reporting Environment 144

Creating a Repository 144
Defining Your Users and User Groups (Supervisor) 145
Creating a Universe (Designer) 145
Creating Documents 147
Updating from Previous Releases 148
Installing Service Pack Updates 153
The Freeware Directory 153

Chapter 6 System Configuration, Sizing, and Benchmarking 155

Overview 157
Performance Factors and Sizing 157
Supported Configurations 175
Example Distributed Deployment Implementation 176
Specific Deployment Guidelines 181
Benchmarking 198

Chapter 7 System Tuning and Troubleshooting 201

Overview 204
Load Balancing 204
Improving Overall Document Computation Time 211
Relieving Overloaded BOManager Machines 212
Determining and Reducing HTML Overhead 212
Enhancing Performance from InfoView 216
Enhancing Performance from Supervisor 216
Enhancing Performance from Designer 217
Enhancing Performance from BusinessObjects 218
Windows NT Server Tuning Tips 220
Tuning NT and Web Servers to Work Together Optimally 226
UNIX Server Tuning Tips 237
Multiple OAD Instances Under UNIX 244
Modifying CORBA Kernel Parameters 247
Tracing System Activity for Tuning Purposes 259

Monitoring WebIntelligence Log Files on UNIX 260
Troubleshooting Your Deployment 262

Chapter 8 The Repository 275

What is the Repository? 276
Repository Domains 277
Which Components Access the Repository? 280
The Repository and Your Data Warehouse 281
Choosing a Repository Database 284
New 5.x Repository Features 287
Working With or Without a Repository 291
Repository Questions and Answers 292

Chapter 9 Access and Security 295

Introduction 297
Security at three levels 297
Business Objects Security 299
Windows NT Authentication Security 302
RDBMS security 311
Extended Security Commands 313
WebIntelligence and the SSL Protocol 314
WebIntelligence and LDAP Authentication 314
WebIntelligence and Firewalls 319

Chapter 10 Document Processing and Distribution with Broadcast Agent 343

Broadcast Agent Features 344
Broadcast Agent Components 348
Broadcast Agent Sizing Guidelines 354
Troubleshooting and Optimization 355
Where to get more information 358

Chapter 11 Windows, UNIX or Both? 359

Overview 360
Why Use UNIX Servers? 360
Version 5.1/2.6 Clusters 361
Migrating Storage from UNIX to Windows servers 365
UNIX Character Set Support 366
Refreshing BusinessObjects documents in Heterogeneous Clusters 367

Appendix A The Repository Tables 369

Overview 370
Broadcast Agent 371
Document Properties 372
Categories 374
Channels 375
Generic table 376

Appendix B Databases and Database Connectivity 377

Overview 378
Database Connectivity Basic Concepts 378
Where To Get More Information 381

Appendix C Understanding Secure HTTP Transactions 383

Introduction to Secure Transactions 384
Public-Key Encryption 385
Client and Server Certification 387
Making a Secure Connection 389
Configuring Your Web Server to Use SSL 390

Index

Finding the Information You Need

In this chapter

- **About this Guide x**
  - Audience x
  - Conventions used in this guide x
- **Business Objects References xi**
  - Deployment and server products xi
  - Query, Reporting and Analysis Products xiii
  - Error messages xviii
  - Developer Suite xviii
  - Database Access xx
  - Rapid Deployment Templates xxii
- **Your Comments Are Welcome xxii**

About this Guide

This guide presents information relating to the deployment of Business Objects products within your company or organization.

Audience

The guide is intended for the person in charge of deploying Business Objects products in a company's computing environment.

This person should be familiar with the hierarchical structure of his or her company, have basic Windows and/or UNIX administration skills and a basic knowledge of the products in the Business Objects product line.

Conventions used in this guide

The conventions used in this guide are described in the table below:

Convention	How Used
SMALL CAPITALS	The names of all Business Objects products: DESIGNER, SUPERVISOR, BUSINESSOBJECTS, etc.
The AvantGarde font	The names of Business Objects classes, objects and conditions. For example, Customer, Sales, Revenue, Service, etc.
The `Courier` font	SQL syntax or @ function syntax. For example: `@Select(Country\Country Id)`

Business Objects References

This section presents the complete list of references for Business Objects documentation. A reference for a guide includes its title, part number, pdf name, and a brief description.

Each guide falls into one of these categories:

- Deployment and Server Products
- Query, Reporting, and Analysis Products
- Developer Suite
- Error Messages
- Database Access
- Rapid Deployment Templates

Deployment and server products

In addition to this guide, the following documentation provides information on how to set up and deploy Business Objects server products within your company or organization.

Installation and Update Guide for Windows	
Part number	**PDF name**
313-10-510-01	InstallationEN.pdf
Gives complete instructions on installing and updating Business Objects software at your site on Windows.	
Installation and Update Guide for UNIX	
Part number	**PDF name**
313-20-260-01	Installationunixen.pdf
Gives complete instructions on installing and updating Business Objects software at your site on UNIX.	

Deployment Guide	
Part number	**PDF name**
383-10-510-01	DeploymentEN.pdf
Presents information relating to the deployment of Business Objects products within your company or organization.	
Supervisor's Guide	
Part number	**PDF name**
311-10-510-01	SupervisorEN.pdf
Describes how to use SUPERVISOR to set up a Business Objects environment: creating and maintaining the repository, defining users and groups, assigning resources to users, controlling security of the repository and its resources, and customizing user access to functionalities.	
Broadcast Agent Administrator's Guide	
Part number	**PDF name**
312-50-510-01	BroadcastAgentEN.pdf
Explains to administrators how to set up and run BROADCAST AGENT, which offers scheduled processing and distribution of documents.	
Broadcast Agent Publisher Administrator's Guide	
Part number	**PDF name**
317-10-100-01	BroadcastAgentPublisherAdminen.pdf
Explains to administrators how to set up and run BROADCAST AGENT PUBLISHER, which enables BROADCAST AGENT users to create publications using a range of powerful subscription-based email distribution and report bursting capabilities.	
WebIntelligence System Administrator's Guide	
Part number	**PDF name**
301-10-260-01	WebISystemAdminEN.pdf
Explains to administrators how to set up and run the WEBINTELLIGENCE system in a Windows and/or UNIX environment. WEBINTELLIGENCE is the industry's leading web-based query, reporting and analysis solution.	

UNIX Diagnostic Tools Guide	
Part number	**PDF name**
387-20-260-01	UNIXTroubleshootingDiagToolsEN.pdf
Describes the tools and methods available to you when you need to troubleshoot performance and functional problems encountered while using your WEBINTELLIGENCE deployment.	
UNIX Troubleshooting Guide: From Symptoms to Solutions	
Part number	**PDF name**
3A1-20-260-01	UNIXTroubleshootingSolutionsEN.pdf
Describes some of the most frequently-encountered issues with Business Objects deployments under UNIX and provides solutions.	
UNIX Performance Optimization and Sizing Guide	
Part number	**PDF name**
313-10-510-01	Eupos10SP301.pdf
Provides a set of guidelines for optimizing the performance of your Business Objects deployment in a UNIX environment.	

Query, Reporting and Analysis Products

This category is made up of end-user documentation relating to querying, reporting, and analysis. This includes the user's guides for BUSINESSOBJECTS, WEBINTELLIGENCE, DESIGNER, INFOVIEW, and all OLAP access packs.

BusinessObjects User's Guide	
Part number	**PDF name**
310-10-510-01	BusinessObjectsEN.pdf
Documents all BUSINESSOBJECTS features and functions. It describes the different ways to access data to use in your documents, how to create and format these documents and analyze the data in them, and how to publish and distribute the finished documents.	
Getting Started with BusinessObjects	

Part number	PDF name
302-10-510-01	BusinessObjectsGettingStartedEN.pdf
Presents a step-by-step introduction to creating, formatting, analyzing, and publishing documents with BUSINESSOBJECTS.	
InfoView User's Guide	
Part number	**PDF name**
384-10-260-01	InfoView_EN.pdf
Covers INFOVIEW, the Business Objects portal that offers an overview of all the documents in the corporate repository. This product is also used to view, manage, and distribute these documents.	

WebIntelligence User's Guide	
Part number	**PDF name**
353-10-260-01	WebIUserEN.pdf
Describes how to use WEBINTELLIGENCE to create, format, analyze, and publish documents. WEBINTELLIGENCE is the first web-based query, reporting, and analysis solution. This guide is available as an Acrobat PDF document, and in HTML, but not as a printed book.	
Getting Started with WebIntelligence Guide	
Part number	**PDF name**
302-10-260-01	WebIGettingStartedEN.pdf
Presents a step-by-step introduction to creating, formatting, analyzing and publishing documents with WebIntelligence.	
Broadcast Agent Publisher User's Guide	
Part number	**PDF name**
319-10-100-01	BroadcastAgentPublisherUseren.pdf
Explains how to use BROADCAST AGENT PUBLISHER, which enables BROADCAST AGENT users to create publications using a range of powerful subscription-based email distribution and report bursting capabilities.	
Designer's Guide	
Part number	**PDF name**
307-10-510-01	DesignerEN.pdf
Contains detailed instructions on creating simple to complex universes, the semantic layer that represents database structure in everyday business terms. It also contains information on managing and distributing universes as well as workgroup issues.	
Metacube OLAP Access Pack User's Guide	
Part number	**PDF name**
351-10-510-01	MetacubeOLAPEN.pdf
Describes how to build BUSINESSOBJECTS documents with data from Informix MetaCube, a multidimensional database solution.	

Essbase OLAP Access Pack User's Guide	
Part number	**PDF name**
308-10-510-01	EssbaseOLAPEN.pdf
Describes how to build BUSINESSOBJECTS documents with data from Essbase, a multidimensional database solution.	
DB2 OLAP Access Pack User's Guide	
Part number	**PDF name**
359-10-510-01	DB2OLAPEN.pdf
Explains how to build BUSINESSOBJECTS documents with data from DB2 OLAP, a multidimensional database solution.	
Express OLAP Access Pack User's Guide	
Part number	**PDF name**
347-10-510-01	ExpressOLAPEN.pdf
Describes how to build BUSINESSOBJECTS documents with data from Oracle PC Express, a multidimensional database solution.	
SAP BAPI Access Pack User's Guide	
Part number	**PDF name**
348-10-510-01	SAPBAPIEN.pdf
Covers the SAP BAPI Access Pack, a data provider that lets BUSINESSOBJECTS users access SAP data using SAP BAPIs (Business Application Programming Interfaces).	
BusinessQuery for Excel User's Guide	
Part number	**PDF name**
305-10-510-01	BusinessQueryEN.pdf
Explains how to bring data from relational databases to Microsoft Excel, using the BUSINESSOBJECTS query technique. It also contains information on managing and updating workbooks containing query data.	

BusinessMiner User's Guide	
Part number	**PDF name**
343-10-420-01	BusinessMiner_EN.pdf
Describes data mining using BUSINESSMINER. Data mining is the analysis of data and the use of software techniques for extracting hidden, comprehensible, and useful information from databases.	
OLAP Connect User's Guide	
Part number	**PDF name**
398-10-510-01	OLAPConnecten.pdf
The OLAP Connect User's Guide is an updated and comprehensive version of the previous Getting Started with OLAP Connect. It is a summarized guide for using BusinessObjects MDX Connect and BusinessObjects BW Connect. New information includes rank filter and sort workflows, as well as specific user interface options for SAP BW.	
WebIntelligence Guide for Olap Data Sources	
Part number	**PDF name**
399-10-260-01	WebiOLAPUserEN.pdf
This guide describes how to use WebIntelligence for OLAP data sources. It includes instructions for each of the features, as well as basic information about OLAP databases. Only Microsoft SQL Server Analysis Services is supported.	
WebIntelligence OLAP Administration and Drill Through Guide	
Part number	**PDF name**
307-10-510-01	DrillThroughAdminen.pdf
The WebIntelligence OLAP Administration and Drill Through Guide contains installation instructions for WebIntelligence OLAP. It also provides the WebIntelligence OLAP administrator with the information necessary to set up and maintain an OLAP-to-relational drill through environment, as well as instructions on using UDS Designer, the application used to the build the translation maps for drill through.	

Error messages

The *Error Message Guide* is a compilation of the error messages that can appear with ordinary use of Business Objects products.

You can open this online guide from any of the Business Objects products by selecting the Error Messages Explained command from the Help menu. From INFOVIEW, click Error Messages in the On-Line Help navigation bar.

Error Message Guide	
Part number	**PDF name**
383-10-510-01	ErrorsEN.pdf
Provides detailed troubleshooting information so that you can determine the reasons for an error and take the appropriate steps to resolve it. It allows you to search for error messages by code. Each error message appears with its probable cause and the recommended course of action.	

Developer Suite

This category is made up of developer's guides for Software Development Kits (SDKs) pertaining to BUSINESSOBJECTS, WEBINTELLIGENCE, and BUSINESSQUERY. It also includes object model diagrams that illustrate the classes, properties, and methods available in the SDKs.

Introduction to Developer Suite	
Part number	**PDF name**
314-10-510-01	IntroDevSuiteEN.pdf
Introduces the concepts, tasks, and procedures required to use the BUSINESSOBJECTS Developer Suite to customize BUSINESSOBJECTS and WEBINTELLIGENCE.	
Putting Developer Suite to Work	
Part number	**PDF name**
315-10-510-01	DevSuiteToWorkEN.pdf
Explains how to get the most from Developer Suite 5.1. It contains advice, tips and code samples for developing BUSINESSOBJECTS SDK and WEBINTELLIGENCE SDK applications.	

BusinessObjects SDK Reference Guide	
Part number	**PDF name**
344-10-510-01	BOSDKEN.pdf
Provides information on the BUSINESSOBJECTS Developer's Suite. It explains how to use Visual Basic for Applications to create custom macros and add-ins for BUSINESSOBJECTS and DESIGNER. In addition, this guide describes the use of OLE Automation for creating BUSINESSOBJECTS scripts.	
BusinessObjects SDK Object Model Diagram	
Part number	**PDF name**
254-10-510-01	BOSDKOMDEN.pdf
These foldout diagrams contain reference information on the classes, properties, and methods available in the BUSINESSOBJECTS Developer's Suite.	
WebIntelligence SDK Reference Guide	
Part number	**PDF name**
388-12-260-01	WIDSKEN.pdf
Describes the various concepts, tasks, and procedures required to develop custom deployments of the WEBINTELLIGENCE system. This includes an overview of how the system operates, and detailed information about the WEBINTELLIGENCE object model.	
WebIntelligence SDK Object Model Diagram	
Part number	**PDF name**
389-12-260-01	WEBISDKOMDEN.pdf
This foldout diagram contains reference information on the classes, properties, and methods available in the WEBINTELLIGENCE Developer's Suite.	
Installing and Configuring Application Servers	
Part number	**PDF name**
316-50-510-01	InstallConfigAppServen.pdf
Describes how to install and configure application servers currently supported by WEBINTELLIGENCE SDK for developing JavaServer Pages using the WIBean object model.	

BusinessQuery SDK Developer's Guide	
Part number	**PDF name**
356-10-510-01	BQSDKEN.pdf
Covers BUSINESSQUERY SDK, which enables application programmers to create applications that query relational databases using the BUSINESSOBJECTS query technique.	
BusinessQuery Object Model Diagram	
Part number	**PDF name**
358-10-510-01	BQSDKOMDEN.pdf
This diagram contains reference information on the classes, properties, and methods available with the BUSINESSQUERY Developer Suite.	

Database Access

This category consists of one guide for each RDBMS supported by the Business Objects product line. Each guide provides information on:

- Running the different Business Objects products with the given RDBMS
- Configuring machines
- Defining a connection to an RDBMS to access data
- SQL syntax specific to an RDBMS

Microsoft SQL Server Database Guide	
Part number	**PDF name**
377-10-510-01	MicrosoftSQLServerEN.pdf
Oracle Database Guide	
Part number	**PDF name**
320-10-510-01	OracleEN.pdf
Sybase Database Guide	
Part number	**PDF name**
321-10-510-01	SybaseEN.pdf

IBM DB2 Database Guide	
Part number	**PDF name**
325-50-510-01	DB2_EN.pdf
Informix Database Guide	
Part number	**PDF name**
322-10-510-01	InformixEN.pdf
Microsoft Access Database Guide	
Part number	**PDF name**
376-10-510-01	MicrosoftAccessEN.pdf
Teradata Database Guide	
Part number	**PDF name**
324-10-510-01	TeradataEN.pdf
Generic ODBC Access Guide	
Part number	**PDF name**
334-10-510-01	GenericODBCEN.pdf
Star SQL Database Guide	
Part number	**PDF name**
381-10-510-01	StarSQLEN.pdf
Ingres Database Guide	
Part number	**PDF name**
330-10-510-01	IngresEN.pdf

Red Brick Database Guide	
Part number	**PDF name**
333-10-510-01	RedBrickEN.pdf
Repository Migration Guide	
Part number	**PDF name**
319-10-510-01	RepositMigrationEN.pdf
Neon Shadow Direct Database Guide	
Part number	**PDF name**
382-10-510-01	ShadowEN.pdf

Rapid Deployment Templates

The guides in this category cover BUSINESSOBJECTS Rapid Deployment Templates (RDTs). An RDT provides users with predefined universes and reports for third party applications; it enables them to begin working immediately with BUSINESSOBJECTS in order to create reports for viewing and analyzing information.

To obtain these guides, contact your Business Objects sales representative.

Your Comments Are Welcome

We encourage you to contact us with any suggestions, ideas or comments that you may have regarding this guide. Your opinions are the most important input we receive as we revise Business Objects documentation.

If you would like to contact us by e-mail, please write to documentation@businessobjects.com.

To find out information about Business Objects products and services, visit our World Wide Web site at http://www.businessobjects.com.

This site also provides information on contacting our technical support services.

Chapter 1 Overview

In this chapter

- **What We Mean by "Deployment" 24**
- **A Standard SQL Reporting Configuration 25**
- **A Secure Business Objects Configuration 26**
  First in e-Business Intelligence 27
- **The Business Objects Product Line 28**
  Enterprise server products 29
  Desktop products 34

What We Mean by "Deployment"

By **deployment** we mean the distribution and arrangement of Business Objects query, reporting and analysis tools within a network infrastructure.

Initially offered for use in a standard client-server environment, the products can now be deployed and used over intranets, extranets and the World Wide Web, in conjunction with WEBINTELLIGENCE. This means that users can access documents from wherever they are, using a standard thin-client web browser. Additionally, documents can be automatically broadcast to users at scheduled intervals using BROADCAST AGENT. These documents can be set up so that the information made available to recipients is determined by their particular user profile.

Each deployment is different. How your deployment is structured and configured depends on the number of users supported, the type of information access that you want your users to have, the number and types of databases you're using, the geographical distribution of your company, and so on.

Before looking at the specific products Business Objects has to offer, it may be useful to review some basics, and to look at the reasons why customers choose Business Objects as their database query, reporting, and analysis tool.

A Standard SQL Reporting Configuration

Before applications such as Business Objects product suite, customers relied on database analysts and SQL specialists to directly access their databases and produce meaningful documents.

This type of information access provides a direct link to the corporate database, and allows a large amount of functionality and customization by building unique combinations of SQL statements. However, it is not particularly flexible or easy to modify, since it requires a sufficient level of expertise to build the SQL statements into a coherent reporting environment. Also, data security is limited to the security mechanisms which are built into the chosen corporate database.

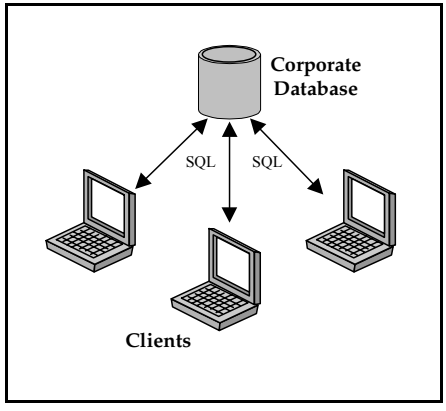

A Standard SQL Reporting Configuration

A Secure Business Objects Configuration

As companies expanded and diversified, it became more and more important that these functions were made accessible to non-specialized users within the company. A simplified interface was required; one that allowed users to execute in-depth database queries, and generate comprehensive and professional documents, without needing to understand the processes that lay underneath. This is the interface that Business Objects delivers with its Business Intelligence software products.

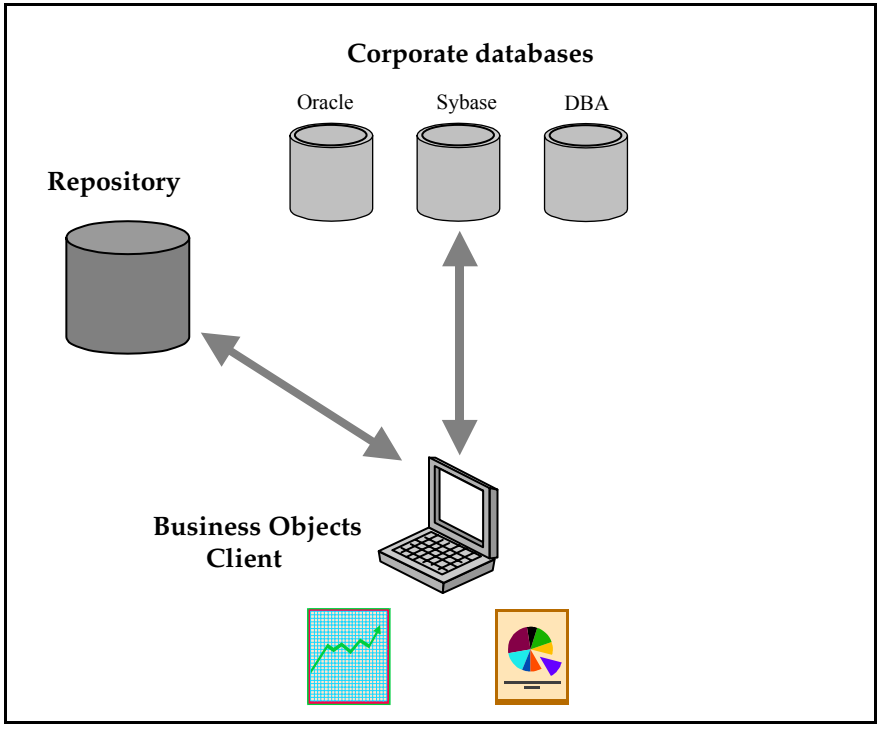

A Secured Business Objects Configuration

The Business Objects product suite uses a **repository** - a database which is stored in a relational database management system. The repository is used by the various Business Objects products and modules to secure access to your data warehouse and to provide a deployment infrastructure for the Business Objects applications.

Your deployment relies on the power of the repository to support the advanced features of the Business Objects solution. For more information, see Chapter 8 "The Repository" on page 275.

First in e-Business Intelligence

Business intelligence technology enables decision-makers in an organization to access, analyze, report and share corporate information. But with the advent of e-business, the Business Objects product line has expanded to take a leadership role in **e-business intelligence**, or e-BI.

With e-BI, organizations around the world can link their customers, suppliers, partners and internal services to share information and work together in new ways. Based on its web-based thin-client product, WEBINTELLIGENCE, and its web-enabled full-client product, BUSINESSOBJECTS, the Business Objects solution provides a robust and secure way to access, analyze and share information, either within an enterprise via an intranet, or beyond an enterprise via an extranet.

What can you get out of it? Reduced cost, greater efficiency, faster transactions and increased revenue.

The Business Objects Product Line

The Business Objects product line can be broken down into:

- Enterprise server products
- Desktop products

The same distinction is used in the installer program.

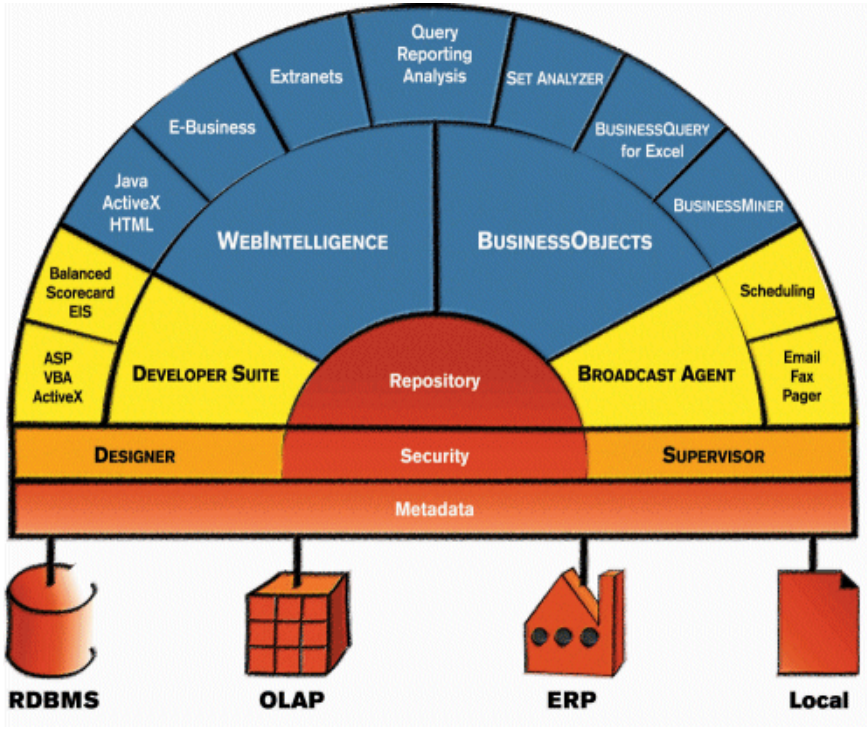

Enterprise server products

There are four main enterprise server products:

- WEBINTELLIGENCE 2.7
- BROADCAST AGENT 5.5
- Zero Admin BUSINESSOBJECTS
- WEBINTELLIGENCE SDK

These products are described in greater detail in "Distributed Architecture Overview" on page 91.

WebIntelligence 2.7

WEBINTELLIGENCE is the leading business solution designed to meet your critical query, reporting and analysis needs over the World Wide Web.

With WEBINTELLIGENCE, you can access the data in your corporate databases or data warehouses from within your office, your home, or around the world, using your corporate intranet, extranet, or the World Wide Web.

WEBINTELLIGENCE provides users with a light-weight, easy-to-use interface called the Web Panel to create and modify documents.

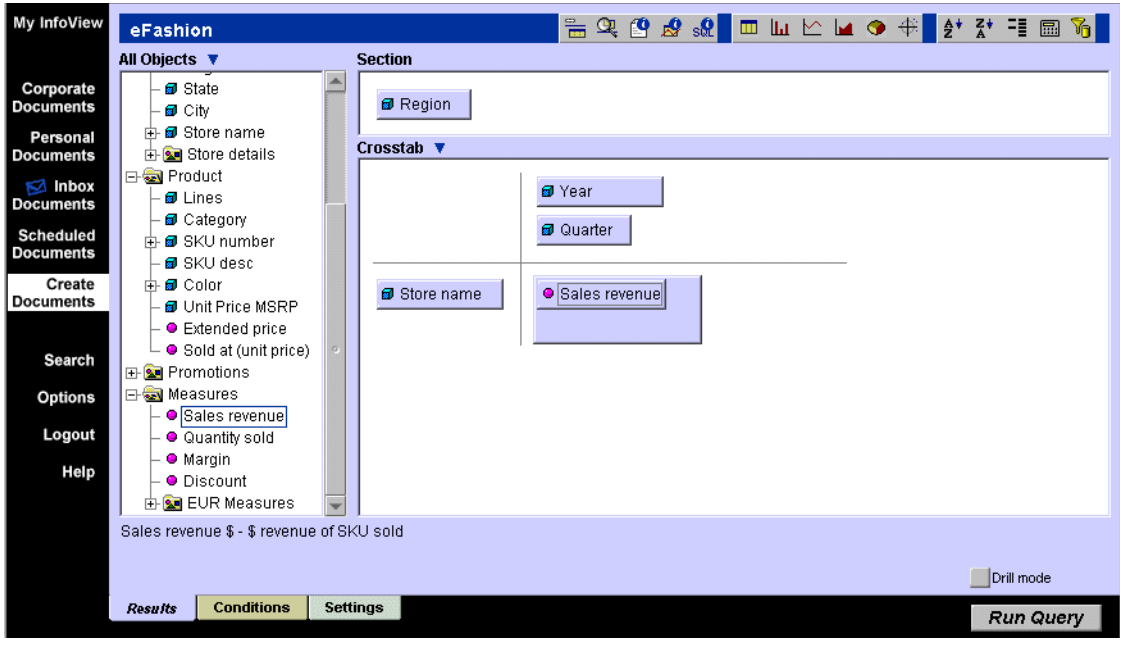

WEBINTELLIGENCE uses the same universes created for full-client users. These universes are the sole data providers for WEBINTELLIGENCE documents.

All users need to do to create and edit WEBINTELLIGENCE documents is to log into a web server and download a Java applet or ActiveX control from the corporate intranet, extranet, or the web. They can then build queries to create documents that display accurate, up-to-date data from their corporate databases.

When users have finished editing the queries, they send them to the server, where they are processed, formatted and sent back as HTML documents.

Once you've used WEBINTELLIGENCE to query a database, you can take the information and turn it into documents as simple as tables, or as complex as dynamic documents with drillable charts.

Example **WebIntelligence document containing a drillable chart**

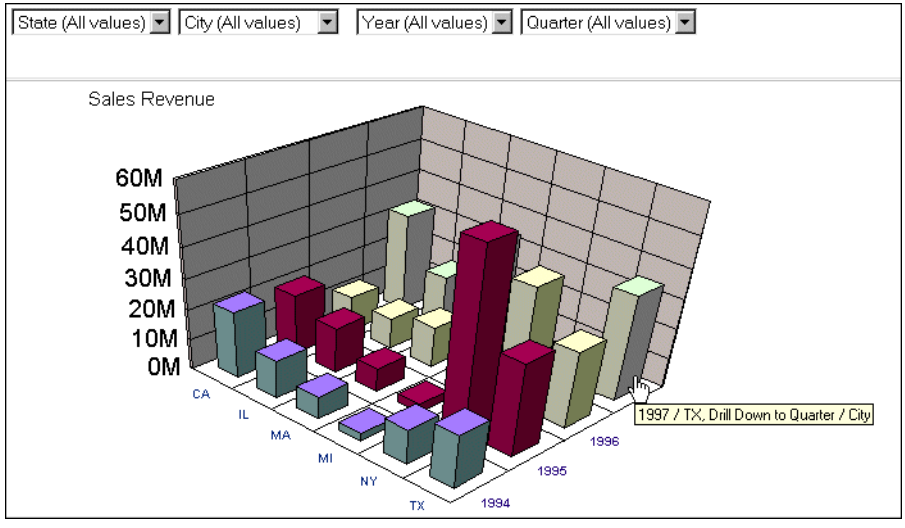

You can then save those documents for your own personal use, send them to other users, publish them to the repository as corporate documents, or schedule them for automatic refresh and distribution through BROADCAST AGENT.

WEBINTELLIGENCE comes in three modules:

- WEBINTELLIGENCE INFOVIEW
  The indispensable web-based entry point for accessing both WEBINTELLIGENCE and BUSINESSOBJECTS documents, and all other types of files that are stored in the repository or in personal document storage areas. INFOVIEW also lets users manage, save, distribute, print and schedule documents for automated processing by BROADCAST AGENT.
  For more information about INFOVIEW, see "Does the system have a single e-Business Intelligence portal?" on page 104.
- WEBINTELLIGENCE REPORTER allows users to create and edit WEBINTELLIGENCE documents.
  The "client" for WEBINTELLIGENCE is a Java applet or an ActiveX control called the Web Panel, which is automatically downloaded for users to build and edit documents.
- WEBINTELLIGENCE EXPLORER allows users to drill in "drillable" WEBINTELLIGENCE documents to analyze data on different levels of detail, as in the example on the previous page.

WEBINTELLIGENCE also comes with the Business Objects Services Administrator for overseeing the entire WEBINTELLIGENCE system.

Broadcast Agent

BROADCAST AGENT allows BUSINESSOBJECTS and WEBINTELLIGENCE users to automatically process and publish their Business Objects documents via the repository, an intranet, extranet, or the World Wide Web.

End users can schedule documents for processing and distribution at times which suit their businesses. This can help reduce network traffic congestion, and enables documents to be printed or refreshed on the web at off-peak times (for example, during the night). End users can even set conditions so that BROADCAST AGENT processes and distributes documents only when events such as increased revenue occur.

Fully integrated with VBA, BROADCAST AGENT also allows users to customize full-client document processing by attaching macros that perform specific tasks such as paging and e-mail distribution.

From an administration point of view, BROADCAST AGENT is part of a distributed solution across a CORBA network. Thus, BROADCAST AGENT offers robust performance, scalability load balancing and failover.

For more information about BROADCAST AGENT, see Chapter 10, "Document Processing and Distribution with Broadcast Agent" on page 343, as well as the *Broadcast Agent Administrator's Guide*.

Zero Admin BusinessObjects

Zero Admin BUSINESSOBJECTS is a new deployment of BUSINESSOBJECTS 5.1. Installed on client machines from WEBINTELLIGENCE via a web browser, Zero Admin BUSINESSOBJECTS is a zero-administration solution:

- There is no configuration on the client side for installation or for middleware.
- There is no need for middleware on the client because the server handles data access.
- A WEBINTELLIGENCE server handles both login and data access.

For more information, see "Zero Admin BusinessObjects" on page 48.

WebIntelligence SDK

WEBINTELLIGENCE SDK (Software Development Kit) is an extension of WEBINTELLIGENCE which enables you to customize the look, behavior, and workflow of a WEBINTELLIGENCE deployment to match the needs of a wide user audience. This type of customization is imperative for any type of extranet deployment, in which your corporate users will be communicating with entire populations of external users through the Web.

From a developer's viewpoint, this means that you have tight control over how your WEBINTELLIGENCE deployment operates and how it appears to end users. From the users' viewpoint, this means they can access documents and perform queries and analysis in an environment with which they are familiar. A common use of WEBINTELLIGENCE SDK is to align WEBINTELLIGENCE deployments with current corporate intranet or extranet solutions.

Customizing the WebIntelligence client

With WEBINTELLIGENCE SDK you can automate many user operations and customize the integration of WEBINTELLIGENCE within your business. For users, this means easy and intuitive access to information.

For example, you can:

- Personalize the user interface.
- Embed prompts in web pages. You can determine which prompts are returned depending on a user's profile and you can determine the results sent back to the user based on their input.
- Add a hyperlink to a document which launches the Web Panel.
- Include ActiveX and other components in web pages. For example, you can display a Microsoft Excel spreadsheet in the same page as a BUSINESSOBJECTS document.

Extending the server

The following are examples of how you can define the means and scope of data access from your WEBINTELLIGENCE server. You can:

- Build a custom web-based document list by editing the existing layout, or by creating a completely new layout. If your business has existing web page templates you can adapt the document catalog to conform to your own styles and standards.
- Personalize the workflow for a group of users and determine the reporting capabilities available, based on user profile.

RDBMS and OLAP modules

In addition to these main server products, there is also a set of RDBMS and OLAP modules.

The server components are linked to each other by the CORBA infrastructure. This means that these components can run on different machines in order to distribute the processing load and improve performance.

Additional server installation options include Data Access modules and OLAP Servers.

Desktop products

Business Objects desktop products include:

- BUSINESSOBJECTS
- BUSINESSQUERY
- BUSINESSMINER
- DESIGNER
- SUPERVISOR
- DEVELOPER SUITE (excluding WEBINTELLIGENCE SDK, which is an Enterprise Server product)
- SET ANALYZER
- The Business Objects Services Administrator
- The Broadcast Agent Console

From the Business Objects Technical Support web site at **http://www.techsupport2.businessobjects.com**, you can also obtain a set of Rapid Deployment Templates, which are pre-defined universes for quick interfacing of Business Objects products with packaged solutions such as SAP and Baan.

Note: These products are only available for Windows environments.

BusinessObjects

BUSINESSOBJECTS is an integrated query, reporting and analysis solution for business professionals that allows you to access the data in your corporate databases directly from your desktop and present and analyze this information in a BUSINESSOBJECTS document.

BUSINESSOBJECTS makes it easy to access this data, because you work with it in business terms that are familiar to you, not technical database terms like SQL. You don't need any knowledge of the database structure or technology.

Once you've used BUSINESSOBJECTS to access the data you need, you can present the information in reports as simple as tables or as sophisticated as dynamic documents with drillable charts.

Example **BusinessObjects document**

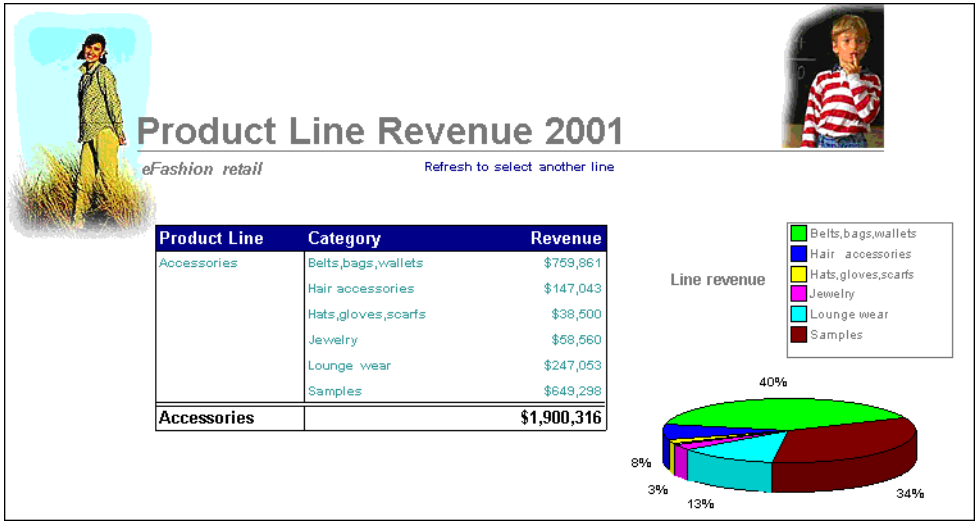

You can then save those documents for your own personal use, send them to other users, publish them to the corporate repository, or send them to BROADCAST AGENT for automatic refresh and distribution.

This product includes three modules:

- BUSINESSOBJECTS INFOVIEW, which connects directly to the repository via the BUSINESSOBJECTS Reader to view, refresh, schedule, distribute and print BUSINESSOBJECTS documents. A full client can also work offline with a local copy of a document.

Tip: Install the full client reader for users who don't need to create documents, but still want to work in off-line mode to view, format and print BUSINESSOBJECTS documents.

- BUSINESSOBJECTS REPORTER
  For reporting and ad hoc queries. REPORTER includes all of the features of Reader, as well as the ability to create queries and documents.
- BUSINESSOBJECTS EXPLORER
  For reporting and analysis

BusinessQuery

BUSINESSQUERY for Excel is an add-in tool that provides Microsoft Excel with fully functional database access. With BUSINESSQUERY, you can access your corporate databases from Excel using familiar business terms. All BUSINESSQUERY commands are available through the BUSINESSQUERY menu and toolbar that appear in Excel. The result is easy and intuitive information access with guaranteed, reliable results.

When you run a query, BUSINESSQUERY automatically places the results into Excel cells. There is no need to copy or export the results to Excel. Nor are the results a static embedded object. Instead, they can be used with the full range of Excel functions, including calculations, charts, and pivot tables.

BUSINESSQUERY provides you with a one-step graphical interface called the Query Panel, that you use to build and run queries using standard Business Objects universes.

BusinessMiner

Data mining is the analysis of data and the use of software techniques for extracting hidden, comprehensible, and useful information from databases. This information can take the form of patterns, trends, or rules that are implicit in data.

BUSINESSMINER works hand-in-hand with BUSINESSOBJECTS. Once you have used BUSINESSOBJECTS to access the information you need from your corporate databases or data marts, you can use BUSINESSMINER to find the trends and patterns hidden in the data.

Designer

DESIGNER allows administrators to create, manage, and distribute universes for particular groups of BUSINESSOBJECTS and WEBINTELLIGENCE users.

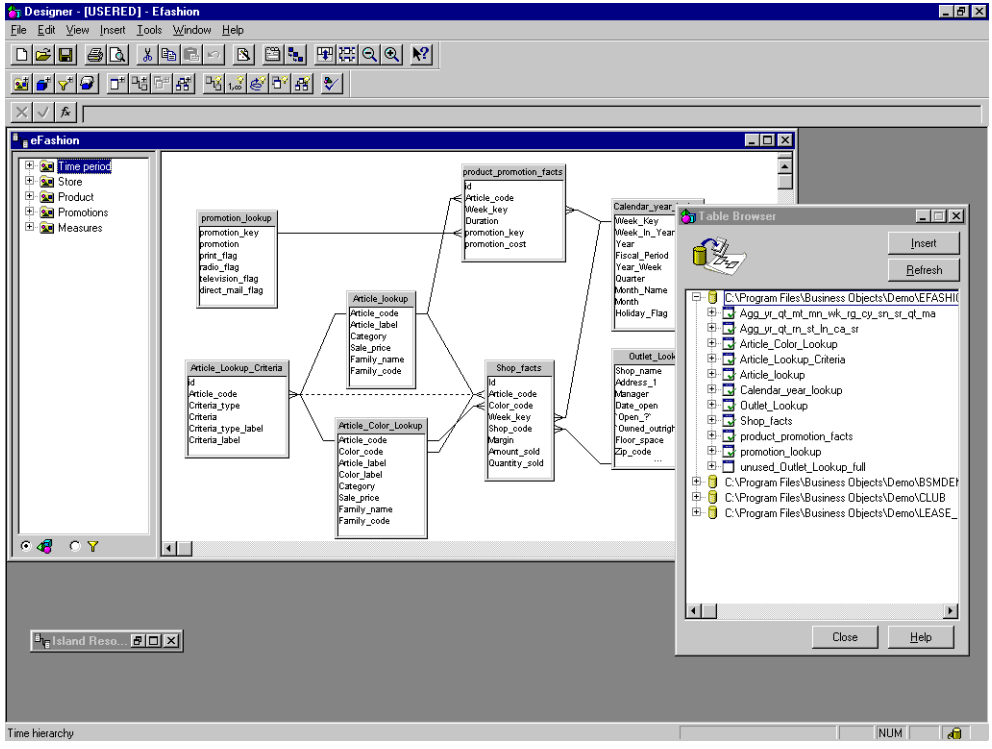

A **universe** is the semantic layer that represents database structures in everyday business terms. Designers distribute universes to users by exporting them to the repository, where end users can access it when they create their documents, or by moving it as a file through the file system.

Supervisor

SUPERVISOR allows you, as supervisor, to set up and maintain a secure environment for the overall Business Objects system. It also provides you with a powerful and easy-to-use structure for distributing information between users. This information is centralized through relational data accounts called repositories.

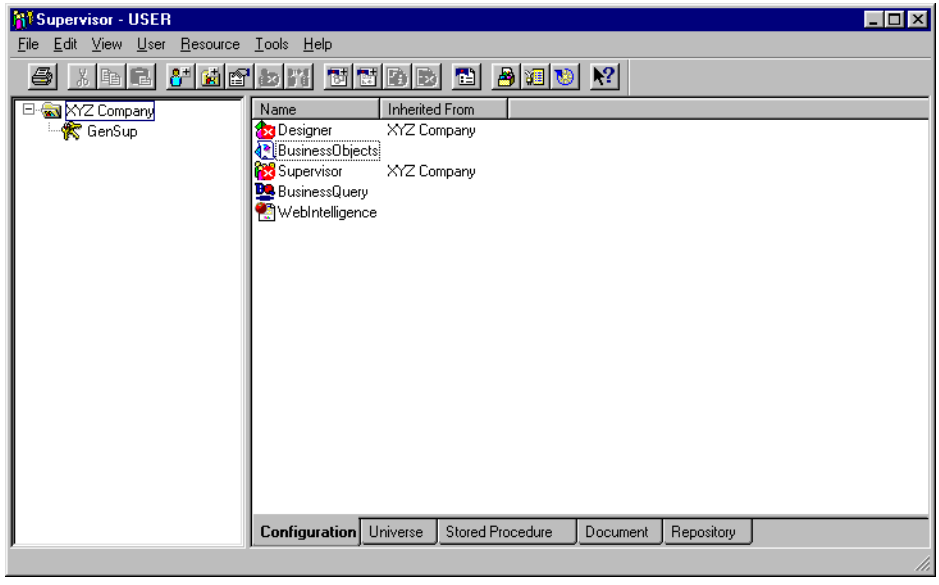

With SUPERVISOR, you define users and user groups, then assign profiles to them. You can control user access to Business Objects products and even the functions they contain, and manage the exchange and distribution of the universes and documents of all users.

SUPERVISOR also lets you manage your company as groups of users with profiles that you determine. User profiles include user identification (user name and password), the products, modules and functions users can work with, the universes they can access, and the documents they can share. Once you have defined these profiles, you can use SUPERVISOR to manage repositories and data accounts.

Developer Suite

The BUSINESSOBJECTS object model provides application developers with tools to customize BUSINESSOBJECTS and DESIGNER. Developers can create internal scripts and OLE automation applications for external applications. These features allow developers to tailor data presentation and retrieval using BUSINESSOBJECTS within their own software environment.

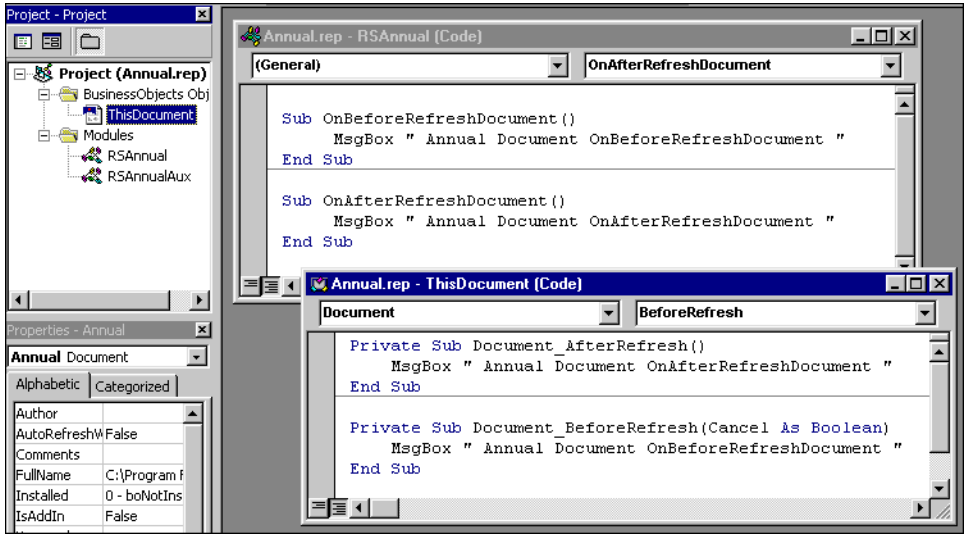

DEVELOPER SUITE provides developers with three key advantages:

- A rich set of business intelligence components
  DEVELOPER SUITE's broad range of customization options run from a slight modification of the BUSINESSOBJECTS interface, to completely integrating BUSINESSOBJECTS or WEBINTELLIGENCE with other applications, such as those used in finance, human resources, sales, marketing and manufacturing.
- A state-of-the-art business intelligence infrastructure
  No matter what kind of customization they're doing, developers will rely on the Business Objects infrastructure. When combined, the centralized repository, powerful semantic layer, and robust security system are the ideal underlying structure for a large business intelligence deployment.

- An easy-to-use and fast development environment
  Develop and deploy customized applications with speed and efficiency, thanks to the use of industry standard technology like Visual Basic for Applications (VBA) for BUSINESSOBJECTS and DESIGNER, as well as Active Server Pages (ASP) and JavaServer Pages (JSP) for WEBINTELLIGENCE. Business Objects components are also available through a standard API.

DEVELOPER SUITE comes with additional documentation containing pre-written code, as well as development examples for both training purposes and accelerating the development process.

Set Analyzer

The bank wants to know the names of all your customers who live in San Jose or San Francisco, who have a mortgage but do not have insurance. How fast, and how easily can you dig that information out of your data warehouse?

Finding meaningful information from ever-increasing amounts of data is one of the biggest challenges in business today. SET ANALYZER extends BUSINESSOBJECTS and WEBINTELLIGENCE functionality, adding more powerful analysis and speeding up queries against very large databases. With SET ANALYZER, you can not only answer your banker's questions in seconds, but you can also tell her what proportion of those customers are male and under 30.

SET ANALYZER allows nontechnical users to query large data stores without the lag time associated with standard relational online analytical processing (ROLAP) tools, thus allowing corporations to make faster and better decisions.

In most ROLAP queries, the number of joins involved can dramatically slow down response time. With SET ANALYZER, you define sets that are actually index tables, which can reduce the number of joins down to as little as one. That means a query using SET ANALYZER can be as much as 100 to 1,000 times faster than before.

SET ANALYZER uses a set-based approach to break down complex business questions into a series of simple steps, giving business users even further autonomy from IT.

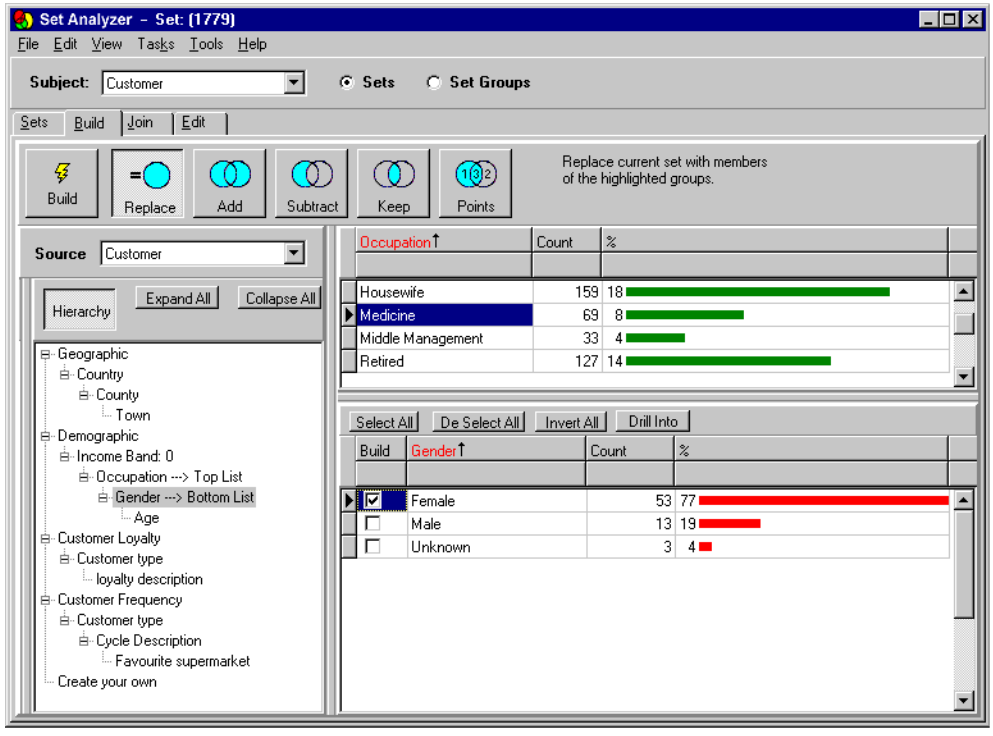

Sets are saved, representing "snap shots" in time of your customers and products.

Set-based analysis is uniquely designed for analysts pursuing solutions to issues such as campaign selections, retention of profitable customers, and cross selling opportunities. The biggest productivity gains with SET ANALYZER will be in sales and marketing implementations. The tool includes several features specifically aimed at improving data analysis, including customizable hierarchies for maintaining the context of an evolving query, and shared set definitions -- eliminating the need to recreate previously used sets.

SET ANALYZER integrates with BUSINESSOBJECTS and WEBINTELLIGENCE at the universe level. Integration with BUSINESSOBJECTS means that all your users can read and refresh production documents with integrated sets.

The Business Objects Services Administrator

The Business Objects Services Administrator gives you an overview of your distributed architecture system:

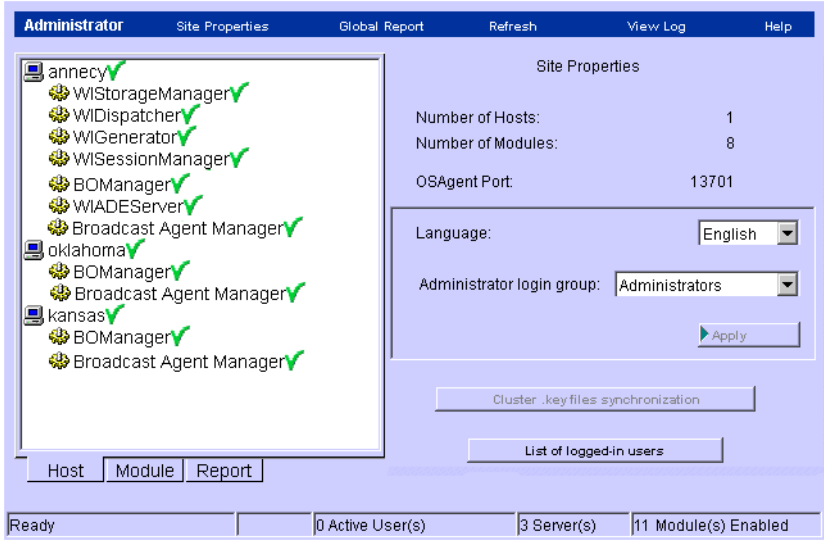

It lets you monitor and control your system's servers and processes, and trace various types of activity in the system. For more detailed information, see "The Business Objects Services Administrator" on page 96.

The Administrator is also automatically available in applet form when you install WEBINTELLIGENCE.

The Broadcast Agent Console

The Broadcast Agent Console allows administrators to track and modify the scheduled automatic processing of documents through BROADCAST AGENT.

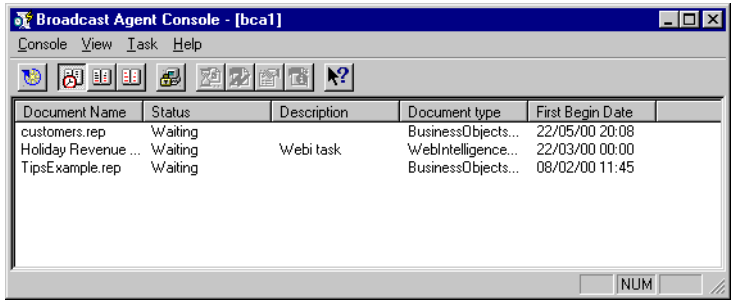

The Console is also available from BUSINESSOBJECTS so BUSINESSOBJECTS users can monitor the documents they have scheduled with BROADCAST AGENT. It can also be installed and used on any Windows machine on the same sub-net as the server cluster.

Chapter 2 What's New in the Latest Release?

In this chapter

- **Overview 48**
- **Zero Admin BusinessObjects 48**
  Working offline 49
  What's the difference between BusinessObjects and Zero Administration BusinessObjects? 50
  Installing the Zero Admin BusinessObjects server software 52
  Installing the client software 53
  Automatically updating versions 54
  Three-tier architecture 55
  Is multiple security domain selection supported? 55
  What about web security? 55

- **Enhanced Audit and Tracing Facilities 56**
- **CORBA Externalization 56**
- **Scheduling BusinessObjects Documents from InfoView 57**
- **Multiple WIGenerators on a Single Server 57**
- **Administration Tools on Windows Platforms 57**
  Business Objects Services Administrator 58
  Broadcast Agent Console 59
- **Managing Default Options for New Users 59**
- **Document Agnosticism 60**
  Uploading documents from InfoView 61

- **Saving BusinessObjects Documents as PDF Files 61**
- **Adding Hyperlinks to BusinessObjects Documents 61**
- **BusinessObjects OLAP Connect 62**
- **The PDF File Cache 63**
- **My InfoView 64**
- **Choosing the Platform for Document Processing 65**
- **Scheduling WebIntelligence Documents With Prompts 65**
- **Improved Date and Time Handling 66**
- **Broadcast Agent in UNIX-Only Clusters 66**
- **Importing/Exporting Universes from SUPERVISOR 67**
- **Printing Repository User and Resource Data to Table Format 67**
- **Purging Inbox Documents 67**
- **New WebIntelligence SDK Object Models 68**
- **Broadcast Agent Publisher 68**
- **What's New in Documentation 68**
  InfoView Quick Tour 69
  Accessing PDFs directly from the InfoView/WebIntelligence Help 70
  Getting Started with WebIntelligence 71
- **Multihomed Configuration Support 72**

Overview

This release introduces several new features enhancing previous releases. The most important of these are introduced in this chapter.

Zero Admin BusinessObjects

This release comes with an important new extension to the functionality of BUSINESSOBJECTS: a zero-administration deployment of BUSINESSOBJECTS.

Zero Admin BUSINESSOBJECTS includes a version of BUSINESSOBJECTS that runs on a Windows PC as a standalone Windows application, but is installed and automatically updated via an Internet browser connected to WEBINTELLIGENCE. Only the minimum required software is installed on the client PC, all middleware remains on the server, and security and access is handled through SUPERVISOR. This means no client-side administration is required.

Once Zero Admin BUSINESSOBJECTS is installed, users can create queries, build, format and analyze documents using the powerful features of BUSINESSOBJECTS 5.1. They can then save the documents locally or in personal storage space on the WEBINTELLIGENCE server or publish or schedule documents to share with other users in the company.

If a zero-administration deployment of BUSINESSOBJECTS 5.1 is set up in the company, BUSINESSOBJECTS users who have installed the full version of BUSINESSOBJECTS 5.1 from a CD or Master installation can also take advantage of middleware stored on the server to build queries and refresh documents. When users log on to BUSINESSOBJECTS, they just have to choose the web portal login.

Working offline

Once BUSINESSOBJECTS 5.1 users have retrieved the data they need, they can save documents locally, then if they have the appropriate rights, choose to work in offline mode, i.e. without a connection to the WEBINTELLIGENCE server.

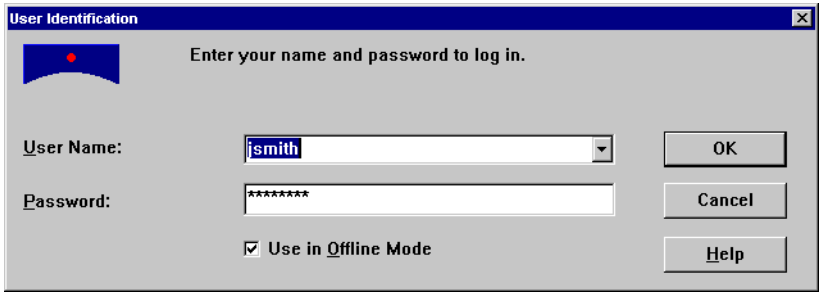

In this case, users indicate they are going to work in offline mode when they log into the product. To authenticate the user, BUSINESSOBJECTS re-uses the information cached in the local .lsi (Local Security Information) file. This file contains the user's security information, profile, command restrictions and authorized documents and universes.

Once they're logged in, users can work on documents stored locally, reviewing and analyzing their data. If users create a new document or refresh an existing one, however, BUSINESSOBJECTS opens a new online session on the WEBINTELLIGENCE server, using the ID and password with which the user originally logged in.

Being able to work offline saves system resources for only the most essential processing.

Note: These BUSINESSOBJECTS users cannot, however, login offline using NT Challenge/Response or Basic Authentication. For more information about authentication, see the *WebIntelligence System Administrator's Guide*.

What's the difference between BusinessObjects and Zero Administration BusinessObjects?

BUSINESSOBJECTS 5.1 and Zero Admin BUSINESSOBJECTS are the same product.

In this and all Business Objects manuals, Zero Admin BUSINESSOBJECTS refers to the server side components of the web deployment of BUSINESSOBJECTS 5.1.

On the client side, the functionality available to you as an end user differs slightly depending on whether you installed BUSINESSOBJECTS 5.1 from the CD, or whether you installed it via a web browser, however, because security and logins are handled in a slightly different way.

- If you install BUSINESSOBJECTS 5.1 from the CD, you will be able to work with both a client/server connection and a web connection to access and refresh data.
  If a zero-administration deployment of BUSINESSOBJECTS 5.1 is set up in the company, BUSINESSOBJECTS users who have installed the full version of BUSINESSOBJECTS 5.1 from a CD or Master installation can also take advantage of middleware stored on the server to build queries and refresh documents. When users log on to BUSINESSOBJECTS, they just have to choose the web portal login. For more information, see the *BusinessObjects User's Guide*.
- If you install BUSINESSOBJECTS 5.1 through a web browser, you can neither access data using stored procedures, free-hand SQL, and OLAP multidimensional servers, nor use data from SAP BW and SAP BAPI applications. On the other hand, you can carry out every action on BUSINESSOBJECTS documents that you can from the traditional desktop product (including document creation, drilling, Slice and Dice, prompts, etc.), and you enjoy the advantages of a true zero-administration solution.

The following table sums up the main similarities and differences:

Feature	BusinessObjects 5.1 installed from a CD	BusinessObjects 5.1 installed via a web browser
Automatically detect and install a new version of the software over the web	No	Yes
Access data using stored procedures	Yes	No
Access data using free-hand SQL	Yes	No
Access data using OLAP multidimensional servers	Yes	No
Access data using Visual Basic for Applications (VBA) procedures	Yes	Yes
Use data from SAP BAPI applications	Yes	No
Access data from personal data files (MS Excel, for example)	Yes	Yes
Access data from queries built on BUSINESSOBJECTS universes	Yes	Yes
Link data from different data sources in one document	Yes	Yes

Feature	BusinessObjects 5.1 installed from a CD	BusinessObjects 5.1 installed via a web browser
Display multiple blocks on a document page (e.g. a mixture of tables, crosstabs and charts)	Yes	Yes
Use VBA macros and add-ins contained in BUSINESSOBJECTS documents	Yes	Yes
Use BUSINESSMINER, the optional desktop data mining product	Yes	No
Send documents to BROADCAST AGENT for automatic scheduling and processing	Yes	Yes

Installing the Zero Admin BusinessObjects server software

To install the server software for the zero-administration deployment of BUSINESSOBJECTS, you must:

- Install the Zero Admin BUSINESSOBJECTS package. This installs on the server the components needed to install Zero Admin BUSINESSOBJECTS via an Internet browser.
- If it isn't installed through the WEBINTELLIGENCE system already, install the required middleware to connect to the data warehouses containing repository and corporate data.

Zero Admin BUSINESSOBJECTS uses the connectivity features of WIQT processes to handle and process BUSINESSOBJECTS documents. It does this by channeling all transactions through the WEBINTELLIGENCE system's WIADE (WEBINTELLIGENCE Application Development Environment) Server.

Can you install Zero Admin BusinessObjects, WebIntelligence and Broadcast Agent on the same machine?

Yes. You can install the Zero Admin BUSINESSOBJECTS, WEBINTELLIGENCE, and BROADCAST AGENT on the same machine. However, in this case, although you can download Zero Admin BUSINESSOBJECTS to client machines, downloading it to the server machine is not supported.

Installing the client software

You install the zero-administration version of the BUSINESSOBJECTS 5.1 client software via an INFOVIEW session in an Internet browser. There are, however, some constraints:

- You must be using a Windows machine.
- You can only install BUSINESSOBJECTS from Internet Explorer 3.x, or Netscape 4.x or above.
- You cannot install BUSINESSOBJECTS via an Internet browser if you have already installed any of the Business Objects 5.x desktop products from the CD. If you want to install BUSINESSOBJECTS via an Internet browser, you must first uninstall all other Business Objects desktop products.
- Make sure that the browser security settings required to download signed ActiveX controls and plug-ins are enabled.
- The first person to download BUSINESSOBJECTS through a web browser on a Windows NT4 workstation must be an NT Administrator.
- The first person to download BUSINESSOBJECTS through a web browser on a Windows 2000 workstation must be an Administrator and must first run the AllowBO51-26-W2000.exe file. For more information, refer to, “Installing Business Objects Desktop Products,” on page 131.

- At the end of the BUSINESSOBJECTS installation, users are prompted to install Visual Basic for Applications (VBA).

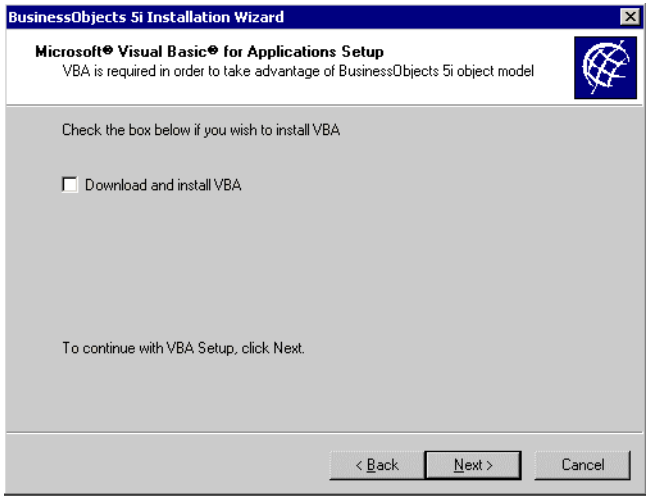

If they have the rights to install VBA from web installer (SUPERVISOR setting) and the correct version has not already been installed, they should do so, otherwise they won't be able to run the macros and add-ins that can be included in BUSINESSOBJECTS documents. If the Installer notes the presence of this version of VBA in the Windows Registry, no prompt appears.

Automatically updating versions

You can automatically update your version of BUSINESSOBJECTS 5.1 only if you installed it entirely through a web browser. As soon as you add a component to BUSINESSOBJECTS via the CD (e.g. add OLAP access), the installation is transformed into a full BUSINESSOBJECTS installation and you will no longer be able to update automatically.

If you installed the product using the Business Objects installation CD, you can update your software from a CD.

Three-tier architecture

Like WEBINTELLIGENCE, Zero Admin BUSINESSOBJECTS has a three-tier architecture:

- Tier 1: the client
  The BUSINESSOBJECTS software is installed on the client machine.
- Tier 2: the Zero Admin BUSINESSOBJECTS server
  BUSINESSOBJECTS 5.1 uses the HTTP processing layer in the WEBINTELLIGENCE system to communicate with the HTTP server. The WEBINTELLIGENCE server manages security and connectivity information and downloads the required information to the client the first time he connects. BUSINESSOBJECTS uses the WIQT for the user session and to communicate with the repository and database.
- Tier 3: data sources

Is multiple security domain selection supported?

Multiple security domain selection is supported for the zero administration deployment of BUSINESSOBJECTS only if you launch BUSINESSOBJECTS from the Windows Start menu. This allows you to choose a remote key file (with an .rkey extension) which points to different portals.

If you have installed BUSINESSOBJECTS from the CD, you can choose either a .rkey or a local .key file found on the client machine when you launch the product using the Start menu.

We do not support multiple .key file selection for WEBINTELLIGENCE and BUSINESSOBJECTS launched by WEBINTELLIGENCE, because installation through a web browser doesn't provide the local SQL BO required to run online against key files.

What about web security?

Like WEBINTELLIGENCE, Zero Admin BUSINESSOBJECTS works with web security standards such as the Secure Socket Layer (SSL), which protects data transferred over the network from unauthorized access through the use of advanced encryption algorithms.

For detailed information, see the *WebIntelligence System Administrator's Guide*.

Enhanced Audit and Tracing Facilities

The Business Objects Services Administrator provides trace facilities that let you monitor the requests and processes passing through the WEBINTELLIGENCE system. This facility now also traces these activities for Zero Admin BUSINESSOBJECTS. The information from this trace can be saved to a log file or to a database.

You can now gather even more precise information using module log files, which trace all the activity for a specific WEBINTELLIGENCE system module on a specific server machine.

CORBA Externalization

Business Objects server products are based on three main CORBA components manufactured by Inprise:

- The OSAgent
- The OAD (Object Activation Daemon)
- The Gatekeeper

They are known as the system's **kernel components**.

With this release, you can modify the command lines used by these three components at run-time. This allows you more flexibility with your system settings, allowing you for example to:

- Give the OAD more time to react to object requests
- Configure your system to use the double-firewall configuration known as a **DMZ** (DeMilitarized Zone)
- Modify the port used by the system's Gatekeeper.

Scheduling BusinessObjects Documents from InfoView

WEBINTELLIGENCE INFOVIEW users can now schedule BUSINESSOBJECTS documents for automatic refresh and distribution.

Multiple WIGenerators on a Single Server

Enabling multiple instances of WIGenerator on a single server machine serves two purposes:

- It allows you to prevent a single WIGenerator from becoming a bottleneck when several WIQT processes have been launched concurrently.
- It prevents UNIX server mutex management problems that may occur on machines with more than 16 processors.

  Mutexes (multiple exclusion objects) are program objects that are created so that multiple program threads can take turns sharing the same resource. Due to resource management constraints under Solaris machines with more than 16 processors, for example, a single WIGenerator can spend more time managing the queue of jobs awaiting processing than actually processing requests. The result is that for the same number of active, concurrent users, you get better performance with a 16-processor machine than with a 24-processor machine, and this therefore blocks the scalability of this node. By multiplying the number of instances of the WIGenerator, you provide more WIGenerators to share the transaction load and therefore minimize this server problem.

Administration Tools on Windows Platforms

The Business Objects Services Administrator and the Broadcast Agent Console can now be installed as a Windows application on client machines as well as on server machines as a Java applet. This allows WEBINTELLIGENCE system administrators to monitor and control WEBINTELLIGENCE and BROADCAST AGENT tasks directly from Windows client machines. Supported Windows operating systems include Windows 95, 98, NT 3.51, NT4 or 2000 workstation, or NT4 Server with no server installation.

Business Objects Services Administrator

The Business Objects Services Administrator is a utility that lets you monitor and control all the servers and modules in your WEBINTELLIGENCE system.

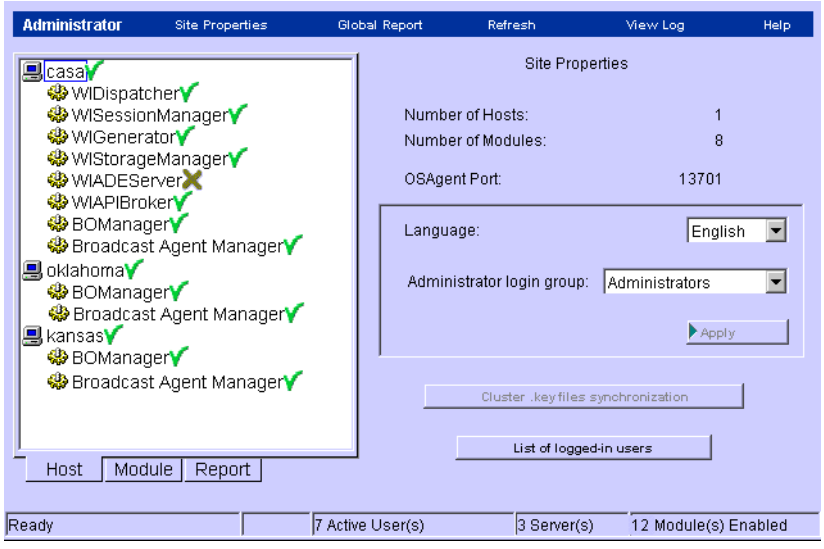

It is now available as:

- A Java applet that runs in a web browser
  The applet is automatically included when you install WEBINTELLIGENCE. It requires an HTTP server.
- A standalone Windows application
  The application is especially useful in BUSINESSOBJECTS deployments with BROADCAST AGENT which do not include WEBINTELLIGENCE and do not use an HTTP server. In this case, the Administrator is installed as a Desktop product on a Windows client machine.
  This installs all the required CORBA components it needs to connect with the BROADCAST AGENT cluster. During installation, you must also provide both the OSAgent Port being used by the cluster, and the IP address of the cluster manager. For more information, see the *Installation and Update Guide*.

Broadcast Agent Console

The **Broadcast Agent Console** is a desktop program that allows administrators to monitor the status of all processed and pending tasks, and to execute actions such as killing a task, changing its distribution list, or running a task immediately.

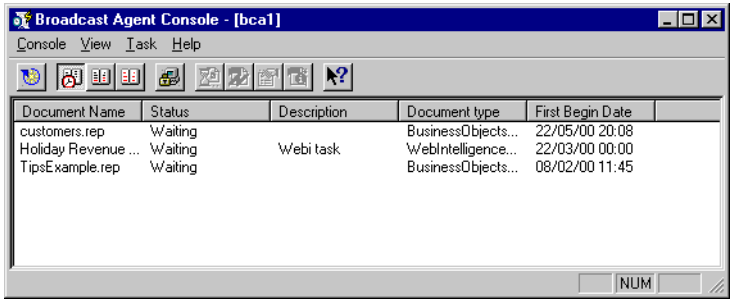

The Console has direct access to both the Scheduler and the repository. This component comes with the full-client BUSINESSOBJECTS product. With the 5.1/2.7 release, it can be installed on any client machine on the subnet.

Managing Default Options for New Users

End users can customize the look, feel and functionality of INFOVIEW and WEBINTELLIGENCE by setting options in the INFOVIEW Personal Options page. They can, for example, define the page that opens when they launch INFOVIEW, and what that page looks like, the type of Web Panel they use for creating and modifying WEBINTELLIGENCE documents, and the type of viewer they use to view BUSINESSOBJECTS documents.

WEBINTELLIGENCE system administrators must have General Supervisor profiles. With this version, they have the right not only to define the default Personal Options page settings for all new WEBINTELLIGENCE/INFOVIEW users, but to use SUPERVISOR to prevent those users from subsequently changing those settings if they wish.

Document Agnosticism

The Business Objects repository now lets WEBINTELLIGENCE INFOVIEW users with the required rights view, send, publish, download and upload not only WEBINTELLIGENCE and BUSINESSOBJECTS documents, but any other type of file, such as Microsoft Word documents, Microsoft Excel spreadsheets, or Adobe Acrobat PDF files.

Icons indicating type of file

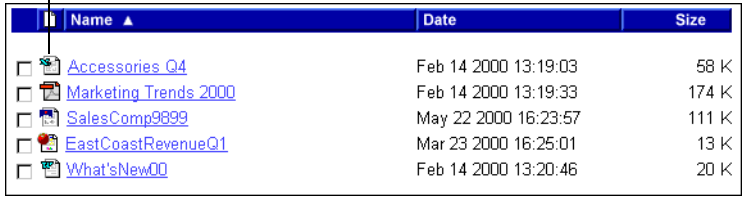

Users can also store all file formats as personal documents.

To view a third-party file, users just click the file name in the INFOVIEW document list. If the relevant application is installed on the user's machine, INFOVIEW automatically launches the application with the file open in it. If the application isn't installed, users can choose to download the file to their machines.

This feature lets users distribute and receive files of all kinds while taking advantage of Business Objects security, controlled via SUPERVISOR.

Uploading documents from InfoView

The Business Objects repository and WEBINTELLIGENCE system can now store documents of all file formats, for example Microsoft Word, Microsoft Excel or Adobe Acrobat PDF files, not just BUSINESSOBJECTS and WEBINTELLIGENCE documents. From INFOVIEW it is now possible to upload documents of any file format except BUSINESSOBJECTS and WEBINTELLIGENCE documents (these must be published or scheduled normally in order to maintain security). Upload options differ according to the documents page:

- From the Corporate Documents page, you can upload documents to publish in the Corporate Documents list.
- From the Personal Documents page, you can upload documents to your Personal Documents list.
- From the Inbox Documents page, you can upload documents to send to other users or groups.

Saving BusinessObjects Documents as PDF Files

BUSINESSOBJECTS 5.1 lets you save BUSINESSOBJECTS documents in Adobe Portable Document Format (PDF), an industry standard for viewing and printing. PDF documents let you view and navigate through the reports in a BUSINESSOBJECTS document and obtain high-quality printed copies. During processing, BROADCAST AGENT can save documents in PDF format before distributing them.

All users need to view PDF documents is the Adobe Acrobat Reader, available free from Adobe's website.

Adding Hyperlinks to BusinessObjects Documents

You can now add hyperlinks such as mail addresses and web site addresses to your documents. BUSINESSOBJECTS automatically recognizes keywords for several different protocols including http, to connect to web sites, electronic mail addresses, FTP and more. These hyperlinks are maintained when you save the document in HTML or PDF format.

BusinessObjects OLAP Connect

BUSINESSOBJECTS OLAP Connect brings users of BUSINESSOBJECTS improved OLAP connectivity for Microsoft SQL Server Services and SAP Business Information Warehouse (BW). Users access MS OLAP and SAP BW OLAP data sources directly via the pop-up OLAP Panel within BUSINESSOBJECTS.

Using the OLAP Panel, users select the data they need using buttons that represent the dimensions of the OLAP data source.

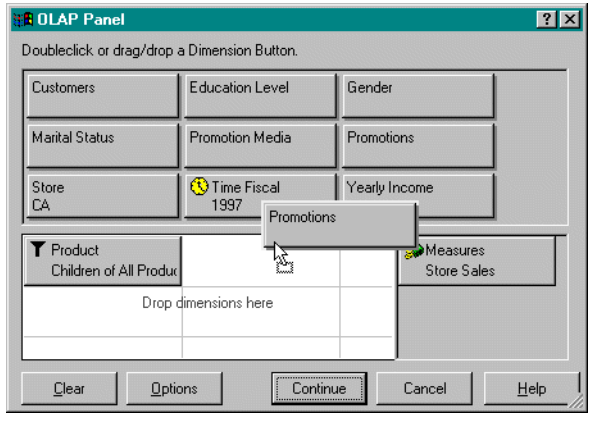

A graphical family tree displays the hierarchies and members of each dimension for easy navigation and data selection. BUSINESSOBJECTS OLAP Connect also includes filters, time series analysis, and a powerful calculation wizard.

Once they generate the BUSINESSOBJECTS report, users are brought back to the familiar reporting and analysis environment of BUSINESSOBJECTS.

The PDF File Cache

When you view a BUSINESSOBJECTS corporate document from INFOVIEW, you can choose to have it first converted to PDF format.

If you choose this option (in your INFOVIEW Personal Options page), WEBINTELLIGENCE will convert the document to PDF and store it in a temporary cache, where it will remain available for you and for other users. This caching of retrieved PDF files saves time and system resources, since the WEBINTELLIGENCE system only needs to generate the PDF file once - the first time it is accessed. This can greatly speed up the retrieval process, as it can take a considerable amount of time to generate PDF files from large files which contain many graphics.

If the source document is modified in the meantime, then the next time the document is downloaded, the PDF will be regenerated again.

My InfoView

My INFOVIEW is a customizable INFOVIEW start page. It lets users divide the browser window into several panes in which to view the information they need.

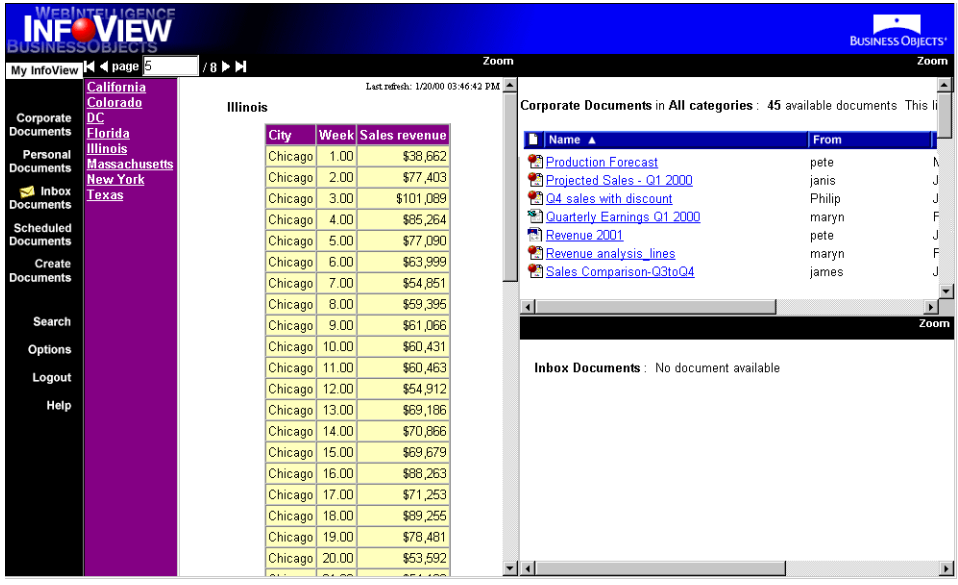

This can include document lists, specific documents, or favorite web pages.

Choosing the Platform for Document Processing

BOManager can now run on either Windows or UNIX machines. If you are using a heterogeneous cluster, you can therefore choose what type of server you want to process interactive BUSINESSOBJECTS document refreshes.

Note: The term "interactive refresh" refers to the specific request for the refresh of a BUSINESSOBJECTS document from an INFOVIEW user.

- The new *Enable Batch Mode* flag in the BOManager parameters allows individual server nodes to be prevented from running BROADCAST AGENT tasks. Setting this parameter to No effectively reserves these machines for processing requests from logged-in users.
- The new *Enable Interactive Mode* flag in the BOManager parameters allows individual server nodes to be prevented from running non-BROADCAST AGENT tasks (i.e., tasks run by interactive users). Setting this parameter to No effectively reserves these machines for use by BROADCAST AGENT tasks.
- The user or the administrator can use the Console to specify whether a scheduled task should be run on a UNIX server or on a Windows server.

Scheduling WebIntelligence Documents With Prompts

INFOVIEW now lets you refresh the prompts contained in a WEBINTELLIGENCE document before sending it to BROADCAST AGENT for scheduled refresh. Previously, BROADCAST AGENT would automatically use the prompts that were used when the document was last refreshed. Now you have the option to choose new prompt values in the Scheduling Options window.

Improved Date and Time Handling

In previous versions of the Business Objects product line, the dates and times displayed for documents which had been scheduled and automatically distributed via BROADCAST AGENT corresponded to the server through which the documents had been processed. Now, all dates and times are shown relative to the client machine which is viewing the document.

For example, a user in New York with a server in Paris will see only New York time, and will not have to convert to Paris time. Scheduled documents will be processed at the correct time without the user needing to know the location of the server.

Broadcast Agent in UNIX-Only Clusters

BROADCAST AGENT can now run on a UNIX-only cluster, with no Windows server nodes. However, you may still require a Windows client machine to use the SUPERVISOR and DESIGNER modules to help set up your system.

The BOManager module (and its associated BUSINESSOBJECTS process) now runs under UNIX as well as under Windows, thus permitting UNIX systems to refresh BUSINESSOBJECTS documents as well as WEBINTELLIGENCE documents.

You must, however, have access to a Business Objects Services Administrator, indispensable for administering the overall system. In this case, the Administrator is installed as a Desktop product on a Windows client machine.

In order to process documents containing VBA macros, however, you still need a Windows node in the custer.

Importing/Exporting Universes from SUPERVISOR

With version 5.0, you could import and export universes from SUPERVISOR as long as the universe was created with a secure connection stored in the target repository. In effect, you could use SUPERVISOR to move universes from one universe domain to another in the same repository.

BUSINESSOBJECTS 5.1 lets you choose the connection with which to store the universe when you export it. This means you can export universes to a repository which were created without reference to that repository: from one repository to another, for example.

This feature lets you fully manage universes in a production environment that does not include DESIGNER.

Printing Repository User and Resource Data to Table Format

SUPERVISOR now lets you print full information on the definition and properties of users and resources to comma-separated values (CSV) files which can be used as a data source to create a BUSINESSOBJECTS document. You can then format the document to optimize the way the data is presented.

Purging Inbox Documents

Inbox documents are documents sent from user to user via the repository, either directly or by means of BROADCAST AGENT, and cached there until their recipients download them. SUPERVISOR now lets you purge all Inbox documents older than a given number of days from the repository with a single command.

This lets you recover the document domain space taken up by documents which have been sent but not yet collected by users. This feature can be especially useful when regularly scheduled documents are sent using the *Refresh according to the profile of each recipient* option.

New WebIntelligence SDK Object Models

WEBINTELLIGENCE SDK now contains two identical object models, WICOM, the COM interface and WIBean, the Java interface, that allow you to customize WEBINTELLIGENCE functionality.

- You can use WICOM to write ASP scripts.
- You can use WIBean to write JSP scripts.

Both interfaces provide equal programming capability.

Each model exposes the same WEBINTELLIGENCE functionality for different environments and allows you to rebuild a customized version of the WEBINTELLIGENCE interface and workflows to maximize the value of your e-business intelligence solution.

Broadcast Agent Publisher

BROADCAST AGENT PUBLISHER is an integrated package of components that empowers BROADCAST AGENT users to create publications using a range of powerful subscription-based email distribution and report bursting capabilities. For more information on BROADCAST AGENT PUBLISHER, see the BROADCAST AGENT PUBLISHER Administrator's Guide and the BROADCAST AGENT PUBLISHER User's Guide.

What's New in Documentation

The documentation for the 5.1/2.7 release is more substantive, complete, and of higher quality than ever before. This section describes just a few new documentation initiatives.

InfoView Quick Tour

The INFOVIEW Quick Tour is an animated, multimedia presentation of all the basic procedures you need to master to successfully get started with INFOVIEW. It shows you how you can access the documents in your company's database, view different types of documents, and share your documents with other users using a web browser.

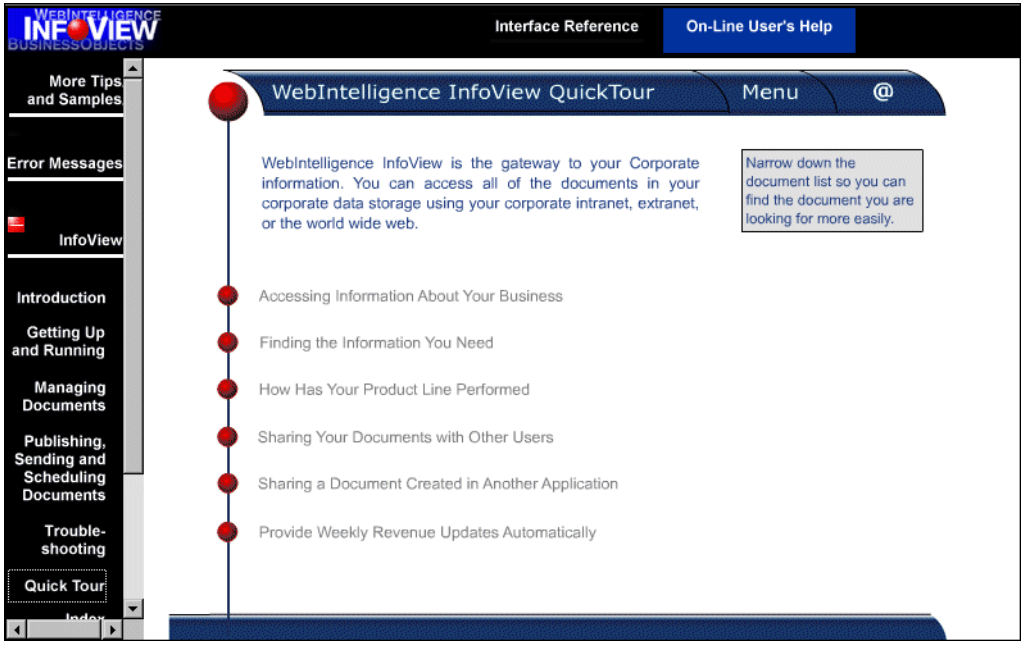

Accessing PDFs directly from the InfoView/WebIntelligence Help

You can now access Business Objects manuals in Adobe Acrobat PDF format directly from the On-line Help for either INFOVIEW/WEBINTELLIGENCE or the Business Objects Services Administrator.

Simply click the Product Documentation link on the Help home page.

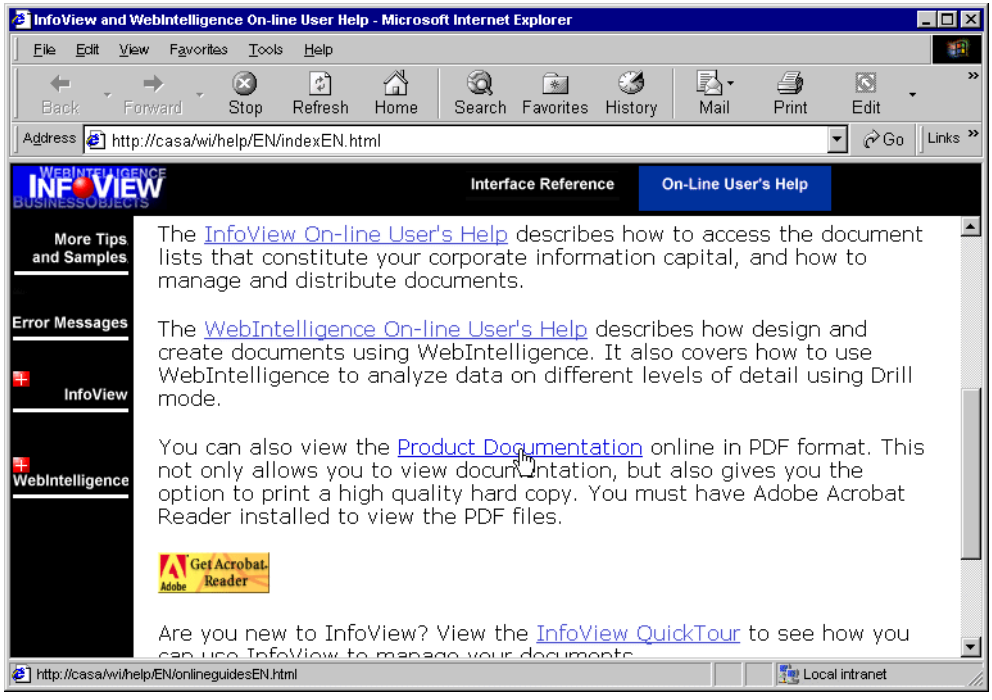

Direct links to all manuals pertaining to use of INFOVIEW or WEBINTELLIGENCE, or administration of the WEBINTELLIGENCE system are immediately displayed.

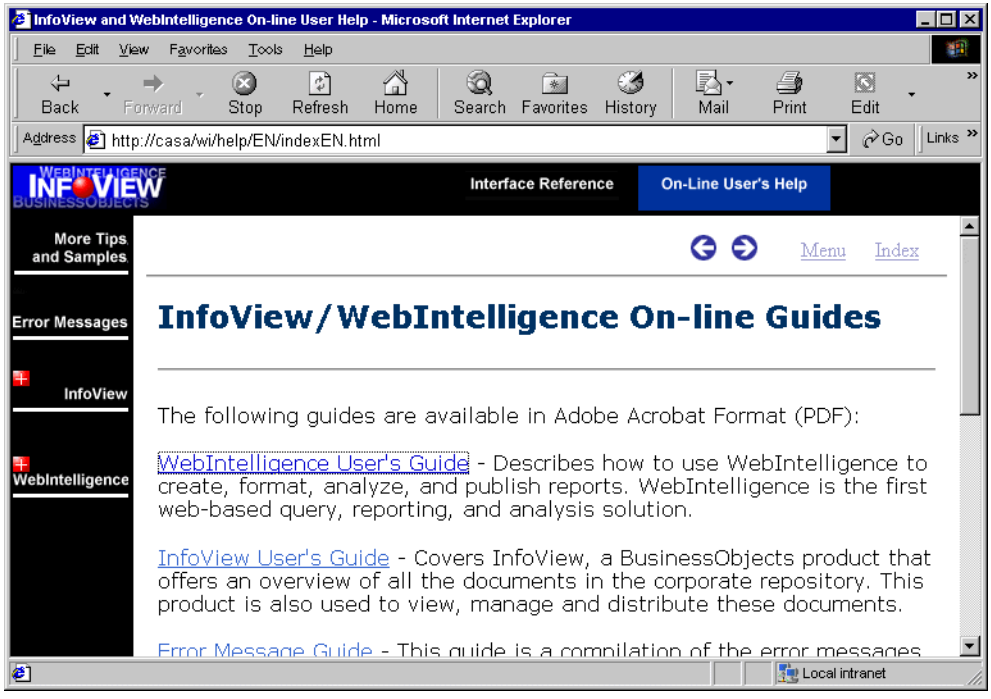

Getting Started with WebIntelligence

Getting Started with WEBINTELLIGENCE is a new step-by-step introduction to creating, formatting, analyzing and publishing documents with WEBINTELLIGENCE. Available in printed or Adobe Acrobat PDF format, this tutorial is the perfect place to begin if you're about to use WEBINTELLIGENCE for the first time.

You can access it in PDF format from the link in the INFOVIEW/WEBINTELLIGENCE On-line Help home page.

Multihomed Configuration Support

With Service Pack 2a, WEBINTELLIGENCE servers under UNIX can use multiple Network Interface Cards (NICs), allowing them to connect to more than one network via multiple physical IP addresses. This is known as a **multihomed configuration**.

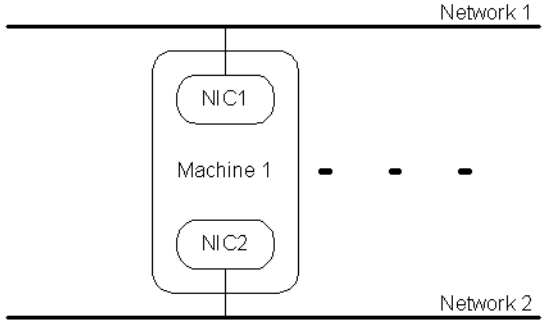

This configuration is useful for many reasons. For example:

- One high-speed network may be highly secured and subject to rigid rules on a firewall, allowing only traffic between the server and an RDBMS. Another network could then be used for other WEBINTELLIGENCE transactions. This not only lightens the transaction load on any particular network or NIC, but allows greater flexibility in adapting networks to specific tasks.
- In some cases, one network is used to back up records of the transactions on another network.
- In some architectures, the nodes in a single server machine communicate via one NIC. In this case, communication with the exterior would have to pass through another NIC.

Implementing a multihomed configuration for WEBINTELLIGENCE is simple. During installation, the installer detects the IP addresses currently assigned to the machine, and asks you specify the address you want to use for the WEBINTELLIGENCE server. You then simply select one or specify a new address.

Chapter 3 Planning Your Deployment

In this chapter

- **Overview 74**
- **Pre-Deployment Checklist 74**
  - Education: providing the tools to be successful 75
  - Data management: making data accessible 76
  - Users: know their numbers and profiles 77
  - Document processing: what documents, and how complex? 77
  - Deployment type: client-server or distributed architecture? 78
  - Server sizing: efficiently supporting peak user activity 78
  - Disaster recovery: what can be done about it? 79
  - Load balancing: matching system processes and server resources 79
- **Planning Your Pilot Project 80**
  - Bring in Business Objects expertise 80
  - Start small 81
  - Think big 82
  - Expand, trace, volume test, benchmark 82
- **Conclusion 83**

Overview

The most efficient way to deploy the Business Objects solution varies from site to site. What is true for all implementations, however, is that planning your deployment carefully, evaluating the needs of your organization and users, and designing a realistic, step-by-step implementation process involving as many representative groups in your organization as possible are critical.

Getting the process right is the key to first-time success. Whether you're a new or an existing Business Objects administrator, you need to consider a number of human and technical issues before you can effectively plan your deployment. This chapter lists the most important deployment issues and tells you where you can turn in the Business Objects documentation for more detailed information on each.

Pre-Deployment Checklist

Before you actually install the Business Objects software and set up your system, you should consider these fundamental deployment issues:

- What do you expect from system performance?
- How can you acquire the technical expertise to deploy Business Objects products, and how will you educate end users?
- How can you make data accessible to everyone?
- Who are your users?
- What types of documents and other files will be processed?
- What types and configurations of servers are needed?
- How can you plan for disaster recovery?
- How can you distribute transaction loads across the distributed system to optimize performance?

In *any* deployment environment, users will expect:

- That the system respond rapidly (login/open a standard document/creating a new document)
- That the system be easy to use and intuitive (finding a standard document/creating a new document)

These expectations should drive the implementation.

Education: providing the tools to be successful

Any deployment requires particular baseline knowledge to be successful.

The development team requires a collection of skills:

- Data warehouse expertise
- Windows and/or UNIX administration expertise
- Generic Business Objects product and deployment knowledge

Don't forget the users, however. No matter how technically fine-tuned a Business Objects deployment is, if the users are ignorant, or even hesitant about using the system, the implementation will take longer and be more costly. Involve user input from the beginning, and give users the knowledge they need to use the system comfortably and efficiently. The more they get out of using the system, the more they'll use it.

Your deployment will probably involve two distinct sets of users:

- The minority will be a small set of "power users," who actually create, edit, and distribute the bulk of the system's documents. These users will require WEBINTELLIGENCE and/or BUSINESSOBJECTS end user skills. To get them up and running most rapidly, training may be required.
- The great majority of users, however, will probably be using the thin-client and full-client INFOVIEW module to view and refresh existing documents. These users will require less in-depth training. A good start would be having them use the multimedia tutorial called the *Quick Tour*, and the *Getting Started with Business Objects* manual.

Data management: making data accessible

Typically, a common goal with a distributed architecture deployment is for everyone in the organization to gain access to data, from top to bottom. In order to meet all your organization's data access needs with one system, the system must be easy to use:

- Universe design must be simple. Each universe object must justify its place in the universe relative to the other objects, so each object must be ultimately linked logically, or in some sense understandable to the user, to all the other objects.
- In general, a universe built against a data warehouse/data mart is easier to use than a universe built against an On-Line Transaction Processing (OLTP) database.

What is your company's geographical distribution?

Are you deploying on a worldwide basis, or locally within a geographically confined area? "Universal" data access also means lots of data movement and data being distributed around the organization.

If your deployment spans more than one time zone, remember that all the servers in a cluster must be in the same time zone. Of course, this doesn't prevent client machines from being in different time zones. And when BROADCAST AGENT is used to refresh and distribute documents, their time and dates will automatically correspond to the client machines' local times.

In geographically distributed deployments, you also need to strategically locate data warehouses, system clusters, and the repository so that they're closest to the heaviest transaction loads. You will also need to decide whether you need more than one of each.

How should you choose a repository database?

Which database will you use to store your repository? You should examine the features of your chosen database - does it allow row-level locking as well as page-level locking? Does it support long BLOB columns? These and related topics are dealt with in Chapter 6, "System Configuration, Sizing, and Benchmarking" on page 155.

Do you need multiple repositories?

Do you need to create multiple repositories and/or document and universe domains to reduce the traffic on the network, and share the load between multiple databases? See "How do multiple repositories work in practice?" on page 289.

Users: know their numbers and profiles

You need to know how many users you will be supporting, as well as the type of users, and their particular requirements. Here, you need to differentiate between "licensed" users (the total number of authorized users) and "actual" users (the users you expect to be simultaneously accessing data).

For more information about user considerations, see "How many users will be using the system?" on page 159.

You also need to identify which users will have access to which services (document creation, document viewing, data analysis, Web access), as well as the users' hierarchical levels or "profiles" (supervisors, designers, users, etc.).

These considerations are greatly impacted by what type of documents users will be viewing and processing (BUSINESSOBJECTS documents can be larger and more complex than WEBINTELLIGENCE documents). For more information, see "What types of documents will be processed?" on page 161. Will they be processing other files, such as Microsoft Word documents, Excel sheets or Adobe PDFs?

If you have a large percentage of "mobile" users, such as a sales force who are constantly on the move, then access to documents via the web from their laptop computers will be critical.

Document processing: what documents, and how complex?

Traditionally, end users have been able to process two types of documents with the Business Objects system:

- **Full-client documents** are those created using BUSINESSOBJECTS REPORTER.
- **Thin-client documents** are those created using WEBINTELLIGENCE REPORTER.

Full-client document processing requires many more system resources than thin-client document processing, but using full-client documents has its own advantages, such as the capacity for much greater complexity, using multiple data sources, and more. For more information about the differences between the two types of documents, see "What's the difference between BusinessObjects and WebIntelligence documents?" on page 162.

Now that WEBINTELLIGENCE lets you access, send, save and publish any type of file, you will need to assess what types of files users are likely to use, of what size and complexity, and with what frequency.

Deployment type: client-server or distributed architecture?

The new distributed architecture allows you to spread out the processing load for WEBINTELLIGENCE, BROADCAST AGENT, and Zero Admin BUSINESSOBJECTS over a larger number of physical servers.

If you decide that the size of your deployment necessitates moving from a client-server architecture to a distributed deployment, you will need to quantify the number and size of the servers you intend to use for the various server products.

This decision will be based on:

- The number of users using the system
- The transaction load being carried by the system

For a description of these two architectures, see "Deployment Architectures" on page 87.

Server sizing: efficiently supporting peak user activity

You need to support not the *average* concurrent use, but the *peak* concurrent use. This can be up to 40% higher than average!

- In a typical system, approximately 15% of the users are logged in at any one time
- 15% of these users are actually using the system's resources

You need to identify how the users are using the system:

- What percentage of processed documents will be WEBINTELLIGENCE documents?
- What percentage will be BUSINESSOBJECTS documents?
- What percentage will be other types of files?
- What percentage will have interactive refresh capabilities for WEBINTELLIGENCE and BUSINESSOBJECTS documents?

There are so many variables in this type of system fine-tuning that you can determine what server configuration best supports your end-user population only by trial and error. We do, however, provide detailed capacity planning guidelines in Chapter 6, "System Configuration, Sizing, and Benchmarking" on page 155.

Disaster recovery: what can be done about it?

Throughout your pilot project and your production deployment, you will need to be able to identify potential critical points of failure:

- The repository
- The data warehouse
- The cluster manager
- Other areas on the cluster nodes

You will also need to have a plan for failure recovery. For example, you might use a multiple cluster deployment, so that if one cluster goes down, the other cluster can take over. You might also implement fault tolerance for cluster managers. For information, see page 184.

Load balancing: matching system processes and server resources

One of the greatest advantages of a distributed architecture is the capacity to distribute system transactions over several -- and indeed, increasing numbers of -- servers. When you're setting up your system, you need to decide exactly how you are going to share the transaction load.

You can balance loads over servers by strategically distributing high-traffic system modules (or processes) across several machines, or in some cases, by directing a greater proportion of processes to higher-capacity servers.

You can dedicate certain servers to a single type of process, or provide for several types of processing on several servers.

For more information, see "Load Balancing" on page 204.

Planning Your Pilot Project

All business intelligence solutions need careful planning, execution and management, but using e-Business Intelligence over the web is exponentially more complex (although given the typically much larger user populations, it requires less deployment and maintenance time per user).

If you are planning a large deployment, set up a pilot project first. The pilot project should be as representative as possible of your entire project, and should be tested before investing time and resources in an actual production project.

A pilot project will allow you to test and tune your deployment, and to iron out any configuration and operational problems that arise from your particular setup and chosen tools. Based on this feedback, you can then realistically plan your actual deployment.

Bring in Business Objects expertise

Web-based delivery is still a relatively new approach. Business Objects skilled input either internal or external is valuable across the life cycle. Business Objects offers a complete range of support and services to ensure you the maximum business benefit from your e-business and business intelligence deployment. A global network of Business Objects and business intelligence technology experts provides customer support, education, and consulting that can be tailored to fit the needs of any organization.

Customer Support

Business Objects operates three global customer support centers staffed with acknowledged experts in our technology. Using an industry-leading call management and problem-reporting system, all engineers share their knowledge to accelerate case resolution time.

The Business Objects online customer support (OCS) web site let you log cases, update case information, and track case progress autonomously via the Internet. Using the OCS, you have online access to the Business Objects knowledge base 24 hours a day, seven days a week. So no matter what time it is, you can access the most up-to-date product and technical information at Business Objects to help answer your questions.

From that site, you can also access the full range of product documentation in Adobe Acrobat PDF format, as well as a variety of tips for using the Business Objects solution to its fullest.

Consulting

Business Objects offers a full range of consulting services to assist customers in all stages of the development life cycle, from initial analysis through to delivery. Depending on your needs, our involvement can range from an advisory status to managing the whole project.

Business Objects consultants have in-depth product knowledge and experience in using and deploying Business Objects products in many companies across different market sectors. In addition, consultants offer specialized experience in relational and multidimensional databases, connectivities, and use of data design tools, as well as embedding product technology using scripting and customization techniques.

Contact your local sales office for more information.

Start small

Start with a single cluster, and incorporate a small group of users, either:

- A couple of users with full functionality
  In this scenario, monitor how users use the system: are they processing more full-client or thin-client documents, how often do they create documents, refresh, schedule them, and when?
- A small set of 10-15 users with restricted functionality, for example the right to view documents but not to create, edit, refresh or schedule them.
  In this scenario, evaluate the system's response times. Experiment with load balancing, for example dedicating servers to processing either thin- or full-client documents, or combining the functionality on the same machines.

Use this restricted system to monitor where the bottlenecks are. Use the Business Objects Services Administrator to monitor user and system activity, and evaluate what "performance" means for your organization.

Above all, don't enter into the pilot with preconceived notions. Accept this stage of development as trial by error. Plan for evolution.

In addition, evaluate not only how the system responds to user transactions, but how users react to the system. Make sure you give the users the knowledge they need to use the system correctly, and incorporate their ideas in plans for future evolution.

Think big

Most systems will necessarily evolve as you discover new ways to optimize performance, and new ways of using the system to work in different ways.

Your system's implementation must be scalable, as more users come on-line with increased functionality.

Consider the possible future extranet requirements from the very beginning. If you interface with the World Wide Web, how can you protect the integrity of your corporate data? Can you scale up to cope with a potentially exponentially-increasing user population? How can you prevent bottlenecks before they happen?

Most systems will be processing both BUSINESSOBJECTS and WEBINTELLIGENCE documents. Later on, however, they may want to migrate entirely to WEBINTELLIGENCE or Zero Admin BUSINESSOBJECTS. And now, what does WEBINTELLIGENCE's support of other types of files mean for your system? Will users be most likely to process Microsoft Word files, Excel sheets, Adobe PDFs, or others? How much disk space will these files require, and how much will they increase transaction loads?

You may be using an exclusively Windows system now. But as the user population and transaction load increase, you may want to switch to a mixed Windows/UNIX environment.

Expand, trace, volume test, benchmark

Once you've designed a preliminary system, you must build it then test it in order to gauge its performance level. Based on the results of the tests, you can modify the system's configuration to ensure acceptable performance levels. For a benchmarking overview, see page 198.

As you run into problems then resolve them, expand your system:

- Turn on more functionality for "power users," then see what they do with it.
- Expand the user population. When are they most likely to use the system's resources for accessing data?
- Force greater workloads through the system. At what point does performance plunge? Where are the bottlenecks? How can the load be better balanced across the cluster?

At every stage, trace your system and user activity using the Business Objects Services Administrator tracing facility. This feature tracks the activity of your entire distributed system. Among other things, it can tell you what universes users are using, what types of documents they're processing, how, and how long the system has taken to respond. It allows you to determine how many users were active at any given time, and thus check the number of concurrent users in the system (that is, users who are making the server work as opposed to the users who are merely logged in).

All this information is critical for tuning your system to perform optimally.

For more information about this facility, see the *WebIntelligence System Administrator's Guide*.

Conclusion

Virtually no organization's needs remain frozen. User populations expand, evolve, or relocate. Technology continues its steady advance, and just as steadily, users demand more out of a system's performance. Your organization may decide to port to UNIX, or extend its reach across the web via an extranet.

No matter how carefully you've planned and managed your Business Objects deployment, chances are you will never really see the end of the deployment procedure.

However your system evolves, keep the same rules in mind:

- Start small.
- Think big.
- Expand, trace, volume test, benchmark.

And don't hesitate to bring in the expertise you need, when you need it.

Chapter 4 Deployment Environments

In this chapter

- **Deployment Architectures 87**
- **Client-Server or Distributed Architecture? 87**
  What products are associated with each type of deployment? 88
- **Distributed Architecture Overview 91**
  Typical features of a distributed architecture deployment 91
  Distributed system server components 92
  An architectural overview of a distributed system 97
  What are the system's administrative layer components? 99
  Does the system have a single e-Business Intelligence portal? 104
  Distributed system performance optimization using caches 113
  Is the Business Objects product suite now server-only? 115

❑ **Distributed Environments and Clusters 116**
The cluster manager 117
Cluster nodes 118
When should you use multiple clusters? 119
What is the OSAgent Port number? 120
Setting up a cluster 121
Clusters and time zones 122

❑ **Intranets and Extranets 122**
Why is e-business important? 122
How does Business Objects fit into e-business intelligence? 123
Customizing your e-BI solution 123

Deployment Architectures

The Business Objects solution can be deployed over one or both of the following:

- A standard client-server environment (the full client version of BUSINESSOBJECTS and BROADCAST AGENT)

  The advantage of having all components running on a single server are that there is only one server machine to maintain and there is less impact on the network bandwidth. The disadvantages are that performance will be lower as the number of users increase, and it is harder to maintain.

- A distributed architecture environment, WEBINTELLIGENCE, BROADCAST AGENT, Zero Admin BUSINESSOBJECTS)

CORBA was already used in previous releases of WEBINTELLIGENCE. With this release, however, the use of CORBA has been extended to become the **standard architecture** for all of the Business Objects enterprise server products. Therefore, this chapter is largely dedicated to distributed architecture deployments.

Client-Server or Distributed Architecture?

A standard client-server deployment has some advantages over a distributed architecture deployment. Among these include faster response time, more powerful document design capabilities, and the creation of OLAP documents (although users can still *access* and *process* these documents using their browsers from WEBINTELLIGENCE).

However, using a distributed architecture deployment allows for much larger deployments. Distributing components over several servers optimizes utilization of the different server resources, reduces the workload on the web server, and increases performance. It can also provide failure recovery if the components are installed on more than one server, as well as help isolate the processes in case of failure. Distributing components is therefore recommended.

A distributed architecture deployment using WEBINTELLIGENCE, BROADCAST AGENT and Zero Admin BUSINESSOBJECTS reduces the total cost of ownership for large-scale deployments.

- A distributed architecture deployment also provides self-service data access, high availability and performance (full 3-tier, n-tier architecture, load balancing, failure recovery), zero-administration client (only the application servers to maintain), security (repository), and unified metadata - Business Objects Semantic Layer (universes).
- It uses Java applets, which ensures platform independence and low cost of ownership as well as maintaining the intuitive drag and drop interface that full client users are used to.
- The compact Java applets are automatically downloaded to the thin-client web browser.
- All processing is carried out on the middle-tier application server.

What products are associated with each type of deployment?

Within each environment, Business Objects products can be deployed in different ways. The following tables list the different deployment types as well as the products needed to create the architecture. DESIGNER and SUPERVISOR are not included in these charts but are needed to deploy all architectures.

Types of Standard Client-Server Deployments and Associated Products

Deployment Type	Product(s)
Standard full client	BUSINESSOBJECTS
Standard full client with scheduling support	BUSINESSOBJECTS, BROADCAST AGENT

Types of Distributed Architecture Deployments and Associated Products

Deployment Type	Product(s)
Thin client with interactive refresh only	WEBINTELLIGENCE
Thin client interactive refresh only, with scheduling support	WEBINTELLIGENCE, BROADCAST AGENT
Full client document interactive refresh only	WEBINTELLIGENCE
Full client document interactive refresh with scheduling support	WEBINTELLIGENCE, BROADCAST AGENT
Full client document refresh and editing	Zero Admin BUSINESSOBJECTS, WEBINTELLIGENCE (to install Zero Admin BUSINESSOBJECTS)
Thin client interactive refresh and full client document refresh	WEBINTELLIGENCE
Thin client interactive refresh and full client document refresh, with scheduling support	WEBINTELLIGENCE, BROADCAST AGENT

Deployment Type	Product(s)
Thin client interactive refresh, full client document refresh, with scheduling support on a **heterogeneous configuration** (Windows and UNIX)	WEBINTELLIGENCE, BROADCAST AGENT
Thin client and full client interactive refresh	WEBINTELLIGENCE, Zero Admin BUSINESSOBJECTS
Thin client interactive refresh, full client document interactive refresh, with scheduling support	WEBINTELLIGENCE, BROADCAST AGENT, Zero Admin BUSINESSOBJECTS

Example **Functional deployment examples**

- A deployment that permits users to access, refresh, and schedule both WEBINTELLIGENCE and BUSINESSOBJECTS from their browsers must include WEBINTELLIGENCE and BROADCAST AGENT.
- To allow users to create and interactively refresh both WEBINTELLIGENCE and BUSINESSOBJECTS, the deployment must include WEBINTELLIGENCE and Zero Admin BUSINESSOBJECTS.
- Any deployment requiring scheduling support must include BROADCAST AGENT.

Distributed Architecture Overview

A distributed architecture is one in which common server resources are shared among the server products, and are distributed over several physical servers.

This section describes the principal features of a distributed Business Objects architecture, and gives you an overview of one of its principal components: WEBINTELLIGENCE INFOVIEW.

It also describes the critical role played by the BOManager and the Business Objects Services Administrator in a distributed system.

Typical features of a distributed architecture deployment

Typically, an enterprise distributed architecture deployment is characterized by:

- A large volume of users (often with different requirements)
- A geographically dispersed implementation
- A large volume of data, complex schemas and data marts
- More than one data source, and dispersed
- A complex physical supporting infrastructure
- Many pre-defined documents and ad-hoc requirements
- A large number of often complex measures

Distributed system server components

The program files are installed on the WEBINTELLIGENCE system server. This server communicates with the web browser via the HTTP server.

Processing layers

This system contains 2 processing layers:

- HTTP server

  The HTTP server components, which include an HTTP Server Abstraction Layer (HSAL).

- WEBINTELLIGENCE backbone

  The common processing layer used by all the Business Objects server products. This layer is an integral part of the CORBA architecture, and is required in all configurations. This layer comprises:

  - The CORBA backbone (all operating systems)
  - The WIOrb process
  - The WIManager process (Windows only)
  - All the modules in the WEBINTELLIGENCE system. For complete information about these modules, see the *WebIntelligence System Administrator's Guide*.

  In a cluster configuration, the MachineManager, which provides the system's module management interface, allocates processes to the cluster manager and cluster nodes. It also handles communications with the Business Objects Services Administrator.

Note: The backbone services must be installed on each of the servers in your configuration. By default, these services are auto-started on Windows servers.

What actually happens when a request is received?

What actually happens in the system depends on the type of request it receives:

- When requests for WEBINTELLIGENCE documents or non-Business Objects files are received, they are passed on to the **WIGenerator**, which starts instances of the WIQT process as and when required.
- When requests for BUSINESSOBJECTS documents are received from INFOVIEW users, they are passed on to the **BOManager**, which starts an instance of the BUSINESSOBJECTS process (BusObj under Windows, `bolightsvr` under UNIX) for each document. On the server, BUSINESSOBJECTS performs the actions the user requested on the document, such as refreshing it.

What role does the BOManager play?

When a request is received from a WEBINTELLIGENCE INFOVIEW user to view or refresh a BUSINESSOBJECTS document, the BOManager launches an instance of the BUSINESSOBJECTS program to process the request.

The BOManager balances tasks between the available resources as further requests are received, and as transactions are completed. This is done so as to optimize resources, and also provides fallback solutions should any of the processes fail.

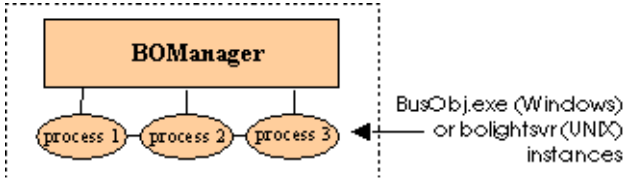

Figure 4.1: How BOManager balances tasks between available processes

The BOManager is an integral part of the WEBINTELLIGENCE system, required for:

- Allowing WEBINTELLIGENCE INFOVIEW users to view and refresh BUSINESSOBJECTS documents in HTML, Adobe PDF, or enhanced formats.
- Running scheduled BROADCAST AGENT tasks. The BOManager receives instructions from the Scheduler to process documents that are due. The BOManager then launches one BUSINESSOBJECTS session on the server for each document that is to be processed.
- Embedding the Report Viewer control which is delivered with the BUSINESSOBJECTS Developer's Suite, in custom applications.

You can deploy several BOManagers running on separate servers.

For more information on deploying BOManagers, see "Deploying Broadcast Agent" on page 181. For both detailed information about how the WEBINTELLIGENCE system works, and BOManager troubleshooting information, see the *WebIntelligence System Administrator's Guide*.

What role does the WIGenerator play?

This module is used the first time a user logs into WEBINTELLIGENCE. It locally launches a WIQT process each time a user logs in. Thereafter, as long as the WIQT hasn't timed out, any of the user's requests requiring a WIQT process are passed on directly to that WIQT.

WIQT processes are used to process both WEBINTELLIGENCE documents and third-party files. In BUSINESSOBJECTS terms, the WIQT (Query Technique) is an SQL Manager, a Query Executor, a microcube engine, and a document (and HTML) generator rolled into one. It is responsible for validating users, checking WEBINTELLIGENCE queries against users' database access rights, generating SQL requests, then sending them to the repository or database. When WIQT receives the results, it does one of the following:

- If WIQT has queried the repository for user or universe information, it passes the information on to the WIGenerator, which dynamically generates either the user's start page, or pages giving the user access to the existing universes or documents for which the user has access rights.
- If WIQT has queried a database for information, it then generates the final document from the query results.

When a request is received from a WEBINTELLIGENCE INFOVIEW user to view or refresh a WEBINTELLIGENCE document or third-party file, the WIGenerator assigns an instance of the WIQT process to handle the request.

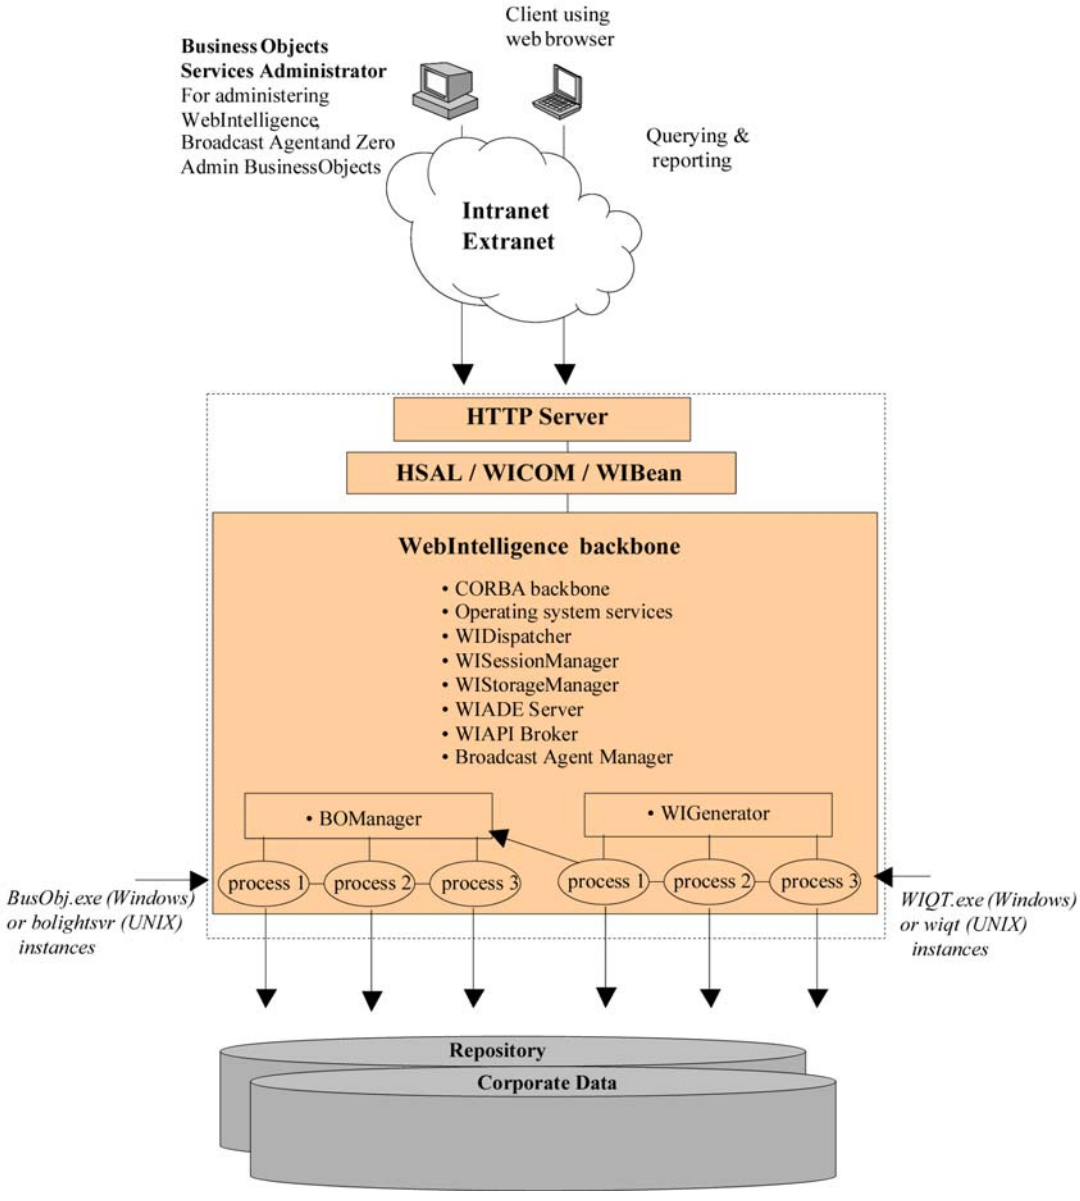

Figure 4.2: Server components

The Business Objects Services Administrator

A distributed solution requires the Java-based Business Objects Services Administrator, which can be used to administer all Business Objects server products, including WEBINTELLIGENCE, BROADCAST AGENT and Zero Admin BUSINESSOBJECTS.

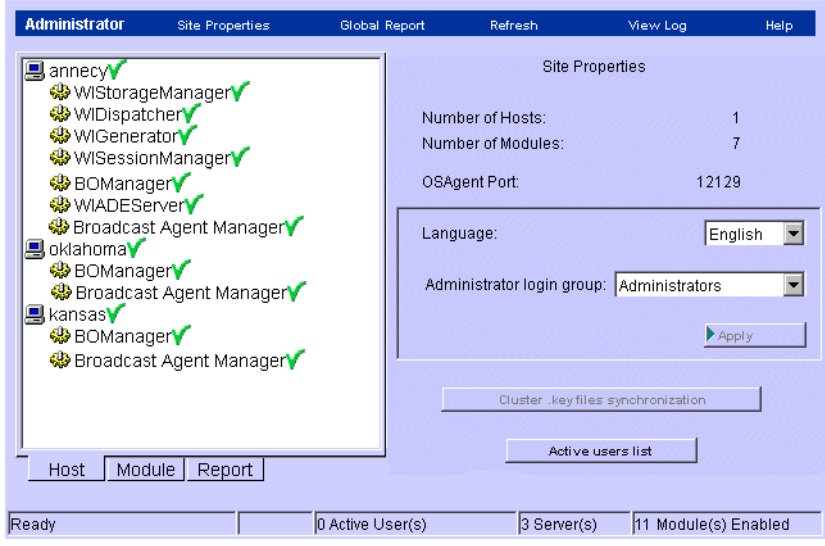

Unlike the Broadcast Agent Console, which provides both a monitoring function for documents which have been scheduled in the BROADCAST AGENT and the ability to modify scheduled tasks, the Administrator lets you actually administer and tune the overall system. For example, it lets you modify Scheduler parameters, enable and disable modules on cluster nodes, and customize BOManagers.

The Administrator provides settings for how and when the BOManager launches processes. The configurable parameters include the maximum number of active document allowed, the number of documents that may be concurrently open, maximum number of processes, and so on.

Types of Administrator

The Administrator is available as:

- An applet that runs in a web browser.
  The applet is automatically included when you install WEBINTELLIGENCE. It requires an HTTP server.
- A standalone Windows application
  The application is used in full client and BROADCAST AGENT-only deployments, which do not use an HTTP server. In this case, therefore, the Administrator must be installed as a Desktop Product on a Windows client machine.
  This installs all the required CORBA components it needs to connect with the BROADCAST AGENT cluster. During installation, you must also provide both the OSAgent Port being used by the cluster, and the IP address of the cluster manager.
  For more information, see the *Installation and Update Guide*.

When do you *not* need an HTTP server in a distributed solution?

If you are deploying BUSINESSOBJECTS and BROADCAST AGENT without WEBINTELLIGENCE, while you need the WEBINTELLIGENCE backbone in order to use the CORBA architecture and Windows services, you do not need an HTTP server.

In any UNIX-only deployment, however, an HTTP server is required in order to access the Business Objects Services Administrator. You can also administrate a UNIX deployment from a Business Objects Services Administrator installed on a PC.

An architectural overview of a distributed system

Figure 4.3 illustrates an architectural model which combines all the server products with a selection of desktop products in a distributed architecture environment.

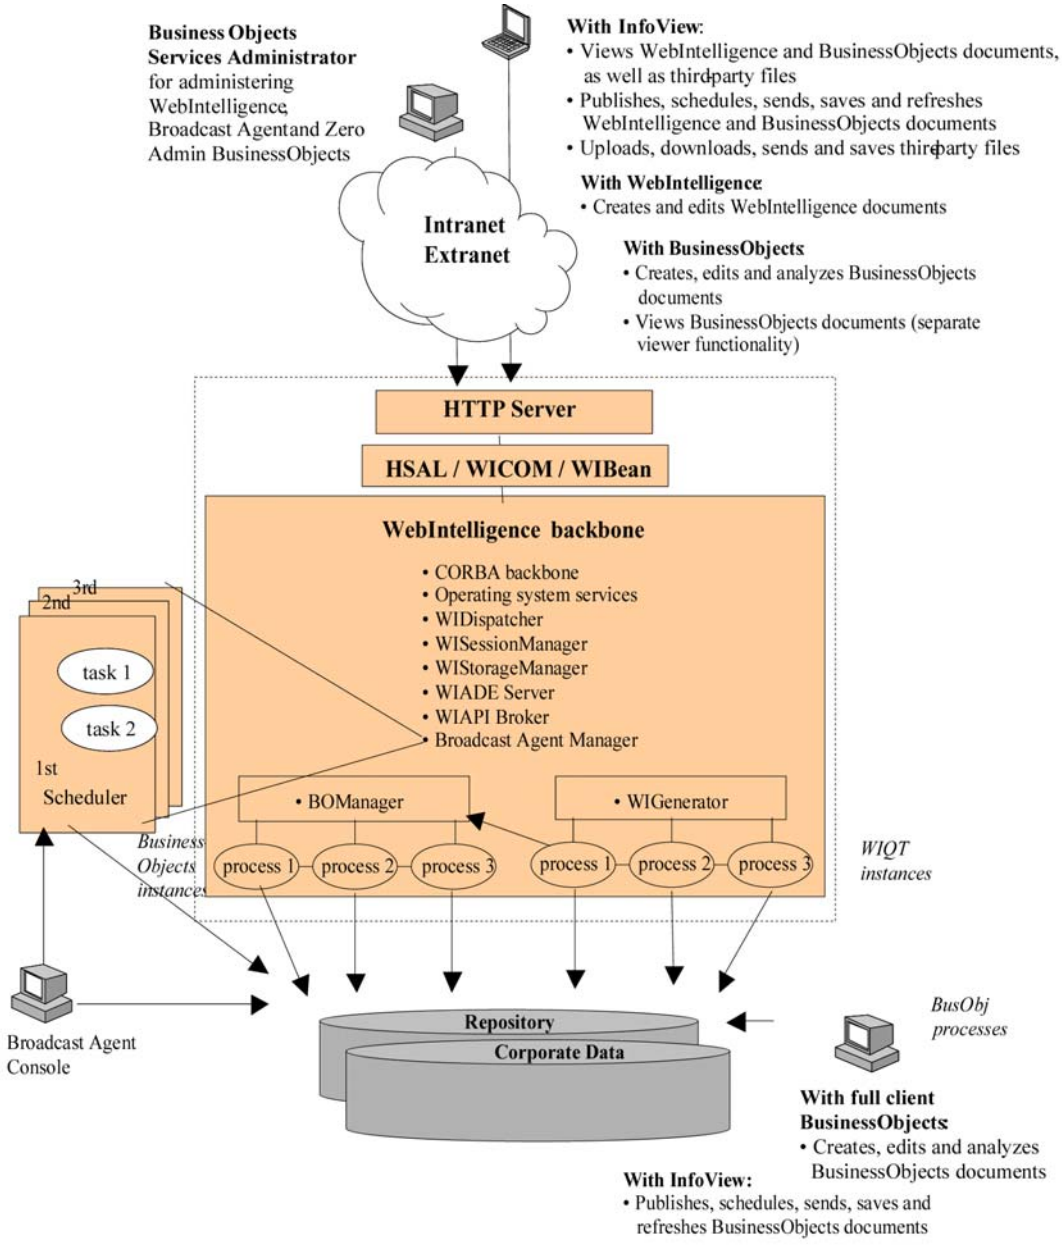

Figure 4.3: Distributed architecture

What are the system's administrative layer components?

The distributed system's applicative layer is provided by specific modules, designed as CORBA server objects. Business Objects CORBA components are licensed from Inprise.

For a complete explanation of how these modules interact, see Chapter 2, "How the WebIntelligence System Works" in the *WebIntelligence System Administrator's Guide*. The information in this chapter is crucial for understanding how the system responds to client requests from INFOVIEW, WEBINTELLIGENCE, BROADCAST AGENT and BUSINESSOBJECTS in a three-tier deployment, and being able to troubleshoot any issues that may arise while using the system.

On each machine, the installer sets a central component to start these modules and provide administration features. The administrative layer is primarily server-side, except for the Business Objects Services Administrator client component.

Although some components are platform-specific, the administration layer is designed as a cross-platform framework.

Here's a quick reference to all the executables in the administrative layer:

Executable	Description
Wiorb (`WIOrb` under UNIX)	Wiorb's main function is to start ORB tools. Wiorb starts the OSAgent and OAD on the cluster manager, and only OAD on cluster node machines. To prevent OAD failure, Wiorb monitors the OAD process. **Notes:** • On Windows platforms, Wiorb runs as a service. • On UNIX platforms, `WIOrb` not only starts Inprise tools but exposes a server object, called resourceServer::registry, which provides Registry emulation features for all components running on the machine. (To start `WIOrb` automatically at boot time, see the *Installation and Update Guide*.)
OSAgent (`osagent` under UNIX)	OSAgent runs only on the cluster manager machine. It runs as a server and waits for potential client broadcasts and locations tentative using a given port value, defined by the OSAgent Port value. All the components (client and server) which initialize the ORB layer with the same OSAgent Port value belong to the same ORB.
OAD (`oad` under UNIX)	Each machine runs one Orb Activation Daemon. Each CORBA object expected to run on a given machine with failover management is registered with the associated OAD. The OAD then works in association with the OSAgent to tell the cluster where server objects are located over the ORB. The OAD maintains an implementation repository, a text file containing all registration information. **Note:** When the OAD crashes, the Wiorb monitoring system kills all processes hosting registered objects. The OAD is then restarted, re-initializing itself from its implementation repository. The system is once again ready to answer new requests.

Executable	Description
Wimanager	This process only exists on Windows platforms and runs as a service on each machine. Wimanager starts and monitors the machine management process, which is different depending on machine type. On UNIX platforms, Wimanager features are provided by the `WIOrb` process.
WIClusterNode	This is the first machine management process. It is started by Wimanager (or `WIOrb` on UNIX) on each cluster node machine and has three main goals. • The first is to provide an object named **MachineManager**. The MachineManager provides an interface for module management on its machine. This means it: --Registers / unregisters servers with OAD --Starts / stops servers --Connects to servers in order to get or set their parameters • The MachineManager manages the Gatekeeper. • When BROADCAST AGENT is installed on the machine, the WIClusterNode process provides another object called **Broadcast Agent Manager**. Functionally, this object appears as other server objects do. It belongs to the application layer and manages Schedulers on the machine (see the *Broadcast Agent Administrator's Guide*). **Notes:** • The MachineManager and Broadcast Agent Manager objects are not registered with OAD. • Broadcast Agent Manager can be enabled on UNIX machines only with the 5.1/2.6 and later product releases.

Executable	Description
WIClusterManager	This is the second machine management process. It is started by Wimanager (or `WIOrb` on UNIX) on the cluster manager machine. Functionally, WIClusterManager resembles WIClusterNode, but does more. It provides the MachineManager and BCA Manager objects on the machine. It also: • Exposes the **SiteManager** object, which is the central point of system administration layer. The site manager maintains information about all MachineManagers running all over the cluster. Each MachineManager register itself with site manager at start up time. The site manager also provides central information like language, product version, or data from the Business Objects repository. • Enables the auditing feature by starting a dedicated process. This means that all auditing actions are grouped on the cluster manager machine.
Gatekeeper	The Gatekeeper is an Internet Inter-ORB Protocol (IIOP) proxy for Java client implementations. It is a Java-based tool designed with strong networking features: subnetting, firewalls, etc. It offers capabilities for forwarding client applet or application requests through the ORB.
WIAdminBOTools	This process is used as an external tool by the site manager to get the list of BROADCAST AGENTs from the repository when needed. It provides all Business Objects security-compliant features.

Executable	Description
WIAdmToolStop	This additional tool is involved when server products have to be stopped on a machine. Depending on the machine type (cluster manager or node), the WIAdmToolStop component connects to the site manager or machine manager and requests shutdown operations. **Note:** On Windows platforms, the Wimanager stops services by starting WIAdmToolStop.
WIKill	WIKill is an additional tool designed to kill server processes. It supports two options: • Processes associated with registered objects only • All processes WIKill is used both when the OAD crashes and registered servers have to be killed, and when the system takes too long during a normal shutdown.
WICpw	This additional tool exists only on UNIX platforms. It helps the site manager to authenticate the Business Objects Services Administrator user as a local administrator on the cluster manager machine.
WISiteLog	This process provides a server object for auditing features, registered with the OAD. It runs on the cluster manager machine. The audit object lets you log events in text files or in a database through a Business Objects security-compliant connection.
tcphsal	Tcphsal resides behind the second firewall in a double-firewall configuration, acting as an HSAL proxy within the DMZ. For more information, see “WebIntelligence and Firewalls” on page 319.

Executable	Description
wihsalfw	Wihsal resides behind the second firewall on the same machine as tcphsal, monitoring it to make sure it doesn't fail. WEBINTELLIGENCE system administrators can administrate this monitor using the Business Objects Services Administrator. For more information, see "WebIntelligence and Firewalls" on page 319.
WINotify	A Windows graphical tool hosted by the task bar, providing: • The immediate status of the Wimanager • Product start / stop on the machine • A link to administration help • A link to the Business Objects Services Administrator • Product version information

Note: With the 5.1/2.7 release, you can customize parameters for the system's three key CORBA components: the OSAgent, OAD and Gatekeeper. For more information, see "Modifying CORBA Kernel Parameters" on page 247.

Does the system have a single e-Business Intelligence portal?

Yes. The base module for both WEBINTELLIGENCE and BUSINESSOBJECTS is INFOVIEW, the personal gateway to an organization's information capital.

With WEBINTELLIGENCE INFOVIEW, however, users can access documents from the office, home, or around the world, using the corporate intranet, an extranet, or the World Wide Web. It is one of the key components of a distributed architecture.

WEBINTELLIGENCE INFOVIEW includes a web-based document catalog and a set of document viewers, including:

- The BUSINESSOBJECTS Reader, which lets users read BUSINESSOBJECTS documents from within BUSINESSOBJECTS
- A viewer for reading WEBINTELLIGENCE and BUSINESSOBJECTS documents in HTML format over the web
- A special enhanced document viewer for reading BUSINESSOBJECTS documents in BUSINESSOBJECTS format over the web

INFOVIEW also includes the document viewer control that application developers can embed in other applications.

INFOVIEW lets users view, refresh, manage and distribute documents, but not create or modify them. To do that, they need the full WEBINTELLIGENCE and/or BUSINESSOBJECTS products (installed either from a CD or via a web browser).

With the 5.1/2.7 release, it also lets users view and distribute any other type of file, such as Microsoft Word or Excel files, or Adobe Acrobat PDF files.

The InfoView document catalog

At the heart of INFOVIEW is the document catalog.

Whether users are looking for a sales report, a financial statement, or an inventory spreadsheet sheet, it provides an instant overview of all the WEBINTELLIGENCE and BUSINESSOBJECTS documents -- or any other types of files in the repository -- that are available to them.

The document catalog includes up to three lists containing files they can access:

- The Corporate Documents page contains files that have been saved or uploaded to a document domain in the corporate repository and made accessible to workgroups within the corporation or organization.
- The Personal Documents page contains documents users have saved for their own personal use.
- The Inbox Documents page contains documents other users have sent each user.

What happens when users click on a document in a list?

What happens depends on the type of document:

- When users click on a WEBINTELLIGENCE document in the document list, the WIGenerator generates an HTML version of the document, which is then displayed.
- If they click on a BUSINESSOBJECTS document, either the HTML reader, the enhanced document reader, the Adobe Acrobat reader, or BUSINESSOBJECTS itself is launched, depending on the reader each user has chosen.
- When users click on a non-Business Objects file, such as a Microsoft Word document or an Adobe Acrobat PDF file, if the associated application is installed on the user's machine, the application is launched with the file open in it. If the application is not installed, users can choose to download the file to their machine.

Icons representing the type of file

Viewing BusinessObjects documents from a web browser

INFOVIEW users can choose to view BUSINESSOBJECTS documents in any of a variety of formats.

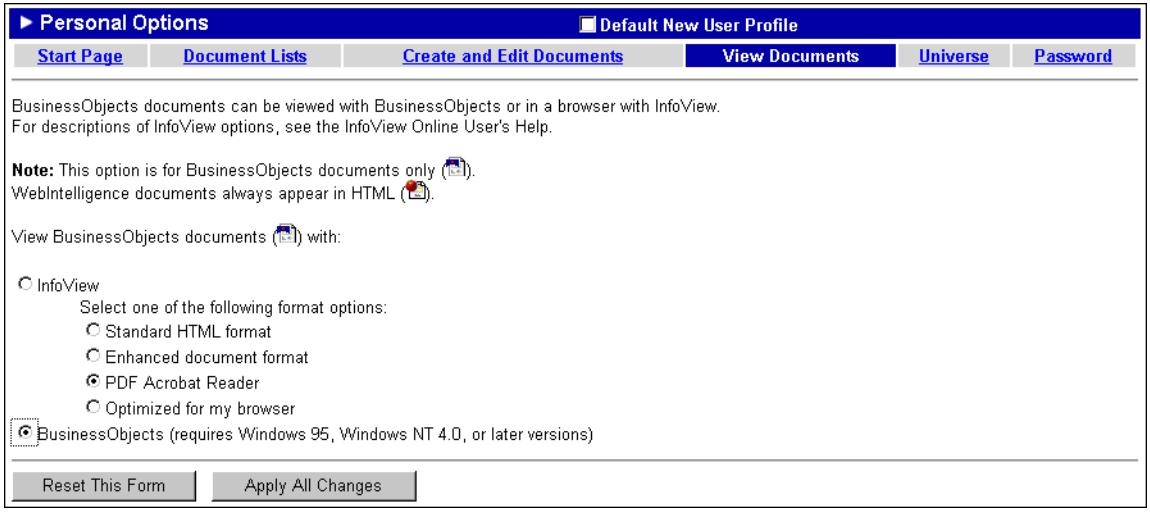

The standard HTML reader

A standard HTML reader allows users to view both BUSINESSOBJECTS documents (.rep) and WEBINTELLIGENCE documents (.wqy) in HTML format.

The enhanced document reader

The enhanced document reader is a lightweight component that is automatically downloaded from the server. This ActiveX reader or Netscape plug-in lets users view and print BUSINESSOBJECTS documents in enhanced format. This means they appear exactly as they were designed and formatted in BUSINESSOBJECTS, rather than relying on HTML formatting's rather limited capabilities.

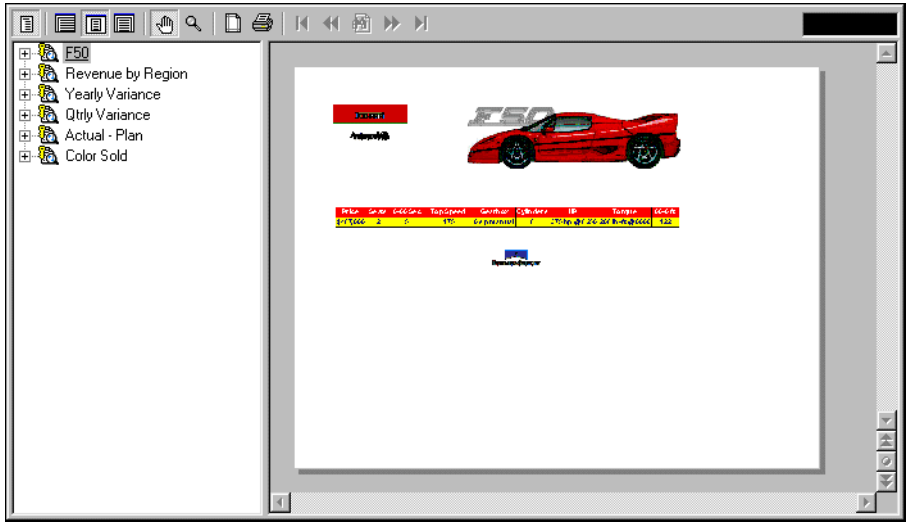

When users print HTML documents directly from a standard web browser, the print quality and formatting is limited to the restrictions of the HTML format. The enhanced reader provides superior WYSIWYG print quality. Users can zoom, pan, and navigate through your data. They can also print out high-quality documents, including headers and footers.

The enhanced reader uses the HTTP protocol, and therefore can be used for extranets (i.e., through firewalls). It requires the INFOVIEW web infrastructure (HSAL, WIDispatcher, etc.).

If users are using a browser that supports this reader, but have not opted to view full-client documents in standard HTML format, the first time they attempt to open a BUSINESSOBJECTS document, the enhanced reader is automatically downloaded from the server and installed on their system. This only needs to be done once. For a list of the browsers supported for this reader, see the ReadMe.

The enhanced reader uses **on-demand paging**. A metafile is created on the server and is transferred to the browser window when users click on a document section displayed in the left part of the window.

Note: The server hosting the WIADE Server module should be a Windows machine if INFOVIEW users in your deployment want to use the enhanced document format viewer for BUSINESSOBJECTS documents, and the deployment is using a web server in CGI mode. In this case, Basic Authentication must be set on both the web server and the cluster node hosting the WIADE Server module in order for the following components to function correctly:

- Netscape plug-in document viewer
- ActiveX document viewer
- ADE document viewer (for use with WEBINTELLIGENCE SDK).

The Adobe Acrobat Reader

An Adobe Acrobat Reader is automatically launched if users have chosen to view the document in PDF format.

The BusinessObjects product

If the user installed BUSINESSOBJECTS via a web browser, they can also choose to view BUSINESSOBJECTS documents by default in BUSINESSOBJECTS itself.

Note: This is most useful for users with rights strictly limited to INFOVIEW functionality, for BUSINESSOBJECTS used in this manner does not enable users to edit the document or create another without first closing the application.

Categories

Users can organize documents into categories, which allow them both to filter unwieldy document lists and make documents they publish or send even more rapidly accessible to recipients. Users can also search for documents using a document's name, author or date of last modification.

Powerful information distribution

INFOVIEW lets users efficiently distribute documents both inside and outside their organization.

It gives them a secure interface for aggregating, managing and sharing critical information. Extranet providers can offer exceptional customer service and reduce costs in their supply chain. Their customers can enjoy self-service access to information tailored to their specific needs.

Using the powerful broadcasting services of BROADCAST AGENT, users can schedule the automatic refresh and distribution of documents to colleagues or partners via the repository, the web, an intranet, or an extranet.

On the server side, what does WebIntelligence InfoView include?

INFOVIEW provides:

- WEBINTELLIGENCE backbone services
- The Business Objects Services Administrator
- One or more WIQTs, which not only control user access rights, but provide the services users need to view and refresh WEBINTELLIGENCE documents and other, non-Business Objects files.
- One or more BOManagers, which control resources, and start instances of the BUSINESSOBJECTS document processors as and when required by users to view BUSINESSOBJECTS documents.
- The software required to interface with BROADCAST AGENT, WEBINTELLIGENCE REPORTER and EXPLORER and BUSINESSOBJECTS, the enhanced document reader for BUSINESSOBJECTS documents, and the Business Objects Services Administrator (for both WEBINTELLIGENCE and BROADCAST AGENT).

What can WebIntelligence InfoView users do with documents or files?

Exactly what each web-based INFOVIEW user can do depends on the user's profile and the type of document being processed.

The following table describes what WEBINTELLIGENCE INFOVIEW users can do with WEBINTELLIGENCE (thin-client), BUSINESSOBJECTS (full-client) documents, and other types of files if they have the appropriate rights:

WebIntelligence InfoView users can do this...	with Webintelligence documents	with BusinessObjects documents	with other types of files
View documents/files	x	x	x (as long as the application associated with the file is installed on the client machine)
Download non-Business Objects files	N/A	N/A	x (when the application associated with the file is not installed on the client machine)
Create documents	N/A	N/A	N/A
Upload non-Business Objects files	N/A	N/A	x
Refresh documents/files	x	x	N/A
Modify documents/files	N/A	N/A	N/A
Publish documents to the repository	x	x	x
Send documents/files to other users	x	x	x

WebIntelligence InfoView users can do this...	with Webintelligence documents	with BusinessObjects documents	with other types of files
Save personal copies of documents/files	x	x	x
Print documents/files	x	x	x
Schedule documents for automated distribution with BROADCAST AGENT	x	x	N/A
Drill in documents	N/A	N/A	N/A

N/A = Not Applicable

To create, modify and drill in documents, however, users need one of the following:

- WEBINTELLIGENCE REPORTER (creation and modification) and WEBINTELLIGENCE EXPLORER (drilling)
- BUSINESSOBJECTS REPORTER (creation, modification, and Slice and Dice functionality) and BUSINESSOBJECTS EXPLORER (drilling)
- Zero Admin BUSINESSOBJECTS (creation, modification, drilling and Slice and Dice functionality)

Distributed system performance optimization using caches

The distribution of the system's functions across several machines creates more complex chains of transaction processes. The Business Objects high security standard creates additional complexity, as user access rights are constantly checked in the repository, to make sure users see only the information to which they are entitled.

When large numbers of users use the system to create, view, refresh and distribute documents, your Business Objects system can end up supporting a heavy transaction load. This release comes with several enhancements aimed at improving the performance of heavily loaded servers, and lightening the load on the repository. These particular enhancements make use of server caches throughout the system.

Login caches

Within the same WEBINTELLIGENCE session, users' requests for displaying or refreshing a BUSINESSOBJECTS document no longer require logging into the repository after the initial request.

At the first request, the system creates a temporary folder and uses it to store the users' security information. The next time users need to access the repository, the system can find their security information without having to connect to the repository again.

Likewise, once WEBINTELLIGENCE users have executed a task involving a BUSINESSOBJECTS document via BROADCAST AGENT, tasks they subsequently submit in the same session no longer have to log in to BOManager again (BOManager is the module that retrieves a requested full-client document from the repository, then refreshes and/or distributes it from WEBINTELLIGENCE).

BOManager caches the login information for each BROADCAST AGENT task and if users have the correct access rights, the session context for the task is restored directly from the cache on the server's local hard disk.

Repository cache

The WIStorageManager manages a cache for repository-related information, such as corporate document and inbox document lists, and lists of available universes. The first time a user logs into INFOVIEW, the WIGenerator generates these pages; they are then saved to the repository cache. Whenever the user needs one of these pages, the system obtains it from the cache instead of querying the repository again.

The presentation cache

To prevent potential overloads at peak transaction periods, users can **pre-load** the cache whenever they process a task. When BUSINESSOBJECTS users schedule a corporate document with BROADCAST AGENT, they can automatically generate the document in enhanced document, standard HTML, or Adobe Acrobat PDF format, then store it in the server cache. These formats correspond to the three possible options for viewing BUSINESSOBJECTS documents within WEBINTELLIGENCE INFOVIEW.

The first time a INFOVIEW user asks to view the document in a certain format, BOManager retrieves the presentation and stores it in a cache on a machine controlled by the system's WIStorageManager(s). When other users subsequently access the document in INFOVIEW, instead of having to recontact BOManager and regenerate the file in the requested format, they access a pre-generated file.

This means that:

- Each INFOVIEW request for a BUSINESSOBJECTS document doesn't require logging into BOManager.
- There are fewer demands on available processes in your cluster.
- The document's presentation doesn't have to be generated -- the document is displayed faster, and having used fewer system resources.
- The response time remains constant and doesn't depend on the document's size or complexity.
- CPU power and BUSINESSOBJECTS processes are therefore made available for refreshing documents (ad-hoc queries).

The PDF file cache

When WEBINTELLIGENCE INFOVIEW users view or download a BUSINESSOBJECTS corporate document, they can choose to have it first converted to PDF format. If they choose this option, the system will convert the document to PDF and store it in a temporary cache, where it remains available for them and for other users. You can customize the length of time documents are retained using the Business Objects Services Administrator.

Caching of retrieved PDF files saves time and system resources, since the system only needs to generate the PDF file once -- the first time it is accessed. This can greatly speed up the retrieval process, as it can take a considerable amount of time to generate PDF files from large files which contain many graphics.

If the source document is modified in the meantime, then the next time the document is downloaded, the PDF is regenerated again.

Note: This feature is available only for corporate documents, and not for personal or Inbox documents. Only corporate documents are shared by all users in the company, and can therefore make use of a shared document cache.

Is the Business Objects product suite now server-only?

This release is not a server-only product suite. Many of the new server module features take advantage of CORBA technology, which greatly facilitates deployment over a distributed network of servers. However, you can still deploy Business Objects desktop products in the traditional way, using a full-client setup on a single system, or a shared installation on a network.

In an enterprise server deployment, the Business Objects server components (such as the WEBINTELLIGENCE server and BROADCAST AGENT) are set up on one or more server machines. The client is the front-end system on which the Business Objects desktop products (such as BUSINESSOBJECTS REPORTER, DESIGNER or SUPERVISOR) are installed, and from which you can access the web-based WEBINTELLIGENCE.

There are different ways in which the products can be set up. Although you will typically set up an Enterprise server configuration, you can also deploy a stand-alone installation on a single or multiple system(s).

Distributed Environments and Clusters

Of course, Business Objects server products can run on a single machine. This machine can host the HTTP server, system services, WEBINTELLIGENCE and BROADCAST AGENT, and can administer all of these. In this case, all these components run in the same environment and have no problem communicating and coordinating. The full client product, BUSINESSOBJECTS, could interact with that system on individual client machines.

However, chances are you will want to scale your system up as demand grows by adding additional servers to the system. They can share the transaction load and provide failover capabilities: if one server fails, the system automatically redirects requests to another functioning server in the system. The system automatically recognizes all the machines in your deployment, as well as the distributed services at runtime.

In this type of distributed deployment, the group of machines over which the server products are deployed is referred to as a **cluster**.

A cluster can include two types of servers:

- The **cluster manager** is the cluster's central machine. It houses the HTTP server, the OSAgent (the Inprise VisiBroker Object Request Broker), an OAD (ORB Activation Daemon), mandatory Business Objects modules, and any optional ones you may choose. There must be one and only one cluster manager in a cluster.
- One or more **cluster nodes**, each of which runs the OAD required to communicate with the cluster manager and start BUSINESSOBJECTS and/or WEBINTELLIGENCE processes on the node, as well as strictly optional services.

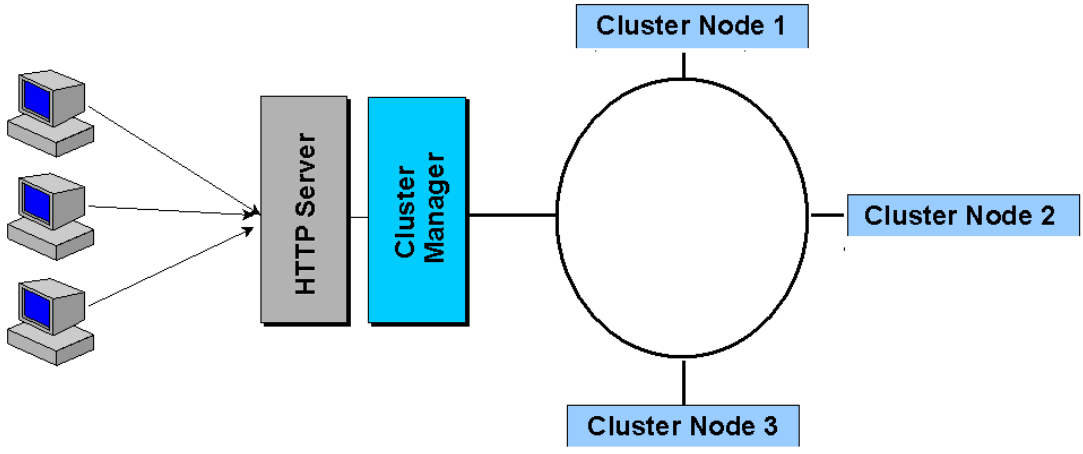

Figure 4.4: Cluster servers

The cluster manager

The cluster manager is the central coordinator between all the servers in the system. It performs the following services:

- Tracks and manages processes throughout the system using WIClusterManager.
- Runs the system session manager, WISessionManager.

  Because of the architecture of a distributed solution, a user logs into the *system* and not a particular server. Even if the log-in procedure is actually handled by another server, WISessionManager maintains the session information.
- Records various trace activities.

The cluster manager tracks and manages the key processes in the cluster. It always hosts the ORB (OSAgent) and the WISessionManager, and generally hosts the HTTP Server Abstraction Layer (HSAL), the Gatekeeper, the WIDispatcher, the WIStorageManager and at least one WIGenerator.

For more information about all these processes, see Chapter 2, "How the WebIntelligence System Works," in the *WebIntelligence System Administrator's Guide*.

Note: The cluster manager occupies a critical position in the system. If the cluster manager server should fail, or if one of its solution-wide modules (such as WISessionManager, or the Inprise VisiBroker OSAgent the WEBINTELLIGENCE system uses as its Object Request Broker) should fail, the entire system needs to be stopped and restarted manually.

You can, however, provide fault tolerance on the cluster manager. See page 184.

Cluster nodes

Cluster nodes are used to host server processes, such as WIGenerator, BOManager and Broadcast Agent Manager.

Cluster nodes rely on the cluster manager to provide the infrastructure required to allow multiple servers to work together. For example, when a cluster node receives a request and needs to validate the user, it does so by checking with the WISessionManager running on the cluster manager.

Using the Business Objects Services Administrator, you can configure a cluster node to run only a subset of the full complement of server modules. You can also tailor the structure of your system and more efficiently serve the needs of your users. If you find that the load on a particular type of server process within the system is high, you can dedicate one or more machines to just that process.

Cluster nodes have failover capability, which means that if one node fails, the system automatically redirects requests to another functioning server in the cluster. If a user session is lost, the user is prompted for authentication, and another available node is invoked to launch another user session.

For information about distributing server processes efficiently within the system, see "Distributed Architecture Overview" on page 91.

When should you use multiple clusters?

You can deploy more than one cluster. Each of the servers in a particular cluster shares the same OSAgent Port number. This ensures that they communicate with servers in the same system only.

A cluster server cannot interact directly with a server from a different cluster, but can exchange information via the repository, which can be linked to all of the configured clusters. Each CORBA network you deploy must have one (and only one) cluster manager. You can add as many cluster nodes as you require.

When too many users spoil a cluster's performance

You should deploy multiple clusters if you have a large number of users that you want to break down into smaller groups, each with an individual cluster manager to improve performance. All users can still access the same repository, and can exchange documents in the usual way.

When you want failover capacities for a cluster manager

By default, there is no cluster manager failover. If you have only one cluster configured, and the cluster manager fails, then all cluster nodes will also go down.

To provide for cluster manager fault tolerance, you can set up multiple clusters, or configure a redundant cluster manager machine. For instructions, see "Providing cluster manager fault tolerance" on page 184.

When you want Broadcast Agent failure recovery

Multiple clusters are also useful for BROADCAST AGENT failure recovery. If you have BROADCAST AGENT Schedulers in one cluster, and Schedulers in another cluster, then if one cluster manager fails, subsequent tasks can be assumed by one of the Schedulers in the second cluster.

When your deployment is geographically distributed

All the servers in a cluster must be in the same time zone. This doesn't, however, prevent *client* machines from being in different time zones. And when BROADCAST AGENT is used to refresh and distribute documents, their time and dates will automatically correspond to the client machines' local times.

When you are using multiple repositories

If you are supporting large numbers of users, you may want to speed up logins and document access by using multiple repositories. When the WEBINTELLIGENCE system is being used in this type of deployment, you must set up multiple clusters, each of which points to a particular repository. Users can select which repository to connect to by using the URL that corresponds to the appropriate cluster.

What is the OSAgent Port number?

If you're configuring WEBINTELLIGENCE and/or BROADCAST AGENT as a cluster solution, all the machines that you want to work together in the distributed system must use the same OSAgent Port value.

By default, the installer sets this to 14999. If an environment variable for this value has already been set on your machine, however, that value is displayed instead. In this case, decide whether you want to install the product(s) on this port, or on another.

The value must be a number between 1024 and 65535, but must not be 15000, which is reserved for the Gatekeeper port. This value must be different from the port value of any other application server running on the machine, whether it be a CORBA, web, database, or other type of server.

If you're configuring WEBINTELLIGENCE as a single-machine system or installing a cluster manager, make sure you're using an OSAgent Port value that is unique on your subnet.

Setting up a cluster

You set up a cluster in two phases:

1. First, you install the server products on each of the machines to be included in the cluster. You must install the cluster manager first, then you can add the cluster nodes.
2. After installation, you use the Business Objects Services Administrator to enable the modules you want to run on each machine in the cluster.

You don't need to tell the cluster manager where its nodes are, and you don't have to tell the nodes where their cluster manager is. Whenever WEBINTELLIGENCE or BROADCAST AGENT is started on a node, the OAD on each node broadcasts throughout the network using a specific TCP/IP port. On the cluster manager, the ORB process, called OSAgent, will use the same port to listen for these broadcasts, and will register the presence of the node along with what processes it can run.

Installing a cluster

First, you install the enterprise server products on the cluster manager machine, by selecting the Cluster Manager option.

Next, you install the server products on each cluster node machine, by selecting the Cluster Node option.

At this time, you assign the OSAgent Port value that you assigned the cluster manager. This shared port will allow the nodes to communicate with the OSAgent on the cluster manager.

For full information, see the *Installation and Update Guide*.

Enabling/disabling server components

After installation, you need to configure precisely how your distributed system functions by using the Business Objects Services Administrator to enable specific modules on specific nodes.

In order for the Business Objects server products to function, all of their modules must be running somewhere within the cluster. Some modules can be enabled on a cluster node, but on one node only. Others can be enabled on several nodes at a time.

For detailed information, see the *WebIntelligence System Administrator's Guide*.

Clusters and time zones

All of a cluster's servers must be in the same time zone.

In addition, the server running the WIStorageManager module must be set to the same time or earlier as the other servers in the cluster. If this server is set to a later time than the other servers, the server's cache mechanisms may not work correctly.

For information on setting time zone information in your Windows or UNIX servers, see the *WebIntelligence System Administrator's Guide*.

Intranets and Extranets

An **intranet** is a network that reaches across all the different entities of an organization. TCP/IP has become a popular protocol since it allows network administrators to easily split heavy network traffic into smaller domains using subnets, which provides links between them using routers. It takes advantage of existing Internet standards such as the standard client-server model. You can picture the Internet as a massive TCP/IP network, where intranets are just smaller, organization-wide TCP/IP networks that provide company-related information to employees.

When a company opens its intranet to selected business partners, the intranet becomes an **extranet**. Suppliers, distributors, and other authorized users can connect to the company's intranet over the web to view the selected subset of data the company makes available.

E-business intelligence represents the move of organizations around the world to link their customers, suppliers, partners and internal organizations, to share information, work together in new ways, and to generate greater revenues.

Why is e-business important?

E-commerce represents the single most important aspect of e-business as it is directly quantifiable, and therefore serves as a useful indicator of the direction and magnitude of e-business.

Forrester Research, a leading Internet research firm, has estimated that $17 billion in transactions were conducted over the internet in 1998. This figure is expected to grow to $1.3 trillion by 2003. This new estimate is a revision of a mid-1997 Forrester report that forecast a $327 billion e-commerce market in 2002.

How does Business Objects fit into e-business intelligence?

As organizations establish their business offerings on the web, several implications have emerged surrounding the role of business intelligence. If you take a look at any e-business application framework, there are more aspects to it than just placing a sales order on the web. Once the sales order is placed customers need to take a look at previous transactions, they need to search on a particular product based on certain criteria, they probably need to define their profile and from time to time get alerted on promotions on a product based on their profile.

The Business Objects distributed architecture perfectly fits these various aspects of customer interaction, enabling customers to communicate and share data over the web with their customers, partners and suppliers.

Customizing your e-BI solution

As the vision of "information democracy" increasingly becomes reality in today's organizations, e-Business Intelligence (e-BI) is being deployed on a large and growing basis. With this massive deployment comes a correspondingly large increase in the diversity of the requirements of users themselves. This inevitably means that e-BI products need to be customized to meet these specific needs.

With version 5.1/2.7, Business Objects continues its commitment to providing the best offering for developers, including:

- Integrated development environments that enable developers to customize tools to meet the requirements of specific categories of users.
- Components that enable developers to integrate specific elements of e-BI technology into their own applications with the lowest possible development investment.
- APIs that open the e-BI tools and technologies to access from third-party environments.

Business Objects offers a comprehensive e-BI platform by including:

- Microsoft Visual Basic for Applications (VBA) for customizing BUSINESSOBJECTS.
- A document catalog and viewer control that can be embedded easily into any application.
- APIs that give access to BUSINESSOBJECTS REPORTER, DESIGNER, and documents stored in the repository, and allow third parties to gain access to new sources of data.
- WEBINTELLIGENCE SDK for customizing WEBINTELLIGENCE.

Microsoft Visual Basic for Applications

BUSINESSOBJECTS embeds Microsoft Visual Basic for Applications (VBA). This replaces the ReportScript 4.1 language (SBL) with a world-class integrated development environment. VBA gives BUSINESSOBJECTS unprecedented customization, extensibility, and integration capabilities.

- Compatibility with ReportScript
  The language syntax of VBA and SBL is similar, so conversion of existing code is smooth. An SBL to VBA converter is provided which will assist in converting existing scripts.
- Customization
  Developers can customize the user interface to meet the needs of any group of users. Toolbar and menu items can be added, removed, or modified. VBA macros can be associated to toolbars and menu items.
- Extensibility
  VBA can be used to extend the capabilities of the core BUSINESSOBJECTS product. Users, third party suppliers, and Business Objects development and consulting services will be able to provide users with product enhancements outside of the regular product release cycle by developing add-in modules. BUSINESSOBJECTS 5.1 is more than a tool; it has become an e-BI platform.
- Desktop integration
  VBA is also used by Microsoft Office products and by an increasing number of other Windows software products. With BUSINESSOBJECTS 5.1, users will be able to leverage their existing investment -- both in desktop software and in development skills -- by integrating and customizing applications already available on the user's desktop, rather than building new systems from the ground up. EIS applications, for example, can be developed using the combined power of BUSINESSOBJECTS and Microsoft Office.

Document catalog and viewer control

An ActiveX control and an object model are provided to give developers access to the document catalog and to BUSINESSOBJECTS document viewing. The control can be embedded in any Windows application -- at a source code level -- to enable users to list, read, and refresh documents. Embedding Version 5 document-reading technology in another application is easy. Minimal software development effort is required to add a fully operational reporting module to an existing application. This control does not require other BUSINESSOBJECTS software to be installed on client PCs, because it makes use of server components.

Extensibility - build a new data provider

Because of its flexible data provider architecture, BUSINESSOBJECTS 4.x already allows users to access, analyze, and share information from a wide range of sources from Excel files on a PC to DB2 databases on a mainframe.

Version 5.1 is built on an extensible architecture so that in addition to offering a wide range of "out of the box" data providers, you can add new data providers. The new data providers can be built as COM objects, following specifications provided by Business Objects, and then integrated with BUSINESSOBJECTS 5.1 using VBA.

Extended object model

- BUSINESSOBJECTS REPORTER object model
  The BUSINESSOBJECTS 4.1 object model has been extended in version 5 to include universes, classes, objects (read-only), and basic document formatting. An event model is also delivered. Access to this object model requires BUSINESSOBJECTS REPORTER to be installed. In addition to increasing the scope for customization, this enables developers to build generic add-in software that will work for any BUSINESSOBJECTS customer.
- DESIGNER object model
  DESIGNER version 5 exposes a new object model for application developers. The DESIGNER object model lets designers create and modify universes in a programmatic, rather than manual, manner. The DESIGNER object model provides access to universes, classes and objects, predefined conditions, strategies, joins, contexts, hierarchies, and connections. Import and export to and from multiple universe domains is also provided. This means that universes can be automatically created, modified, imported, and exported, which will facilitate metadata exchange.

Chapter 5 First Steps: Installing and Setting Up the System

In this chapter

- **Introduction 129**
- **Fresh Installation or Update? 129**
- **System Requirements 130**
- **Business Objects Products and Windows 2000 130**
  Installing Business Objects Desktop Products 131
  Business Objects Enterprise Server Products 132
  Is Windows 2000 "privatization" supported? 135
- **Installing Business Objects Software 136**
  Segregating desktop products and enterprise server products 137
  Business Objects products 137
  First step for all types of installations from the CD 137
  Where are the programs installed? 137
  What products for what users? 138
- **Installing and Setting Up Enterprise Server Products 139**
  Server installation profiles - cluster managers and nodes 139
  Enabling modules on cluster servers 140
  Starting a cluster 141
  Configuring Business Objects services to start automatically 142
  Configuring a web server 142

- **Installing Desktop Products (Master Installation) 143**
  Choosing a client installation profile 143
  Choosing the modules you want to install 143

- **Creating Your Database Reporting Environment 144**

- **Creating a Repository 144**

- **Defining Your Users and User Groups (Supervisor) 145**
  The significance of user groups 145

- **Creating a Universe (Designer) 145**
  The quick design wizard 145
  Creating a universe manually 146
  Defining a secured connection to a universe 146
  Exporting universes to the repository 146

- **Creating Documents 147**

- **Updating from Previous Releases 148**
  From BusinessObjects version 3.1 to version 5.x 148
  From BusinessObjects version 4.1 to version 5.x 148
  From WebIntelligence version 2.0 to version 2.5/2.6/2.7 152

- **Installing Service Pack Updates 153**

- **The Freeware Directory 153**

Introduction

Once you have given the critical deployment issues discussed in the previous chapters adequate consideration, you are ready to actually get your Business Objects system up and running.

This chapter outlines the first concrete steps in the deployment procedure:

1. Installing Business Objects server and desktop products, either for the first time, or updating from previous releases
2. Setting up server products in a distributed system to make a functioning cluster
3. Creating your database reporting environment
4. Creating a repository
5. Defining users and user groups using SUPERVISOR
6. Creating universes using DESIGNER
7. Generating the documents that will be used in your Business Objects system

Fresh Installation or Update?

If this is the first time you're deploying Business Objects products in your company, you will be making a fresh installation.

If you're already using BUSINESSOBJECTS version 3.1 or 4.1, you will need to migrate your current version to the current version. This process varies depending on the version you are migrating from.

If you're already using WEBINTELLIGENCE 2.0, you will need to update to version 2.7.

If you're already using DOCUMENT AGENT SERVER, you will need to update to BROADCAST AGENT 5.5.

These topics are summarized in "Updating from Previous Releases" on page 148, and are thoroughly covered in the *Installation and Update Guide.*

System Requirements

System requirements may change with service packs.

- To make sure you have the latest information on supported platforms and products, read this release's ReadMe file.
- For a complete list of supported browsers, simply click the Supported Browsers button in the Browser Check page before logging into WEBINTELLIGENCE.

Business Objects Products and Windows 2000

Windows 2000 is based on Windows NT 4.0 technology. All Windows NT 4.0 features and functionalities are included or enhanced in Windows 2000.

The correspondences between the 4.0 and 2000 families are:

Windows NT 4.0	Windows 2000
Windows NT Workstation	Windows 2000 Professional
Windows NT Server	Windows 2000 Server
Windows NT Server + Options Pack	Windows 2000 Advanced Server

Installing Business Objects Desktop Products

BUSINESSOBJECTS (including its zero administration server deployment that can be installed through INFOVIEW) and all applets running inside a browser are supported on Windows 2000 Professional and Server platforms.

We advise against using BUSINESSOBJECTS via Terminal Server Services when many clients use the machine concurrently.

Installing BusinessObjects

To install BUSINESSOBJECTS for the first time on a Windows 2000 client machine, you must log on as a local administrator and run the AllowBO51-26-W2000.exe file. You must do this for each machine on which you wish to install BUSINESSOBJECTS.

Once you have run the AllowBO51-26-W2000.exe file on client machines, you can then run the installer. You must still be logged on as a local administrator. When you install BUSINESSOBJECTS, you must either keep the default installation folder hierarchy (i.e. `C:\Program Files\Business Objects\`) or select a folder for which the person who will use the installed product has write privileges.

Note: You must run this file only on client machines and not on the machine on which you installed your WebIntelligence server. This can cause unreliable operation of Enterprise Server products.

You are warned about the process failure with appropriate error messages if you try to install the product before running the AllowBO51-26-W2000.exe file.

You can find the AllowBO51-26-W2000.exe file on the installation CD in `Setup\Win2000`. It is also copied to the *`installation-folder`*`\Business Objects\Setup 5.0` folder if you select the Setup option when you install Enterprise Server or Desktop products.

An efficient method of accessing the AllowBO51-26-W2000.exe file in a deployment is to copy the files from the installation CD to a file server, then access *`file-server`*`\Setup\Win2000` over the network.

To run the AllowBO51-26-W2000.exe file:

1. Double-click the AllowBO51-26-W2000.exe file.
2. Before you install BUSINESSOBJECTS from INFOVIEW, allow 15 to 20 seconds while AllowBO51-26-W2000.exe is running.

Tip: To check if AllowBO51-26-W2000.exe is running, right-click on the blank area of the task bar, select Task Manager and click the Processes tab. If AllowBO51-26-W2000.exe is listed on the Processes tab, it is still running. If it is not listed, you can install BUSINESSOBJECTS.

Business Objects Enterprise Server Products

Although this version is Windows 2000 compliant, it is not Windows 2000 certified. WEBINTELLIGENCE can only run on Windows 2000 Server and Advanced Server platforms.

WEBINTELLIGENCE uses none of the new features introduced in the Windows 2000 family, such as Microsoft Installer (MSI) or of the Active Directory, which gives network users access to permitted resources anywhere on the network using a single logon process. Users can access WEBINTELLIGENCE resources from anywhere on the network using a single logon to the NT domain.

Thus unless there is a real need for the Windows 2000 Advanced Server features, we strongly recommend you deactivate them. WEBINTELLIGENCE will get the most out of the machine and system resources if Windows 2000 Server is installed on the machine, with basic services activated.

WebIntelligence SDK

When using WEBINTELLIGENCE SDK to customize WEBINTELLIGENCE:

- WICOM: ASP applications are supported with WEBINTELLIGENCE 2.7 on Windows 2000.
- WIBEAN: For a full list of JSP servers supported on Windows 2000, refer to the Busines Objects Installer Readme.

Web servers

Windows 2000 Server and Advanced Server contain Internet Information Server (IIS) 5. WEBINTELLIGENCE 2.7 supports this web server, and the install configures it automatically on setup.

For more information on web servers supported on Windows 2000, refer to the Installer readme.

Windows Components	Interaction with WebIntelligence	Recommendation
Certificate Services	None	Disable
Cluster Service	Not supported	Not supported
Indexing Service	None	Disable
FTP Server	None	Disable
FrontPage Server Extensions	May conflicts with WebI	Disable
NNTP Service	None	Disable
SMTP Service	None	Disable
RAD Remote Deployment Support	None	Disable
WWW Server	Connection Manager Components	None
Network Monitor Tools and Management protocol	None	

Windows Components	Interaction with WebIntelligence	Recommendation
Message queuing Services	None	Disable
COM Proxy	None	Disable
DNS	None	Disable
DHCP	None	Disable
Internet Authentication Service	Not supported	Disable
WINS	None	Disable
Simple TCP/IP Service	None	Disable
QoS Admission Control Service	None	Disable
Site Server ILS Services	None	Disable
File/Print Services for UNIX or Macintosh	Allows UNIX nodes to use the Windows 2000 machine printers	Enable
Remote Installation Services	None	Disable
Remote Storage	None	Disable
Script Debugger	None	Disable
Terminal Services / Licensing	Unacceptable performance	Disable
Windows Media Services	None	Disable

Is Windows 2000 "privatization" supported?

With Windows 2000 privatization, all applications install their own version of common system DLLs in the application's directory. The result is that many popular DLLs will exist in many system folders. They must be 'side-by-side' aware so that each of the applications will run properly without provoking system crashes.

BusinessObjects is not side-by-side aware. This is not an issue because Windows 2000 comes with system DLLs that are more recent than the ones we provide. Therefore, when you install Business Objects software on a Windows 2000 platform it won't install any system DLLs.

Installing Business Objects Software

In a Windows environment, you use the Business Objects installer to install both Business Objects desktop products and enterprise server products.

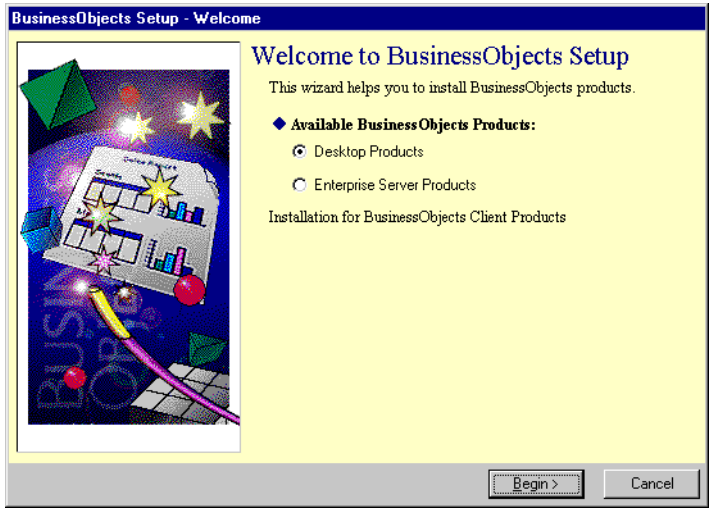

Install first the enterprise server products, then the desktop products on selected workstations, or vice-versa. The Setup Wizard installs both desktop and server products at once.

Many product components and mechanisms are shared between Desktop and Enterprise Server products. The installer manages the inter-relationships between common components at a file level, and also at a Registry level when installing or updating products (same entries on the registry, product keys, etc.).

To install Enterprise Server products on a UNIX server, see the *Installation and Update Guide*.

Segregating desktop products and enterprise server products

We strongly recommend that Desktop Products and Enterprise Server Products not be installed on the same server machine.

In particular, using BusinessObjects on the same server machine as Enterprise Server Products can cause unreliable operation of BusinessObjects and is not supported.

In addition, any Master installation of Desktop Products launched from a machine also housing Enterprise Server Products will not install the Setup utility, making subsequent uninstallation on the client machine impossible.

For these reasons, you should not install Desktop Products and Enterprise Server Products on the same server machine.

Note: Should you choose to ignore this recommendation at your own risk, it is essential that you install Enterprise Server Products first, and only then install Desktop Products.

Business Objects products

For a list of the supplied Desktop and Server products, see "The Business Objects Product Line" on page 28.

First step for all types of installations from the CD

Before installing or updating any product from the Business Objects CD, you must choose between installing or updating Desktop products or Enterprise Server products. Install Enterprise Server products, then Desktop products, because the installer for the server products dtects existing any existing Business Objects installation on the machine and will stop the installation process.

You must have administrator access rights for the machines on which you are installing the Business Objects products.

Where are the programs installed?

The Setup program manages the installation of all Business Objects products. All products are installed inside the main Business Objects folder. This is the folder in which all other folders will be created.

What products for what users?

Within your Business Objects deployment, not every user will need the same products. The products users use depends on their needs. The following table presents the Server product that should be used for each type of decision support requirement:

Example Group Profile	Recommended Server Products
Document consumers only	INFOVIEW For more information, see the *InfoView User's Guide*.
Standard ad hoc query, reporting and analysis users	WEBINTELLIGENCE For more information, see the *WebIntelligence User's Guide*.
Power users with a need for access to OLAP servers	BUSINESSOBJECTS For more information, see the *BusinessObjects User's Guide*.
Power users accessing RDBMSs only	Zero Admin BUSINESSOBJECTS
Remote power users with no direct database access	Zero Admin BUSINESSOBJECTS
Power users with no middleware installed on their PCs	Zero Admin BUSINESSOBJECTS
Users who need to perform offline analysis of existing documents	Zero Admin BUSINESSOBJECTS

Installing and Setting Up Enterprise Server Products

Enterprise server products must be installed on a Windows or UNIX server, on a computer which is neither the Primary Domain Controller (PDC) nor the Backup Domain Controller (BDC) of its local domain. For a complete list of supported Windows and UNIX platforms, see the ReadMe file for this version.

Under Windows, you can see if the computer you plan on installing to is a PDC or a BDC by choosing Start | Programs | Administrative Tools (Common) | Server Manager. Server Manager is a Windows administrative tool that lists the computers in your domain and gives their type.

In addition to Database Access Packs and OLAP Server Access Packs, there are four Enterprise Server Products that can be installed via the setup program:

- WEBINTELLIGENCE
- BROADCAST AGENT
- Zero Admin BUSINESSOBJECTS
- WEBINTELLIGENCE Software Development Kit (SDK)

Server installation profiles - cluster managers and nodes

You can install each of the server products on one server, or you can distribute your installation over several servers to balance the workload. The type of distribution you select is referred to as your **installation profile**.

You can choose either of the following installation profiles:

- Cluster Manager
  This is the default setting, which allows you to set up the main server in your network. This option also lets you set up a single-server deployment.
- Cluster Node
  When you have installed one cluster manager, the subsequent servers that you install are cluster nodes.

In all situations, you must first install a cluster manager. This deploys the product on one server machine. If you want to deploy the various product components over two or more server machines in order to balance the workload, you must add nodes to the cluster.

During installation, you must assign the cluster manager a unique OSAgent Port value. If you configure additional nodes, you must also assign the same OSAgent Port value to each of them to allow them to communicate with each other.

By default, the installer sets the OSAgent Port value to 14999. If an environment variable for this value has already been set on your machine, however, that value is displayed instead. In this case, decide whether you want to install WEBINTELLIGENCE on this port, or on another.

The value must be a number between 1024 and 65535, but must not be 15000, which is reserved for the Gatekeeper port. This value must be different from the port value of any other application server running on the machine, whether it be a CORBA, web, database, or other type of server.

If you're configuring a single-machine system or installing a cluster manager, make sure you're using an OSAgent Port value that is unique on your subnet.

Note: You must define an installation profile regardless of the server product you are installing.

What is installed when your cluster is complete?

When you have completed the installation process on the last cluster node in your system, *all* the machines in the cluster will contain *all* the server products. The installer also installs the OSAgent on the cluster manager, and the OAD on the cluster manager and all cluster nodes.

For more information on cluster managers and nodes, as well as setting up multiple cluster networks, see "Distributed Architecture Overview" on page 91. You can also find more detailed information in the *Installation and Update Guide.*

Enabling modules on cluster servers

If you have installed a distributed system, you need to configure precisely how the system functions by enabling specific modules on specific nodes. You do this using the Business Objects Services Administrator.

By strategically orchestrating where each process functions, you can optimize performance on your system. For more information, see "Load Balancing" on page 204.

In order for the system to function, all of its modules must be running somewhere within the cluster. Certain modules can be enabled on one cluster node only. Others can be enabled on several nodes at a time.

Module enablement recommendations

By default, after installation the WISessionManager and WIStorageManager are automatically deactivated on all cluster nodes.

- Although you can enable the WIStorageManager on a cluster node, we highly recommend that you enable it on the cluster manager, then enable it on nodes as well if required. For more information, see "The WIStorageManager" on page 206.
- Once you have enabled these modules, they must never be deactivated then enabled elsewhere.
- To change module enablement, you must respect the following steps:
  - Start the cluster manager and nodes.
  - Enable/disable the cluster's modules using the Administrator.
  - Stop and restart first the cluster manager, then each cluster node.

You define the distribution of functions across the machines in the system very simply, using the Administrator. For complete information on module enablement and the Administrator, see the *WebIntelligence System Administrator's Guide*.

Starting a cluster

In order for a cluster to function, the installed Business Objects products must be started on the cluster manager before being started on cluster nodes. This ensures that:

- The **OSAgent** is ready to register which modules are running on each cluster node and then provide the WIDispatcher with the addresses for machines running requested processes.
- The **WISessionManager** is ready to authenticate users as they try to log in, and track each user session until it ends.
- The **HTTP server** and **HSAL** are running, permitting communication between client browsers and the WIDispatcher.

The **WIStorageManager** must also be running in order for users to login and create sessions.

For more detailed information, see the *WebIntelligence System Administrator's Guide*.

Configuring Business Objects services to start automatically

Server products include two Windows services/UNIX processes:

- WEBINTELLIGENCE ORB (Windows and UNIX)
- WEBINTELLIGENCE Manager (Windows only)

Under Windows, by default, these services/processes are started automatically when you boot your server system. You can also configure your system so that these services must be started manually.

Under UNIX, the default installation requires that you start and stop the WEBINTELLIGENCE system manually. Two example scripts are provided, however, which allow you to set WEBINTELLIGENCE to automatically start at boot time and stop at shutdown. For complete information, see the *Installation and Update Guide*.

Configuring a web server

If you are installing WEBINTELLIGENCE, you must select a web server. You can choose to have the web server configured automatically by the Installer, or to configure it manually.

Installing Desktop Products (Master Installation)

To install desktop products from a CD for a shared Windows installation, you must choose the *Business Objects Desktop products* option from the installer Welcome dialog, and then select Master Setup (Shared).

Note: You can install one desktop product, BUSINESSOBJECTS, not just from the Business Objects installation CD, but through an Internet browser. This type of deployment of BUSINESSOBJECTS provides a slightly different functionality.

For information about this type of deployment, see "Zero Admin BusinessObjects" on page 48. For information about installing and using this type of deployment, see the *InfoView User's Guide*.

Choosing a client installation profile

You then determine the options that will be given to users when they install or update Desktop Products from the Master Setup:

- Standalone Installation: All application files are copied to the user's hard disk. The user will have no options during client installation.
- Shared Installation: All shareable application files remain on the server and run remotely. The user will have no options during client installation.
- User's Choice: The user is prompted to choose a shared or standalone installation. The user can customize the setup during client installation.

You must then select an installation profile (Express, Minimal, or Custom) and installation folders for the Master Setup, and default client installation folders for installations made from the Master Setup.

Choosing the modules you want to install

Finally, if you selected the Custom profile, you can select exactly which products and modules you want to install by checking the corresponding options.

Creating Your Database Reporting Environment

Once you have installed all the necessary server and desktop products, you can begin creating your own database reporting environment.

One of the most important tasks in any deployment is the creation of your repository using SUPERVISOR. The repository is used by the various modules to:

- Secure access between your Business Objects deployment and your data warehouse
- Provide a centralized storage location for the Business Objects applications

Your deployment relies on the power of the repository to support the advanced features of the products.

To set up your repository, you need to run both SUPERVISOR and DESIGNER, and perform specific actions in each.

The steps involved are:

1. Start SUPERVISOR, and create a repository.
2. Define your user groups, and set your user security levels.
3. Start DESIGNER, and create one or more universes and database connections based on your particular corporate database.
4. Define a secured connection between the universe and the corporate database.
5. Export your universes to the repository.

Each of these steps is described below. For more detailed information, see the relevant chapters of the *Supervisor's Guide* and the *Designer's Guide*.

Creating a Repository

You create a repository using SUPERVISOR. A Setup wizard is automatically started when you launch SUPERVISOR for the first time. This wizard guides you through the process of creating the repository.

Note that you must also have a suitable database available in which to create your repository (see Chapter 8 "The Repository" on page 275).

For more information on creating your repository, see the *Supervisor's Guide*.

Defining Your Users and User Groups (Supervisor)

After creating your repository, the next step is to define your users and user groups. You need to have these defined before you export a universe to the repository, since during the export you also need to specify which groups will have access to the universe.

The significance of user groups

When you export a universe, you must link it to one or more user groups. Similarly, when you configure a BROADCAST AGENT, you must link it to a specific user group. This is a Business Objects security feature since, by default, any user belonging to a group has access to the associated group resources (universes, documents, scripts, etc.). The supervisor can therefore implement security and allocate resources at a group level rather than at an individual user level.

You can find complete information on creating users and user groups in the *Supervisor's Guide*.

Creating a Universe (Designer)

Your next step is to create a universe based on your corporate database. To do this, you need to start DESIGNER.

The quick design wizard

For a universe based on a simple relational schema, DESIGNER provides Quick Design, a wizard for creating a basic yet complete universe. You can use the resulting universe immediately, or you can modify the objects and create new ones. In this way, you can gradually refine the quality and structure of your universe.

The Quick Design wizard assists you throughout the creation of a universe. It guides you in establishing a connection to the database and then lets you create simple classes and objects. The wizard also provides built-in strategies for the automatic creation of objects, joins, and tables.

Creating a universe manually

Alternatively, you may choose to create an entire universe manually. You will then need to build the underlying database structure of a universe. This involves selecting the initial tables and columns which form the classes and objects of your universe.

With its graphical and intuitive interface, DESIGNER makes it easy to create and organize a working universe in a short amount of time. From the basic model of a universe, you can then refine the classes and objects so as to obtain the degree of complexity or sophistication you desire.

For more information on creating universes, see the *Designer's Guide*.

Defining a secured connection to a universe

The final step in making the information in your corporate database available to your specified user groups is to define a secured connection between the universe and the corporate database. Only the Supervisor can do this, or the Supervisor can assign a Supervisor profile to the Designer so that the Designer can do it himself.

For information on how to do this, see Chapter 6, "Managing Resources," in the *Supervisor's Guide*.

Exporting universes to the repository

After it has been created, a universe must be exported (stored in the repository) and made available to specific groups. It can then be imported for use by users.You export and import universes from DESIGNER or SUPERVISOR.

Export

You can only export universes that are present locally on your disk drive, and to which you have been given access (with a Supervisor or Designer profile). Also, you can only export to groups you belong to (with a Supervisor or Designer profile).

Import

You can only import universes to which you have been granted access (i.e., a universe that is assigned to one of the groups to which you belong).

Creating Documents

Once you have set up your reporting environment as described in the preceding sections, you are ready to start generating documents. For complete information on document creation, see the *BusinessObjects User's Guide* and the *WebIntelligence User's Guide*.

In addition to basic document generation, you can also schedule the automatic generation and distribution of documents to specific users and groups. This feature is discussed in greater depth in Chapter 10 "Document Processing and Distribution with Broadcast Agent" on page 343.

You can also allow users to access documents via a web browser, using WEBINTELLIGENCE INFOVIEW. See "Does the system have a single e-Business Intelligence portal?" on page 104 for more detailed information.

Updating from Previous Releases

Can you update from one Business Objects release to the most recent one without risking the considerable investment you made in your original deployment? Absolutely, but only if you follow the procedures described in detail in the *Installation and Update Guide*.

To give you an idea of what's involved, however, here's a summary for each type of update.

From BusinessObjects version 3.1 to version 5.x

BUSINESSOBJECTS deployments using version 3.1 cannot be updated directly to version 5.x. Only version 4.1.x can be updated directly. To go from version 3.1 to version 5.x, you must first update to version 4.1.x. Customers going from version 3.1 to version 4.1 should have received BUSINESSOBJECTS 4.x and instructions for this update.

For more information on updating universes and users from BUSINESSOBJECTS version 3.1 to 4.x, see the Appendix of the *Supervisor's Guide*, "Converting Universes and Users for Update from Version 3.1 to Version 4.x."

From BusinessObjects version 4.1 to version 5.x

Installing the software

With the exception of a stand-alone installation of version 4.x Desktop products, previous installations of Business Objects products must be uninstalled before installing the latest release of Business Objects Desktop or Enterprise Server products.

When installing Desktop products on top of a 4.1 Standalone installation, the Setup wizard detects the 4.1 Standalone setup and allows you the choice of replacing your 4.1 version or installing the 5.x Desktop products in a separate Standalone installation on the same machine.

If you choose to replace the 4.1 version, it is uninstalled and the 5.x release installed in its place. The new release will use the existing 4.1 folders for documents, universes, scripts, user libraries, and security and administration files. You will be able to reuse your 4.1 files directly from the new release (see below for details on updating these files).

If you choose to keep your 4.1 version, a separate file structure will be created to install the 5.x Desktop products. To use the 4.1 files, you will first have to transfer them manually to the appropriate 5.x folders. You will be able to uninstall your 4.1 release at any time using the 4.1 Setup program, without affecting the 5.x installation.

The repository

The repository has undergone few changes. A few tables have been added or expanded to account for new features and enhancements, and no tables, rows or columns have been deleted.

To update a BUSINESSOBJECTS 4.1 repository, you must first install the 5.x software. Then log into SUPERVISOR as the General Supervisor who created the repository to be updated. Any other login will be refused and an error message will be returned. Upon login you will be asked if you want to update the repository to version 5.x:

- Choosing Yes updates the repository immediately.
- Choosing No exits the application.

Once a repository has been updated to version 5.x, it can no longer be used in version 4.1.

Universes

BUSINESSOBJECTS 5.x universes have exactly the same format as version 4.1. Universes created with version 4.1 can continue to be used with no changes.

Documents

The following table describes what 4.1 and 5.x users can do with specific versions of BUSINESSOBJECTS documents:

	From version 4.1	From version 5.x
Users can read 4.1 documents	Yes	Yes
Users can read 5.x documents	No	Yes
Users can access and publish 4.1 documents in a 5.x repository	Yes	Yes When a 5.x user saves a 4.1 document or publishes it to the repository, it is converted to 5.x format. This means 4.1 users can no longer read the document.
Users can schedule 4.1 documents with Broadcast Agent 5.x	Yes BROADCAST AGENT detects the format of documents it processes and saves them in their original format.	No
Users can schedule 5.x documents with Broadcast Agent	No	Yes

Note: 4.1 users prior to Service Pack 6 can successfully export (send or publish) a 4.1 document to a 5.x repository that already contains a 5.x version of the same document. Afterward, however, those users can no longer re-import the document.

With Service Pack 5, 4.1 users get a warning message before exporting the document. With Service Pack 6, users can actually re-import the document.

BusinessObjects internal scripting language

BUSINESSOBJECTS now uses Microsoft VBA (Visual Basic for Applications) instead of Mystic River SBL (Softbridge Basic Language) as its internal scripting language. BUSINESSOBJECTS 5.x includes a module that assists you in converting your existing SBL ReportScript scripts into VBA projects.

To access the module, in BUSINESSOBJECTS choose Tools / Macro / Convert from ReportScript.

For further information, see the *BusinessObjects SDK Reference Guide*.

Document Agent Server to Broadcast Agent

BUSINESSOBJECTS 4.1 used DOCUMENT AGENT SERVER (DAS) to provide scheduled or batch processing of documents. BUSINESSOBJECTS 5.x has improved this component, now called BROADCAST AGENT.

If you are updating a Business Objects deployment including one or more DOCUMENT AGENT SERVERs, you must use the following sequence of operations:

1. Stop any DAS that is running.
2. Run the Setup Wizard to update SUPERVISOR to version 5.x.
3. From SUPERVISOR, update the repository to 5.x format.
4. Run the Business Objects Setup Wizard to install BROADCAST AGENT.

The settings and scheduled tasks of the DOCUMENT AGENT SERVERs are inherited without change by the BROADCAST AGENTs.

If your updated deployment is running a BROADCAST AGENT and still includes DOCUMENT AGENT SERVERs, each DAS then functions as a console for the running BROADCAST AGENT: it can be used to monitor the operation of the BROADCAST AGENT, but no interaction is possible.

From WebIntelligence version 2.0 to version 2.5/2.6/2.7

If you have previously installed WEBINTELLIGENCE 2.0 on your system and you are updating to WEBINTELLIGENCE 2.5, 2.6 or 2.7, you will need to update:

- The cluster manager
- All the cluster nodes
- The repository the cluster uses

You can update your cluster in two ways:

- By shutting down your WEBINTELLIGENCE 2.0 system completely, then migrating to version 2.5/2.6/2.7
- By keeping your 2.0 cluster manager running as long as possible until the 2.5, 2.6 or 2.7 cluster is complete or nearly complete

For complete information, see the *Installation and Update Guide*.

Documents

WEBINTELLIGENCE INFOVIEW users can read WEBINTELLIGENCE 1.0 and 2.0 documents, but if they save, send or publish these documents, they are automatically converted into 2.5/2.6/2.7 format.

The same users can view and refresh BUSINESSOBJECTS 4.1 documents, but can save nor publish them.

Installing Service Pack Updates

When you install Service Pack updates on top of an existing installation, the installer detects if there are already products installed.

When there are, the wizard goes straight to a custom installation screen allowing you to choose which products to install (either Enterprise Server or Desktop products depending on which Installer option you chose).

Although we strongly discourage this type of configuration, you may have installed Enterprise Server and Desktop products on the same machine for evaluation or demonstration purposes. In this case, you must update both Enterprise Server and Desktop products, by launching the Setup wizard for each set of products. This requires running the Setup wizard separately for each type of product.

Reminder: You cannot simply replace BUSINESSOBJECTS 4.1 with version 5.1, or WEBINTELLIGENCE 2.0 with 2.7! You will find a detailed procedure for updating in the *Installation and Update Guide*.

The Freeware Directory

A Freeware directory is provided on the Business Objects CD. This subdirectory contains supplementary information and utilities, which are for use at your discretion, but are not supported elements of the software.

Chapter 6 System Configuration, Sizing, and Benchmarking

In this chapter

- **Overview 157**

- **Performance Factors and Sizing 157**
  Factors impacting system performance 158
  How many users will be using the system? 159
  What types of documents will be processed? 161
  How big are your universes? 163
  What type of database or data warehouse will be used? 164
  How fast is the repository database engine? 164
  How much space do you need to allocate for the repository? 164
  How fast is the network? 168
  What are the optimal server machine configurations? 170
  How should you set Windows heap settings? 173
  How can you get around UNIX thread per process constraints? 174
  How should system components be distributed? 174

- **Supported Configurations 175**

- **Example Distributed Deployment Implementation 176**
  Number of registered users 176
  Maximum number of users 176
  Types of documents that will be processed 177
  Type of architecture and number of nodes 177
  Required memory 179

- **Specific Deployment Guidelines 181**
  Deploying Broadcast Agent 181
  Deploying WebIntelligence 183
  Providing cluster manager fault tolerance 184
  Cisco LocalDirector 189
  Deployment limitations 195
  Deploying multiple clusters 195
  Installed server components 197

- **Benchmarking 198**
  Basic methodology 198

Overview

This chapter describes the major performance factors for each type of deployment environment, and provides guidelines on how to deploy the server components of a Business Objects architecture.

Each deployment is different. It depends on the number of users supported, the type of information access that you want your users to have, the number and types of documents and other files, as well as the databases you are using, the geographical distribution of your company, and so on.

There are many variables to consider when measuring performance as well. These include the network speed between the clients and the middle tier, and the network speed between the middle tier and the databases.

The factors and formulas described in this chapter give you a starting point for deploying your system. We suggest you implement one or more clusters with a supported architecture, containing as many CPUs, memory and disk space as needed to support the system constraints.

Once you have set up your cluster, you'll need to plan a series of benchmark tests that will allow you to see if your system meets performance requirements. For information, see "Benchmarking" on page 198.

To ensure high availability, performance and maximum scalability, the system's components can be distributed across different servers to take full advantage of the server's resources, and to enable load balancing. For this type of information, see Chapter 7, "System Tuning and Troubleshooting" on page 201.

Performance Factors and Sizing

Some performance factors affect all deployment types, while others are specific to the deployment environment. Before setting up your system, you need to know what type of deployment environment best meets your requirements, as well as how many users will be using the system. Afterward, you can compare your needs to the specific deployment types while taking into consideration the performance factors for each type.

Factors impacting system performance

Performance Factor	Standard Client-Server Deployment	Distributed Architecture Deployment
The database or data warehouse	x	x
The repository database	x	x
The network	x	x
The number of concurrent active users	x	x
The types of documents	x	x
The size and complexity of the documents	x	x
The number of documents	N/A	x
Universe size and complexity	N/A	x
The server machines (CPU, memory, disk space)	x	x
The server system settings (CPU usage, memory usage, I/O and disk access, number of threads per process)	x	x
User activity	N/A	x
The distribution of CORBA components	N/A	x
Intranet vs. extranet	N/A	x

N/A = Not Applicable

The following sections describe these factors in more detail.

Note: These factors and formulas give you a starting point for deploying your system.

The optimal solution is a cluster with the recommended architecture containing as many CPUs as needed to support the system constraints.

Once you have set up your cluster, you need to plan a series of benchmark tests (see "Benchmarking" on page 198) that will allow you to see if your system meets performance requirements.

How many users will be using the system?

There are two important ratios you need to determine before actually deploying a distributed architecture deployment:

- The user activity ratio at peak usage times
- The total supported users ratios

User activity ratio

Only needed in a distributed architecture deployment, this ratio is defined as:

logged active users : logged users

This ratio particularly impacts the server's CPU usage.

A **logged active user** is an active user who performs an action that invokes server CPU, such as creating or refreshing a document. These users consume both processor and memory resources.

A **logged user**, however, is a user that is only logged into the system, but is not currently using CPU. These users consume memory resources and one database middleware connection, but no processor resources. Typically, this corresponds to a user who has performed an action such as refreshing a document, but is now viewing the document without invoking server CPU.

Each time a user logs into WEBINTELLIGENCE, a session is created and a WIQT instance is launched. By default, it stays in memory for one minute after the last user activity (the time-out can be set using the Business Objects Services Administrator).

Note: If a user is logged in but inactive for a duration exceeding the *Max WIQT active time* parameter on the Business Objects Services Administrator, the WIQT instance is deactivated and the user no longer consumes these resources.

Similarly if the user is inactive for a duration exceeding the time-out for the database connection, the database connection is deactivated and the user no longer uses the database middleware connection.

Total supported users ratios

All logged users have their own WIQT process running (from the time they login until the time they logout, or until the process times out). Therefore, **the system must be able to support as many running WIQT processes as logged-in users**. The following formula tells you how many WIQT processes you need:

number of registered users * ratio of logged users at peak usage time

For example, your system has 1000 registered users, but only 50% of them are actually logged in at any one time. You therefore will need 500 running WIQT processes.

How many BOManagers should you enable? This depends on how many BUSINESSOBJECTS processes you need in the cluster. As a general rule of thumb, allow for:

- One BUSINESSOBJECTS process for each user using full-client documents
- A maximum of 5 BUSINESSOBJECTS processes per processor

You can set the maximum and minimum number of these processes you want to run on any BOManager-enabled node.

The following formula can help you estimate how many users are actually processing BUSINESSOBJECTS documents:

number of registered users * ratio of logged users * ratio of logged active users * ratio of BUSINESSOBJECTS documents

Example **How many running BusinessObjects processes will you need?**

Your system has 1000 registered users, 60% of whom are logged in at any given time. Only 25% of all logged in users are actually generating or refreshing documents. In general, 40% of the documents being processed by the system are BUSINESSOBJECTS documents. Applying these figures to the formula, you get:

1000 * 60% * 25% * 40% = 60 BUSINESSOBJECTS processes

What types of documents will be processed?

The Business Objects system provides full processing of two types of documents:

- Full-client documents (created using BUSINESSOBJECTS REPORTER)
- Thin-client documents (created using WEBINTELLIGENCE REPORTER)

There are different advantages in using each type of document within each environment. After reviewing the differences between the two types of documents, you can decide what deployment is best suited to your user's needs.

Note: With version 5.1/2.7, the Business Objects system can provide access to all types of files, and allow users to upload, download, send and save them. These files cannot, however, be edited or scheduled. See "What other types of files can the system process?" on page 163.

What's the difference between BusinessObjects and WebIntelligence documents?

The most significant difference between full-client and thin-client documents is that full-client documents can be much more complex than thin-client documents due to additional reporting features in the full client product. Thin-client documents, however, can allow for much larger deployments. Other differences include:

- Refreshing a full-client document in the distributed architecture deployment requires more resources than refreshing a thin-client document, because of the ability to handle more complex documents built into BUSINESSOBJECTS.
- BUSINESSOBJECTS users can use a scheduling option for full-client documents called *Faster Document Viewing Over the Web* that generates either a Windows metafile, HTML or PDF version of the document. Click Options to select the format.

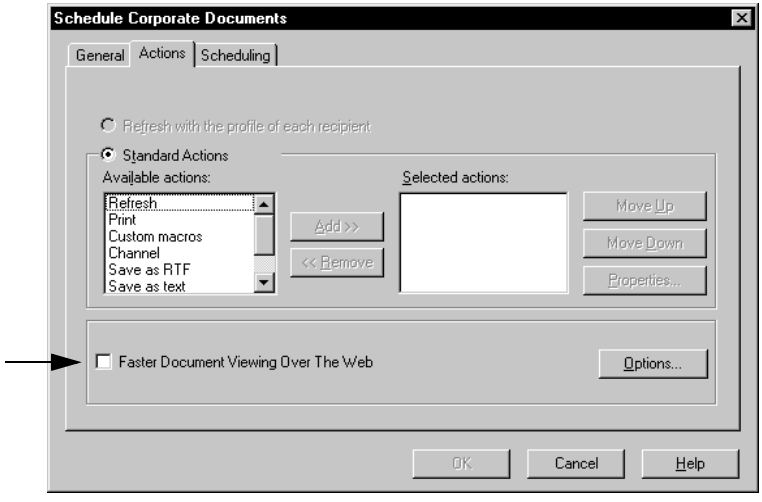

  Once a recipient of the document opens it in WEBINTELLIGENCE, the document's file is added to the system's cache mechanism. When the document is subsequently requested in that particular presentation, the file is retrieved directly from the cache, thereby reducing time and system load.
- Formatting is lost when a thin-client document is opened in BUSINESSOBJECTS.

What other types of files can the system process?

With this release, Business Objects users with the appropriate rights can upload, view, send, publish and download solution can process any type of file, whether it be a Microsoft Word document, Excel sheet, Adobe PDF file or other. These files are included in the document lists which form the core of INFOVIEW.

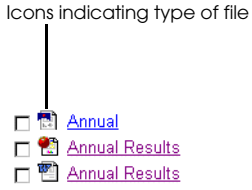

When they are downloaded, or viewed, the associated application is automatically launched if it resides on the user's computer.

BUSINESSOBJECTS users can even save their full-client documents in Adobe Acrobat PDF format once they've been automatically processed by BROADCAST AGENT.

How big are your universes?

In some large enterprises, universes can contain literally thousands of objects. Typically, when a universe contains this many objects, the time it takes to refresh a simple document is considerably lengthened. Here's why:

When the WEBINTELLIGENCE INFOVIEW user clicks Refresh, the WEBINTELLIGENCE system runs its SQL engine to re-generate the document's SQL query. To do this, the engine must load the entire universe into memory. If the universe contains multiple joins and contexts, the SQL generation is that much more complex and time-consuming.

Therefore, we suggest you make your universes as simple and manageable as possible. It's easier for the users to find the objects they need, and your system's document computation time will be optimized.

What type of database or data warehouse will be used?

The speed of the database engine can play a significant role in overall system performance. The number of open middleware connections may also be limited for a given database. If this number is limited, the number of users logged into the system will also be limited.

Make sure you verify with database vendors how many open connections the database or data warehouse can support per session, both for the repository and corporate data databases.

How fast is the repository database engine?

The speed of the database engine containing the repository can play a very significant role in overall system performance. Therefore, if possible, the repository should be run on a dedicated machine. If the choice of the repository database has not been made, be sure to compare available vendor benchmark numbers before making the final decision.

If you have the opportunity to choose a database platform for your repository, you should consider choosing databases that support:

- Row-level locking
- BLOBs (Binary Large OBjects)

How much space do you need to allocate for the repository?

A repository contains three parts:

- The **security domain** contains the characteristics of the other domains, as well as user definitions and access rights.
- **Universe domains** are meta-models of the database containing a description of the data to be accessed.
- **Document domains** contains the structures for storing shared documents and for executing tasks according to a timestamped definition.

When initially installed, the repository takes up roughly 1 MB. Its size increases as data is added. The security and universe domains expand relatively little per user, group or universe. It is the document domain that typically takes up the most space.

The security domain

The security domain is the core of the repository and is used to store all information about users, the groups they belong to, the applications and features they can use, the universes they have access to, the documents they have shared, and the list of tasks submitted to BROADCAST AGENT.

Space requirements per user

When they are created by the supervisor, all users require 4 KB of disk space in the security domain. If they belong to more than one user group, however, users should be counted as many times as they appear in the repository. For example, if particular users belongs to five users groups, they should be counted five times.

Later on, users will need 1.5 KB per constraint (link between the users and documents, accessible universes, etc.) + 0.4 KB per document + 1.5 KB per task submitted to BROADCAST AGENT, of disk space in the security domain.

To be safe, we consider that an average user will have 100 constraints, 100 documents and 100 tasks. **Thus the average user will need up to 344 KB in the security domain. By providing a safety margin, we round this figure up to 400 KB, or 0.4 MB, for each user.**

You can therefore calculate the initial required disk space for the security domain using this formula:

number of registered users * 0.4 MB

Security complexity as a performance factor

Because of the quantity and complexity of information in the security domain, it may grow and become large in a way that may not necessarily be linked to the number of users. The compexity of the security being used becomes another important performance factor, as the system must check these security measures during login or while processing documents. For instance, assigning a user to more than one user group has an impact on performance, since the system must check the user's rights by executing SQL to compare and compile the rights for the user in each group.

Another example is the SUPERVISOR feature which allows supervisors to secure certain commands. Use this feature with care, since it too requires extra processing.

You can now also set up several security domains, possibly distributed in different geographical locations. WEBINTELLIGENCE does not technically support security domain selection, but there is a way to connect to different repositories. For information, see "WebIntelligence and security domain selection" on page 279.

The universe domain

The universe domain is used to store the contents of each shared universe. This is the area to which designers can export the universes that they create. Each universe is a meta-model of the related corporate database.

For every universe domain, a document domain must also exist in the same database account. This document domain will store the list of values files.

You can calculate the initial required disk space for the universe domain using this formula:

number of universes * 0.5 MB

Note: The 0.5 MB figure is twice the size of the average .unv file, but it is advisable to include a sizable margin for safety.

The document domain

The document domain is used to store the contents of each shared document and file. This is the area to which users can publish Business Objects documents, or upload other types of files, to be shared by all users. This is also where documents such as templates, scheduled documents, and documents or files sent from users to other users are stored. BROADCAST AGENT accesses the document domain to run automated Business Objects document-processing tasks.

To begin with, you should allow 150 MB for the document domain. **Be sure, however, to monitor how fast the content grows, and add more if necessary!** As BUSINESSOBJECTS 5.1 and WEBINTELLIGENCE 2.7 can process all other types of files, you will need to be especially vigilant about allowing enough space for those files as well.

Note: Before any type of file or document other than those with add-ins (BUSINESSOBJECTS, WEBINTELLIGENCE or third-party) is sent to a 5.x repository, it is compressed. BUSINESSOBJECTS also has extended support for BLOBs, yielding better performance (results also depend on the BLOB-handling capabilities of your database).

In many deployments, administrators choose to create multiple repository domains to support the requirements of their organizations. For instance, it is possible to create a deployment that features a single security domain, two universe domains and five document domains. See “How to set up multiple repositories” on page 290.

Reducing the size of a document domain

To reduce the size of a document domain, run the Scan and Repair utility on the document domain in SUPERVISOR. Deletions in the repository are "logical" deletes; i.e. the records are still in the database. Scan and Repair should clean out the deleted records.

You can also use purge all Inbox documents older than a given number of days from the repository with a single SUPERVISOR command. This lets you recover the document domain space taken up by documents which have been sent but not yet collected by users. This feature can be especially useful when regularly scheduled documents are sent using the *Refresh according to the profile of each recipient* option.

Note: You may find that users are not deleting documents that have been stored in the repository. You should encourage users to delete old or outdated documents on a regular basis.

Initial space required for the repository as a whole

To calculate the initial required repository disk space in MB, use the following formula:

(number of registered users * 0.4) + (number of universes * 0.5) + 150

This corresponds to:

(space required for the security domain) + (space required for the universe domain) + (minimal initial space required for the document domain)

How fast is the network?

Network bandwidth has a significant role in performance for all deployments. You should pay particular attention to the following factors:

- The speed of the network between the server products and the repository, and between the servers and the corporate databases
- The speed of the network between cluster nodes in a load-balancing configuration
- The speed of the network between the client machine (browser) and the servers

Since TCP/IP (Transmission Control Protocol/Internet Protocol) is independent of any specific physical network hardware, it can run over Ethernet, Token Ring, a dial-up line, etc. Ethernet is recommended because it was designed to be easily expandable and meet the networking needs of a given site. Although the majority of computer vendors today equip their products with 10-Mbps Ethernet LANs, the 100-Mbps is becoming more widely adopted and is recommended. For system optimization, we also recommend the use of fiber channels.

Extending systems by using hubs

To help extend Ethernet systems, networking vendors sell **hubs** to provide multiple Ethernet ports.

There are two major kinds of hubs: repeater hubs and switching hubs.

We recommend switching hubs because each port of a packet switching hub provides a connection to an Ethernet system that operates as a separate Ethernet LAN. Unlike a repeater hub, whose individual ports combine segments together to create a single large LAN, a switching hub makes it possible to divide a set of Ethernet systems into multiple LANs that are linked together by way of the packet switching electronics in the hub. The round trip timing rules for each LAN stop at the switching hub port. This allows you to link a large number of individual Ethernet LANs together. While an individual Ethernet LAN can typically support anywhere from a few up to several dozen computers, the total system of Ethernet LANs linked with packet switches at a given site may support many hundreds or thousands of machines.

In Token Ring systems, switching can double the bandwidth of a single power-user or server by suspending the token passing protocol. Switching can transfer packets from one ring to another with minimal delay. Through increased aggregate bandwidth, reduced transfer delay and the ability to direct numerous parallel conversations, the productivity of a Token Ring network can rise to remarkable heights. It can handle about 80 percent utilization before performance degradation, as opposed to 30 to 60 percent utilization in Ethernet.

Therefore use 100-Mbps LANs with switched hubs (supporting duplex capability) if possible since it enables the maximum bandwidth efficiency.

Networks in web solutions

In a clustered server environment, the network can become a major performance bottleneck. The selection of the network architecture directly influences the network performance relative to the server. A powerful Network Interface Card (NIC) like 32-bit Bus Mastering Cards should be used.

The network load on the server should also be balanced. This means you should distribute the heavily used network segments between two or more NICs and configure the cluster nodes NICs to Duplex operation (simultaneous send and receive effectively doubling the network bandwidth).

We suggest you place each node of the cluster on a dedicated segment. Place the RDBMS on one segment of the switched hub if possible, or else make sure the cluster has the fastest link possible to the RDBMS. This minimizes network traffic and maximizes the cluster's bandwidth.

How can you identify potential network bottlenecks?

To monitor the network traffic and identify potential bottlenecks you can use:

- For Windows, the Performance Monitor or the Net Mon command.

  You can measure traffic a product generates for common actions. You can also test these actions for a given number of active users. Then you can extrapolate these measures for the projected number of active users.
- For UNIX, use the `netstat` command. Overall, it displays network statistics such as the number of packets in and out, the number of collisions, erros, etc. You can use the following flags, however, for more detailed information:

Flag	Information
`-a`	Details for active connections
`-i`	Activity on local machine interfaces
`-s`	Protocol statistics
`-m`	Streams allocation

  You can also use the `vmstat` command for statistics about disk, processes, virtual memory, trap and CPU activity.

You can also use the auditing feature to monitor Zero Admin BUSINESSOBJECTS, WEBINTELLIGENCE and/or BROADCAST AGENT activity as well. For more information, see the *WebIntelligence System Administrator's Guide*.

What are the optimal server machine configurations?

Hardware

WEBINTELLIGENCE, BROADCAST AGENT and ZERO ADMIN BUSINESSOBJECTS are server applications and therefore need server configurations. Of course, the type of server machine you need depends entirely on the type of deployment you have, the number of concurrent active users, and all other performance factors. However, as a general baseline, Windows servers should be at least of a class equivalent to Dual 500 MHz Xenon CPUs with 2 MB L2 cache. UNIX servers should be at least of a class equivalent to Sun Enterprise 450.

CPU

WEBINTELLIGENCE and BUSINESSOBJECTS document generation can be very processing-intensive.

The best performance is obtained with two CPUs. Therefore, the most limiting factor for greater performance is CPU power. A document that is large because a large graphic or video has been embedded in it will be processed much more quickly than a document of the same size containing pages and pages of HTML.

The number of supported modules on a CPU depends on your performance requirements.

- An average Windows 350 Mega Hertz CPU can support 5 BUSINESSOBJECTS processes or 15 WIQTs.
- An average UNIX 400 Mega Hertz CPU can support 20-25 BUSINESSOBJECTS processes or 25-60 WIQTs. These numbers are valid for up 16 CPUs per node.

Hard disk space

Installation requires from 300-400 MB depending on the number of languages installed.

Distribution over multiple disks on a single UNIX server

If you distribute your WEBINTELLIGENCE system over several disks on a single UNIX server, we recommend you use 4 distinct partitions (each partition to be placed on a separate disk; a single partition may be defined within one, or over several stripped disks):

- Disk 1: the WEBINTELLIGENCE software, including documentation
- Disk 2: the `WIDATA` directory and storage, containing user data
- Disk 3: the `WEBITEMP` directory, a temporary dirctory used by all concurrent users (whether this requires an entire disk, however, depends on how much user activity the system is supporting)
  This partition should be the most efficient for read/write access (for example, in the case of a write caching mechanism)
- Disk 4: the directories containing log and audit information.

Memory

Memory usage per user is linked directly to the size and type of the document being manipulated -- larger documents and full-client documents use more memory. Using small thin-client documents conserves memory.

System performance scales upward by adding more servers, rather than keeping the same number of servers and adding extra memory and/or processors on existing machines.

Minimum system requirement

The minimum system requirement for RAM (in MB) is calculated by:

(100 * number of cluster nodes) + (50 * ratio of BUSINESSOBJECTS documents * ratio of logged active users) + (30 * ratio of logged users * number of registered users)

Minimum requirement for each machine

For documents smaller than 3 MB, the RAM (in MB) required by each machine is calculated by:

For NT	**100 + (25 * number of BUSINESSOBJECTS processes supported by the machine) + (15 * number of WIQT processes supported by the machine)**
For UNIX	**180 + (14 * number of BUSINESSOBJECTS processes supported by the machine) + (8 * number of WIQT processes supported by the machine)**

For documents larger than 3 MB, the RAM (in MB) required by each NT machine is calculated by:

100 + (3 * average size of documents * number of BUSINESSOBJECTS processes supported by the machine) + (30 * number of WIQT processes supported by the machine)

For examples of these formulas, see "Example Distributed Deployment Implementation" on page 176.

Virtual memory

Virtual memory is related to the server's memory. On a Windows server, use a fixed size for the Paging file size rather than a dynamic one (Control Panel | System | Performance tab).

Disk access

Business Objects server products use the disk for a number of interactions. We recommend you use the fastest disk possible on the server running the WIStorageManager and WISessionManager, as the throughput on this server is one of the critical aspects for performance.

Under UNIX, the `WEBITEMP`, `<Application Data>` (this directory's location is defined at installation) and `<Storage>` directories are the most resource intensive.

How should you set Windows heap settings?

Due to limits in Windows resources, the maximum number of active processes per machine is 180. We strongly recommend, however, that you plan for only 150 active processes per machine.

In Windows NT 4.0, memory that can be allocated to the desktop is limited to 48 MB. Each BUSINESSOBJECTS process started in a virtual desktop consumes an equal part of this memory.

The *desktop* variable is critical for IIS 4.0. A desktop represents the context in which a process is running. Several processes can be running at the same time in different desktops:

- The interactive desktop is associated with the interactive, logged-on user (if any)
- Any number of non-interactive desktops can exist to run service processes and COM server processes.

Each time an application creates Windows objects—such as windows, menus, strings, etc.— it uses some memory from its desktop heap. If for some reason system memory runs out for a given desktop, applications will fail in various and unpredictable ways.

Entering a higher heap size for a given desktop type enables more processes to run concurrently in this desktop, but will limit the number of separate desktops that can be created simultaneously (due to the 48 MB limit). Conversely, lowering the heap size limits the number of processes that can run concurrently in a given desktop type, but enables the creation of a larger number of separate desktops.

There are three Windows settings which concern heap size:

- Global Heap
  Defines the amount of system memory that Windows reserves globally for all desktops. This value cannot be modified.
- Interactive
  Defines the amount of system memory that Windows reserves for the interactive desktop. The default value is 3072 and should generally not be modified.
- Non-Interactive
  Defines the amount of system memory that Windows reserves for each non-interactive desktop. If this value is empty, non-interactive desktops share the same setting as interactive desktops. We recommend that you set this value to 512.

  The default value for this setting in a standard Windows installation will be empty. However, installing the Windows NT Option Pack—or specific applications such as Internet Information Server 4.0— will automatically set this value to 512, which is the setting currently recommended by Microsoft. This variable may be changed by using the servconf utility.

How can you get around UNIX thread per process constraints?

Under UNIX, the maximum number of threads per process is limited to 1024. You can get around this constraint by using multiple instances of a module such as the WIGenerator on a single machine. For information about how to do this, see the *WebIntelligence System Administrator's Guide*.

How should system components be distributed?

To ensure high availability, performance and maximum scalability, the WEBINTELLIGENCE system uses a distributed component architecture (DCA). Their components can be distributed across different servers to take full advantage of the server's resources, and to enable load balancing. For more information about load balancing, see page 204.

Supported Configurations

We support three types of architectures:

- **Architecture A**

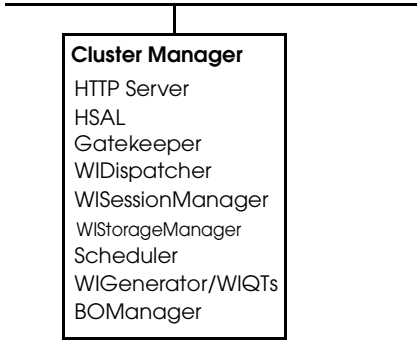

- **Architecture B**

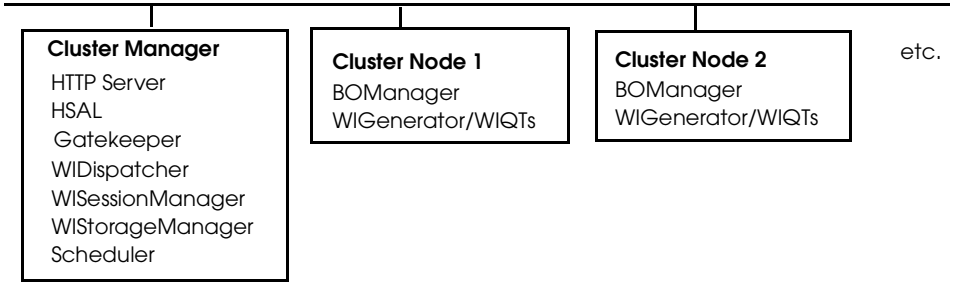

- **Architecture C**

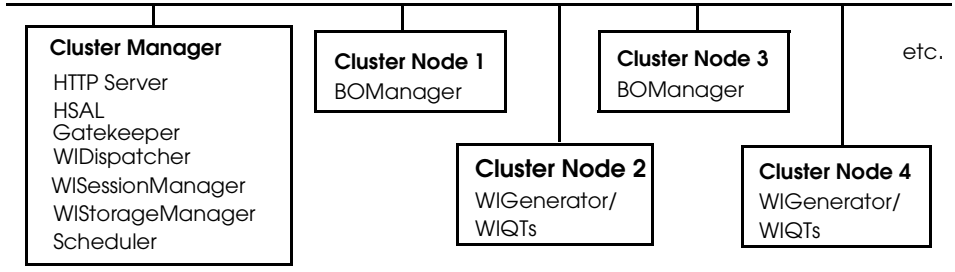

In these configurations, Architecture A illustrates a single server hosting all components. Both Architectures B and C require the same number of nodes (CPUs) to cope with the same constraints, but we recommend architecture C for better performance.

Example Distributed Deployment Implementation

This section provides a step-by-step solution implementation. Each subsection represents the type of information the administrator must have before beginning the concrete deployment of the solution.

Number of registered users

The number of registered users: 900

Maximum number of users

- 40% of the registered users will be logged at any time.
- 40% of those logged users will be actively processing documents at peak time.

Types of documents that will be processed

- 30% of the users will process BUSINESSOBJECTS documents
- 70% of the users will process WEBINTELLIGENCE documents

Of these documents:

- 40% contain 30 rows
- 40% contain 300 rows
- 20% contain 3000 rows

Type of architecture and number of nodes

The system will be supporting:

- 360 logged users at peak times
- 144 active logged users
- 44 users accessing BUSINESSOBJECTS documents

Therefore a maximum of 360 WIQT and 44 BusinessObjects processes need to be supported at any given time by the system.

- To cope with this demand we will use a type C architecture for a windows deployment:

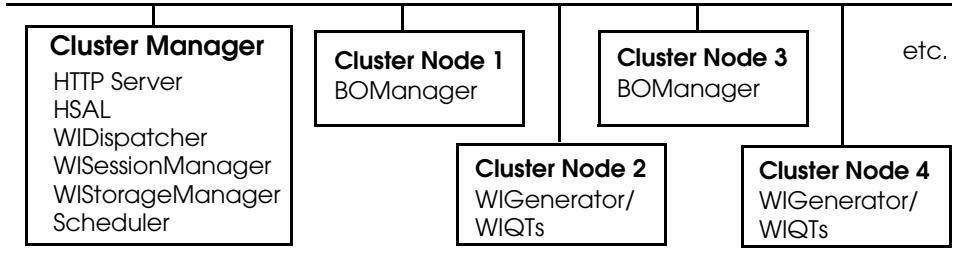

and a type A architecture for a UNIX deployment:

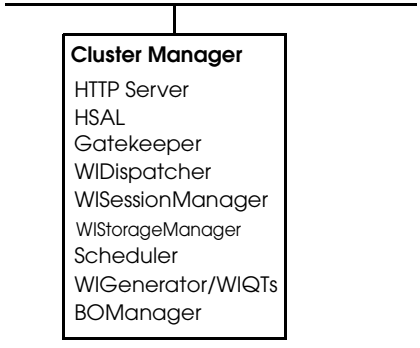

BusinessObjects processes

For Windows servers:

- As we suggest 4 active BUSINESSOBJECTS processes per CPU, the 44 full-client users will require 11 CPUs.
- If you have quadri-processor machines as your nodes, you will need at least 3 nodes to run the BUSINESSOBJECTS processes.
- The maximum number of active BUSINESSOBJECTS processes on each machine should be set to 15.

For UNIX servers:

- With 20-25 BUSINESSOBJECTS Light processes per CPU, the 44 active users will require 2 CPUs.

WIQT Processes

The WIGenerators and their WIQT processes can be distributed over several nodes to optimize performance. For UNIX you can have up to 16 CPUs on one machine.

For Windows servers:

- With 15 WIQTs per CPU, the 360 active users will require 24 CPUs.
- You will therefore need 6 nodes of the same configuration to run the WIQTs.

For UNIX servers:

- With 25-60 WIQTs per CPU, the 360 active users will require 8-14 CPUs.
- You can run the WIQTs on an 8-14 processor machine.

The deployment configuration

Thus the Windows cluster will contain a cluster manager plus:

- 6 QT nodes with a 4 * processor machine on each.
- 3 BO nodes with a 4 * processor machine on each.

The UNIX deployment will contain an 8-14 processor machine.

Required memory

NT

Using the formulas for RAM requirements for NT machines, you get:

(25 * 4 * 4 + 100) MB, or 500 MB of RAM for each BO node

(As there are 4 * 4 BUSINESSOBJECTS processes on each BO node, each BUSINESSOBJECTS process will need up to 25 MB of RAM.)

(15 * 15 * 4 + 100) MB, or 1.3 GB of RAM for each WIQT node

(As there are 15 * 4 WIQTs on each node, in general, each WIQT will need less than 15 MB of RAM.)

Using the formulas for RAM requirements for NT machines if the average document size is 10 MB, you get:

(3 * 10 * 4 * 4 + 100) MB, or 580 MB of RAM for each BO node

(As there are 4 * 4 BUSINESSOBJECTS processes on each BO node, each BUSINESSOBJECTS process will need up to 30 MB of RAM.)

(30 * 15 * 4 + 100) MB, or 2.5 GB of RAM for each WIQT node

(As there are 15 * 4 WIQTs on each node, in general, each WIQT will need less than 30 MB of RAM.)

UNIX

Using the formulas for RAM requirements for UNIX machine, you get:

(14 * 44 + 180) MB, or 796 MB of RAM for the BO processes

(As the maximum number of BUSINESSOBJECTS processes is 44, each BUSINESSOBJECTS process will need up to 14 MB of RAM.)

(8 * 360 + 180) MB, or 3 GB of RAM for the WIQT processes

(As the maximum number of logged on users is 360, in general, each WIQT will need less than 8 MB of RAM.)

Specific Deployment Guidelines

This section provides more guidelines, rules, and tips for deploying specific components or configurations in your server installation.

Deploying Broadcast Agent

This section addresses:

- The number of BROADCAST AGENTs and Schedulers you need in your deployment
- BROADCAST AGENT server sizing

How many Schedulers do you need?

How many BROADCAST AGENTs should you deploy, and how many Schedulers?

You can configure multiple BROADCAST AGENTs in your configuration, and multiple Schedulers within each BROADCAST AGENT. The advantage of having several BROADCAST AGENTs is that you can have one for each department of your company (Sales, Finance, etc.). It is then a simple matter to associate specific user groups with specific BROADCAST AGENTs.

The advantage of having two or more Schedulers for each BROADCAST AGENT is for failover purposes. Schedulers periodically query the repository to determine which documents are due for processing; when a scheduled task is due, the Scheduler instructs the BOManager to process it. Without a Scheduler, no jobs will be processed. Since an additional Scheduler does not use much in the way of resources, it is always a good idea to configure a second Scheduler.

A Scheduler does not tie up many system resources. You can therefore afford to put a BOManager on the same machine as the Scheduler.

When configuring BROADCAST AGENT, you must specify the connection between the Scheduler and the repository. Otherwise, when users send jobs to the repository for processing, nothing will happen.

Using the Business Objects Services Administrator, start the Scheduler, then enter an authorized name and password. For more information on the Administrator, see the *WebIntelligence System Administrator's Guide*.

Broadcast Agent server sizing

The main factors in BROADCAST AGENT sizing are the number and size of the documents that you want to process per hour.

The number of required BUSINESSOBJECTS processes is equal to the number of documents being processed simultaneously. You can estimate this number using the following formula:

(number of documents to be refreshed * average time it takes to refresh a document) / time available to run the schedule

The average time it takes to refresh a document depends on the CPU being used, and the number and size of the documents being processed. You can only determine this by experimenting with your system.

RAM required for each BusinessObjects process

The amount of RAM the system requires depends on the size of the documents being processed. Most documents use from 3 to 15 MB. For documents of 3 MB and larger, the amount of RAM used follows this equation:

RAM required for each BusinessObjects process = (3 * average size of a document) + 4 (MB)

Let's take a system regularly processing very large documents of 15 MB as an example: (3 * 15) + 4 = 49 MB. Rounding it up for safety, **we recommend that you allocate 50 MB of RAM for each BUSINESSOBJECTS process**, unless your documents' average size is greater than 15 MB.

We recommend that you have no more than 4 active BUSINESSOBJECTS processes per CPU. Therefore, if 4 * number of CPUs is less than the number of active BUSINESSOBJECTS processes, simply add more CPUs until 4 * the number of CPUs is greater than the number of BUSINESSOBJECTS processes.

The RAM required by a node is given (in MB) by:

100 + (50 * number of BusinessObjects processes on the node)

Deploying WebIntelligence

Here are a few WEBINTELLIGENCE deployment notes:

- The type of request executed by a user, the size of the document, and the amount of documents all influence the number of active users that the system will support, which in turn influence the number of servers needed.
- Viewing corporate documents and building ad-hoc queries require more resources than viewing personal documents.
- A BusObj.exe (Windows) or `bolightsvr` (UNIX) process will be started for every BUSINESSOBJECTS document that is viewed or refreshed from WEBINTELLIGENCE. Each process should be allocated between 15 and 30 MB of RAM.
- A login mechanism for interactive users improves performance in all deployments. This mechanism creates temporary .key and .lsi files so that when the user's process is reassigned to another user because of inactivity, and the user becomes active again, no login against the repository is done. Instead, the user logs in against the temporary .key and .lsi file.
- For information about supported browsers, and which Web Panel can run on each browser, see the ReadMe file for this release.

WIStorageManager deployment notes

The bulk of what is stored in the WEBINTELLIGENCE system Storage is the documents INFOVIEW users have saved as personal documents. No overall advice can be given about how much space to allot for this function, as the space required for personal documents will depend on the number of users in the deployment, how often they will save documents for their personal use, and the size of the documents they're saving.

You may find that users are not deleting old or outdated personal documents. You should encourage them to do so on a regular basis. If the space allotted for Storage becomes saturated, users will no longer be able to save new documents!

In cluster configurations, we suggest you place the WIStorageManager drive outside of the cluster configuration onto a fast file server drive, such as a RAID drive. This facilitates sharing in the case of cluster manager failure.

Providing cluster manager fault tolerance

You can provide fault tolerance for your cluster's cluster manager by deploying two of them, only one of which would be active at any given time. You can do this two ways:

- By deploying main and backup cluster managers.

  In this scenario, because the backup cluster manager uses a different OSAgent Port, it is up and running continuously.

- By deploying a redundant cluster manager

  In this scenario, because the backup cluster manager uses the same OSAgent Port as the regular cluster manager, it is started only when the regular cluster manager fails.

Deploying main and backup cluster managers

To deploy a main and backup cluster manager, WEBINTELLIGENCE must be installed on a specific OSAgent Port on one or several Windows servers (for example, the cluster manager is on Server1, with several cluster nodes).

You then install a backup cluster manager on another server, Server2, and several cluster nodes, **using a different OSAgent Port**. This cluster manager runs continuously, whether Server1 is functioning or not.

When Server1 goes down...

Users logged into WEBINTELLIGENCE on http://Server1/wi receive a Server Time Out error in their browsers. They can no longer access the login page to log back into WEBINTELLIGENCE.

If they need to use WEBINTELLIGENCE, however, they can immediately access the backup cluster manager using the URL http://Server2/wi or the second login button on the corporate intranet home page.

Tip: You can also write a Visual Basic or Java Script program for the initial web page that tests if Server1 is running and if not, it redirects the link to Server2.

During this time, the system administrator for Server1 stops the current WEBINTELLIGENCE system by:

1. Stopping the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on all cluster nodes.
2. Stopping the WEBINTELLIGENCE Manager and WEBINTELLIGENCEOrb Windows Services on Server1.

The administrator then gets the Server1 cluster manager and its nodes back up and running by:

1. Rebooting Server1. (In case of disk failure, the administrator must plug a copy of the hard drive into Server1).
2. Restarting the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows services on Server1.
3. Restarting the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows services on all the cluster nodes.

The main WEBINTELLIGENCE system is now ready to be used again on Server1.

The administrator should send an email to the end user community saying that they can now use WEBINTELLIGENCE again on http://Server1/wi.

Tip: To allow users to access personal documents from either cluster manager, you could use Microsoft Site Server software to synchronize the files between the cluster managers' separate WIStorageManagers. For information, see http://mspress.microsoft.com/prod/books/1335.htm.
You could also put the Storage on a RAID disk.

The following illustrates main and backup cluster manager implementation:

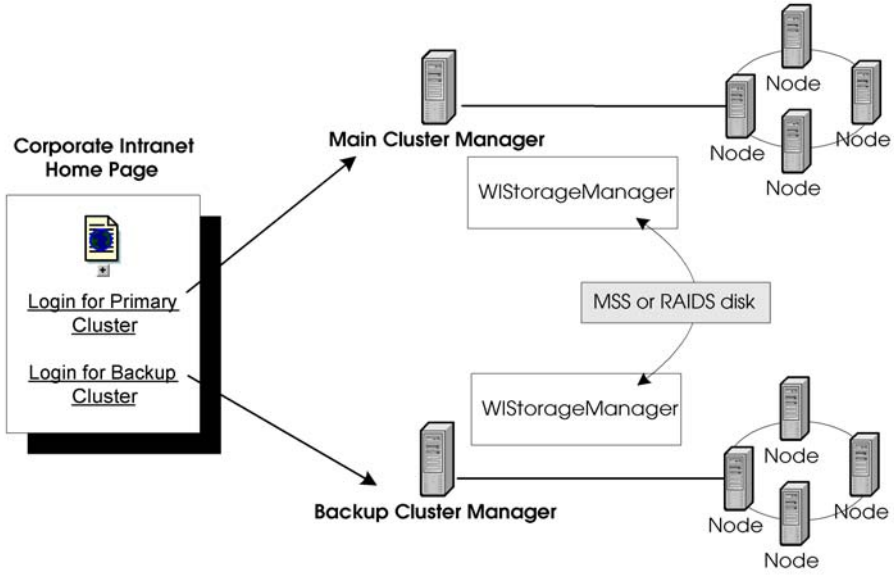

Deploying a redundant cluster manager

In this scenario, WEBINTELLIGENCE is installed on a specific OSAgent Port on one or several Windows servers (with Server1 as cluster manager, and several cluster nodes).

The redundant cluster manager is installed on Server2, **using the same OSAgent Port**.

When Server1 goes down...

End users logged into WEBINTELLIGENCE using the URL http://Server1/wi receive a Server Time Out Error in their browsers. They can no longer access the login page to log back into WEBINTELLIGENCE.

At this time, the system administrator shuts down the Server1 cluster by:

1. Stopping the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on all the cluster nodes.
2. Stopping the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on Server1.

The administrator then starts up Server2 as the new cluster manager for the same cluster nodes by:

1. Booting Server2, then starting the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on it.
2. Restarting the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows services on all the cluster nodes.

The main WEBINTELLIGENCE system is now ready to be used on Server2.

End users can then use WEBINTELLIGENCE either by using the URL http:// Server2/wi, or by clicking the second login button on the corporate intranet home page.

What happens to Server1 once Server2 is being used?

In this scenario, once Server1 is up and running again, it can then be used as a redundant cluster manager for Server2 if desired, although the corporate internal home page will have to be modified.

If you decide not to use Server1 not as the redundant server but as the main server, you will need to:

1. Stop the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on all cluster nodes.
2. Stop the WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on Server2.
3. Restart WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on Server1.
4. Restart WEBINTELLIGENCE Manager and WEBINTELLIGENCE Orb Windows Services on all the cluster nodes.

You should send all users an email informing them of the imminent shutdown of the Server2 system and the re-starting of the Server1 system.

The following is an illustration of the redundant cluster manager implementation:

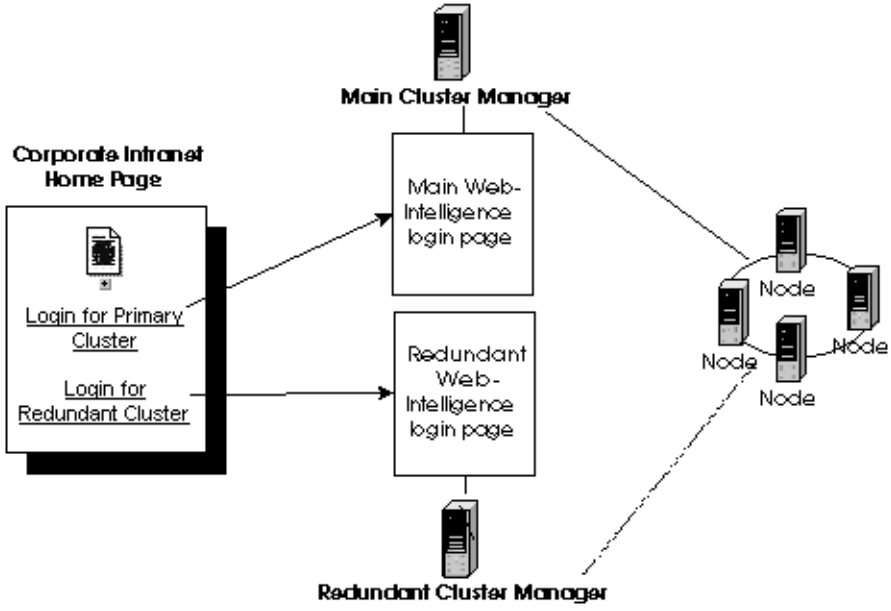

Cisco LocalDirector

The Cisco LocalDirector is a device that load balances TCP/IP traffic across multiple servers. It receives HTTP requests and redispatches them to different web servers. The following section describes how you can use the Cisco LocalDirector to provide load balancing and failover features for WEBINTELLIGENCE clusters.

Full information on the Cisco LocalDirector is available on the Cisco website: http://www.cisco.com

What Business Objects software/tasks can use the LocalDirector?

The following Business Objects products and tasks take advantage of the features of the Cisco LocalDirector:

- Zero Admin BUSINESSOBJECTS in 3-tier mode
- Navigating in INFOVIEW
- Viewing and refreshing full-client documents in INFOVIEW
- Viewing and refreshing thin-client documents
- Creating thin-client documents.

What Business Objects software/tasks can't use the LocalDirector?

The following Business Objects products and tasks do not take advantage of the features of the Cisco LocalDirector:

- Administration tool

  In order to administer several clusters, the Administrator Tool needs to connect directly to the web server.
- Full-client products (BUSINESSOBJECTS, DESIGNER, SUPERVISOR)

  These tools query the repository directly, without accessing the web server or the WEBINTELLIGENCE cluster.

- BROADCAST AGENT

  This tool does not go against the web server. However, it does have its own load balancing mechanism, and for failover, the same Scheduler can be started on each cluster.

- BROADCAST AGENT console

  The BROADCAST AGENT console queries the repository directly, without going against the web server. It can also connect directly to a single cluster for administration purposes.

- Audit facility

  The audit facility may not work properly as the session ID can be the same for two different clusters.

Configuring the Cisco LocalDirector

When configuring the Cisco LocalDirector, you have a number of options:

- One virtual web server and multiple real web servers

  Users always point to the same virtual web address and their requests are then sent to each of the real web servers.

- Multiple virtual web servers and one real web server

  Users specify different virtual web addresses on a single port and their requests are then sent to the same real web server on a different port.

- Multiple virtual web servers and multiple real web servers

  A number of virtual web servers access a number of real web servers.

When using the Cisco LocalDirector in a WEBINTELLIGENCE cluster, you should also consider configuring the following features. For details of how to do so, see your Cisco documentation:

- Client-assigned load balancing

  This allows specific clients to be load balanced to different real servers, according to their IP address. There is no failover feature in this case. If the web server or the cluster manager is down, the client will not be able to connect to any machine.

- Maximum connection and Weighted configuration

  This tells the Cisco LocalDirector how to load balance to the real servers by specifying the maximum number of connections to each server and/or the weight of each server. See the section "Calculating the maximum connections and weighted configuration" on page 192 for details of how to determine these amounts.

- The `sticky` command

  This LocalDirector command **must** be enabled for WEBINTELLIGENCE connections. It ensures that once a connection is opened to a real server, requests from this client will always go to the same real server until the timeout is reached, thus maintaining the statefulness of the client/server connection.
  There are two ways you can implement the `sticky` command, one based on the IP address of the client PC and the other using a cookie sent from the web server to the client PC. When using the Cisco LocalDirector in a WEBINTELLIGENCE cluster, you need to implement the `sticky` command based on the IP address of the client PC.
  For more information, see "Using a non-transparent proxy policy" on page 191.

Tip: We recommend making the timeout for the `sticky` command the same as or slightly less than the WEBINTELLIGENCE session timeout.

Using a non-transparent proxy policy

If your network is set up so that the connections go through a proxy server, the Cisco LocalDirector may see the proxy IP address and redirect all requests to the same cluster manager if the `sticky` command has been enabled. To prevent this, you need to use a non-transparent proxy policy.

If a proxy is "transparent", this means that it hides the IP address of the client PC. In this case, the requests are sent to the Cisco LocalDirector showing the proxy IP address. Although the LocalDirector can tell that it is talking to a proxy, it cannot differentiate between the calls and sends them all to the same web server and the same WEBINTELLIGENCE cluster.

For this reason, you must use a non-transparent proxy policy if you want to take full advantage of the load balancing features offered by the Cisco LocalDirector with BUSINESSOBJECTS. If a proxy is "non-transparent", this means that it performs IP translation so that calls to the Cisco LocalDirector are made as if they came from the client PC itself. The LocalDirector can then redirect each call to a different web server and the different WEBINTELLIGENCE clusters as appropriate.

Calculating the maximum connections and weighted configuration

You can tell the Cisco LocalDirector how to load balance to the servers by specifying the maximum number of connections to each server and/or the weight of each server.

To do this:

1. Determine the maximum number of concurrent active users for each of the clusters connected to the Cisco LocalDirector. The sum of these users will determine the maximum number of connections.
2. Next, determine the ratio between the clusters in order to calculate the weighted configuration of the Cisco LocalDirector.

Example **Maximum connections and weighted configuration**

Take for example a set-up with three clusters.

- Cluster1 (three machines) has 300 active users.
- Cluster2 (two machines) has 200 active users.
- Cluster3 (one machine) has 500 active users.

In this case, the maximum number of connections is 1000 (300 + 200 + 500).

We can then determine the ratio between the different clusters, as follows:

- Cluster1: Weight = 300/1000 = 30%
- Cluster2: Weight = 200/1000 = 20%
- Cluster3: Weight = 500/1000 = 50%.

This gives us the weighted configuration for each cluster.

Note: For further information on load balancing in your WEBINTELLIGENCE deployment, see the *UNIX Performance Optimization and Sizing Guide*.

Deployment schema

The following diagram shows an example of how you can incorporate the Cisco LocalDirector in the network topology.

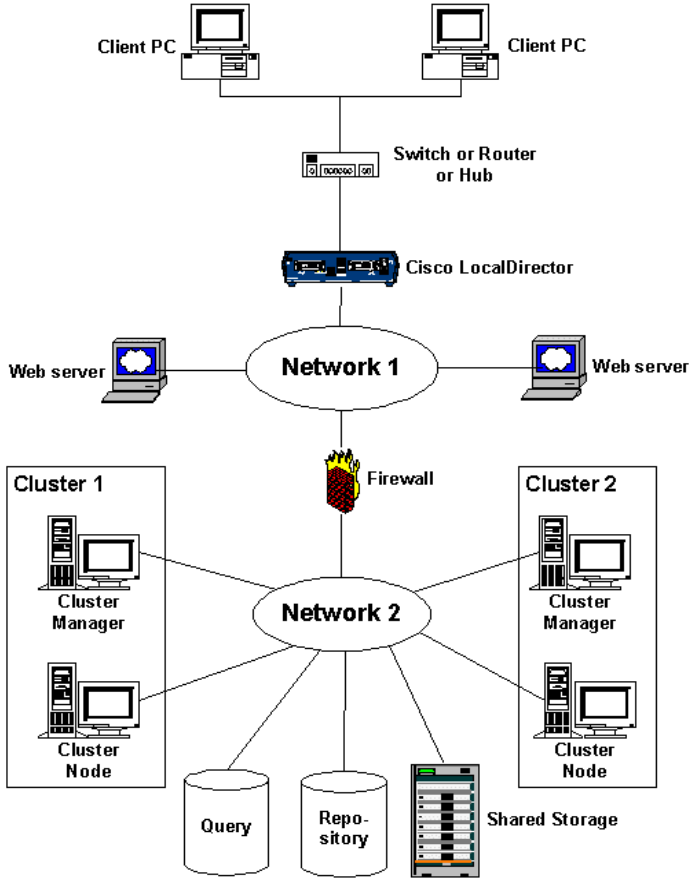

Note: Each of the servers in the cluster **must** be running exactly the same version of WEBINTELLIGENCE and BROADCAST AGENT. The same Business Objects modules must also be installed and enabled. If this is not the case, a user may need a specific module that is present on one server and not the others.

When the web server fails

What happens when the web server fails?

For new connections

If the Cisco LocalDirector can detect that no data is sent, it redirects the request to one of the other web servers.

For existing connections

If the Cisco LocalDirector can detect that no data is sent, it redirects the request to one of the other web servers. If you are working with WEBINTELLIGENCE, the session context is lost and all work that has not been saved will be lost. This is not the case if you are working with Zero Admin BUSINESSOBJECTS.

When the cluster manager fails

What happens when the cluster manager fails and the web server is still running?

If the request were to go to the same web server, the connection would fail because the cluster manager was down. In order to avoid this, the Cisco LocalDirector needs to know when the cluster manager has failed, even though the web server may still be running.

For new connections

If the Cisco LocalDirector can detect that no data is sent, it redirects the request to one of the other web servers.

For existing connections

If the Cisco LocalDirector can detect that no data is sent, it redirects the request to one of the other web servers. If you are working with WEBINTELLIGENCE, the session context is lost and all work that has not been saved will be lost. This is not the case if you are working with Zero Admin BUSINESSOBJECTS.

However, the web server or the Cisco LocalDirector may not know that the cluster manager has failed. In this case, the web server may need to be stopped to force failover to take place.

There are a number of ways you can do this. For example, you can:

- Create a script that checks if the cluster manager is running. On UNIX, you can see if the cluster manager is running by checking for the following compulsory processes on the cluster manager machine: WIClusterManager, oad, osagent, jre, Wiorb, WISitelog. For this, you can use the WIMonProc script, which you can find in the `Tools/bin` directory.
- Test the login page to make sure new users can log in. You can also run the login test from a JSP script that connects directly to the web server (without going through the LocalDirector).

If you discover that the cluster manager has failed, you have to:

- Stop the web server (to avoid new connections during the time needed to start the cluster manager)
- Reinitialize all the configuration files to ensure that you start from a clean environment
- Restart the cluster manager and the web server.

Deployment limitations

Running BUSINESSOBJECTS and BROADCAST AGENT, or BUSINESSOBJECTS and WEBINTELLIGENCE on the same machine is not recommended. When BROADCAST AGENT and/or WEBINTELLIGENCE are running, they automatically launch and stop BUSINESSOBJECTS instances. On the same machine, these "system" sessions can be mixed up with the BUSINESSOBJECTS user session.

You can, however, run BROADCAST AGENT and WEBINTELLIGENCE on the same machine. They both make calls to the BOManager component that manages a pool of BUSINESSOBJECTS instances. In fact, installing these two products on the same machine makes the system easier to administer.

Deploying multiple clusters

When should you deploy multiple clusters?

You should deploy multiple clusters if you have a large number of users that you want to break down into smaller groups, each with an individual cluster manager to improve performance. All users can still access the same repository, and can exchange documents in the usual way.

Multiple clusters are also useful if you need BROADCAST AGENT failure recovery. If you have BROADCAST AGENT Schedulers in one cluster, and Schedulers in another cluster, then if one cluster manager fails, the tasks will be taken over by one of the Schedulers in the second cluster.

Note: The cluster manager occupies a critical position in the system. If the cluster manager server should fail, or if one of its solution-wide servers (such as WISessionManager, or the Inprise VisiBroker OSAgent the WEBINTELLIGENCE system uses as its object request broker) should fail, the entire system needs to be stopped and restarted manually.

You can, however, provide fault tolerance on the cluster manager. See page 184.

Some rules about multiple clusters

- Each cluster must have a unique OSAgent Port value.
- Each cluster must have ONE cluster manager.
- Each cluster may have one or more cluster nodes on which you can run your chosen server components.

How many machines will you use to host your server components?

The more machines you can deploy, the more evenly the processing load will be distributed, and the better the overall performance. Installing additional processors in your existing machines will also improve performance.

Installed server components

Begin by installing all of the server components on each of your systems, and then activate or deactivate components, as required.

Broadcast Agent

A single BROADCAST AGENT comprises the following server components:

- One or more Schedulers
- Backbone services:
  - The CORBA backbone (all operating systems)
  - WIOrb
  - WIManager (Windows only)
- One or more BOManagers

WebIntelligence

WEBINTELLIGENCE comprises the following server components:

- An HTTP server
- Backbone services:
  - The CORBA backbone (all operating systems)
  - WIOrb
  - WIManager (Windows only)
- One or more BOManagers

Benchmarking

Basic methodology

Benchmarking tools work in the following way:

First users record scripts that they want to use during the system performance measurement. This is usually done via a recording tool that registers traffic through a given port (in this case, the HTTP port) while the client (web browser) is running the desired actions.

Once the traffic is registered, it is converted into a C-like script, and the benchmarking tool emulates the recorded actions of the client by running the registered script. This script can be run using many virtual users from many machines:

- The master computer is the machine that controls the machines from which the load is submitted. It also gathers data from them, performs analyses, and creates statistics.
- Agent computer(s) is (are) the machine(s) from which the load is submitted.

Here is a script you can use to test your system:

1. Open the WEBINTELLIGENCE Login page.
2. Click the Login button.
3. Enter the login and password, then click OK.
4. Click Corporate Documents in the navigation bar to the left.
5. Click Refresh in the toolbar.
6. Click on the document you want to open (open a different document each time you repeat this step, you will need to set two types of scripts, one dealing with BUSINESSOBJECTS documents, the other with WEBINTELLIGENCE documents).
7. When the document opens, click Refresh.
8. Repeat steps 4-7 five times.

9. Click Logout in the navigation bar.

10. Confirm the logout.

11. Set up three types of user groups:
    - Logged users
    - Logged active thin client users
    - Logged active full client users

12. Each of these user groups should contain as many users as the system is expected to support for that particular group.

13. Run this scenario 100 times.

Once the benchmark is done, the tool usually provides you with statistics files regarding the response times, success rates and the performance of the servers.

The two major criteria for judging the system are:

- The response times for the different script steps, which should be equal or less than what you specified when you designed the system.
- The success rate, which should be at least 94%.

Chapter 7 System Tuning and Troubleshooting

In this chapter

- **Overview 204**
- **Load Balancing 204**
  On the cluster manager 205
  On cluster nodes 205
  Thin-client vs. full-client load balancing 208
- **Improving Overall Document Computation Time 211**
  Improving full-client document processing (Windows) 211
- **Relieving Overloaded BOManager Machines 212**
- **Determining and Reducing HTML Overhead 212**
  Measuring HTML overhead for a BOManager server 212
  Determining other performance issues related to HTML overhead 214
  Optimizing HTML generation from BusinessObjects 215

- **Enhancing Performance from InfoView 216**
- **Enhancing Performance from Supervisor 216**
- **Enhancing Performance from Designer 217**
- **Enhancing Performance from BusinessObjects 218**
  Making full use of the presentation cache 218
  Logging in in offline mode 219
- **Windows NT Server Tuning Tips 220**
  Use an NTFS partition 220
  Use a disk defragmentation utility 220
  Tune foreground versus background application tasking 221
  Configure display properties 221
  Log out unused servers 221
  Disable unused Windows services 222
- **Tuning NT and Web Servers to Work Together Optimally 226**
  Optimizing the NT network subsystem 227
  Tuning the web server software engine to use all server resources 229
  Designing the disk subsystem layout for low latency 235
  More information 236
- **UNIX Server Tuning Tips 237**
  Setting the number of file descriptors allowed per process 237
  Maximum size of a core dump 238
  Swap file management 238
  Tuning web servers 239
  Optimizing the WebIntelligence Login 240
  Optimizing the Temp directory 240
  Optimizing the Storage directory 241
  Optimizing the Audit directory 241
  Optimizing the var/tmp directory 242
  Monitoring CPU use 243
  Where should you place database files? 243

- **Multiple OAD Instances Under UNIX 244**
  What Exactly Is the OAD? 244
  How Does Multiple OAD Enablement Work? 245
  Starting Multiple OADs 246
  When Should You Use this Feature? 246

- **Modifying CORBA Kernel Parameters 247**
  Modifying the timeout for object implementation 247
  Setting the OAD, WIDispatcher and WIAPI Broker server ports 249
  Modifying Gatekeeper Settings 251
  Modifying the Gatekeeper's Default Port 252
  Changing the Java VM Used by the Gatekeeper 256

- **Tracing System Activity for Tuning Purposes 259**
  The Audit kit 259

- **Monitoring WebIntelligence Log Files on UNIX 260**

- **Troubleshooting Your Deployment 262**
  Check your middleware 263
  Check your network configuration 263
  Tune your server configuration (Windows) 264
  Run the BOManager interactively (Windows) 269
  Trace system activity for debugging purposes 269
  Check your WebIntelligence server processes 272
  Check your WebIntelligence client software 273

Overview

You can obtain the best overall results from your Business Objects system through load-balanced configurations. This chapter describes how you can distribute your deployment over a series of servers, and demonstrates how you can progressively scale up your system without reducing performance. For example, some system modules are more processing-intensive than others. By enabling these modules on additional, dedicated servers, they can work more efficiently, with less risk of failure.

This chapter also gives you advice on how to improve overall system performance to speed up document processing times, and eliminate the system's web server as the cause of system bottlenecks. It ends with a troubleshooting section covering common problems and their solutions.

Load Balancing

A distributed deployment can scale to a greater number of users by automatically redirecting requests to the machines that are less "loaded" in the system. This capability is called **load balancing**. To scale up, simply add machines. Multiple machines allow overloaded components to redirect to less-used ones.

You use the same type of strategy with WEBINTELLIGENCE's system components, or modules. By having more than one of them activated in the cluster, they can share the transactions loads between them.

WEBINTELLIGENCE's system components are called modules. Each module provides a particular functionality. For a full description of each module, see the *WebIntelligence System Administrator's Guide*.

The load for various modules varies greatly. As the mechanisms for querying the repository and corporate databases, for example, the BOManagers and WIGenerators can carry heavy loads.

If particular modules are taking too long to respond to user requests, you can balance the load at strategic points by distributing high-traffic modules across several machines.

In cluster configurations, we suggest you put hardware in a place that can intelligently load balance WEBINTELLIGENCE requests across different clusters, and which has the ability to have WEBINTELLIGENCE servers automatically and transparently placed in or out of service, and which has a host standby failover mechanism. Some organizations use a product from Cisco systems called Cisco LocalDirector to provide this capability. For more information on how Cisco LocalDirector can fit in to your WEBINTELLIGENCE deployment, see page 189.

On the cluster manager

The cluster manager controls the entire server system. It handles all user requests and is the most critical node in the cluster.

Since the WIStorageManager stores users' personal documents, cached documents, and full-client presentations on the server which hosts it, it should be on the machine with the greatest available hard disk capacity.

We highly recommend that the WIStorageManager run on the cluster manager because it needs to be one of the first modules started. For more information, see "The WIStorageManager" on page 206.

On cluster nodes

Here are some tips for enabling modules on cluster nodes.

HTTP server, HSAL and other components

- As start pages, and document and universe access pages are generated by the WIGenerator and sent to the client through the HTTP interface, it is also advisable to keep a WIGenerator on the same machine as the HTTP server and HSAL.
- As a general rule, you should keep the HTTP server and HSAL on the same machine as the WIDispatcher. This will speed up the transmission of requests to the appropriate processes.

WIGenerator and BOManager

Because they are so processing-intensive, we recommend that you activate BOManager modules and WIGenerators on separate, dedicated cluster nodes. WIGenerator, however, should also be enabled on the cluster manager. These modules use much more RAM and CPU capacity than the other components, and can degrade the cluster manager's ability to handle and distribute user requests.

If only one module can be put on a separate cluster node machine, and the deployment processes both full-client and thin-client documents, it should be the BOManager module. A BUSINESSOBJECTS process is started for every full-client document that users want to view or refresh, and these processes take more resources than the WIQT processes used for thin-client documents. Full-client refreshes under Windows also use OLE automation, which is CPU-intensive.

Several BOManagers can exist in a single cluster. Free to function on machines hosting few other processes, these modules can work more efficiently, and with a reduced risk of failure. If one BOManager does fail, the system automatically and seamlessly redirects requests for the repository and corporate databases to another BOManager module on another machine within the same cluster.

For more information about WIGenerator and BOManager load balancing, see "Thin-client vs. full-client load balancing" below.

The WIStorageManager

When you install Business Objects server products, you can choose which server or servers on which you want to actually store the WEBINTELLIGENCE system's cached, temporary and personal documents and files. For performance reasons, it is often a good policy to locate application files, temporary files and personal documents on separate disks or machines. You can thus better control where the heavier transaction loads are likely to occur, and avoid bottlenecks before they happen.

In cluster configurations, the WIStorageManager drive should be placed outside the cluster configuration on a fast file server drive, such as a RAID drive. This facilitates sharing if the cluster manager fails.

With this release, the same Storage can be shared by multiple servers in the same cluster, or by multiple clusters, provided that they use the same repository.

Implementing shared storage

In order to enable shared Storage between WIStorageManagers in the same cluster, or in different clusters using the same repository, you must do the following on each of the servers hosting this module:

- Activate the Enable shared storage parameter for the WIStorageManager.
- Define Storage's address using the WISTORAGE environment variable, which points to the system's Storage area. This variable in each WIStorageManager must point to the same area. In Windows, you define system environment variables using Control Panel | System | Environment tab. In UNIX, you define them in the `WebiEnv.sh` file.

Ensuring access to shared Storage

To implement this sharing capability under Windows, however, the WEBINTELLIGENCE Orb and WEBINTELLIGENCE Manager services on every machine on which you install WEBINTELLIGENCE or BROADCAST AGENT using UNC paths or network disk mapping must run on a different account than the system account. This account must belong to the Administrator group and have *Act as part of the operating system* rights.

To configure these services for this:

1. Click Start | Programs | Administrative Tools (Common) | User Manager for Domains.
2. Choose User | Select Domain, then select a private domain.
3. Choose a user in the Administrator group, then choose Policies | User Rights.
4. Enable advanced User Rights and select *Act as part of the operating system*.

The account can now be used by the WEBINTELLIGENCE Orb and WEBINTELLIGENCE Manager services, which in turn will understand the UNC paths and the network disk mapping required to manage distributed Storage space.

To enable access to shared Storage under UNIX, you must ensure that administration account on each machine has sufficient privileges to read/write shared storage. You can do this by defining a group to which each account belongs, giving the group the required access rights.

Thin-client vs. full-client load balancing

Load balancing is generally performed in a simple "round robin" way, with the system automatically assigning a process to the next available node in the cluster. This version of the server products, however, has brought substantial improvements to how load balancing is handled.

Remember that there are two major processes that are affected when load balancing a system:

- WIQT processes managed by the WIGenerator, which:
  - Are started at login of each user instance
  - Generate the Corporate Documents and Inbox Documents lists
  - Are used to create and refresh thin-client documents
  - Generate LOVs (list of values) for full-client documents
  - Process scheduled thin-client documents
- BUSINESSOBJECTS processes managed by the BOManager, which:
  - Are used to refresh full-client documents
  - Process scheduled full-client documents

Thin-client process load balancing

The WEBINTELLIGENCE system includes an enhanced means of distributing the WIQT processes in a cluster.

The WIGenerator module has a parameter, called Node Load Factor, which allows you to "weight" the transaction load of the WIQT processes for a specific server dynamically. For example, if a particular cluster node machine is much more powerful than other nodes in the cluster, you can make sure it receives more WIQT processes by giving it a higher Node Load Factor than the other machines in the cluster. Likewise, if one machine is particularly restricted, by giving it a low Node Load Factor, you can make sure it does not penalize the entire system.

Now, when a new user session starts, instead of automatically sending the new transaction to the node running the fewest transactions, the WISessionManager sends it to the machine whose current session count/Node Load Factor ratio is the smallest. This is called the **load ratio**.

Example **Where does the WISessionManager send a new transaction?**

Take for example a small cluster with two machines:

- Machine 1 is currently handling 30 user sessions, and has a Node Load Factor of 1000. Its load ratio is therefore 30/1000 (or 3/100).
- Machine 2 is currently handling 40 user sessions, and has a Node Load Factor of 2000. Its load ratio is therefore 40/2000 (or 2/100).

Even though Machine 2 currently has a greater transaction load, its load ratio is smaller. The WISessionManager therefore sends the new session to it for processing.

Tip: When the Node Load Factor is set to 0, the WIGenerator will perform all other tasks except register WIQT processes on that particular machine. Instead, the WIGenerator registers WIQT processes on other cluster machines that have a Node Load Factor of greater than 0. Since the Node Load Factor is dynamic, this can be useful when shutting down cluster nodes to redirect new WIQT processes to another machines.

For information about setting a server's Node Load Factor, see the *WebIntelligence System Administrator's Guide*.

Full-client process load balancing

The BOManager modules on the clusters execute scheduled and existing full client documents as BUSINESSOBJECTS processes. Each BOManager may be assigned parameters of minimum and maximum number of processes. The OSAgent assigns the processes to the first available BOManager module.

Load balancing is managed in a "round robin" way, with the system automatically assigning a process to the next available node within the cluster.

Under Windows, full-client document refreshing uses OLE automation, which is CPU intensive. Each process should have at least 15 to 20 MB of RAM allocated to it.

For more information, see "Deploying Broadcast Agent" on page 181.

Improving Overall Document Computation Time

Document execution time can be significantly improved by arranging WHERE clause tables from largest to smallest (in terms of number of rows). For some documents, rearranging these tables can reduce execution time from hours to seconds.

Many other factors can also affect document computation speed:

- The number of rows in your queries
- The number of variables in your document (BUSINESSOBJECTS documents only)
- Multiple queries being "linked" into a single document (BUSINESSOBJECTS documents only)
- One query with multiple select statements being "joined" by BUSINESSOBJECTS (BUSINESSOBJECTS documents only)
- Your client configuration
- The size of the universe being queried

Here are some important recommendations:

- Restrict the amount of data to be transferred as much as possible by using conditions.
- Restrict the number of groups to which individual users are assigned.
- Restrict the number of universes in a single repository.

Improving full-client document processing (Windows)

You can considerably speed up the processing of BUSINESSOBJECTS documents by installing a printer driver on all Windows servers on which the BOManager module is activated.

Whenever the system needs to build the document's formatting for display, such as when a user refreshes the document, the driver gives the BusObj.exe process a jump start in formatting the document exactly as it was designed.

You can use any printer driver; no physical printer actually needs to be connected to the WEBINTELLIGENCE system.

Relieving Overloaded BOManager Machines

System performance can be diminished when the server machine(s) running BOManager is (are) overloaded, as when three to five users are refreshing a (different) document at the same time on a single mono-processor machine.

You can determine whether the server machine(s) running BOManager is (are) overloaded by using the Performance tab in the Windows Task Manager, or the Performance Monitor utility. If this is the case, you can improve performance by adding more machines running BOManager to your cluster and taking advantage of the system's load-balancing capability.

If performance does not improve significantly, some component in your configuration that is external to the Business Objects architecture may be overloaded. Overloading any of the following components, for example, will limit the performance of your configuration:

- The web server
- The database server(s) hosting the repository
- Network bandwidth

You should verify periodically that all these components perform reliably and are not adversely affecting the performance of your configuration.

Note: For suggestions for optimizing your web server, see page 226.

Determining and Reducing HTML Overhead

The HTML generated for typical BUSINESSOBJECTS documents is large, typically more than 500 KB. The resulting HTML overhead can result in a significant performance handicap.

Measuring HTML overhead for a BOManager server

A large proportion of BOManager performance problems arises from documents that result in large HTML presentations.

You can measure the HTML overhead for your particular configuration on the server running BOManager by doing the following:

1. On the server running BOManager, launch BUSINESSOBJECTS and open the document for which you want to measure the HTML overhead.
2. Measure the time it takes to open and display the document.
3. Save the document in HTML format in a temporary folder on the computer's local file system.
4. Measure the time it takes to generate the HTML document.
5. Measure the size of the generated HTML, making sure that you include subfolders.
6. If the web server is not running on the same computer, define a network share on the temporary folder.
7. On the computer running the web server, configure the web server to define an Internet alias to the folder created in step 3 that contains the saved HTML.
8. On the client machine, point the web browser to the Internet alias defined above, and measure the time it takes to open the HTML presentation in the browser.

In this instance, the HTML is already present on the hard disk. Therefore, performance has nothing to do with BUSINESSOBJECTS, BOManager, or WEBINTELLIGENCE. The overhead comes instead from:

- The web server reading the HTML from the hard disk
- Streaming the HTML through the network from the web server to the client computer
- Interpreting the HTML and displaying the page in the web browser on the client computer

Determining other performance issues related to HTML overhead

To gain an even better understanding of the scope of the performance issues that are related to HTML overhead, you may want to take additional measures.

1. On the machine running BOManager, with BUSINESSOBJECTS already running, measure the total time required to interactively complete the following operations without pausing (*tBusObj*):
   - Log in as the same user trying to view the BUSINESSOBJECTS document in the web browser.
   - Import the document for which you want to measure the HTML overhead from the repository.
   - Open the same document in BUSINESSOBJECTS.
2. On the client machine, measure the time it takes from clicking the name of the BUSINESSOBJECTS document in the corporate document list, until the HTML presentation is fully displayed in the browser (*tUserExp*).
   For this measure to be meaningful:
   - The BUSINESSOBJECTS document should not contain embedded HTML—i.e., the *Store HTML in document* option should not have been selected when the user sent the document to the repository.
   - The document's HTML presentation should not be in the presentation cache.

You can gain a good understanding of the nature of your performance issues by comparing three values, *tUserExp*, *tBusObj*, and *tPureHTML*, and by computing and comparing the percentages:

%BusObj = 100 * (tBusObj / tUserExp) and %PureHTML = 100 * (tPureHTML / tUserExp)

Optimizing HTML generation from BusinessObjects

If this performance is unacceptable compared to the equivalent operation in BUSINESSOBJECTS, we suggest that you redesign the document(s) to generate a smaller HTML presentation.

You can reduce the time it takes to open the BUSINESSOBJECTS document and generate the HTML by breaking down the document into smaller documents and/or by breaking the report into separate report tabs. We suggest that you:

- Create a document per section and/or per report tab.
- Reduce the number and/or the complexity of the formulas in your documents, where possible.

Note: With WEBINTELLIGENCE and BOManager, the HTML is generated for each report tab on an on-demand basis. Therefore, breaking up a large document into separate report tabs will improve performance by generating (and transmitting to the client browser) only the HTML that the users actually select, as opposed to generating the entire HTML for all the report tabs at once.

Optimizing HTML generation is only useful when the document will be refreshed from within WEBINTELLIGENCE, or it is sent to the repository without the Store HTML in Document option having been activated.

If the document was sent with that option activated, the document already contains the HTML as part of the .rep file. This can result in an extremely large .rep file, which means that importing and exporting the document can take a long time. If this is the case, you can improve the performance by making several smaller documents.

Enhancing Performance from InfoView

Ask your system's WEBINTELLIGENCE INFOVIEW users to select the *Enhanced Document Format* viewing option rather than HTML in the Personal Options page.

The enhanced presentation format is superior to HTML in most situations:

- It provides a much more precise rendering of the document, complete with headers and footers, precise alignment of blocks and graphics, and several navigation aids, such as page numbers and a hyper-linked navigation map.
- The presentation is generated and streamed to the client browser page by page, not per tab or section.
  For example, imagine a document with a tab containing several thousand rows. In HTML, the page for the whole tab, with the associated graphics, must be streamed across the network from the server machine to the client browser, which must then parse and lay out the incoming HTML before displaying the page to the user. This means several hundreds of kilobytes of data must be transferred across the network and parsed by the client browser before the user will see his page. On the other hand, using the enhanced document format, only the first page of the report tab will be streamed across the network and displayed. Additional pages will be generated and streamed only if and when the INFOVIEW user asks for them.
- The server compresses the enhanced presentation before it is streamed to the client browser, where it is decompressed. This conserves network bandwidth and further improves performance.

Enhancing Performance from Supervisor

Regularly delete old or outdated documents from the repository. You can do this:

- For corporate documents, one document at a time, using the Delete Document command in the Tools menu in SUPERVISOR.
- For documents that have been saved by or sent to users, using the Purge Inbox Documents command in SUPERVISOR. This clears out all "user" documents that have been in the repository for more than a specified number of days.

See the *Supervisor's Guide*.

Enhancing Performance from Designer

- Use Linked Universes.

  Linking two or more universes allows you to access one database per connection but may be deployed over different servers.

- Use External Strategies.

  A strategy is a script that you declare in an RDBMS folder that reads structural information from a specific database or flat file. External strategies limit the number of tables that are loaded into your repository's table browser.
  When you run a query, Business Objects products run the resulting SQL statement against the system catalog. If your database contains many tables, the catalog can be huge, and the SQL request will take longer to process.
  By using external strategy files, you can make your universes much more efficient. For example, for Oracle databases, you can use a strategy to retrieve the tables from a single account containing only a few number tables.
  Two external strategies are already supplied on the Business Objects CD.

- Optimize database performance by using aggregate awareness and shortcut joins. Do all aggregation on the server side instead of in the document (this can result in huge time savings).

See the *Designer's Guide* for more information.

Enhancing Performance from BusinessObjects

You can enhance your overall system's performance from BUSINESSOBJECTS by:

- Making full use of the distributed system's presentation cache
- Having all BUSINESSOBJECTS users log in in off-line mode.

Making full use of the presentation cache

One powerful way to enhance performance is to schedule corporate documents using the *Faster Document Viewing Over The Web* options in BUSINESSOBJECTS.

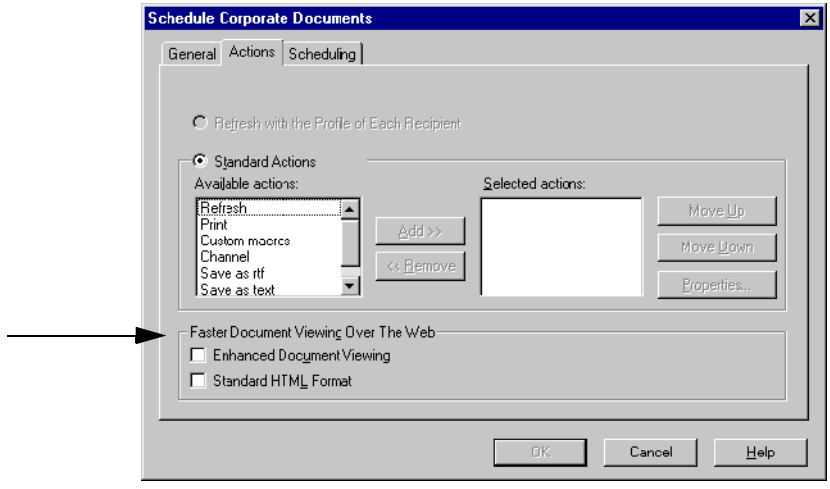

These options have been specifically designed to improve performance by coordinating the work of BROADCAST AGENT with the presentation cache maintained by BOManager.

To use this feature for maximum benefits:

- In BUSINESSOBJECTS, schedule the frequent refresh of the corporate documents that you expect will be accessed by a significant number of WEBINTELLIGENCE users (for example, every 10 minutes). This will ensure that these documents are always up-to-date and that users won't need to refresh the documents themselves.
- Check the F*aster Document Viewing Over The Web* option corresponding to the presentation formats you expect WEBINTELLIGENCE users will be requesting.

  This will ensure that the corresponding presentation formats are stored in the presentation cache maintained by BOManager after BROADCAST AGENT has processed the document. WEBINTELLIGENCE users will then be able to view these presentations straight from the cache, thereby eliminating the need to log into the BUSINESSOBJECTS server process, import the document, open it into the BUSINESSOBJECTS server process, and finally compute the required presentation.
- Disable full-client refresh rights for all users. Because the documents are refreshed so frequently, users are automatically guaranteed up-to-date documents anyway. WIQT processes are therefore used for user validation only, and the interaction between the server, the repository and the RDBMS is kept to a minimum.

For more information, see "The presentation cache" on page 114, as well as the *InfoView User's Guide*.

Logging in in offline mode

When starting any Business Objects applications, users can save time if they normally login in offline mode. Off-line mode establishes a direct connection with the database(s) containing report data, and bypasses the need for frequent transactions with the repository -- a frequent source of system bottlenecks.

For more information about off-line mode, see the *BusinessObjects User's Guide*.

Windows NT Server Tuning Tips

If you are using Windows NT machines as application servers, here are a few basic optimization tuning tips. You should spend some time evaluating each of these settings to ensure that your server's basic performance settings are properly configured.

Use an NTFS partition

We recommend installing WEBINTELLIGENCE on an NTFS (NT File System) partition. In addition to its security and reliability, NTFS was designed for speed. NTFS provides impressive disk I/O performance on large volumes such as those typically found in file servers or advanced workstations, and is therefore usually a good file system choice for these systems.

Use a disk defragmentation utility

One of the most important practices in maintaining a well-optimized disk subsystem under any operating system is minimizing the level of file fragmentation on the system's drives. You should therefore regularly run a disk defragmentation utility -- the more a server is used, the more often you should run the utility on it.

Tune foreground versus background application tasking

Another basic performance-related setting you should evaluate is the foreground background tasking option. The configuration is located in the Control Panel under System | Performance.

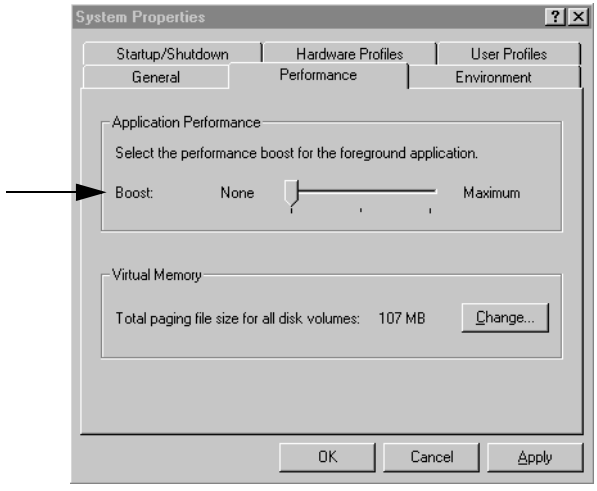

It controls the process and thread priorities Windows NT automatically assigns to foreground versus background applications.

In the Application Performance section, set the performance boost for foreground application to None.

Configure display properties

Set the Display Settings to 16 Colors and 800*600 pixels to improve the performance of the server.

Log out unused servers

Keep your servers logged out whenever possible. A locally logged-in user on a server wastes a significant amount of CPU time.

Disable unused Windows services

Windows comes with a number of services that may be running and consuming system resources even though you don't need them. Disable any services that are not indispensable to your deployment.

Here's a list of common Windows NT server services:

Service	Description
Alerter	Notifies administrators what alerts are occurring on the machine where alert is running. To receive these alerts, the Messenger service must be enabled.
ClipBook Server	Allows a user to share the contents of the ClipBook. **Not useful** for a server in most environments.
Computer Browser	Provides a list of servers and server resources to clients that request it. **Not useful** for a server in most environments.
DHCP Client	Establishes and connects to DHCP (Dynamic Host Configuration Protocol) servers that lease IP addresses. **Not useful** for a server in most environments if you do not use DHCP.
Directory Replicator	Allows you to transfer files and directories between defined Windows systems. Not enabled by default. **Not useful** for a server in most environments.
Event Log	Records all system, security and application events.
License Logging Service	Establishes the number of client and server licenses and the method of counting them. **Not useful** for a server in most environments.
Messenger	Provides pop-up messaging for Windows systems. **Not useful** for a server in most environments.

Service	Description
Net Logon	Provides account authentication, keeps the LSA/SAM up to date and synchronized with all other domain member servers.
Network DDE	Provides both a transport driver and security for Dynamic Data Exchange communications. Set to Manual by default.
Network DDE DSDM	Manages DDE communications.
Windows NT LM Security Support Provider	Provides security for Remote Procedure Call (RPC) services which use transports other than named pipes. Set to Manual by default.
Plug and Play	Allows such things as network adapters to be automatically configured at boot. **Not useful** for a server in most environments, but you will not be able to open Control Panel if you disable it.
Remote Procedure Call	Provides the endpoint mapper and miscellaneous RPC functions and allows distributed programs to use RPC name service. Set to Manual by default.
Schedule	Enables the capability to run AT and WinAt to schedule a particular action to take place at a particular time. Set to Manual by default.
Server	Provides RPC, file, print and named pipe sharing.
Spooler	Controls the print spooler. **Not useful** for a server in most environments, especially if there is no printer.
UPS	Uninterruptible power supply management utility. Disabled by default.
Workstation	Provides network connections and communications.
World Wide Web Publishing Service	Microsoft HTTP Server Service.
FTP Publishing Service	Microsoft FTP Server Service. **Not useful** for a server in most environments.

Service	Description
IIS Admin Service	Provides the administration tools for the Microsoft web server.
Microsoft SMTP Service	Microsoft SMTP Server Service. Set it to Manual.
Content Index	Microsoft Index Server Service. **Not useful** for a server in most environments.
WEBINTELLIGENCE Orb	WEBINTELLIGENCE CORBA backbone service.
WEBINTELLIGENCE Manager	WEBINTELLIGENCE service.

Detect and alleviate memory bottlenecks

Because memory and CPU performance are so critical to good application server performance, you need to know all you can about identifying bottlenecks in these subsystems.

To determine if a memory shortage exists, monitor the system's total virtual memory usage over time (including physical RAM and paging file usage), using the Performance Monitor. If the paging file is constantly being resized, there is insufficient physical memory.

Another useful Performance Monitor counter is the Memory object's Available Bytes. It tells you the total amount of free memory on the system. Make sure that this monitor always indicates a healthy buffer zone even during peak system usage.

Another method for examining the amount of available physical memory on the system is to use the Windows Task Manager and examine the Available figure in the Physical Memory (K) section of the Performance tab.

Detect and alleviate paging file bottlenecks

A paging file bottleneck occurs when the system's paging file is inadequately sized for the paging activity the system experiences. By default, the size of the system paging file is equal to the amount of physical memory present when Windows is installed, plus 11 MB.

You can monitor the paging file's size by using the Paging File object in the Performance Monitor tool.

If the paging file's usage is over 85, increase the file's size to accommodate the memory usage on the system. To do this:

1. In the Control Panel, click System, then click the Performance tab.
2. Click the Change button in the Virtual Memory section.

Detect and alleviate processor bottlenecks

The processor's performance and usage levels are critical to system performance. The easiest way of verifying these levels is to check the Windows Task Manager.

You can also use the Performance Monitor to monitor the key objects and counters concerning CPU usage on the system:

- The Processor object's % Processor Time
- The System object's Total Processor Time

Tuning NT and Web Servers to Work Together Optimally

You can dramatically improve the performance of your distributed system by optimizing NT's use of server resources (CPU, network, disk, memory) and tuning the web server software engine itself.

Your HTTP server functions as a part of the WEBINTELLIGENCE system by providing the main communication channel from users to WEBINTELLIGENCE. This includes sending requests from the user to WEBINTELLIGENCE, and sending back to the client interface generated HTML pages.

Because this solution constantly generates dynamic content, CPU, memory, network, and disk resources are generally in greatest demand.

Effectively tuning your HTTP server involves:

- Optimizing the NT network subsystem
- "Tuning NT and your web application to stay in RAM".
- "Tuning IIS memory usage".
- "Tuning your web application to use available CPU resources".
- "Grouping similar disk I/O activities and files together on fast disk subsystems".

Caution: Many of the tuning options outlined below require editing the NT registry. Before executing any of them, make sure you backup the registry and understand how to recover it!

Optimizing the NT network subsystem

Optimizing the NT subsystem for web-based workloads involves:

- Making sure the network path to your web server is not creating a bottleneck
- Tuning specific Registry values to optimize network response times
- Tuning NIC card usage to optimize web server response times

Eliminating network bottlenecks

Web server activity places a particularly heavy strain on the servers' network subsystem, NT's TCP stack, and Network Interface Card (NIC). Your first step in alleviating that strain is to ensure that the network path to your web server is not a bottleneck.

Under NT, the most common Performance Monitor counters to use to detect a network bottleneck are:

- Network Segment object: %Network Utilization
  This counter determines the total amount of bandwidth your web server is using. For an Ethernet environment, if %utilization is greater then 70%, your network connection is limiting your performance. If your server is connected to the network via a shared Ethernet connection (non-switched), consider the network connection saturated at 40-50% use.
- TCP object: Connections Established
  This counter measures the number of simultaneous TCP connections last detected in the Established or Closed-wait State.
  Over time, if this counter hits a level at which it stops dynamically changing, you should increase your bandwidth and investigate whether other NT resources are restricting your web server's performance. The web server solution will run only as fast as your slowest utilized resource.

You can also remove NT network bottlenecks by increasing the bandwidth of the technology you're using (10Baset, 100BaseT, Gigabit Ethernet) or by using multiple NICs (e.g. LAN connections) in a trunking fashion.

Trunking LAN connections is not provided natively by NT servers, so for this option, consider using a third-party package such as Adaptec's Duratec or NSI's Balance Suite.

Tuning TCP-related registry values to improve network response times

A badly-tuned TCP stack will prevent your NT web server from responding to as many connections as your resources will allow.

Transmission control blocks store data for each TCP connection. A control block is attached to the TCB hash table for each active connection. If there aren't enough control blocks available when a web request (http) arrives at your server via TCP/IP, there is added delay while it waits for additional control blocks to be created.

By increasing the TCB timewait table size, you can allow more web connections to be serviced more quickly. To adjust this value, add the following registry value:

Key path	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters
Key name	MaxFreeTcbs
Value	0xFA0

This example increases the TCB timewait table to 4,000 entries from the default of 2,000.

Now that you've reduced the time overhead incurred by TCP for the web server, you must adjust the corresponding hash table, where the TCBs are stored.

Do this by adding the following registry value:

Key path	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters
Key name	MaxHashTableSize
Value	0x400

This increases the TCB hash table size from 512 to 1,024, allowing more room for connection information. TCB information is stored in the non-paged memory pool. If the web server is experiencing memory bottlenecks and more memory cannot be allotted to the server, lower the above values. See "Tuning NT and your web application to stay in RAM" on page 230.

Optimizing NIC card usage

Another way you can you can improve web server response time is by optimizing the NIC and CPU relationship for web server workloads.

Each web request received over the network generates an “interrupt” to the processor requesting service. If the request doesn’t have a high enough interrupt level, the processor defers the request. This deferred interrupt request becomes a Deferred Procedure Call (DPC).

As the web server receives more and more requests, the number of interrupts and DPCs increases. When an interrupt is sent to a particular CPU then deferred to another CPU in the server, additional server overhead is incurred. This default NT behavior results in a costly performance handicap.

To stop this from happening, set the ProcessorAffin registry value to 0.

This forces the CPU handling the interrupt to handle associated DPCs as well. It also means that the network interface card or cards are not associated with a specific CPU. This improves the CPUs servicing of interrupts and DPCs generated by the network interface card(s).

Tuning the web server software engine to use all server resources

Every web server engine such as Apache, Microsoft Internet Information Server 4 (IIS), Netscape, etc. can be tuned to optimize performance. In this section, we discuss how to optimize the Microsoft IIS web server engine.

Note: You can use the same strategies for other web server engines. Talk to your vendor or check out their documentation to implement these tactics.

Tuning NT and your web application to stay in RAM

Networks introduce enough latency to slow down web requests to your users, so avoid adding any more where ever possible in your web server solution. Accessing memory is much faster then retrieving what your web server might need from the disk drive.

IIS is a multi-threaded application that runs as a single process instance, referred to as Inetinfo.exe under NT Server. Included in this process, IIS provides an object cache frequently used by Active Server Pages (ASP), Web Services (http), FTP and SMTP services. The object cache is part of IIS's **working set**, e.g. the area in physical RAM that Inetinfo.exe occupies.

The more you keep in Inetinfo's working set, and consequently as many frequently-used objects in RAM, the better performance will be. Unfortunately, other activities occurring on your web server take up memory space too, including:

- Every web connection provided by your web server
- NT's file system cache
- NT itself (approximately 24Mb)
- Perl scripts, database engines nd other applications running on web servers

Tuning NT's management strategies for web server workloads

To do this:

1. Tune NT's memory strategy to favor the application working sets by selecting Start | Program Files | Control Panel | Network | Services | Server, then clicking the *Maximize Throughput for Network Application* option.

2. Check how much memory is available, and block your web server from paging on a regular basis. To do this, monitor the following information using Performance Monitor:
   - Process Object
     Select the Inetinfo.exe process, and watch the working set counter. This shows you how much memory IIS is trying to use.
     If you see this lowering over time, you must retune your web server (see below) or add additional memory to insure that no parts of inteinfo.exe are ever paged out to disk. If this happens, web server performance will slow to a crawl.
     If you are running other applications on your web server, you will want to monitor their working sets too.
   - Memory Object: Available Bytes
     This displays how much memory is available. If you are using the NT memory tuning strategy above, this value should always be above 4Mb. If it is lower, and the Memory object's pages/sec counter begins to increase, your server is running low on memory, which hurts web server performance.

Tuning IIS memory usage

Now that you know how much memory is available on your web server, you can tune IIS to take advantage of that memory.

Note: After each tuning change, watch your memory usage closely. If there is still memory available, you have room to tune. But remember to take into consideration your estimated peak web workloads, as each web connection adds approximately 10Kb of memory usage to the Inetinfo.exe's working set.

You can set Inetinfo's object cache to be larger, and then tune its cache characteristics in an effort to optimize its cache usage for your specific environment.

The default value of the IIS Object Cache is 10% of memory, but you can set it to any value you want. With a larger cache, you create more room in cache to support other services too.

To adjust the cache size, add following the registry entry:

Key path	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters
Key name	MemoryCacheSize
Value	0x 1E84800

Note that this value was added as a REG_DWORD using hexadecimal notation.

To determine if the IIS cache is large enough to be effective, check the Performance Monitor counter called:

Internet Information Services Global: Cache Hits %

How effective the cache is depends on the server's workload. Typically if the cache hit rate percentage is less then 90%, you should increase the IIS Object Cache size, provided of course that there is sufficient available memory.

Any cache is more effective if the relevant data is already in the cache when the request is processed. Based on end user usage, you can tune the IIS Object Cache to keep related web page information in its cache for longer periods of time. To do this, edit the following registry entry:

Key path	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters
Key name	ObjectCacheTTL
Value	0x8CA0

This example keeps the IIS open descriptors active for at least 10 hours – an average person's work day. The default is 30 seconds.

You can monitor the effectiveness of this tactic by looking at the Performance Monitor counter:

Internet Information Services Global: Cache Flushes

If the rate of cache flushes is a high value and is associated with a high value of cache misses and page faults, the cache may be flushing too quickly. In this case, repeated requests for the same data do not result in a cache hit when they should.

If this occurs, increase the ObjectCacheTTL and MemoryCacheSize even more. Increasing the ObjectCacheTTL keeps data in the cache for a longer period of time, which increases the chances that the required data is found in the cache. The converse can also happen, in which case you should increase the rate in which the IIS Object Cache is flushed by lowering the ObjectCacheTTL value.

Tuning your web application to use available CPU resources

Tuning the IIS queue size

Under heavy web workloads, keeping IIS properly fed with data is important.

First, determine if you have additional CPU capacity available. You can do this by monitoring Performance Monitor's Processor object's Processor Utilization counter for each CPU in your web server.

Not all workloads are distributed well across all processors in an NT environment, so watch each processor's usage closely. If your processor usage looks relatively well distributed, and only 30% of processor resources are being used, consider increasing the workload that IIS will accept.

To increase the workload IIS will accept, add the following registry value:

Key path	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters
Key name	ListenBacklog
Value	0x1E

This increases the maximum active connections held in the IIS queue from the default of 15, to 30. Whenever this queue length limit is reached, the IIS will reject any new connections until the queue length shortens.

Note: Set this queue length lower if a system bottleneck begins to occur; you can subsequently increase this value if web requests are being rejected and you have the CPU resources to support a greater workload.

Setting the ListenBackLog value high when you do not have the CPU resources available to support the increased workload will cause longer waiting periods for your web server's end users. This is sometimes worse then no response at all.

Similarly, you can also increase the number IIS threads used to respond to web requests, by adding the following registry entry:

Key path	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\InetInfo\Parameters
Key name	MaxPoolThreads
Value	12

This registry entry specifies maximum network request threads per processor. The default is 10. Here we increased the value to 12.

Use extreme caution when setting these CPU-specific options: more is not necessarily better. You must always insure that you have sufficient processor resources, as it is possible to increase these values too high and actually lower your servers' overall performance for your end users. If your CPUs are becoming overworked by Inetinfo threads (you can monitor this using Performance Monitor's Processor object Processor Usage counter), lower the MaxPoolThreads value.

Designing the disk subsystem layout for low latency

To make web server performance as fast as possible, tune your web server engine to keep as much of the web-server provided content located in the web server cache and RAM as you can.

For some environments, such as intense file server and dynamic web application environments, this is not realistic. For these environments, optimize your RAM usage anyway, but also configure your disk subsystem for the lowest latency possible.

Grouping similar disk I/O activities and files together on fast disk subsystems

IIS relies on NT's file system cache to provide caching services when it needs information from the disk subsystem. The file system cache leverages a read-ahead algorithm when working with the disk subsystem.

Whenever possible, group web server files together on their own physical partition. Place these files on the fastest disk subsystem you have, such as a single Fast SCSI-3 Ultra-Fast Wide disk or a small RAID array.

By laying out your disk subsystem in this manner, you leverage NT's file system cache and reduce system latency. If this isn't possible, at least run a good disk defragmenter from time to time to increase the chances that your files are grouped together on the disk subsystem.

More information

A vast quantity of web sites can provide you with more generic NT and UNIX optimization information. As a starter, we can recommend the two following sites:

- http://www.microsoft.com/technet/winnt/default.asp
- http://www.tuningandsizingnt.com

UNIX Server Tuning Tips

You can dramatically improve the performance of your distributed system by optimizing the use of UNIX server resources (CPU, network, disk, memory) and tuning the web server itself.

You should regularly monitor virtual memory using the `vmstat` and `swap` commands.

Minimize the number of remote file systems to reduce transaction times.

Setting the number of file descriptors allowed per process

The default value for the maximum number of file descriptors allowed per process is 64 (this is considered a **soft limit**, or the default value for each user). You can, however, set a higher value, as long as it falls below the default hard limit of 1024. A **hard limit** is the maximum authorized number. Although with Solaris 2.6, the root user can surpass the hard limit or redefine it, we advise against it.

You can set this parameter by configuring either of the following:

- The user environment
- The kernel

Configuring the user environment

Add the following command into the WEBINTELLIGENCE administrator's user `.profile` file. This sets the soft limit for the maximum number of file descriptors to the specified <VALUE>:

- For sh: `ulimit -n <VALUE>`
- For csh: `limit descriptors <VALUE>`

Any user can set a soft limit to any value below the hard limit. Only a super-user can raise a hard limit.

To determine what the hard limits are, use the following command:

- For sh: `ulimit -H -a`
- For csh: `limit -h`

To determine what the soft limits are, use the following command:

- For sh: `ulimit -S -a`
- For csh: `limit`

Note: For more information, see `man` for the `ulimit` and `limit` commands.

Configuring the kernel

To increase the maximum number of file descriptors:

1. Define the hard limit (maximum authorized limit), then the soft limit (default value for each user) in the `/etc/system` file. You do this by adding the following lines:

```
set rlim_fd_max = <VALUE>
set rlim_fd_cur = <VALUE>
```

2. Reboot the server in order to take into account these settings.

For example:

```
set rlim_fd_max = 1024
```

Maximum size of a core dump

The maximum size of a core dump can be set using:

- For sh: `ulimit -c <VALUE>`
- For csh: `limit coredumpsize <VALUE>`

The `VALUE` must be specified in 512-byte blocks.

It may be helpful to prevent core dumps from happening to avoid disk saturation. This could be done by setting the previous value to `0`.

For more information, see `man` for the `ulimit` command.

Swap file management

The `swap` file size should be three times the size of RAM. To check how much swap space you have, type:

```
# swap -s
```

If the file size is too small, increase the file size, or create a second `swap` file.

For more information, see `man` for the `swap` and `mkfile` commands.

Tuning web servers

You can optimize web server performance by:

- Increasing the number of processes and transactions allowed for the server. This is governed by the following two parameters:
  - MaxProcs, which determines the number of processes allowed for the server.
  - RqThrottle, which determines the number of simultaneous transactions supported by the server.
- Regularly monitoring the size of the access log.

Netscape web servers

Increasing authorized processes and transactions

By default, the MaxProcs parameter is set to 1.

To increase Netscape server capacity, we recommend adding the following lines to the `netscape/suitespot/http-P24/config/magnus.conf` file:

```
MaxProcs <VALUE>
RqThrottle <VALUE>
```

For example:

```
MaxProcs 8
RqThrottle 4000
```

Monitoring the Netscape access log

You should check the size of the access log on a regular basis, by typing:

```
<NES install directory>/suitespot/<https server name>/logs/
access
```

You can disable this log using the web server's administrator.

IPlanet web servers

Increase the number of concurrent connections to IPlanet web server, by stopping it then increasing the value for the RqThrottle parameter in the `magnus.conf` file.

If you are using a multi-processor machine, you can also add the following line:

```
MaxProcs = <nb proc / 2>
```

For a more detailed explanation on tuning the Iplanet web server, see http://docs.iplanet.com/docs/manuals/enterprise.html

Optimizing the WebIntelligence Login

By following a few guidelines when creating groups or users in Supervisor, you can improve the response time of the WEBINTELLIGENCE login.

You should try to minimize the number of groups to which a user belongs. When setting up your repository hierarchy, don't mimic your organization's hierarchy, as it may result in the creation of levels and user groups which are functionally superfluous and degrade performance. Think carefully about the functions each user will realistically require, not the user's hierarchical status in the organization.

Unless necessary, deselect the Enable Real Time User Rights Update option on the Definition tab of the User Properties dialog box for the user.

On the Command restrictions dialog box, set as many commands as possible to Inherit status for the group or user.

Optimizing the Temp directory

The WEBINTELLIGENCE `Temp` directory (by default `<WebI Install Directory>/Temp`) is used to store information relative to each current session.

Mounting it on a separate file system

If you plan to have a high number of simultaneous sessions, we suggest you use a separate file system for this directory. **This disk should have an extremely rapid I/O rate, especially for write operations.**

You can do this any of the following ways:

The space you need to allocate to this directory depends on the repository you're using. To get an idea of how much space you need for this directory, log into the system as a typical user:

```
#cd WEBITEMP/sessions/
#du -kas Enxxxxxxxxx
```

The result is the size in kilo octets of that directory. Multiply this result by the number of registered users, and you have the space needed for the `Temp` directory.

- Mount this partition on a dedicated disk.
- Mount this partition on multiple disks. This can be done via the Volume Manager. That way disks work in parallel, which avoid bottleneck on the disks accesses.

- Mount the partition on the `tmp` file system. This way accesses occur in memory and not on physical disks, which is faster. But you will need a large amount of memory.

Mapping the system's temporary directory into memory

Mapping the system's directory into memory by setting up a temporary file system allows files to be placed in a memory cache. While this can enhance performance considerably, make sure you limit the total size of the temporary file system, or else the performance of other software may deteriorate.

Sun Solaris

To map the temporary directory into memory, we recommend that you use the `wupdate` script.

Use the `wupdate` script to map the `$WITEMPDIR` directory to the following file system mounted in memory:

```
/tmp/$WITEMPDIR
```

For more information about the `wupdate` script, see the *Installation and Update Guide for UNIX*.

Optimizing the Storage directory

The WEBINTELLIGENCE Storage directory contains user information and personal document storage. You set its location during installation (referred to as `<WEBI STORAGE DIRECTORY>`).

In large-scale deployments with large numbers of users or intensive personal document use, we suggest you use a separate file system.

Optimizing the Audit directory

By default, audit files are stored in the `<WebI Install Directory>/Server/System2.5/bin/` directory. You can, however, set the location of these files using the Business Objects Services Administrator (see the *WebIntelligence System Administrator's Guide*). We recommend you use a separate directory for these files, regularly monitoring the disk space they use whenever the Audit utility is enabled.

Optimizing the var/tmp directory

Whenever WEBINTELLIGENCE or INFOVIEW users process BUSINESSOBJECTS documents, temporary files are created in the `/var/tmp` directory, which is where UNIX systems creates temporary files by default. With more than 300 users carrying out identical, concurrent actions, high peaks of disk use can be attained rapidly.

An alternative is to create a link from `/var/tmp` to a high-capacity disk, or to set the TMPDIR environment variable before starting the WEBINTELLIGENCE system, with a directory mounted on a disk with lots of available space.

For example:

- For csh or tcsh:

  ```
  setenv TMPDIR /ma/grande/poubelle
  ```

- For sh:

  ```
  TMPDIR=/ma/grande/poubelle;export TMPDIR
  ```

Monitoring CPU use

The CPU system should not dominate CPU user. A good ratio is 80% CPU user for 20% CPU system.

Run Queue: the run queue or load average should not be more than 4 times the number of CPUs. This monitoring can be made using the `vmstat` command (with procs [r, b] and cpu [us, sy, id]). See `man` for `vmstat` for more information.

Where should you place database files?

To prevent any conflicts within the file system, every database file should be on different physical disks. If this not the case, try to organize the files according to how they are used. Take the following scenario:

Example **Example user scenario**

A large user population does the following:

1. They login.
2. They refresh the Corporate Document list.
3. They view a small document, then refresh it.
4. They log out of the system.

In this type of typical INFOVIEW activity, redo logs and rollback segments aren't used, as the database are not updated. Therefore, these files can be located with other database files.

Multiple OAD Instances Under UNIX

With the 5.1/2.6 release, WEBINTELLIGENCE system performance was improved with support for multiple WIGenerators on single cluster nodes. Multiple WIGenerator enablement serves two purposes:

- It can prevent a single WIGenerator from becoming a bottleneck when several WIQT processes have been launched concurrently.
- It prevents UNIX server mutex management problems that may occur on machines with more than 16 processors. By multiplying the number of instances of the WIGenerator, you provide more WIGenerators to share the transaction load and therefore minimize this server problem.

Nonetheless, the 5.1/2.6 release is still limited to 1024 concurrent active users (users actively using system resources to process documents by viewing, refreshing, sending and publishing them) per machine. Why? Because each time a user logs in and remains active, the system launches a wiqt process, which is registered with the machine's OAD. As the OAD can only handle a maximum of 1024 open files, it cannot process more than 1024 users (wiqts).

This type of problem prevents us from supporting an extranet deployment on a single machine.

With the 5.1/2.6 Service Pack 1 release, multiple instances of the OAD are supported on a single WEBINTELLIGENCE server, thereby removing this constraint.

What Exactly Is the OAD?

The ORB Activation Daemon, or OAD, is a CORBA component that runs on each server in a WEBINTELLIGENCE cluster, monitoring server components for failover purposes. It makes sure objects are activated whenever they are required.

At WEBINTELLIGENCE system start-up, a component called Wiorb starts the OSAgent on the cluster manager and the OADs on all the cluster's machines.

How Does Multiple OAD Enablement Work?

1. When the WEBINTELLIGENCE system is started, Wiorb starts the first OAD on the server machine, keeping the IOR (Interoperable Object Reference, which defines how to access a server object by means of its IP address, server port, etc.) in an environment variable called BO_OADIOR.
   When WEBINTELLIGENCE is stopped, or if the system administrator wants to disable a module, the administration framework uses this variable to unregister the module on the right OAD.
2. If other OAD instances have been defined for the server, Wiorb then starts the other OADs.
3. Wiorb then starts WIClusterManager (if the server is the cluster manager) or WIClusterNode (if the server is a cluster node).
4. WIClusterManager or WIClusterNode binds itself to the OAD identified by the BO_OADIOR variable and registers with it the main WEBINTELLIGENCE modules such as WIDispatcher, WIGenerator, WISessionManager, etc. Afterward, the other OADs are launched, so that when the WIGenerator actually starts, all of the OADs are available to it. Using a round robin algorithm, each WIGenerator that is launched binds itself to the next available instance of the OAD.

The OAD and WIGenerator

When a WIGenerator binds itself to an OAD, it stores the OAD's IOR in a global variable. Whenever a user clicks the WEBINTELLIGENCE login button, the WIGenerator registers the wiqt with the OAD identified by that variable. The IOR is then stored at the session level.

This way, even if the WIGenerator times out, the system is able to unregister the wiqt with the correct OAD. This means that the wiqts are bound to the OADs to which the WIGenerator itself is bound.

How many OADs should you enable?

Ideally, you should enable as many OADs as the number of WIGenerators enabled on the machine. That way, the frequent registering and unregistering of wiqt processes required by the WIGenerator will not tie up a single OAD.

Starting Multiple OADs

If you decide to enable multiple instances of the OAD on a single machine, you must define the same number of OAD instances on each machine in the cluster.

To start several instances of the OAD on a single machine:

1. Stop WEBINTELLIGENCE.
2. Edit the boconfig.cfg file, adding the following line to the [Business Objects\Shared\Administration Services\5.0\Properties] section:

   ```
   "NbOad"=string:"<NumberOfOads as int>"
   ```

   For example:

   ```
   "NbOad"=string:"2"
   ```

3. Make sure the -f option is included in the OAD command line, as in the following example:

   ```
   "OadCmdLine"=string:"/$WI/Server/ORB2.5/bin/oad -k -f -t 60"
   ```

4. Restart WEBINTELLIGENCE.
   The number of OADs specified for NumberOfOads in Step 2 are started.

When Should You Use this Feature?

The WebiServer.log file logs all the events that occur on a machine during WEBINTELLIGENCE execution. It is stored in the WEBINTELLIGENCE Log directory defined at installation.

If the message "oad: too many files opened" appears in this file, you should probably start more oads.

Note: This error message can also appear if the user starting WEBINTELLIGENCE has a hard limit (maximum authorized limit) on the number of file descriptors per process. For more information on setting the number of file descriptors allowed, see the *Deployment Guide*.

Modifying CORBA Kernel Parameters

Business Objects server products are based on three main CORBA components manufactured by Inprise:

- The OSAgent
- The OAD
- The Gatekeeper

They are known as the system's **kernel components**.

The Business Objects server backbone manages the startup of these three tools. With the 5.1/2.7 release, you can modify the command lines used by these three components at run-time:

- For NT servers, use the Registry Editor (regedit.exe).
- For UNIX servers, use the `boconfig.cfg` file, which provides Registry emulation.

Modifying these parameters can be useful in specific circumstances, such as those described below.

Modifying the timeout for object implementation

To be automatically activated, CORBA objects must be registered with the OAD. Registration is composed of:

- The object name
- The object interface
- The path for the program which provides the object instance

When a client component needs to establish a connection with a particular object, the OSAgent tries to locate the object directly. As it has no reference for the object, it then asks the OAD for it. Because the object is registered with the OAD, the OAD can locate the required process using specified program path.

After a specific duration, the object is supposed to be ready to respond to requests. If it isn't, the OAD kills the process it has just started and the client sees a "CORBA::No_Implement" exception.

To modify the timeout duration, change the following flag in the OAD command line:

Key path	• Under Windows: HKEY_LOCAL_MACHINE\SOFTWARE\Business Objects\Shared\Administration Services\5.0\Properties • Under UNIX: `\BusinessObjects\shared\Administration Services\5.0\Properties\`
Key name	OadCmdLine
Value	`-t <number of seconds>`

For example:

```
-t 90
```

The OAD default for Business Objects server products is 60 seconds.

When would you modify this parameter?

- When you are using a low-capacity server
  This will give the server more time to respond to the object request.
- When a specific server will be receiving requests for too much information at start-up time, such as a specific server running the WIQT.

Note: Setting the timeout value too much higher may generate other timing issues, or hide other issues related to product time tuning.

Setting the OAD, WIDispatcher and WIAPI Broker server ports

When the OAD, WIDispatcher and WIAPI Broker listen for requests from the system, they use a specific port. By default, CORBA dynamically sets these ports randomly at run-time, choosing them from among a set of free values.

Using specific flags, you can actually set the server port value used by these components. These flags are supported by all Business Objects server modules through the command line in the Registry.

When would you modify this parameter?

When you are deploying WEBINTELLIGENCE with two firewalls, such as in the following diagram:

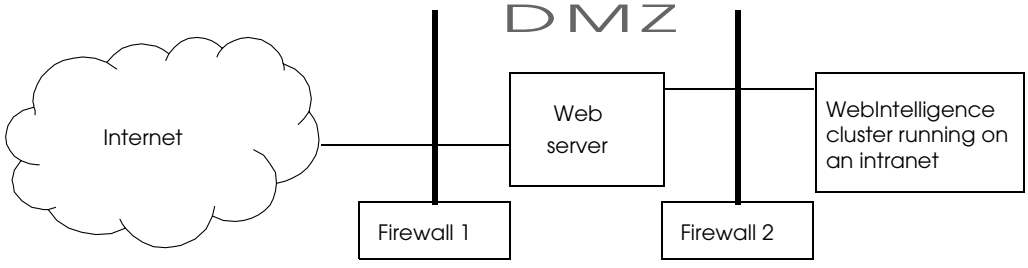

This double-firewall configuration creates what is known as a **DMZ** (**D**e**M**ilitarized **Z**one). The DMZ is a protected buffer zone between an organization's intranet and the external network or Internet. It is designed to keep outside users from accessing a server containing company data.

One of the key agents in WEBINTELLIGENCE communication is the HTTP Server Abstraction Layer, or **HSAL**. This is a CORBA-based component which acts as the interface between the HTTP server and WEBINTELLIGENCE. It receives user requests for WEBINTELLIGENCE services and converts them from the gateway protocol used by the gateway into a form WEBINTELLIGENCE understands.

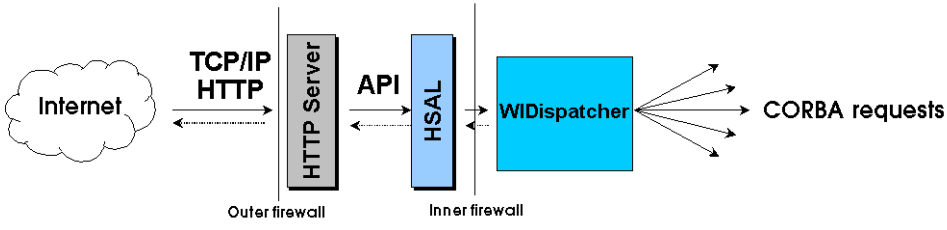

HSAL then passes the request through the firewall to the WIDispatcher, which decides which WEBINTELLIGENCE process the request should be sent to. The WIDispatcher converts the textual URLs into CORBA requests that can start the required WEBINTELLIGENCE process, then transmits the requests to that process via the OSAgent and the OAD.

The HSAL is always located on the web server. In a DMZ configuration, therefore, the HSAL is in the DMZ, separated from the WEBINTELLIGENCE cluster by the inner firewall.

In order for the HSAL to communicate with the WIDispatcher, the OSAgent and the OAD in the cluster, the ports used by those components must be open in the firewall. By default, however, these ports are not set, but dynamically allocated by the CORBA system at run-time.

For this type of configuration, therefore, you therefore need to set the server ports for these components by modifying specific CORBA kernel components. For instructions for deploying this configuration properly, see page 320.

Modifying Gatekeeper Settings

The Gatekeeper is a Java-based Internet Inter-ORB Protocol (IIOP) proxy that is an integral part of the WEBINTELLIGENCE system.

- Used with a firewall configuration, the Gatekeeper acts as an intermediary between the web server and the Internet to ensure security, administrative control, and caching services.
- Used with a system that spans more than one subnet, it receives requests from one subnet and passes them on to another.
- In a more basic WEBINTELLIGENCE system which uses neither firewalls nor multiple subnets, the Gatekeeper transmits requests between the Business Objects Services Administrator and the system modules it administrates.

The Business Objects server products setup includes a basic Gatekeeper command line which defines among other things the Java Virtual Machine (VM) the Gatekeeper uses, and the file containing various Gatekeeper properties.

When Must You Modify Gatekeeper Parameters?

You must modify Gatekeeper settings in the following cases:

- If you want to change the Gatekeeper's default exterior port (for example if another application in your CORBA system is already using the current port and that port cannot be changed). By default, the Gatekeeper's exterior port is always set to 15000.
- You want to change the Java VM (Virtual Machine) used by the Gatekeeper.

For instructions for both of these procedures, see below.

Note: We highly recommend that you *not* modify these settings unless absolutely required. In the event that you do make changes, make only those discussed here.

Modifying the Gatekeeper's Default Port

When you install any Business Objects server product, the Gatekeeper Configuration Manager, the Inprise graphical tool which allows you to set Gatekeeper properties, is also installed.

This tool generates a configuration file containing all the Gatekeeper settings. By default, this file is called gatekeeper.properties, and it is generated in the directory from which you started the Gatekeeper:

\<Business Objects Installation directory>\Server\ORB 2.5\bin

Example **Typical gatekeeper.properties file**

Here, the Gatekeeper's exterior port, for example, is set to the default value of 15000 (in bold).

```
#GateKeeper Properties
#Tue Jun 29 10:30:29 CEST 1999
disable_location_service=true
exterior_port=15000
disable_http_service=true
enable_callbacks=true
interior_port=15000
ior_file_name=C:\\\\gatekeeper.ior
exterior_address=199.10.10.2
log_level=normal
interior_proxy_address=
interior_address=199.2.2.4
```

Ensuring that the Gatekeeper Configuration Manager works

The Gatekeeper Configuration Manager (gkconfig) is a Java-based tool which therefore requires a Java virtual machine, or VM, to work (a VM is an interface between compiled Java binary code and the machine that actually performs the program's instructions). You must specify which Java VM it will use by adding a parameter to the gkconfig command line.

This involves first unsetting the machine's CLASSPATH variable if one is already set; when this variable is defined, the system expects it to provide the required paths. If the variable has not been defined, the tool launcher builds the CLASSPATH variable it needs using the installation path and tree.

Under Windows.

1. On Windows, open the system prompt by doing one of the following:
   - Click Start | Run.
   - Click Start| Programs | Command Prompt.
2. Unset any existing CLASSPATH variable by typing the following command:

   ```
   set CLASSPATH=
   ```

3. Change the current directory to <ORB 2.5\bin>. You can do this directly using the command

   ```
   cd /d %WIORB%
   ```

4. Type:

   ```
   gkconfig -VBJjavavm <the complete path to jre.exe>
   ```

   For example, if you have run a default installation, the command should be:

   ```
   gkconfig -VBJjavavm "C:\Program Files\Business          (...)
   Objects\Server\WebIntelligence 2.5\Jre\bin\jre.exe"
   ```

Under UNIX

1. Unset any existing CLASSPATH variable by typing the following command:

   ```
   unset CLASSPATH
   ```

2. Change the current directory to <ORB 2.5\bin>. You can do this directly using the command

   ```
   cd <WebIntelligence installation directory>/Server/ORB2.5/bin
   ```

3. Type:

```
gkconfig -VBJjavavm <the complete path to jre.exe>
```

For example, if you have run a default installation, the command should be:

```
gkconfig -VBJjavavm "<WebIntelligence installation     (...)
directory>/Server/JDK/bin/jre"
```

It may take half a minute, but the tool should start.

Tip: You can use the Gatekeeper Configuration Manager's Save option to write the gatekeeper.properties file wherever you want.

Modifying the Gatekeeper port number in its configuration file

1. Launch the GateKeeper Configuration Manager by following Steps 1 and 3 above and then typing:

```
gkconfig.exe [property_file]
```

For example:

```
gkconfig.exe gatekeeper.properties
```

2. Click the Exterior tab in the GateKeeper Configuration Manager.
3. Enter an unused, valid port number in the Exterior Port box.
4. Choose File | Save to save the changes.
5. Now add the correct argument in the Gatekeeper command line to take this file into account:

Key path	HKEY_LOCAL_MACHINE\SOFTWARE\Business Objects\Shared\Administration Services5.0\Properties
Key name	GateKeeperCmdLine
Value	`-props <complete path to Gatekeeper properties file>`

Modifying the Gatekeeper port number manually under UNIX

Under UNIX you can also modify the Gatekeeper's exterior port by editing the boconfig.cfg file located in the <WebIntelligence Installation Directory>/SetUp/ directory. You must add the following option to the GatekeeperCmdLine:

```
-VBJProp exterior_port=<port number>
```

For example, to set 15001 as the Gatekeeper port number (instead of the default 15000):

```
"GatekeeperCmdLine"=string:"/export/home/users/webi/webi2/
Server/ORB2.5/bin/gatekeeper -VBJjavavm /export/home/users/webi/
webi2/Server/JDK/bin/jre -VBJprop disable_http_service=true -
VBJprop ior_file_name=/export/home/users/webi/webi2/Server/
WebIntelligence2.5/http-server/classes/gatekeeper.ior -VBJprop
exterior_port=15001"
```

How can you check whether a port number is currently being used?

The procedure varies depending on the platform you're using.

Under Windows

You can use either the services file or the Visispy utility:

- The services file lists the known services running on the machine, together with the port number for each. You can view this file to check that the OSAgent Port value you choose does not conflict with those being used by other services.
  To do this, use a text editor such as Notepad to view the file:
  <Windows Installation Directory>\system32\drivers\etc\services
- The Visispy utility, visispy.exe, is automatically installed when you install Business Objects server products in the <WebIntelligence Installation Folder>\bin folder. It listens to network traffic on a given port and lists the sources of any traffic it detects.

  To run it, open an MS-DOS window and type:

  ```
  visispy -p <value> -t <time>
  ```

  where <value> is the port you are considering using for the OSAgent, and <time> specifies the number of milliseconds it should be monitored (30000, for 30 seconds, should be adequate).
  The utility waits for the specified time before returning to the prompt. If successful, it displays no text. If it lists one or more services on your local machine or on other machines in your subnet, these services are generating traffic on this port, so you must choose another port number for the OSAgent.

Under UNIX

Use the `netstat -a` command and look for the required port number:

```
netstat -a | grep <port number>
```

This command returns the state of the socket associated with the specified port number. If the socket is listening for incoming connections (and therefore in use), its state is "LISTEN".

For example, if port 15001 is in use:

```
netstat -a | grep 15001
```

the command returns something like

```
*.15001  ....                        LISTEN
```

Changing the Java VM Used by the Gatekeeper

The Gatekeeper is a Java application. Starting the Gatekeeper means starting a Virtual Machine (VM), which, again, is an interface between compiled Java binary code and the machine that actually performs the program's instructions.

For various reasons, system administrators may not want to use the Gatekeeper's default. They may want, for example, to standardize the Java level throughout their network.

To set the VM used by the Gatekeeper, enter the complete pathname and the desired VM in the VBjavavm parameter in the Gatekeeper command line:

Key path	• Under Windows: HKEY_LOCAL_MACHINE\SOFTWARE\Business Objects\Shared\Administration Services5.0\Properties • Under UNIX: `\BusinessObjects\shared\Administration Services\5.0\Properties\`
Key name	GatekeeperCmdLine
Value	`-VBjavavm [complete pathname]`

The VM must be at least JDK 1.1.6 compliant.

Parameter	Set in...	Comments
Limit Execution Time to (also, Limit Size of result set to)	DESIGNER	Sets a default value for all queries related to a specific universe.
Universe Properties - Controls	SUPERVISOR	Overrides DESIGNER universe settings.
Group Level: Limit Execution Time to (also, Limit Size of result set to)	SUPERVISOR	Overrides DESIGNER universe settings.
User Level: Limit Execution Time to (also, Limit Size of result set to)	SUPERVISOR	Overrides DESIGNER universe settings and the SUPERVISOR Group Settings
Maximum Fetch Time	WEBINTELLIGENCE ActiveX and Java Web Panels	Can be raised above the duration set in DESIGNER and SUPERVISOR
Maximum rows fetched	WEBINTELLIGENCE ActiveX and Java Web Panels	

Parameter	Set in...	Comments
Max. WIQT Active Time	WIGenerator module in the Business Objects Services Administrator	
Timeout for Interactive actions	BOManager module in the Business Objects Services Administrator	
Timeout for Batch Actions	BOManager module in the Business Objects Services Administrator	
Max Duration	Broadcast Agent Manager module in the Business Objects Services Administrator	

Tracing System Activity for Tuning Purposes

In order to effectively tune your system, you need detailed information about how the system is being used, and which modules on which machines are experiencing the heaviest transaction load.

The BusinessObjects Services Administrator provides an audit facility that lets you monitor the requests and processes passing through the WEBINTELLIGENCE system. These can come from WEBINTELLIGENCE, BROADCAST AGENT, Zero Admin BUSINESSOBJECTS, and WEBINTELLIGENCE SDK.

The Audit facility tracks critical information relating to user activity and places it in files or a relational database where it can be accessed with Business Objects products. This information includes the duration of a user's request as well as the time the system takes to respond. In this way, you can determine how many users were active at any given time, and thus check the number of concurrent users in the system (that is, users who are making the server work, as opposed to users who are merely logged in).

The Audit facility also traces vital system information such as when cluster nodes start and stop, as well as when module executables start and stop.

For more information, see the *WebIntelligence System Administrator's Guide*.

The Audit kit

To make it easier for you to sort and analyze the information you obtain, a User Activity Audit Kit is included in the Freeware directory of the 5.1/2.7 installation CD-rom. It contains:

- A universe that maps to the data in the database
- Three sample queries based on that universe
- Complete instructions for using the kit, in Adobe Acrobat PDF format

These queries should give you an idea of how to use the universe, and even serve as a launching point for building your own personalized queries in WEBINTELLIGENCE or BUSINESSOBJECTS to create detailed system activity charts, tailored to your particular activity analysis needs.

Monitoring WebIntelligence Log Files on UNIX

You should monitor the following log files (particularly the audit log, when Audit is "On") as, over time, the files can take up a great deal of disk space.

There are three types of log files:

- The **Audit log** files. These can be stored in a database of several files. You specify the files and their location by means of the Administrator tool. The following example is the one set by default, and is stored in `<Audit Directory>` :
  - `mm.log` -- contains all logs initially destined for the audit feature, if the bind to WISiteLog object fails. If there are no problems with the audit feature, the file remains empty.

Tip: Keep an eye on the mm.log file. If it is not empty, this is a sign of a CORBA communication problem.

  - `sm.log` -- contains audit system events, such as “WIClusterManger started” or “WIGenerator disabled.” Also created in NT installations.
  - `user.log` -- contains audit users events, such as “Logon” and “Refresh docs.”
  - `userdet.log` -- contains audit users events details that are associated to main details. For example, when a "Login" event is audited, one of the associated details is "Browser IP".
- The **WEBINTELLIGENCE Server** log files.These log files are stored in the *<Log Directory>* directory:
  - `internal.log` -- Contains internal trace information generated by modules when the parameter "Internal Trace" is set to “On”.

Note: Setting WebI Install Generator Internal trace to “On” causes logging of all WIQT's internal events. This can cause the `internal.log` file to quickly grow to several MB in size.

  - `WebiServer.log` -- logs all events that occur during WEBINTELLIGENCE execution (note that file size can get quite large)

- The **Visibroker** log files. These log files are stored in the `<WebI Install Directory>/Server/ORB2.5/adm/log/ directory:`
  - `oad.log` -- produced by `oad` when `oad` is started using the -v flag.
  - `osagent.log` -- produced by `osagent` when `osagent` is started using the `-v` flag.
  - `viserr.log` -- contains any output to the `cerr` that is produced by either the client or the server.
  - `visout.log` -- contains any output to the `cout` that is produced by either the client or the server.
  - `vislog.log` -- contains any output to the `clog` that is produced by either the client or the server.
  - `WebiServer.log --` logs all events that occur during WEBINTELLIGENCE execution (note that file size can get quite large).

Note: Setting WIGenerator Internal trace to "On" causes logging of all WIQT's internal events. This can cause the `internal.log` file to quickly grow to several MB in size. You can customize the wstart and wstop scripts to clean up, or back up the log and audit files each time WEBINTELLIGENCE is started or stopped. See "Starting WEBINTELLIGENCE at Boot and Stopping at Shutdown" in the UNIX *Installation and Update Guide*.

Troubleshooting Your Deployment

If you experience problems with your deployment, try to isolate the problem by running through the following checklist and procedures.

For detailed troubleshooting information concerning:

- BOManager and WEBINTELLIGENCE, see the *WebIntelligence System Administrator's Guide*.
- BROADCAST AGENT, see the *Broadcast Agent Administrator's Guide*.

For detailed error message descriptions and resolutions, see the *Error Message Guide*.

We recommend you check out the Business Objects ongoing documentation service on the World Wide Web. From here, you can obtain the latest in troubleshooting information, as well as updates, tips, and samples. To get there:

- From desktop products, select the More Tips and Samples command in the Help menu.
- From an Internet browser, go to http://www.businessobjects.com/services/infocenter.
- From WEBINTELLIGENCE INFOVIEW, click the Tips and Samples link in the On-Line Help navigation bar.

We also recommend checking the Online Customer Support web site. To access this site, you must be a registered user of Worldwide Customer Support. For information about registering with Worldwide Customer Support, go to http://www.techsupport.businessobjects.com/Register/Register.asp.

Once you're registered:

1. In a web browser, go to www.techsupport.businessobjects.com
2. Click the Enter button.
3. Enter the login name and password provided by Worldwide Customer Support.
4. Search through the Knowledge Base for the troubleshooting information you need.

Check your middleware

In most deployment situations, problems are often associated with middleware configurations. Check that the middleware is set up and working correctly. Also, make sure you're using the correct middleware version. For more information, see "The Repository" on page 275. Check the repository and corporate database connections.

Check your network configuration

There are a number of network issues that you should check:

- The OSAgent Port value must be the same on each of the machines within the same CORBA network cluster. If you're using two clusters, you should also be using two distinct OSAgent Port values for the machines of each cluster. You can check these values using the Inprise Smart Finder.

  To do this under Windows:

  - Open a DOS command prompt window.
  - Change the directory to <Installation Directory>\Server\ORB 2.5\bin.
  - Type `osfind`, then press Enter.

  To do this under UNIX:

  - Set the WEBINTELLIGENCE environment variables:
    If you are using `sh` or `ksh`, type: `./<WebI Install Directory>/SetUp/WebiEnv.sh`
    If you are using csh or tcsh, key a space then type: `/<WebI Install Directory>/SetUp/WebiEnv.csh`
    Then type `osfind`.
- Each of your server machines must have a fixed TCP/IP address.
- If you're using ODBC, then your Data Sources must be declared as System DSNs, not as User DSNs.

Tune your server configuration (Windows)

Servconf is a Windows utility you need to run on each computer hosting BOManager in order to work around some limitations in the way Windows handles some shared memory allocations. By default, it is located in the <Installation Directory>\Business Objects\Server\BusinessObjects Manager 5.0 folder, but you can set another location during installation.

This tool lets you view and modify two types of settings, both of which affect the context in which Windows runs the BusObj processes under the control of BOManager:

- System heap size settings
- The identity for the BusObj process automation server

Defining the appropriate settings is essential to ensure the correct behavior and performance of your BOManager system, especially regarding the number of BusObj processes that can be launched simultaneously.

Recommended settings

In the System Heap Size group:

- Either keep your system's current value for the Interactive setting or reset this value to the Windows system's default. For example, the Windows NT default is 3072.
- In the System Heap Size (KB) group, set the Non-Interactive setting value to 512 (try 1024 if you encounter problems later).

In the Identity For BusObj.exe Automation Server group:

- Select the This User option.
- Type the name of a valid user account (using the "domain_name\user_name" format, without the quotes but with the backslash).
  This account must be in the group 'Administrators' and have the special rights *Log on as a batch job* and *Act as part of the operating system*.

The following sections describe in a little more detail what these settings are and how they can affect the behavior and performance of your system.

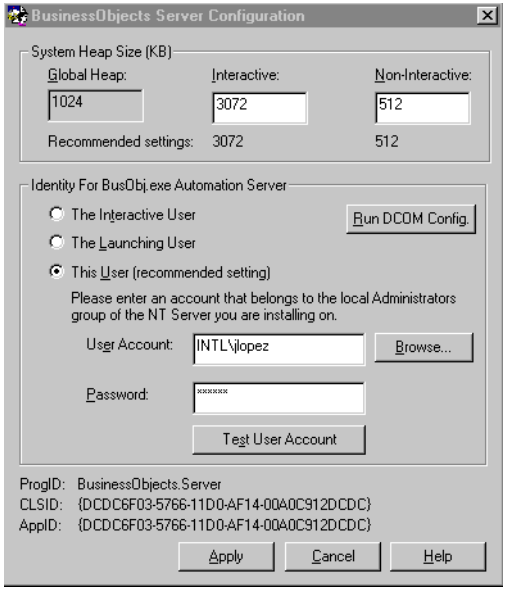

System heap size

These values control how Windows allocates system memory to the various server systems which have been configured.

Each system can be in one of two states:

- Interactive: The system is associated with an interactive, (logged-on) user.
- Non-interactive: Any number of non-interactive systems can exist to run service processes and COM server processes.

At any given time, several BusObj processes can be running simultaneously from one or more BOManagers. If for any reason there is insufficient system memory available for a given BOManager, then processes will fail in various and unpredictable ways.

Entering a higher heap size for a given BOManager enables more processes to run concurrently, but will limit the number of separate BOManagers that can be created simultaneously. Conversely, lowering the heap size limits the number of processes that can run concurrently in a given BOManager, but allows the creation of a larger number of separate BOManagers.

The values displayed are extracted from the Registry. Note that all values are in kilobytes (KB):

Setting	Description
Global Heap	Displays the amount of system memory that Windows reserves globally for all BOManagers. This value is shown here only for reference purposes, and cannot be modified.
Interactive	Defines the amount of system memory that Windows reserves for interactive BOManagers. The default value is 3072 KB and should not be modified.
Non-Interactive	Defines the amount of system memory that Windows reserves for each non-interactive BOManager. If no value is specified, non-interactive BOManagers use the same setting as interactive BOManagers. We recommend that you set this value to 512 KB. In a standard Windows NT installation, for example, no Non-Interactive value is specified. However, if you install the Windows NT Option Pack, or a specific application such as the Internet Information Server 4.0, this value is automatically set to 512 KB, which is the setting currently recommended by Microsoft

Identity for BusObj.exe automation server

These options control howBusObj processes are run when launched from the BOManager. They also determine which Windows machine will start the BusObj process.

Option	Description
The Interactive User	The BusObj process executes in the security context of the end user currently logged on (on the current server machine). This option is not recommended for the following reasons: • A user must always be logged on in order for the BusObj process to execute. • The BusObj process may have different rights depending on who is logged on, which may cause unexpected failures.
The Launching User	Selecting this option causes all BusObj processes to execute from the same BOManager. This option can adversely impact the number of BusObj processes that can run concurrently. The maximum number of processes depends on the *non-interactive* heap setting, and is around 10-12 processes for a non-interactive value of 512.
This User	This is the option that potentially allows the largest number of BusObj processes to run concurrently. The maximum number of processes depends on the *non-interactive* heap setting, and on the total number of BOManagers deployed. This is the recommended option. You should specify a user account using the "domain_name\user_name" format (without the quotes, but with the backslash). To ensure correct behavior, that account must belong to the local Administrators group on the current machine.

Run the BOManager interactively (Windows)

When you run BOManagers in a Windows server environment, the processes are running in background. This makes it difficult to identify any problems which may occur, since you won't see the error messages that are generated.

If you suspect that you have a problem with the BOManager on one of your servers, you should:

1. Run the Servconf utility (by default located in the <Installation Directory>\Business Objects\Server\BusinessObjects Manager 5.0 folder) to change the server identity to *Launching User* (see previous section).
2. Stop the BOManager service.
3. Run the following command on the same server:

   BOMgr.exe -visible

   This opens a DOS window. You should then restart the WEBINTELLIGENCE services. This effectively starts the BUSINESSOBJECTS full client interactively, and allows you to view any error messages which are flagged.
4. Run the Servconf utility again, and change the server identity back to *This User.*

Trace system activity for debugging purposes

Beyond the Audit facility, several methods of tracing sytem activity are available to you to analyse that's going on in the system, and where problems are occurring.

For complete information on all these methods, see the *WebIntelligence System Administrator's Guide*.

Internal tracing

Internal tracing records all the requests passing through a module, including generated SQL and detailed connectivity information.

Except for the WIStorageManager, all system modules contain a parameter called Internal Trace. When this is set to On, it activates the internal trace for the module by recording its requests.

This information is for debugging purposes only:

- Under Windows, it's stored in the log file called internal.log in <Installation Directory>\Server\System 2.5\bin.
- Under UNIX, it's stored in the log file called `internal.log` in the directory you specified as the WEBINTELLIGENCE Log directory at installation.

This setting is not persistent, which means it is not remembered if you stop and start the module on the server.

Note: Internal tracing activity (especially that of the WIGenerator) adds a great deal of overhead to the server processes and can cause significant performance degradation. You shouldn't try to trace system activity on a production system unless there is urgent need or unless instructed to do so by technical support staff.

Module debug logs

Module logs trace the activity of a particular module on a specific server machine. As opposed to WEBINTELLIGENCE system Audit files, these traces are most useful in explaining abnormal module behavior, and are therefore used primarily for debugging purposes.

You can also generate debugging logs for any of the following WEBINTELLIGENCE modules on a specific server machine:

- WIQT
- WISessionManager
- WIDispatcher
- WIStorageManager
- WIGenerator
- WIAdminBOTools
- WISiteLog
- WIClusterManager
- WIClusterNode
- Wiorb
- BOManager

You can also generate a debugging log for the system's Schedulers.

UNIX-specific log files

UNIX log files trace all events occurring during Setup or WEBINTELLIGENCE execution. Stored in the WEBINTELLIGENCE Log directory defined at installation, they include:

- `WebiSetup.log`
  Logs all events that occur during Setup execution.
- `WebiServer.log`
  Logs all events that occur during WEBINTELLIGENCE execution (note that file size can get quite large). For troubleshooting purposes, you can also add oad and osagent trace details to this file.

VisiBroker log files

Visibroker is the manuacturer of the CORBA components included in the WEBINTELLIGENCE system. VisiBroker log files record all ORB-related activity and errors. By default, they are stored in the following directories:

- Under Windows, in <Installation directory>\Server\ORB 2.5\Adm\log
- Under UNIX, in the the `<WebI Install Directory>/Server/ORB2.5/adm/log` directory.

These files include:

- oad.log
  Produced by the OAD when it is started using the -v flag.
- osagent.log
  Produced by the OSAgent when it is started using the -v flag.
- viserr.log
  Contains any output to the `cerr` produced by either the client or the server.
- visout.log
  Contains any output to the `cout` produced by either the client or the server.
- vislog.log
  Contains any output to the `clog` produced by either the client or the server.

Check your WebIntelligence server processes

Server problems may be due to processes that are not being started correctly. For complete information, see the Troubleshooting chapter in the *WebIntelligence System Administrator's Guide*.

Check your WebIntelligence client software

In order to log into WEBINTELLIGENCE, client browsers need to be cookie-enabled.

If users are using WEBINTELLIGENCE REPORTER, their browsers must also be both Java and JavaScript enabled. If they aren't, some parts of WEBINTELLIGENCE may not work. If there are any problems of this nature, run the Browser Configuration Check by clicking its link in the Login page.

Note: If you also want to use the ActiveX Reader to view BUSINESSOBJECTS documents, you must install a supported browser. For complete information, see the ReadMe file for this release.

If you're using a proxy server, check that you are allowed to download signed ActiveX controls.

Chapter 8 The Repository

In this chapter

- **What is the Repository? 276**
- **Repository Domains 277**
  Security domain selection at login 278
- **Which Components Access the Repository? 280**
- **The Repository and Your Data Warehouse 281**
- **Choosing a Repository Database 284**
  Row-level locking 284
  BLOB support 285
- **New 5.x Repository Features 287**
  Security domain selection 287
  Importing/exporting user and group lists 290
- **Working With or Without a Repository 291**
- **Repository Questions and Answers 292**
  How has the repository structure changed? 292
  How is the repository updated to release 5.x? 292
  Are there any compatibility issues I should know about? 293
  How can I reduce the size of the document domain? 293
  How much space does the repository require? 293

What is the Repository?

One of the most important tasks in any Business Objects deployment is the creation of your repository using SUPERVISOR.

The repository is a database that is stored in a relational database management system. The repository is used by different Business Objects applications to:

- Secure access to your data warehouse
- Provide a deployment infrastructure for the applications

Your deployment relies on the power of the repository to support the advanced features of the Business Objects solution.

For complete and up-to-date information concerning the database systems you can use to create your repository, see this release's readme file.

Note: It is important to understand that the repository is a separate entity from the database that your users will access to build documents and perform analysis.

Repository Domains

The repository consists of several domains. At a minimum, each deployment includes one security domain, one universe domain, and one document domain.

Each domain serves a specific purpose:

- The security domain is the core of the repository and is used to store all information about users, the groups they belong to, the applications and features they can use, the universes they have access to, and the documents they have shared.

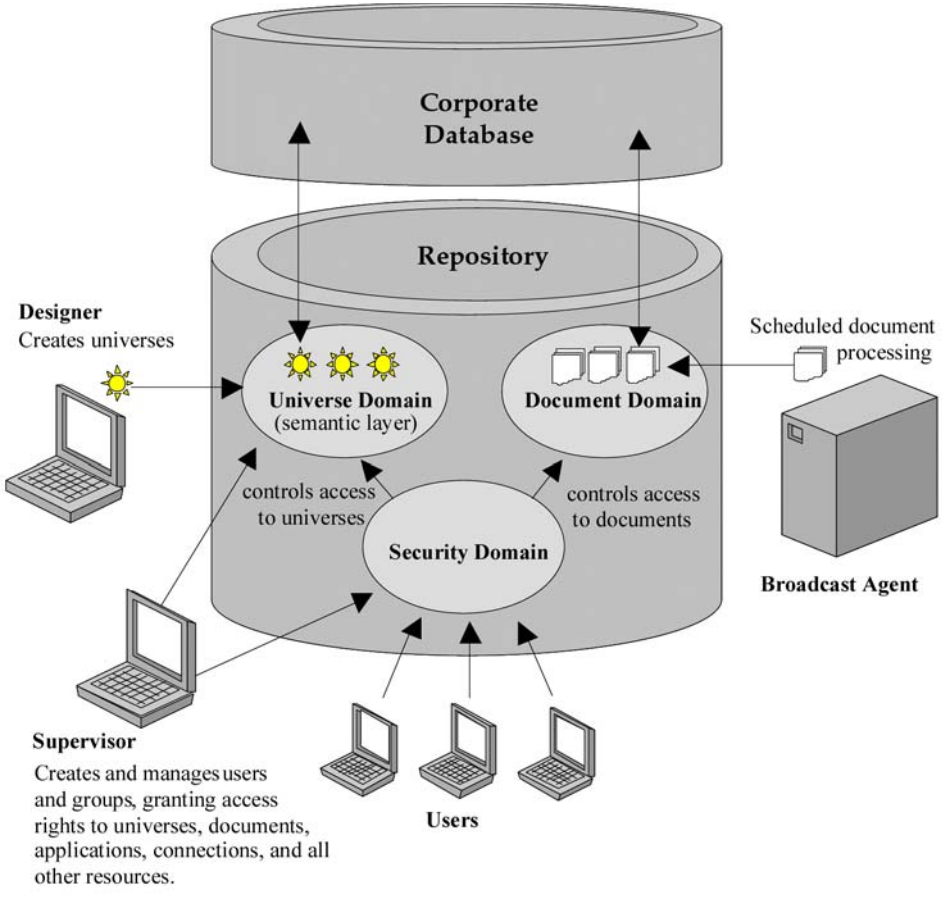

Figure 8.1: The repository

- The universe domain is used to store the contents of each shared universe. This is the area to which designers can export the universes that they create. Each universe is a meta-model of the related corporate database.
- The document domain is used to store the contents of each shared document. This is the area to which users can send documents that are to be shared by all users. BROADCAST AGENT can also access the document domain to run automated document processing tasks.

In many deployments, administrators choose to create multiple repository domains to support the requirements of their organizations. For instance, it is possible to create a deployment that features a single security domain, but two universe domains and five document domains.

In this release, you can also set up several security domains, possibly distributed in different geographical locations.

Security domain selection at login

If the supervisor sets up multiple repositories and grants users access to more than one of them, when users attempt to log on to BUSINESSOBJECTS, they can select the security domain to which they want to connect for the current session.

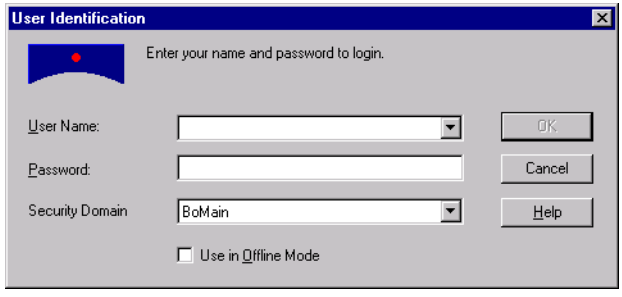

The use of multiple repository security domains can speed up network response times for BUSINESSOBJECTS, since it circumvents the restrictions of a single security domain. If you have a large number of users based in different locations, it may be more efficient to divide them between two or more repositories, each with its own unique security domain. You can also create a series of document and universe domains. This feature is especially interesting for international deployments, where it is advantageous to provide rapid repository access and response for the main body of users.

There are, however, some important restrictions that you should be aware of before setting up multiple repositories, since no interaction between repositories is supported.

WebIntelligence and security domain selection

The web has a different paradigm than a client/server application. A different approach to security domain selection is therefore required for the web-based WebIntelligence.

When WEBINTELLIGENCE users log in using their web browsers, WEBINTELLIGENCE refers to the .key file stored on the cluster manager. The users are unaware of the .key file, and in effect select the security domain by selecting the URL of the WEBINTELLIGENCE system.

Instead of defining .key files with different names, administrators can set up multiple WEBINTELLIGENCE clusters, each of which points to a different repository. Users can select which one to go to by picking the appropriate URL. As bookmarks are available in every browser, it is even easier to let users choose the right domain. This approach preserves the single .key file for each WEBINTELLIGENCE site, a more efficient approach for the administrator and for the end user.

Changing a cluster's repository

When users log into a cluster's repository using its current BOMain.key file, all their identification and access data is retrieved from the repository and cached in a file named BOMain.lsi (for Local Security Information). This file serves as the cluster's security cache concerning that repository.

The Business Objects system also dynamically caches documents and universes that have been accessed by users in its storage areas, located in:

- Under Windows, the <Installation Directory>\Servers\Storage folder and subfolders.
- Under UNIX, the `*_WiData` directories and the Storage directories. Their location was set at installation.

When the system processes subsequent requests for these items, instead of requesting the same information from the repository again, it simply uses its BOMain.lsi file to determine the user's identity, and whether the user has the appropriate access rights.

If you change a cluster's repository, the information in the cluster's new .lsi file no longer corresponds to the repository associated with the documents and universes in the Storage folders. To avoid incorrect results, you must either flush the cluster's Storage folders or archive their contents in a different location.

Which Components Access the Repository?

The repository is the processing center of your deployment. It is in the repository domains where user access is secured, and where shared universes and documents are stored.

Because of its central role in any deployment, the repository is integral to most of the product components, and is accessed by:

- The Broadcast Agent Console
- The Scheduler
- BusObj.exe (Windows) or `bolightsvr` (UNIX) processes
- WEBINTELLIGENCE processes
- All desktop applications (SUPERVISOR, DESIGNER, BUSINESSQUERY, etc.)

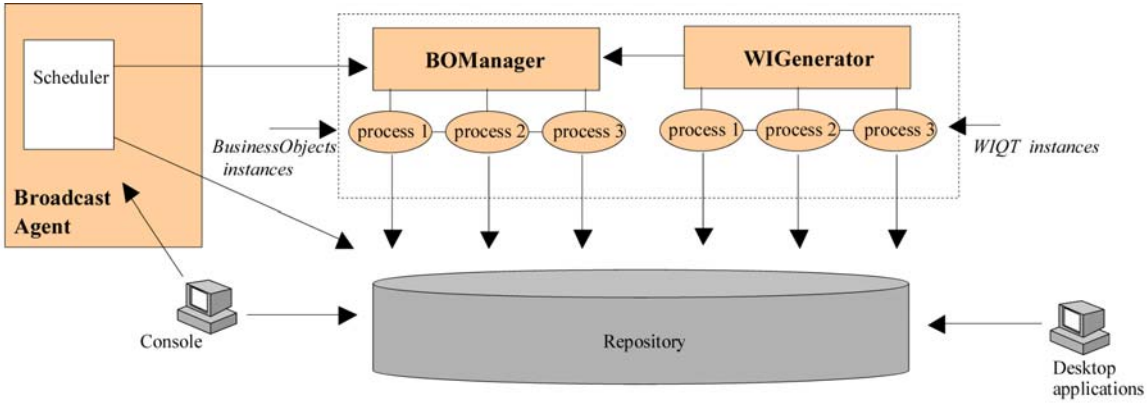

Figure 8.2: Components which access the repository

The Repository and Your Data Warehouse

When you install and configure Business Objects products, you will actually be deploying two databases:

- The repository
- Your corporate database(s) or data warehouse(s)

When you first run SUPERVISOR, you must select a database on which you will store your repository.

In some deployments, the repository will be created in the same relational database management system used to host the data warehouse. In other cases administrators will choose to develop their warehouse in one database while hosting their repository on another system.

For example, you could have:

- A company that hosts both its data warehouse and its repository in an Oracle database (see Figure 8.3).
- A company that hosts its data warehouse in DB2 but chooses to host its repository in a Sybase database (see Figure 8.4).
- A company that hosts its repository in a DB2 database, with connections to three different data warehouses based on different database types (see Figure 8.5).
- A company that hosts different parts of its repository (the security, document and universe domains) in separate databases, which connect to each other and to the corporate data warehouses, as needed (see Figure 8.6).

These choices are left to the administrators responsible for the deployment, and are based on the requirements of the organization and the end users.

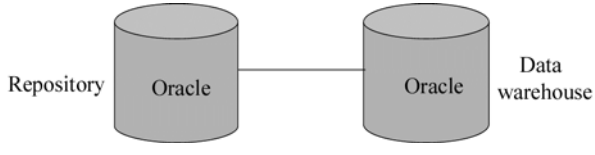

Figure 8.3: Company A hosts both data warehouse and repository in an Oracle database

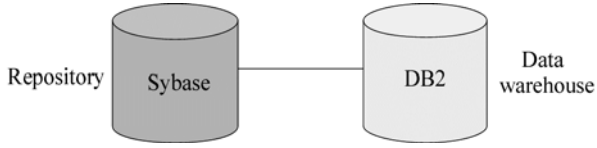

Figure 8.4: Company B hosts its data warehouse in DB2 but chooses to host its repository in a Sybase database

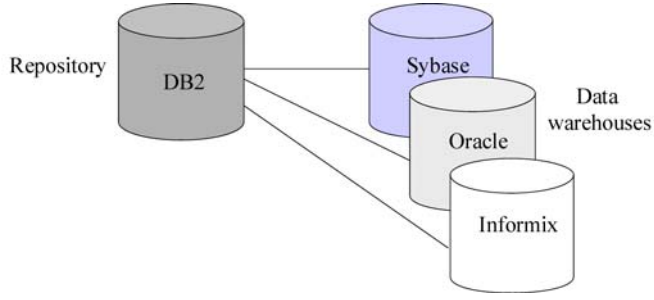

Figure 8.5: Company C hosts its repository in a DB2 database, with connections to three different data warehouses

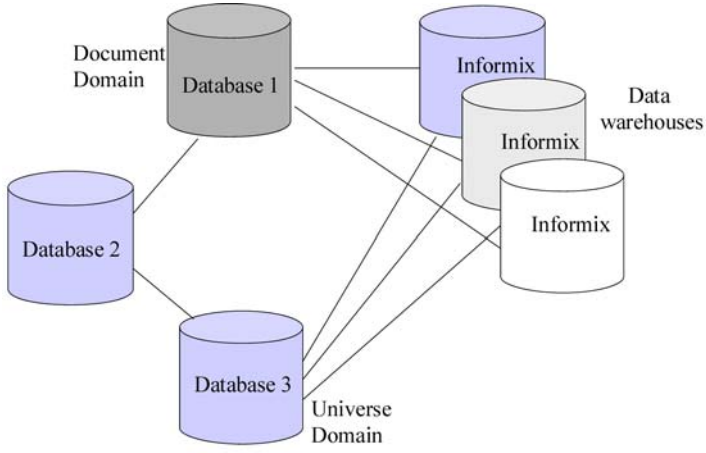

Figure 8.6: Company D hosts different parts of the repository on different databases

Choosing a Repository Database

Administrators who have the opportunity to choose a database platform for their repository are advised to consider the following issues:

- Databases which support row-level locking
- Databases which support BLOBs

Row-level locking

The repository is an SQL application. Many of the day-to-day activities of your users, designers, and supervisors involve database transactions (updates for changed passwords, inserts for exported documents, deletes for removing obsolete documents, etc.). As with any database transaction, these activities invoke the locking mechanisms of the RDBMS to ensure that only one user can update a given value at a given time.

Some database versions have more advanced features than others for handling concurrent user access to the same data. In general, for the repository it is best to select a database system that features row-level locking. This mechanism allows for the highest degree of concurrency and minimum conflicts between multiple users accessing and updating data in the same repository domain(s).

The following list gives some examples of databases that feature row-level locking:

- IBM DB2 Universal Database 6
- Informix Dynamic Server 7
- Microsoft SQL Server 7
- Oracle 7 or 8
- Sybase Adaptive Server 11.9

Often, row-level locking is not "turned on" by default and must be activated at either the database or table level. If such configuration is necessary, you should activate row-level locking for the database prior to creating your repository tables with SUPERVISOR.

See your database documentation for more information on configuring row-level locking.

BLOB support

A BLOB is a Binary Large Object - a common data type in many relational database systems.

When users exchange documents via the repository, the documents are stored in the document domain. The binary content of a document is sliced into units of a defined length and stored in slices in the OBJ_X_DOCUMENTS table of the document domain. Depending on its size and the length of each slice, a single document might be stored in one or more rows in that table.

For each database version, this data is stored in a column of the OBJ_X_DOCUMENTS table that is defined either as a BLOB or a VARCHAR column. If the database supports BLOBs, the driver will use a BLOB type. If not, the driver will convert the binary data to character data and use a VARCHAR type.

In this release, some changes have been made to the management of BLOB data in the document domain. It is now possible to vary the length of the BLOB slice to take advantage of the optimal BLOB size for each database system. The larger the BLOB slice, the fewer rows will have to be inserted for each document stored in the repository. This factor can have an enormous positive impact on performance and network traffic.

We therefore recommend that you select a repository that supports long BLOB columns.

BLOB management: how did it work in BusinessObjects 4.1?

In version 4.1, when you export a document to a repository database that supports BLOBS, the document is sliced into the default BLOB size of 1000 bytes.

This means that if you export a 30 KB document, 30 records are inserted in the OBJ_X_DOCUMENT table in the document domain. If you export a 300 KB document, 300 records are inserted in the document domain table.

The problem here is twofold:

- Importing/exporting a document can take a long time.
- Because importing/exporting documents can take so long, locking problems can occur (see "Row-level locking" on page 284).

Automatic BLOB management in BusinessObjects

Because of the differences in the various RDBMS that may be used, the default BLOB size used is optimized for each RDBMS, and will therefore differ from one to another.

Automatic compatibility between versions 4.1 and 5.x

Documents are not downwardly compatible between version 5.x and 4.1. If a BROADCAST AGENT is used to process a 4.1 document, the document is automatically exported using its original BLOB size, otherwise 4.1 users will not be able to open the document.

Which RDBMSs support BLOBS?

Not all RDBMSs support BLOBs. Where BLOBs are not supported, Varchars are used instead to store your documents. For more information, see your database documentation.

New 5.x Repository Features

Apart from additions and modifications to the repository tables (described in Appendix A), you can now take advantage of this release's distributed architecture in order to set up multiple repositories within your DSS environment. You can also directly import lists of users, without having to create them one by one, and you can copy or move universes using only SUPERVISOR.

Security domain selection

The larger your deployment, the more information must be channeled through the single repository security domain, which can quickly become the bottleneck in your configuration.

With this release, Business Objects products have moved from simple client-server architecture to CORBA. This allows the distribution of the processing tasks over several different physical machines. Instead of having a single repository, or a single primary BROADCAST AGENT, you can now have multiple repositories, and as many BROADCAST AGENT machines as you need.

The more users you have, the more useful it may be to provide two or more repositories, each with its own unique security domain, as well as a series of document and universe domains.

Deploying multiple repositories can be particularly useful in international configurations, where you can considerably speed up your network response times. This is because you are avoiding the security domain bottleneck. For example, rather than having a single, centralized repository located in Europe that serves large numbers of users in both Europe and America, a second repository can be created in America for those users, reducing network congestion for all.

You can select the security domain of the repository you want to work with from a list when you log in:

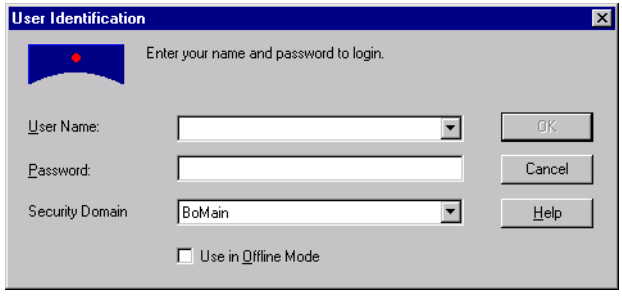

Each item in the list corresponds to a security key file, which points to the security domain controlling access to the associated repository. In previous releases, this key file was named BOMain.key. In this release, you can create multiple repositories, each with its own security domain and key file. You can define any names you want for these key files, although the default name remains BOMain.key.

Note: The WEBINTELLIGENCE system doesn't support security domain selection. see "How do multiple repositories work in practice?" on page 289.

How can you synchronize data between multiple repositories?

It is important to realize that even though you can select your security domain at login, no allowance is made for transparently sharing secured resources (universe, connections and documents) between repositories.

Suppose you start out with a single repository and security domain, and a very large number of users. To improve network performance, you duplicate your configuration and allocate half your users to the new repository. As time passes, your repositories and database information will change, and the two configurations will slip out of synchronization.

Trying to share secured resources by simply moving them from one repository to another may seem the obvious solution, but it will not work.

For example, if you import a universe and the documents that use it from one repository to a local machine, then re-export them all to a second repository, you will not be able to refresh the documents in the second repository. Similarly, if

you have exported a universe separately to two repositories, and created a document that uses that universe in one repository, you will not be able to refresh the document by means of the universe in the second repository.

If it is important that users of both repositories have access to exactly the same data, then you have two options:

- You can periodically log into the remote repository as a General Supervisor, import the documents you require, then log into your local repository and export them there. Users of the local repository will be able to read the documents but not refresh them.
- You can periodically copy the entire corporate database and repository tables from one location to the other. This means that only the repository used as the source of this copy can evolve safely - any changes to the other repository, to which the copy is made, will be lost with the copy.

These operations could be automated. However, when executing either of the above, you will have to prohibit all user access until the operation is completed.

How do multiple repositories work in practice?

When you log in to a Business Objects product (other than WEBINTELLIGENCE), if more than one repository is available, you can select the repository that you want to work with from the displayed listbox. If you also check the *Save for all users* option in BUSINESSOBJECTS when saving your documents or universes locally, then users to whom you send the documents (and this applies to sending documents by email as well as through the system) will be able to access them even if they do not have access to the repository the document was created in.

How does this work? Normally, for security reasons, a document that is stored in the repository also contains the repository ID. This means that the document is only accessible to users who have access to this particular repository. In version 5.x, the *Save for all users* option clears the repository ID from the document. The document is then available to all users to whom it is sent, no matter what repository they work with.

The WEBINTELLIGENCE system doesn't support security domain selection. On system start-up, it looks for a file called BOMain.key (without case-sensitivity, even under UNIX). If it doesn't find the file, it uses the first file found using the command `dir *.key` under Windows, and `ls *.key` under UNIX.

How to set up multiple repositories

You can create additional repositories from SUPERVISOR, by clicking the Admin button to start the wizard, and then specifying a different key file for the repository connection. The procedure is the same as for creating your first repository.

For more information, see the *Supervisor's Guide*.

Importing/exporting user and group lists

Groups of users can be batch imported in text file format. You can view the format of the text file by performing an export operation from SUPERVISOR. User passwords are encrypted and are also saved in the exported file.

To import users in this way, you must first format your data into the same format in a text file, using the import command set (NU = New User, NG = New Group, and so on). For more detailed information, see the *Supervisor's Guide*.

Groups of users can also be imported and exported in interactive mode.

Working With or Without a Repository

There are two modes in which you can work with BUSINESSOBJECTS:

- Workgroup mode - where you work without a repository
- Enterprise mode - where you work with a repository
  In this mode, you can work either online or offline.

The mode in which you save your files determines whether other users are able to access them. By default, files are saved in the mode in which you are already working. For example, if you launched a session in enterprise mode, any file you save is automatically stored in that mode. However, if you want to make a document accessible to another user working without a repository, then you must check the *Save for all users* option in the Save as document dialog box.

Online/offline modes

Online and offline modes are options that apply when you are working in enterprise mode, which means an environment with a repository.

If you are not sure whether the general supervisor has set up a repository, you can check to see whether there is a repository key file located in:

- The LocData or ShData subfolder of your Business Objects folder under Windows
- The `<Application Data Directory>/LocData/` directory for UNIX (you define the <Application Data Directory> at installation).

This encrypted file contains the security domain's connection parameters.

Specifying a mode

When you launch a BUSINESSOBJECTS session, a check box in the User Identification dialog box lets you indicate whether you wish to work in offline mode.

The choice of mode determines whether you can work with shareable or personal files, and the ways in which you can distribute them.

Online mode

Online, the default mode, is appropriate for a networked environment in which the general supervisor has set up a repository. In online mode, you can share resources with other users.

Offline mode

Working in offline mode means essentially that you work with resources that are stored locally on your computer.

You may wish to use this option whenever access to the repository is not required. Alternatively, this mode may be useful when you need to work away from your site, such as with a laptop computer, or whenever your network goes down.

Note: You must have already logged on at least once in on-line mode before you can work in offline mode.

Repository Questions and Answers

This section provides a set of frequently asked questions and answers about the repository and repository domains.

How has the repository structure changed?

New tables have been added with version 5.x releases to store new document attributes and categories. Also, several columns have been added to an existing table for specific BROADCAST AGENT features. For more detailed information, see Appendix A.

How is the repository updated to release 5.x?

You must first install BUSINESSOBJECTS 5.x. To update an existing 4.1 repository, you must log into SUPERVISOR 5.x as the General Supervisor who created the repository to be updated. You are then prompted to update the repository or exit the application. For more information, see the *Installation and Update Guide*.

Note: You should not use this release alongside your 4.1 production repository. You should either start from scratch, or you should make a copy of your 4.1 repository and update it in a test environment. Note that you cannot test version 5.x on top of an existing 4.1 repository without first updating the repository.

Are there any compatibility issues I should know about?

Both BUSINESSOBJECTS 4.1 and 5.x users can access and use the updated repository. However, WEBINTELLIGENCE 2.0 users cannot use the 5.0 repository, since WEBINTELLIGENCE 2.0 is not compatible with BUSINESSOBJECTS 5.x. WEBINTELLIGENCE 2.0 users must therefore first migrate to WEBINTELLIGENCE 2.5.

For more information, see the *Installation and Update Guide*.

How can I reduce the size of the document domain?

Start SUPERVISOR, and run a Scan and Repair on the document domain. Deletions in the repository are "logical" deletes, i.e. the records are still in the database. Scan and Repair should clean out the deleted records.

SUPERVISOR now lets you purge all Inbox documents older than a given number of days from the repository with a single command. Inbox documents are documents sent from user to user via the repository, either directly or by means of BROADCAST AGENT, and cached there until their recipients download them. This lets you recover the document domain space taken up by documents which have been sent but not yet collected by users. This feature can be especially useful when regularly scheduled documents are sent using the *Refresh according to the profile of each recipient* option.

You may also find that users are not deleting corporate documents that have been stored in the repository. You should encourage users to delete old or outdated documents on a regular basis.

How much space does the repository require?

When initially installed, the repository takes up roughly one megabyte. Its size increases as data is added. The security and universe domains expand relatively little per user, group or universe. It is the document domain that typically takes up the most space, and its size is simply equal to the sum of the documents it contains, rounded off to the highest 10 KB.

For more detailed information, see "How much space do you need to allocate for the repository?" on page 164.

Chapter 9 Access and Security

In this chapter

- **Introduction 297**
- **Security at three levels 297**
- **Business Objects Security 299**
  Security and administration files 300
- **Windows NT Authentication Security 302**
  NT Authentication benefits for end users 303
  Restrictions when using particular authentication methods 304
  Setting up the driver 305
  Launching Business Objects applications with the driver installed 307
  Trusted NT domains 307
  Preventing users logging in from previous releases 309
  Frequently asked questions 310
- **RDBMS security 311**
- **Extended Security Commands 313**
  Disabling logins from 4.x 313
  Improved timestamp feature 313
  Script execution disablement for scheduled documents 313

❑ **WebIntelligence and the SSL Protocol 314**

❑ **WebIntelligence and LDAP Authentication 314**
LDAP server authentication of WebIntelligence users 315
Importing user information from the LDAP directory 315
How do the web and LDAP servers and WebIntelligence interact? 318

❑ **WebIntelligence and Firewalls 319**
What is a DMZ Configuration? 320
Communication through a Firewall 320
How WebIntelligence can work through firewalls 322
Installing the WebIntelligence Cluster 322
Configuring the CORBA Environment 323
Using the HSAL in TCP Mode Only 334
Downloading Three-Tier BusinessObjects 339
Troubleshooting Your DMZ Configuration 340
Supported Platforms for WEBINTELLIGENCE 342

Introduction

Security is an increasingly important issue in today's organizations. Implementing and maintaining security for a large number of users can be a costly operation, and has implications for long-term deployment requirements and user functionality.

One of the main strengths of the Business Objects product suite is the range of security options and capabilities provided by the centralized repository. Because security can be applied within universes at the user and group levels, Business Objects can meet the security needs of most types of organizations.

If you are deploying a particularly large-scale distributed solution, such as an extranet, however, you will probably want to build on these basic security options by adding firewalls, which filter the requests leaving and coming into your system from the outside. For more information, see "WebIntelligence and Firewalls" on page 319.

Security at three levels

Security can be applied by the supervisor at three levels. Each level of security has its own advantages and disadvantages:

- Business Objects security: This is provided through the security domain, which is set up when a repository is first created. Each repository that you create has its own security domain, which contains references to the universe and document domains.
- Windows NT authentication security: Checks user names against Windows user names. If the names are the same, users can then launch a Business Objects application simply by double-clicking its icon, without entering any user identification.
- RDBMS security: Allows users to access Business Objects documents through their own individual accounts using only their RDBMS user names and passwords

When deploying, you need to choose carefully the security mechanisms you're going to use, bearing in mind that Windows NT authentication security is not compatible with RDBMS security.

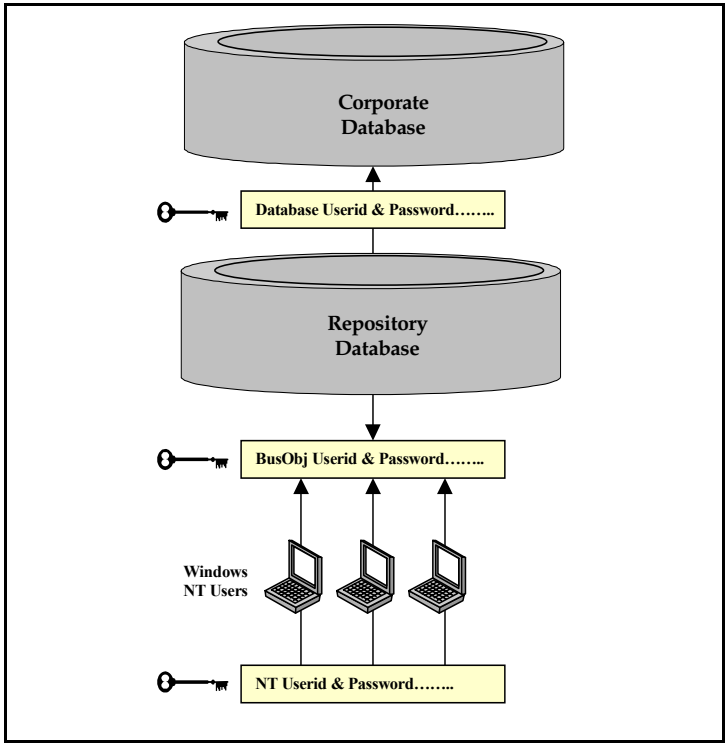

Figure 9.1: Three levels of security

Business Objects Security

Business Objects security is provided through the security domain, which is set up when a repository is first created. Each repository that you create has its own security domain, which contains references to the universe and document domains. The security domain also contains information on the identification of Business Objects users, and on the management of the different applications.

The address of the security domain must be recognized by all workstations using BUSINESSOBJECTS in on-line mode, so that all users can communicate with the other domains of the repository in a transparent manner. This address is contained in the key file, which is created at the same time as the security domain, and which must be distributed to all authorized users.

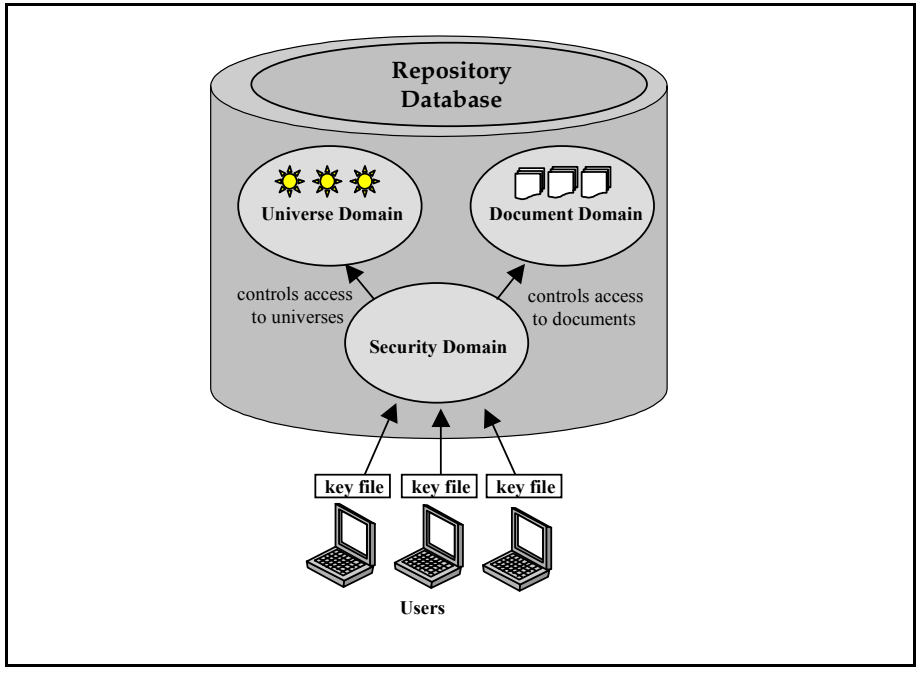

Figure 9.2: Business Objects security

To access documents, users must select a repository to which they have been granted access (via a key file), and must enter a valid user name and password.

Furthermore, user access to objects can be restricted using SUPERVISOR and DESIGNER. For more information, see both the *Supervisor's Guide* and the *Designer's Guide*.

Security and administration files

The security and administration files are not installed with Business Objects products. They are generated dynamically as you create and use repositories and connections. Their location in the file structure is important to the way the products are used.

The security and administration files, stored in either the LocData or ShData folders, include:

- PersonalData Account file (pdac.lsi or pdac.ssi)
- SharedData Account file (sdac.lsi or sdac.ssi)
- User Login Information file (*.lsi or *.ssi)

Personal Data Account file

The Personal Data Account file stores security information concerning the user's personal connections to the database. The pdac.lsi file is created and stored in the LocData folder on the client machine, by default. In that case, a user can use his personal connections only when he works on his own PC.

If you store the Personal Data Account file in the ShData folder on the server (in which case its filename is pdac.ssi), a user can use his personal connections on any PC that has been configured to use the ShData folder to connect to the database.

Shared Data Account file

The Shared Data Account file stores security information concerning shared type connections to the database. These connections are shared among users of Business Objects products. When you install from a Master Setup, by default the Shared Data Account file (sdac.ssi) is set to be created and stored in the ShData folder on the server. Several users can use the connections, as long as their PCs have been configured to use the ShData folder to connect to the database.

If you install from CD-ROM to the server or in a Standalone setup, by default the Shared Data Account file is set to be created and stored in the LocData folder with filename sdac.lsi. In that case, any user working on that PC can use the connections.

Note: For more information about connection types and creating new connections to the database, refer to the *Database Guide* concerning the RDBMS at your site.

User Login Information file (*.lsi or *.ssi)

The User Login Information file stores security information required to run Business Objects products without connecting to the repository. This security information is taken from the repository when the user is working on-line. It is created the first time a user connects to the repository and is updated with each new session. In previous releases, it was called objects.lsi or objects.ssi, depending on its location in the LocData or ShData folder respectively.

Because the current release makes it possible to use multiple repositories with the Desktop products, the name of the file now takes the name of the repository the security information is taken from, so there is one User Login Information file per repository.

- When you install a shared installation from a Master Setup, by default the User Login Information file (with filename *.ssi) is set to be created and stored in the ShData folder on the server, so that all login data for the installation is centralized.
- When you install a Standalone installation from a Master Setup, by default the User Login Information file (with filename *.lsi) file is set to be created and stored in the LocData folder on that PC.

Where are security and administration files located?

In a Standalone Setup, the security and administration files are located by default in the LocData folder. In a Master Setup, they are located by default in the ShData folder. With the exception of the key file, whose filename extension is always .key, the extension of the security files is .lsi when they are stored in the LocData folder, and .ssi when they are stored in the ShData folder.

When updating your configuration, you can choose to change the location of these files from the Setup wizard. In this way you can fine-tune the access possibilities of installed products to suit individual users. See the *Installation Guide* for more information.

Windows NT Authentication Security

The Windows NT authentication driver is a Dynamic Link Library (DLL) delivered with Business Objects desktop products.

The authentication driver avoids the redundant use of login dialog boxes in Business Objects full-client applications. Normally, when you log into Windows NT, you have to specify your user name and password. Once logged in, if you want to run a Business Objects application, you will be prompted again to enter a user name and password, which in many situations will be the same as the ones used to log on to the system.

The NT authentication driver simplifies this process, since it can authenticate Business Objects users by reusing information from the system login. When users launch Business Objects applications, the driver checks (authenticates) Business Objects user names against NT user names. If the user names are the same, users do not have to enter their Business Objects identification.

Besides checking the user name, the driver also verifies the domain the user is logging in from. A list of trusted domains, defined by the General Supervisor, specifies from which domains users can log via the authentication driver. At login time, the Windows driver first verifies if the repository accepts NT authentication; if it does, the driver retrieves the list of trusted domains, otherwise it doesn't allow the user to login. The driver also gets the domain name and the user name for the current user.

Only if the domain is a trusted one, and the user is defined in the repository, will access be granted.

Note: Windows authentication drivers can be used by WEBINTELLIGENCE and BROADCAST AGENT only when Microsoft Internet Explorer 4.01 or later browsers are being used.

For complete information about selecting an authentication method for the WEBINTELLIGENCE system, see the *WEBINTELLIGENCE System Administrator's Guide*.

Using the authentication driver, you have only one password and one login name to maintain and remember. The Supervisor must define a user in the security domain that has exactly the same name as the Windows user. When this user then runs a Business Objects product or module, he/she will not be prompted for authentication.

NT Authentication benefits for end users

The following table shows this driver's benefits to the supervisor and end users:

Benefit	Description
Reduced administration costs	The supervisor has to manage only one set of passwords, i.e., NT passwords, as none are required in Business Objects applications.
Simplified workflow for end users	Users save time and effort when launching Business Objects applications.

Restrictions when using particular authentication methods

NT Challenge/Response and Basic Authentication are subject to two constraints:

- They do not work with a distributed architecture when the *Use Business Objects user name and password* option has been checked, either for the user or user group in SUPERVISOR, or for the universe in DESIGNER.

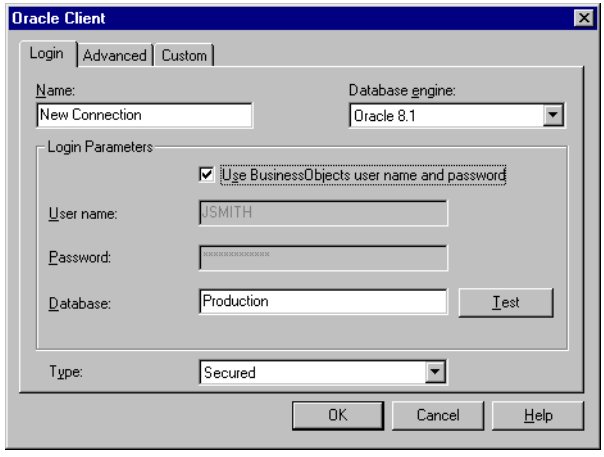

  In this case, the authentication password is not transmitted by the web server.

- For security reasons, Zero Admin BUSINESSOBJECTS users cannot log into the system in offline mode if either of these two authentication methods is used.

Setting up the driver

This section describes what to do before you can use the NT authentication driver, provides information on the drivers installed with desktop products, and explains how to activate and deactivate the driver.

Before you begin

First, ensure that you create users with the same user names as in Windows.

When you work with the Windows NT authentication driver, the user passwords are ignored: once users have logged into Windows NT, they can start the application directly, without entering a user name.

Note: For security reasons, you should always assign passwords to users, whether or not you are using Windows NT authentication. Users' Business Objects passwords do not have to be the same as their Windows NT passwords.

Example **Creating a user**

John Doe's user name is JOHN in NT. You create the user John Doe in SUPERVISOR, user name JOHN, and enter a password. When John Doe logs into NT as JOHN, he can then launch a Business Objects application simply by double-clicking its icon, without entering any user identification.

Drivers installed with Business Objects products

Three drivers are installed with Business Objects products, in the \BusinessObjects 5.0 folder:

Driver	Description
KGLUI50.DLL	**This is the driver called by the application, by default a copy of KBOUI50.DLL.**
KBOUI50.DLL	This is the BUSINESSOBJECTS standard authentication driver. By default, this driver is inactive.
KNTUI50.DLL	This is the Windows NT authentication driver *for the full client product only*. By default, this driver is inactive.

WEBINTELLIGENCE installs the following driver in the same folder:

Driver	Description
KBWUI50.DLL	This is the WEBINTELLIGENCE Authentication driver. This driver supports any WEBINTELLIGENCE authentication option you can select using the Administrator. By default, this driver is inactive.

All these drivers are based on the BUSINESSOBJECTS Plug & Play security driver mechanism. As the application always calls the file KGLUI50.DLL, you simply give that name to the DLL that corresponds to your chosen authentication method.

To activate the NT Challenge/Response driver

1. Delete or move KGLUI50.DLL from the \BusinessObjects 5.1 folder.
2. In the \BusinessObjects 5.0 folder, copy the KNTUI50.DLL file then rename the copy KGLUI50.DLL.

Note: If you are using this method for a distributed system, be sure you don't activate the KNTUIDRV.DLL driver: it is used by the full client only!

To activate the Basic Authentication driver

1. Delete or move KGLUI50.DLL from the \BusinessObjects 5.1 folder.
2. In the \BusinessObjects 5.0 folder, copy the KBWUI50.DLL file then rename the copy KGLUI50.DLL.

To activate the Business Objects standard driver

This driver is activated by default.

To go back to the default driver

1. Delete or move KGLUI50.DLL from \BusinessObjects 5.1.
2. Rename make a copy of the KBOUI50.DLL file then rename the copy KGLUI50.DLL.

Launching Business Objects applications with the driver installed

When the user runs a Business Objects product and an authentication driver is installed, the program checks if there are multiple repositories available. If there are, the program displays a dialog box to select the security domain to log to, then it tries to authorize the users with that specific domain.

If only one security domain is available, no dialog will be displayed and the program will proceed directly to authorization. In both cases, if NT authentication is not allowed for the current domain or if the user is not registered in the repository, a standard error message is displayed.

Trusted NT domains

For a given repository, there can be a list of domains from which users can log using the Windows authentication drivers. Any attempt to use the drivers from a domain that is not included in the list will fail and will deny access to the repository.

There are three typical repository configurations:

- No trusted domains (NT authentication is disabled)
- Any domain is a trusted one (NT authentication ignores the domain)
- Only selected domains are trusted (a list of domains is specified)

This configuration is defined in SUPERVISOR, by the General Supervisor:.

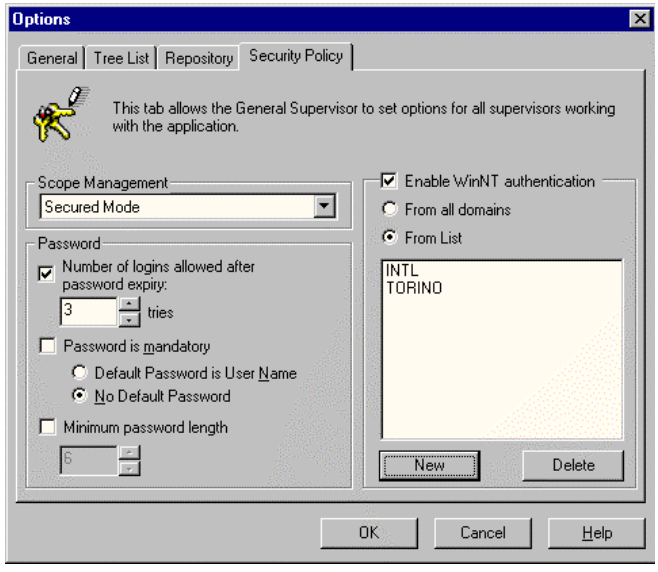

You activate NT authentication by checking the *Enable WinNT authentication* option. This activates the whole group in the dialog.

By default, the option is checked.

Preventing users logging in from previous releases

Because of backward compatibility, users can still access the 5.1 repository from previous product releases, using the old security mechanism. However, you can disable this type of access at the group level, in the Properties tab, by checking the *Disable login from 4.x release or earlier* option (the setting applies to all the users in the group).

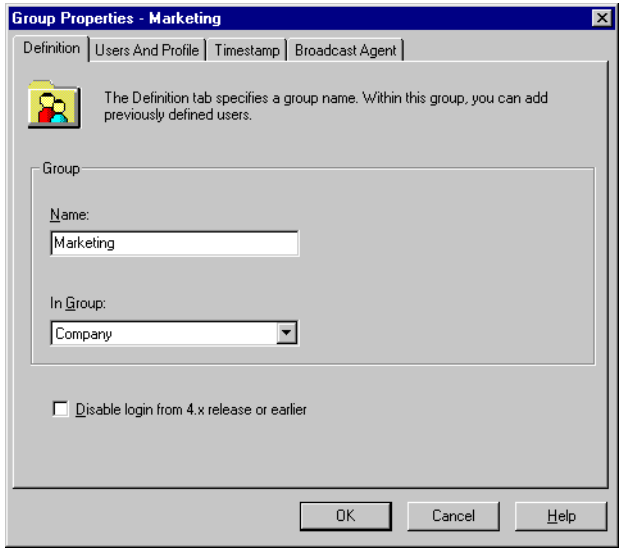

Inheritance applies to this option: You can set it on a group and it will automatically affect all subgroups. Similarly, by setting it at the root of the hierarchy, you can enable or disable logins from earlier versions for all users.

Note: This feature requires enormous system resources, and may take a long time to fully activate. We suggest you use this option only if absolutely required.

For more information on these security options, see the *Supervisor's Guide*.

Frequently asked questions

This section provides answers to the questions that people frequently ask about the NT authentication driver.

Is the driver compatible with BROADCAST AGENT?

Yes. You can send documents from BUSINESSOBJECTS to BROADCAST AGENT from a machine running the NT driver, even if BROADCAST AGENT is running on NT without the driver.

Note: You can install the NT driver on the BROADCAST AGENT machine. For more information, see "Setting up the driver" on page 305.

Does WEBINTELLIGENCE 2.7 support the driver?

No. WEBINTELLIGENCE does, however, support Windows NT authentication for WEBINTELLIGENCE documents.

Can users log on in offline mode when the driver is activated?

Yes, as from BUSINESSOBJECTS 4.1.2. The off-line option that users can add to the Run command is available when the NT driver is installed.

Can the driver be a passthrough to change NT passwords?

No. The driver does not enable you to change NT passwords from BUSINESSOBJECTS.

Can you use the driver to import NT user profiles to the repository?

No. You have to set up NT user profiles and Business Objects profiles in their native applications.

Is the driver compatible with the BOPASS variable?

No. Here's why:

With the NT driver installed, users need to enter their password only when logging onto NT, not when launching Business Objects applications. When the user then tries to connect to the database, for example by running a query, or to the repository, and in either case the connection calls the BOPASS variable, the connection fails and an error message is returned. This occurs because:

- The password is empty in BUSINESSOBJECTS.
- Neither the database nor BUSINESSOBJECTS can retrieve the password from NT.

Workarounds

You can get around this problem by doing one of the following:

- By redefining the database connection without using the BOPASS variable, i.e. by hard-coding the connection password
- By uninstalling the NT driver (see "To go back to the default driver" on page 306)

  Users will then have to enter user names and passwords for both NT and BUSINESSOBJECTS, but they will not experience problems when connecting to their database.

RDBMS security

Business Objects allows users to connect to their database using their own data account for queries and refreshes. Supervisors can choose whether users of a given universe access the database through the universe's data account, or through their own individual accounts.

When using their own data account to access a universe's data, users log into Business Objects products with their database user name and password. In this way, they benefit from the database security mechanisms, and have only one password to remember.

This type of security can be selected by checking the *Use BusinessObjects username and password* option on the Login tab of the connectivity dialog box. When this option is selected, the User name and Password fields are disabled, and the name and password of the currently connected user are used instead.

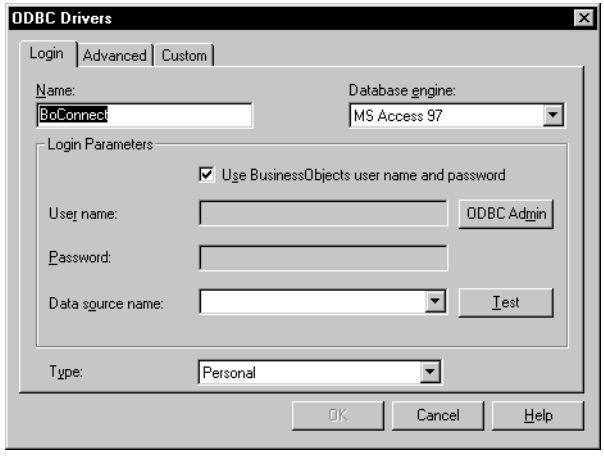

Note: If you select the *Use BusinessObjects user name and password* option in a deployment, ensure that you also select the *Disconnect after each transaction* option on the Advanced tab.

For more information, see the *Supervisor's Guide* or Business Objects *Database Guide* for your database.

Note: This feature is incompatible with Basic or NT Authentication, described above.

Extended Security Commands

SUPERVISOR 5.x features a remodeled and enhanced way of working with security commands, allowing a finer degree of control than earlier versions. Compatibility between version 4.1 and version 5.0 security commands is provided.

For example, a new security command called *Restrict SQL to 'Select' Only* restricts freehand SQL scripts to the Select statement only.

There are several new security commands for WEBINTELLIGENCE.

Disabling logins from 4.x

This new SUPERVISOR option, available from the Group Properties dialog box, lets you prevent users from logging into the repository from previous releases. This can be useful to force users to migrate, for example when you want to use the enhanced Windows NT authentication driver, which doesn't work with version 4.1.

Improved timestamp feature

SUPERVISOR now lets you restrict user or group access to the repository with greater control than was possible in Version 4.1. You can now limit access according to date, time or periodic criteria, or combine these three for flexible and powerful control of access.

Script execution disablement for scheduled documents

To prevent scripts which may cause damage to the system from being executed on the server, administrators can now disable script execution within BUSINESSOBJECTS documents when opened by interactive users.

WebIntelligence and the SSL Protocol

The Secure Socket Layers protocol(**SSL)**, is a means of encrypting or encoding data as it passes between a Web server and a browser. Most popular free and commercial Web servers provide SSL encryption. SSL provides three major benefits:

- Public-key encryption provides a means of exchanging data without transmitting unsecured data or data that would permit snoopers to decrypt the transmissions.
- Authentication certificates permit positive identification of the server and optionally of the client through the exchange of certificates that have been validated by a certificate authority.
- Transaction validation prevents data from being modified between the sender and the receiver. Changes in the data invalidate its authentication code and alerts the receiver that the data has been modified.

You can apply SSL to all WEBINTELLIGENCE transactions simply by enabling the SSL option for each path in your WEBINTELLIGENCE system mapping on your Web server. For more information on SSL, see Appendix C"Understanding Secure HTTP Transactions" on page 383.

WebIntelligence and LDAP Authentication

In recent years, Lightweight Directory Access Protocol (**LDAP**) has become a de facto standard for accessing directory information over the Internet or intranets.

LDAP is a client-server protocol for accessing a directory service. Initially used as a front-end to X.500, it can also be used with standalone and other kinds of directory servers.

LDAP extracts directory information more quickly than a standard directory access protocol, and requires a smaller client. With LDAP, developers can use platform-independent APIs to create applications that leverage a directory. In addition, as LDAP is under the jurisdiction of the Internet Engineering Task Force (IETF), it can more easily evolve to meet Internet requirements.

LDAP server authentication of WebIntelligence users

You can use LDAP authentication to validate users and allow access to WEBINTELLIGENCE while maintaining robust management of user rights and data access offered by Business Objects through the repository.

To do this:

- Set your system's web server to authenticate users via the LDAP server. Netscape Enterprise Server can be configured for this. User IDs will subsequently be authenticated by the web server via LDAP.
- Set WEBINTELLIGENCE to Basic Authentication. This will allow users to log into WEBINTELLIGENCE only if they have been authenticated by LDAP. Each user must still be defined in the Business Objects repository using SUPERVISOR.

Importing user information from the LDAP directory

To simplify user management, BUSINESSOBJECTS/WEBINTELLIGENCE has introduced an import capability that can be used to import users from an LDAP directory. The IT department can develop a script that will export new users from LDAP into the Business Objects security layer on a regular basis. Several fields can be imported, such as:

- Group (I.e., the group the user belongs to)
- User name
- Password
- Profile (user, designer, supervisor…)
- Password options (can change the password, need to change it at the first login…)

Notes:

- The system administrator is responsible for writing the script to extract user information from the LDAP server.
- You are not required to import user passwords for Basic Authentication security. We highly recommend, however, that you not leave the password field empty. We suggest that you use the same secret password for all imported users: nobody will know it except you, the administrator. Be sure to change that password on a regular basis.
- You can streamline assigning user rights by defining security command settings at the group level. This means that when you associate new users with a given group, the users automatically inherit the profile defined for the group.
  For example, if you define an INFOVIEW profile for a group, all the users you add to that group are automatically INFOVIEW users.

For detailed information concerning importing users and groups, see Chapter 8 in the *Supervisor's Guide*.

To import user definitions from an LDAP server:

1. Export the user definitions from the LDAP server. See your LDAP server's documentation for help writing the required script.
2. Package the extracted data in a simple text file, in the syntax required by SUPERVISOR.
3. Import the file to SUPERVISOR, using the Import Users and Groups command in batch mode. This allows you to automate the import with no intervention from the user interface.

Example **Sample text file for a new user, for import in batch mode**

```
BATCH
NU, Sales, TSmith, 3n4n5n6, S, Y, Y, Y, Y, Y, CD, F, Y, 30, N
```

This file creates a new user called TSmith, belonging to the Sales group, with a Supervisor profile. It also sets his password and a series of specific rights. See the *Supervisor's Guide* for more information.

How do the web and LDAP servers and WebIntelligence interact?

The following diagram illustrates how the web server, the LDAP server and WEBINTELLIGENCE interact to provide LDAP authentication of WEBINTELLIGENCE users.

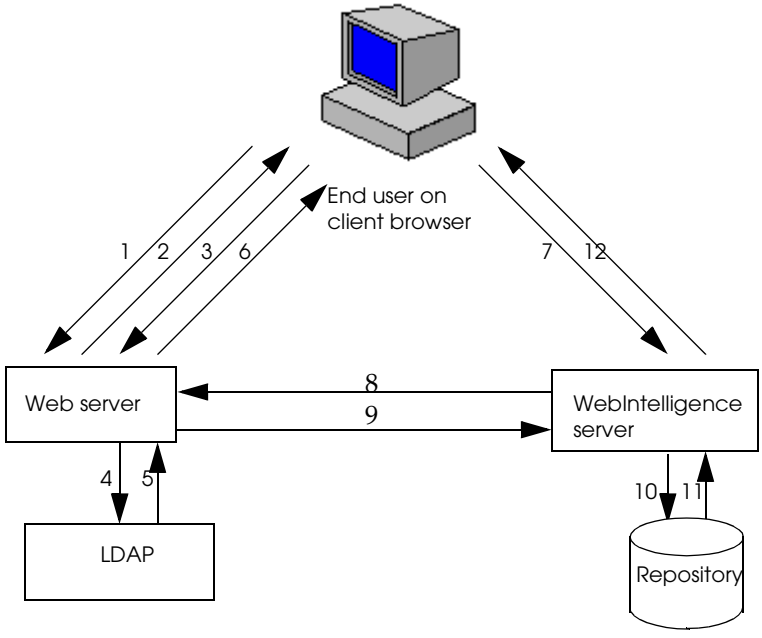

1. The end user enters an HTTP address on the client browser.
2. The web server sends a login dialog box.
3. The end user enters login information.
4. The web server submits that information to the LDAP server.
5. The LDAP server returns OK to the web server.
6. The web server returns the default HTML WEBINTELLIGENCE login page to the client browser. You can easily change the default login page under the /wi directory.
7. The user clicks the blue login button. No login dialog box is displayed.

8.,9. WEBINTELLIGENCE retrieves the user name from the web server.

10. WEBINTELLIGENCE requests user information from the repository.

11. The repository returns the user's rights.

12. If the user has the appropriate rights, the WEBINTELLIGENCE server sends (via the web server) the WEBINTELLIGENCE Welcome page (or another, user-defined start page).

Note: You can easily change the default HTML page to reflect your organization's environment or needs, using WEBINTELLIGENCE SDK.

WebIntelligence and Firewalls

A **firewall** is a frequently used security mechanism in large-scale, generally extranet deployments. A firewall system can be a router, a personal computer, a host, or a collection of hosts, set up specifically to shield a site or subnet from protocols and services that can be abused from hosts outside the subnet. It implements a network access policy by forcing connections to pass through the firewall, where they can be examined and evaluated.

A firewall can greatly improve network security and reduce risks to hosts on the subnet by filtering inherently insecure services. As a result, the subnet network environment is exposed to fewer risks, since only selected protocols will be able to pass through the firewall.

A firewall also provides the ability to control access to site systems. For example, some hosts can be made reachable from outside networks, whereas others can be effectively sealed off from unwanted access. A site could prevent outside access to its hosts except for special cases such as mail servers or information servers.

Lastly, but perhaps most importantly, a firewall provides the means for implementing and enforcing a network access policy. In effect, a firewall provides access control to users and services. Thus, a network access policy can be enforced by a firewall, whereas without a firewall, such a policy depends entirely on the cooperation of users. A site may be able to depend on its own users for their cooperation, however it cannot nor should not depend on Internet users in general.

What is a DMZ Configuration?

The term **demilitarized zone**, or **DMZ**, describes a network topology where the application server is behind a firewall, and separated from the web server, which runs on a different subnet. The web server is then in turn protected by a firewall from the external network.

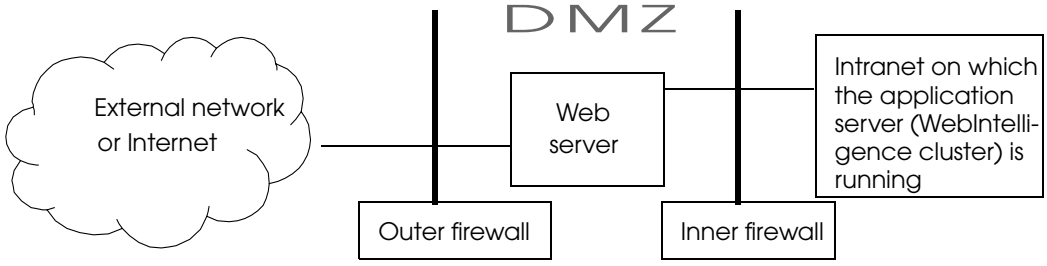

The DMZ is thus a protected buffer zone between an organization's intranet and the external network or Internet. It is designed to keep outside users from accessing a server containing company data.

Communication through a Firewall

One of the key agents in WEBINTELLIGENCE communication is the HTTP Server Abstraction Layer, or **HSAL**. This is a CORBA-based component which acts as the interface between the HTTP server and WEBINTELLIGENCE. It receives user requests for WEBINTELLIGENCE services and converts them from the gateway protocol used by the gateway into a form WEBINTELLIGENCE understands.

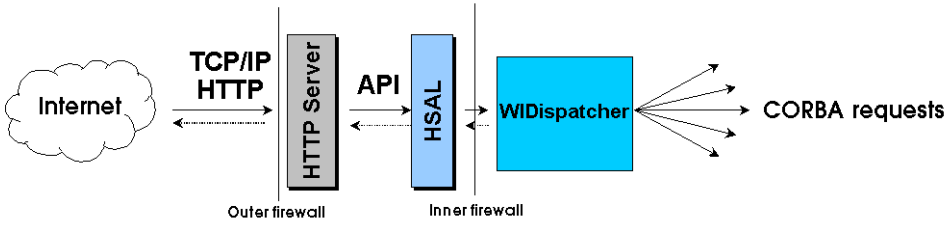

HSAL then passes the request through the firewall to the WIDispatcher, which decides which WEBINTELLIGENCE process the request should be sent to. The WIDispatcher converts the textual URLs into CORBA requests that can start the required WEBINTELLIGENCE process, then transmits the requests to that process via the OSAgent and the OAD.

The HSAL is always located on the web server. In a DMZ configuration, therefore, the HSAL is in the DMZ, separated from the WEBINTELLIGENCE cluster by the inner firewall.

In order for the HSAL to communicate with the WIDispatcher, the OSAgent and the OAD in the cluster, the ports used by those components must be open in the firewall. By default, however, these ports are not set, but dynamically allocated by the CORBA system at run-time.

If you're using WebIntelligence SDK...

When you use WEBINTELLIGENCE SDK, requests are performed by the WIAPIBroker through either the WICOM or WIBean object. In this case, the WICOM or WIBean object replaces the HSAL, and WIAPIBroker runs with the WIDispatcher..

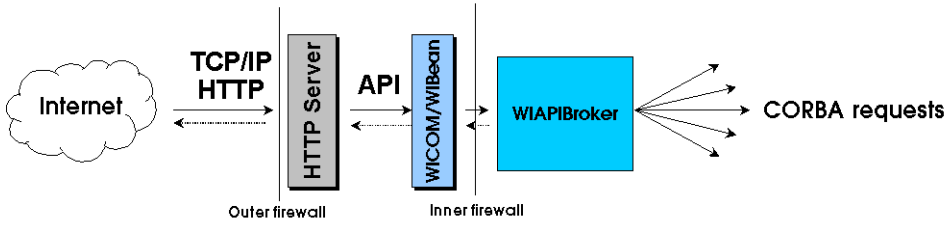

How WebIntelligence can work through firewalls

You can set up the WEBINTELLIGENCE system to work a double-firewall configuration in either of two ways:

- **In CORBA mode**, which requires configuring the TCP/IP and ORB environment to use the correct TCP and UDP ports for each relevant server component.
- **In TCP mode only**, a dramatically simplified configuration which was new with the 5.1/2.6 Service Pack 1 release. For more information, see "Using the HSAL in TCP Mode Only" on page 334. This mode may somewhat slow down communication, depending on your configuration.

Both of these configurations require installing the WEBINTELLIGENCE cluster beforehand. All of these steps are described in detail in the sections below.

Note: One and only one WIDispatcher or WIAPIBroker must be enabled in either of these configurations.

Installing the WebIntelligence Cluster

Installing WEBINTELLIGENCE in a double-firewall configuration is simple:

1. Install your WEBINTELLIGENCE cluster normally, beginning with the cluster manager, then adding cluster nodes if required. It doesn't matter whether you choose to have the web server configured automatically or manually.

2. On the web server, install WEBINTELLIGENCE as a cluster node.
   - Enter the same OSAgent Port value for this machine as you did for the rest of the WEBINTELLIGENCE cluster (Step 1).
     In this document, the variable PORT_B represents the OSAgent Port value.
   - Under Windows, do *not* check the automatic start-up option for Business Objects services. This prevents the WebIntelligence Manager and WebIntelligence Orb services from running but allows the node's HTTP server abstraction layer (HSAL) and the ORB DLLs to function normally.

Configuring the CORBA Environment

If you have chosen to modify CORBA kernel components in order for your system to communicate through a DMZ, you must set the ports that your deployment will use, on both the web server and in the WEBINTELLIGENCE cluster.

Configuring the TCP/IP and ORB environment consists of the following steps, each of which is described in detail below:

1. Get the TCP/IP information for all the machines in the cluster. You will need this information for the following steps.
2. Set the port information on the cluster manager and the cluster server running the Dispatcher if it isn't running on the cluster manager.
3. Set the port information on the web server.
4. Set the port information on the firewall to permit communication with the WEBINTELLIGENCE cluster manager, as well as cluster servers running the Dispatcher and WIAPIBroker (if you're using WEBINTELLIGENCE SDK).

Note: Throughout this section on configuration, the symbol `( . . . )` at the end of a line of code, configuration text file entry, etc., indicates that the following line should be entered continuously with no carriage return. The line following this sign is indented.

As each relevant component in the WEBINTELLIGENCE system must communicate through a different port, these steps will define the following seven ports:

Port*	Direction	Port's Function and Protocol**
PORT_A	intranet > DMZ	Allows the components write to the web server using TCP.
PORT_B	DMZ > intranet	Lets the web server communicate with the OSAgent using UDP.
PORT_C	DMZ > intranet	Lets the web server communicate with the OSAgent using UDP.
PORT_D	DMZ > intranet	Lets the web server communicate with the OSAgent using TCP.
PORT_E	DMZ > intranet	Lets the web server communicate with the OAD on the cluster manager using TCP.
PORT_F	DMZ > intranet	Lets the web server communicate directly with the Dispatcher process on the cluster manager using TCP.
PORT_G	DMZ > intranet	Allows the web server to communicate directly with the WIAPIBroker process on the cluster manager, using TCP (for WEBINTELLIGENCE SDK use only).

* The values in this column are to be used as variables. When you are configuring TCP/IP and ORB information, you will substitute a numeric port address.

** UDP and TCP are two different Internet communication protocols used by CORBA.

Note: If you're using the web server as a standalone server (i.e. in workgroup mode as opposed to being part of a domain), you'll need to follow these steps and another. See "If you're using a standalone web server..." on page 330.

Obtaining TCP/IP information

You must obtain the IP address and subnet mask for each machine in the cluster.

- The **IP address** is often depicted as four groups of decimal digits separated by periods, such as 130.5.5.25.
- The **subnet mask** is a network prefix which identifies a separate part of an organization's network in the same format.

To obtain this information, type one of the following commands at the DOS or command line prompt:

- For a Windows NT server, `ipconfig`
- For a UNIX server, `ifconfig -a`

Example **Sample TCP/IP information for cluster manager**

```
Ethernet adapter Elpc5751:
            IP Address . . . . .  : 192.168.0.1
            Subnet Mask  . . . . .: 255.255.0.0
            Default Gateway . . . : 192.168.255.254
```

Setting port information in the cluster

By setting the port information on the cluster manager and all the cluster nodes, you tell the entire cluster precisely which ports are available for TCP and UDP communication.

Add required environment variables.

- On an NT cluster manager, add two new variables (Control Panel | System | Environment tab):

Variable	Value
OSAGENT_CLIENT_HANDLER_UDP_PORT	PORT_C
OSAGENT_CLIENT_HANDLER_TCP_PORT	PORT_D

- On a UNIX cluster manager, insert the following lines at the beginning of the `WebiEnv.sh` file:

```
OSAGENT_CLIENT_HANDLER_UDP_PORT=PORT_C ; export   (...)
        OSAGENT_CLIENT_HANDLER_UDP_PORT
OSAGENT_CLIENT_HANDLER_TCP_PORT=PORT_D ; export   (...)
        OSAGENT_CLIENT_HANDLER_TCP_PORT
```

- Add the OSAGENT_ADDR variable to all cluster node machines. This variable is set using the cluster manager's IP address.
  - On NT machines (Control Panel | System | Environment tab):

Variable	Value
OSAGENT_ADDR	<Cluster manager's IP address>

  - On UNIX machines, insert the following line after the line

    `OSAGENT_CLIENT_HANDLER_TCP_PORT`

    in the `WebiEnv.sh` file:

    `OSAGENT_ADDR=<Cluster manager IP address>;export OSAGENT_ADDR`

Edit specific CORBA kernel parameters.

You need to change specific CORBA kernel parameters containing the OAD and WIDispatcher command lines on the cluster servers to start them using a specific port.

If you're using WEBINTELLIGENCE SDK, you will also need to change the WIAPIBroker parameter.

- For NT servers, use the Registry Editor (regedit.exe).
- For UNIX servers, use the `boconfig.cfg` file, which provides Registry emulation.

Add the text in bold to the default value for the OAD key:

Note: The initial command names might differ from the Registry to the `boconfig.cfg` file due to platform-specific settings. The text you must add, however, is the same regardless of the platform.

Key path	• Under Windows: HKEY_LOCAL_MACHINE \ SOFTWARE \ Business Objects \ Shared \ Administration Services \ 5.0 \ Properties • Under UNIX: `/Business Objects/Business Objects/Shared/Administration Services/5.0/Properties`
Key name	OadCmdLine
Value	`"<Installation directory>\Server\ORB2.5\bin\oad.exe" -OAport PORT_E -ORBisNTService 1 -C -k -t 60`

Add the text in bold to the default value for the WIDispatcher key:

Key path	• Under Windows: HKLM \ SOFTWARE \ Business Objects \ Shared \ Administration Services \ 5.0 \ WIDispatcher • Under UNIX: `/Business Objects/Shared/Administration Services/5.0/ WIDispatcher`
Key name	Args
Value	`-ORBisNTService 1` **`-OAport PORT_F`** `-OalocalIPC 0 -ORBconnectionMaxIdle 0 -ORBcacheDSQuery 0`

If you're using WebIntelligence SDK, define a port for the cluster machines.

If you're using WEBINTELLIGENCE SDK, add the text in bold to the default value for the WIAPIBroker key:

Key path	• Under Windows: HKLM \ SOFTWARE \ Business Objects \ Shared \ Administration Services \ 5.0 \ WIAPIBroker • Under UNIX: `/Business Objects/Shared/Administration Services/5.0/WIAPIBroker`
Key name	Args
Value	`-ORBisNTService 1` **`-OAport PORT_G`** `-OalocalIPC 0 -ORBconnectionMaxIdle 0 -ORBcacheDSQuery 0`

Setting the port configuration on the web server machine

You need to define two environment variables so that the OSAGENT_ADDR variable points to the address of the WEBINTELLIGENCE cluster manager.

This will make sure the HSAL or WICOM/WIBean component inherits these settings.

- For NT servers, use Control Panel | System | Environment tab to add the following variables:

Variable	Value
OSAGENT_ADDR	<IP address of cluster manager>
WIHTTPArgs	-ORBisNTService 1 -ORBconnectionMaxIdle 0 -OAlocalIPC 0 -ORBcacheDSQuery 0 -OAPort PORT_A

- For UNIX servers, add the following to the Netscape server's start file in the `<Installation Directory>/https-<name of web server>` directory, or to your web server's dedicated configuration file:

```
OSAGENT_PORT=<OSAgent Port Number>;export OSAGENT_PORT
WIHTTPArgs ="-ORBconnectionMaxIdle 0 -OAlocalIPC 0  (...)
        -ORBcacheDSQuery 0 -OAPort PORT_A";export WIHTTPArgs
OSAGENT_ADDR=<Cluster manager IP address>;export OSAGENT_ADDR
```

- For Apache servers on UNIX, add the following to the apachectl file in the `<Apache Install Directory>/bin>` directory.

```
OSAGENT_PORT=<OSAgent Port Number>;export OSAGENT_PORT
WIHTTPArgs ="-ORBconnectionMaxIdle 0 -OAlocalIPC 0  (...)
        -ORBcacheDSQuery 0 -OAPort PORT_A";export WIHTTPArgs
OSAGENT_ADDR=<Cluster manager IP address>;export OSAGENT_ADDR
```

Setting port information on the inner firewall

How you actually set port information depends on the type of firewall you're using. NT firewalls come with a graphical interface which should allow you to do this easily. For UNIX firewalls, you must write a script.

To correctly set the port information, you must:

- Open the following ports from the web server to the cluster manager:
  - PORT_B
  - PORT_C

- PORT_D
- PORT_E
- PORT_F

The ports PORT_B and PORT_C use the UDP protocol, and the rest use TCP.

- Open the port PORT_A from the cluster manager to the web server (TCP).

If you're using a standalone web server...

If you're using the web server as a standalone server, you have no Domain Name Server (DNS) to locate IP addresses or provide name resolution for communication between the web server and the cluster manager. In this case you will need to add two additional entries to the "hosts" file on both the cluster manager and the web server machines:

- Under Windows, in the \winnt\system32\drivers\etc directory
- Under UNIX, in the `/etc` directory

```
IPADDRESS OF WEBSERVER              name of webserver
IPADDRESS OF CLUSTERMANAGER         name of clustermanager
```

For example:

```
10.2.0.1                            mywebserverindmz
192.168.0.1                         myclustermanager
```

An example configuration with UDP and TCP communication

This section describes a sample WEBINTELLIGENCE deployment using a firewall configuration. First, a diagram represents the sample deployment configuration. Below it, you'll find the firewall rules.

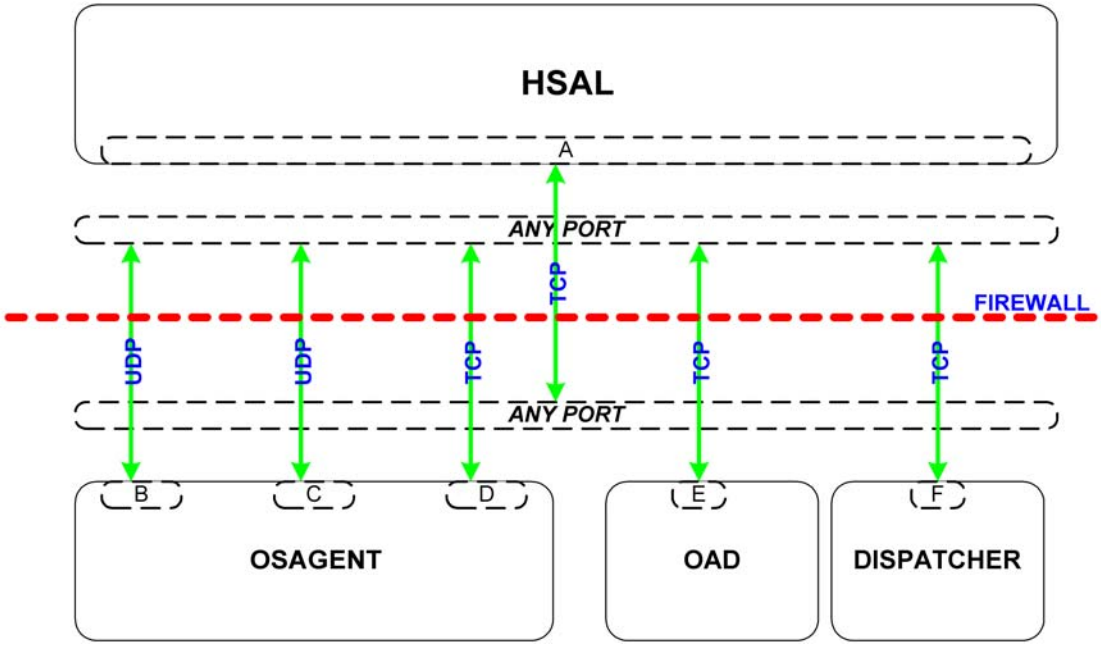

Firewall rules

The firewall rules are the following:

Source Address	Source Mask	Destination Address	Destination Mask	Protocol	Source Port	Destination Port
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	TCP	Any	A
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	UDP	B	Any
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	UDP	C	Any
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	TCP	D	Any
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	TCP	E	Any
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	TCP	F	Any
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	TCP	A	Any
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	UDP	Any	B
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	UDP	Any	C
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	TCP	Any	D
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	TCP	Any	E
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	TCP	Any	F

If you have a router or gateway in your network

If you have a router or gateway in your network (either a router/gateway from DMZ into the Intranet, or a router/gateway from the INTRANET into the DMZ, that is NOT the firewall machine), you need to add the following rule on your firewall:

- When the firewall machine *is* the router/gateway, open all traffic "to" and "from" the gateway (if the gateway is the firewall).
- When a machine *other than the firewall* is the router/gateway, open the specific ports (A - E/F) required for communication between the DMZ and the intranet.

Note: Without this opening, the traffic will always die at the firewall, without receiving requests from or routing them to, the router/gateway machine.

If you're using WebIntelligence SDK...

This example is based on an HSAL/WIDispatcher configuration. If you are using WEBINTELLIGENCE SDK, you must allow WICOM or WIBean and the WIAPIBroker communicate correctly by adding the following rules:

Source Address	Source Mask	Destination Address	Destination Mask	Protocol	Source Port	Destination Port
192.168.0.0	255.255.0.0	10.2.0.0	255.255.0.0	TCP	G	Any
10.2.0.0	255.255.0.0	192.168.0.0	255.255.0.0	TCP	Any	G

Using the HSAL in TCP Mode Only

Until now, the implementation of a double firewall required the use of both the UDP and TCP communication protocols between the HSAL on the web server and the WEBINTELLIGENCE modules in the cluster. Several ports were required for this type of CORBA communication, which opens the system to greater security risks. This type of communication is described in the preceding sections.

With the 5.1/2.7 release, the WEBINTELLIGENCE system's HSAL can operate in TCP mode only, thereby reducing the number of open ports required between the web server and the WEBINTELLIGENCE cluster to one.

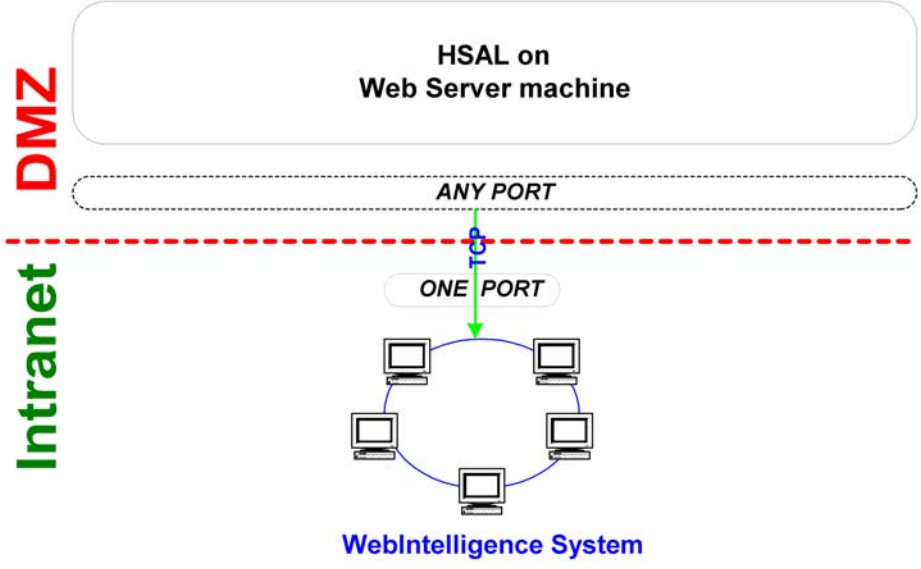

In this type of configuration, the web server recognizes one or more WEBINTELLIGENCE servers by their IP addresses, and knows which port number to use for each of them using the value defined for the web server's **WIHTTPArgs variable**.

On the WEBINTELLIGENCE server side, two new modules simulate an HSAL which allows only a single inbound TCP port to be open, as well as corresponding outbound TCP ports between one or more servers in the cluster and the web server.

If you're using WebIntelligence SDK...

Currently, you cannot configure WIAPIBroker to run through a single TCP port. If you want to configure a WEBINTELLIGENCE SDK implementation (ASP or JSP), you will have to use the six TCP/UDP ports through your firewall.

Be aware, however, that if you configure WIAPIBroker using six ports, you have no failover fuctionality!

New modules for enabling TCP mode

This type of configuration requires the standard HSAL as well as two modules that are new with this Service Pack. They are automatically installed whenever WEBINTELLIGENCE is installed on a cluster server:

Module	Description	Installation Directory
TCPHSAL	Acts as an HSAL proxy within the DMZ. • Under Windows, this module is called tcphsal.exe. • Under UNIX, this module is called `tcphsal`.	Both modules are installed in the same directory. • Under Windows: (...)\Server\System 2.5\bin • Under UNIX: `<Installation Directory>/Server/System2.5/bin`
WIHSALManager	Resides on the same cluster machine as TCPHSAL, monitoring it to make sure it doesn't fail. WEBINTELLIGENCE system administrators can administrate this monitor using the Business Objects Services Administrator. • Under Windows, this module is called wihsalfw.exe. • Under UNIX, this module is called `wihsalfw`.	

Enabling WIHSALManager on a cluster manager

In order for this type of DMZ-enabled configuration to be activated, the WIHSALManager module must be enabled on at least the cluster's cluster manager. If you are running a system with multiple cluster managers using an IP/firewall redirector or Cisco Local Director, you can enable this module on multiple cluster managers.

You enable the module using the Business Objects Services Administrator:

1. In the Administrator's Host or Report page, make sure WIHSALManager is followed by a green checkmark.

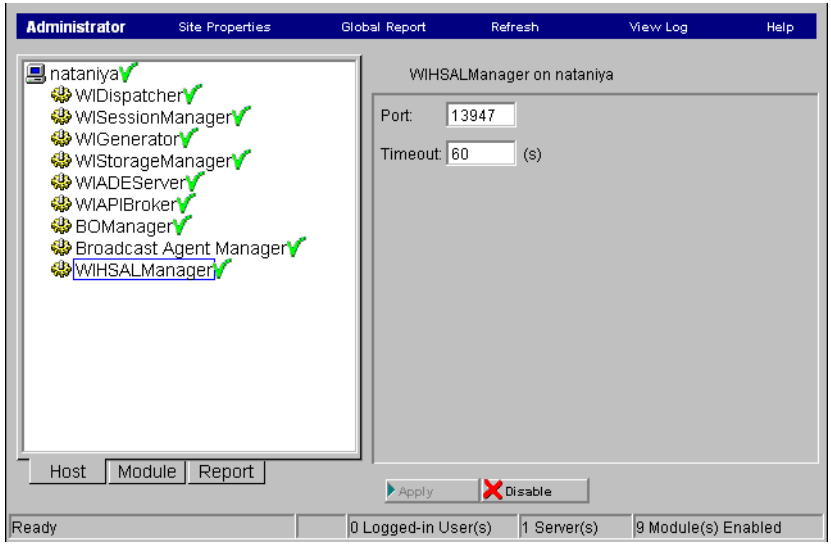

   If it isn't, select the module name then click the Enable button at the bottom of the right pane.

2. With the module selected, enter a valid number for a port that is not being used by any other application.

3. In the Timeout field, enter the number of seconds the process is given to respond to a request before it times out.

Repeat this procedure for each cluster manager on which you want this module enabled. If you enable this module on multiple servers, you will provide failover for this feature (see the following section).

You can define the same port number for each server if you want.

Defining the WIHTTPArgs variable on the web server

Once the WIHSALManager module has been activated on at least one cluster manager, you must add an extra variable called WIHTTPArgs to the web server's configuration file with the following values:

```
-tcpmode 1 -tcpservers [<server1 ip address>:<port1>;   (.../...)
     <server2 ip address>:<port2>[...]]
```

in which `<server1 ip address>:<port1>;<server2 ip address>:<port2>[...]` represents the IP address and port number for any cluster managers running the WIHSALManager.

For example:

```
-tcpmode 1 -tcpservers 192.168.0.0:13947
```

To enter this information for the WIHTTPArgs variable:

- Under Windows, set the variable manually in the list of system environment variables (Control Panel | System | Environment tab).
- Under UNIX, how you set the WIHTTPArgs variable depends on the web server you're using:
  - For Apache in CGI mode, you add the following line to the `webi26.conf` file:

    ```
    SetEnv WIHTTPArgs "-tcpmode 1 -tcpservers [<server1 ip (.../...)
    address>:<port1>;<server2 ip address>:<port2>[...]]
    ```
  - For Apache in DSO mode, you must modify the `wstart` script (in the WEBINTELLIGENCE Setup directory) and `wapachectl` script (in the Apache `bin` directory):
    After the following line:

    ```
    OSAGENT_PORT=<port number>; export OSAGENT_PORT
    ```

    add the line:

    ```
    WIHTTPArgs= "-tcpmode 1 -tcpservers [<server1 ip (.../...)
    address>:<port1>;<server2 ip address>:<port2>(.../...)
    [...]];export WIHTTPArgs
    ```
  - For Netscape, you add the following line to the web server's start file (usually in the `https-<webservername>` directory):

    ```
    WIHTTPArgs="-tcpmode 1 -tcpservers (.../...)
    [<server1 ip address>:<port1>; (.../...)
            <server2 ip address>:<port2>[...]]; export WIHTTPArgs
    ```

The web server will connect to only one of the specified servers at a time. If it gets no response from the first server in the list, it tries the next. If the next server fails to respond, it tries the next, and so on.

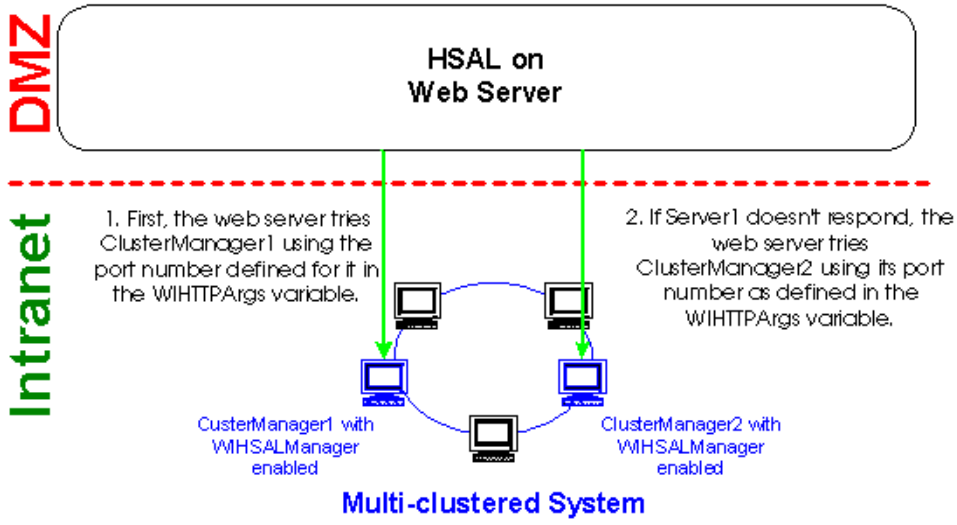

Note: If the -tcpmode option isn't present in the variable, or if the variable isn't initialized, the HSAL will function in CORBA mode.

To deactivate TCP mode, simply replace the "1" following `-tcpmode` with a "0":

```
-tcpmode 0 ...
```

Firewall configuration

In TCP-only mode, you must configure the firewall so that:

- The web server can access the port specified on the WEBINTELLIGENCE server.
- The WEBINTELLIGENCE server can use the same port to open a socket to the web server.

Example **Firewall configuration using the Windows 2000 routing utility**

Input filters:

Source Address	Destination Address	Source Port	Destination Port
WEBINTELLIGENCE server	Web server	Port number	

Output filters:

Source Address	Destination Address	Source Port	Destination Port
Web server	WEBINTELLIGENCE server		Port number

Downloading Three-Tier BusinessObjects

You should have no problem downloading the three-tier version of the BUSINESSOBJECTS product through INFOVIEW.

If you are using a single firewall, you need no special configuration.

In a double-firewall configuration, open the HTTP port 80 in the outer firewall and the TCP port 13794 in the inner firewall with the option "Drop all packets except those that meet the criteria below" option enabled.

Troubleshooting Your DMZ Configuration

This section describes both the tools you can use to troubleshoot your configuration, and solutions to problems you may encounter with your WEBINTELLIGENCE configuration.

Troubleshooting tools

The following tools are indispensable for finding out what is causing problems in your WEBINTELLIGENCE system:

Tool	Description
OSFind	You can launch the Inprise Smart Finder by typing at the DOS or UNIX prompt: `osfind [-a] [-o]` • The `-a` option returns all the OSAgents in your domain • The `-o` option returns all the OADs (Object Activation Daemons) in your domain Use the results to check that you have only one OSAgent running in the cluster, and that only one OAD process is running per cluster node.
OADUtil list	The `oadutil list` command returns all ORB object implementations registered with the Object Activation Daemon. Each OAD as its own implementation repository database where the registration information is stored. To find out which objects are registered with the OAD, see the Troubleshooting chapter in the *WebIntelligence System Administrator's Guide*.
TRACERT	The `tracert <ipaddress>` command lists the IP addresses involved in completing the communication and identifies any routers or gateways. This may "die" at the firewall, but at least it will identify all the components involved from the beginning of the transaction to the firewall.

The first login request times out

The first UDP call from the web server to the cluster manager is a broadcast. Therefore:

- Check that the cluster manager and all the required processes are up and responding. If they are, try accessing WEBINTELLIGENCE from the cluster manager.
- Check for the correct oadcmdline parameter on the cluster manager (see "Edit specific CORBA kernel parameters." on page 327).
- Run `osfind` from the web server to verify that the OSAgent is listening at the location and on the port expected.

WebIntelligence does not respond and times out

Several factors could be causing this problem:

- Windows environment variables are either incorrect, or have not been defined as *system* variables.

  Run `osfind` as the user under which the application is running. This will indicate if / where the OSAgent is defined.
- Hostname entries have not been defined on all machines.

  Visibroker uses the CORBA naming service as well. Ping each machine by name from each of the other machines. If pinging produces no response from any machines, register hostnames for those machines on all the other machines in the cluster in their "hosts" file:

  - Under Windows, in the \winnt\system32\drivers\etc directory
  - Under UNIX, in the `/etc` directory
- The firewall rules are incorrect.
  - Make sure traffic is allowed through the firewall in both directions.
  - Allow traffic between routers / gateways.

WebIntelligence stops working

This could have several causes:

- The cluster manager was restarted but the web server (including the HSAL on it) was not rebooted.
  - Frequently web servers are managed by different groups in order to support all web applications. This means that while you may be able to stop and restart the WEBINTELLIGENCE system whenever you want, you may need to give various groups 24-hour or longer notice before stopping and restarting the web server.
  - The cluster node (HSAL) was bound to a previous instance of OSAgent.
- The network configuration has changed.

  The move from test to production configurations can introduce new routers / gateways. If this is the case, each new router or gateway needs to be registered with the system.
- A cluster node has been added.

  Firewall rules are required to allow all traffic with the subnet (allow the cluster to communicate). Rules on the cluster manager side need to be at the subnet level to allow the cluster nodes and the DMZ to communicate.

Supported Platforms for WEBINTELLIGENCE

WEBINTELLIGENCE supports the following platforms and configurations:

Cluster Manager/ Application Server	Cluster Node/ Web Server	Supported?
Windows	Windows	Yes
UNIX	Windows	Yes
UNIX	UNIX	Yes
Windows	UNIX	No

Chapter 10 Document Processing and Distribution with Broadcast Agent

In this chapter

- **Broadcast Agent Features 344**
  - Automated document scheduling and distribution 344
  - A CORBA solution 346
  - Managing Broadcast Agent activity and errors 346
  - Access and security 347
- **Broadcast Agent Components 348**
  - The Business Objects Services Administrator 349
  - The Scheduler 350
  - The BOManager 351
  - The WIGenerator 352
  - The Broadcast Agent Console 352
  - Broadcast Agent workflow 353
- **Broadcast Agent Sizing Guidelines 354**
- **Troubleshooting and Optimization 355**
- **Where to get more information 358**

Broadcast Agent Features

BROADCAST AGENT features fall into the following categories:

- Automated document scheduling and publishing - how end users submit documents, and the options they can use
- On the server side, CORBA-enabled features such as load balancing
- Tracing BROADCAST AGENT activity
- Access and security

Automated document scheduling and distribution

BROADCAST AGENT provides scheduled or batch processing of BUSINESSOBJECTS and WEBINTELLIGENCE documents. All the end user has to do is send a document to BROADCAST AGENT using BUSINESSOBJECTS or WEBINTELLIGENCE, and specify scheduling and distribution information. No understanding of the server configuration is required.

BROADCAST AGENT lets users not only automate simple actions such as refreshing and printing documents, but control processing with conditions which trigger distribution when pre-defined events occur.

Users can choose to publish documents to the repository, on a web server or on the server file system. Users with access to the repository can view documents that BROADCAST AGENT has processed in BUSINESSOBJECTS or WEBINTELLIGENCE. BROADCAST AGENT supports HTML publishing on the web, on channels and on your intranet or extranet.

The following table gives an at-a-glance overview of what BROADCAST AGENT can do, and indicates which features are new:

Broadcast Agent at a Glance		New?
Broadcast Agent can do all this with documents...	• Refresh and print	
	• Use a printer that the end user selects, not just the default	
	• Report bursting, i.e., refresh the document with each recipient's user profile	
	• Save in .rtf, .txt or .pdf format	
	• Perform custom tasks with VBA macros	
	• Publish on channels	
	• Begin processing only if a specific condition is met	
Based on a schedule...	• At regular times or intervals, for example every Friday at midnight	
	• At start and stop dates that you set	
And/or when another file is present...	• Process your document only when another file is present	
Then distribute the processed document...	• Via the repository	
	• Over the World Wide Web (or intranets and extranets)	
	• On the server file system	
	• Or on all of the above	
For viewing on client machines...	• With dates and times corresponding to the client machine's time zone	*x*

A CORBA solution

From the administrator's point of view, BROADCAST AGENT is a collection of server components that can be deployed over a distributed CORBA network. The distributed architecture enables scalability, load balancing and failover.

Scalability

CORBA enables server solutions to be deployed on multiple server machines. Thus, as the document scheduling and distributing needs of your organization grow, you can manage the extra workload by adding BROADCAST AGENT machines to your network.

Load balancing

BROADCAST AGENT allows you to distribute your document processing over several different machines, in order to balance the workload. When one Scheduler is unable to process tasks, it can pass tasks to another Scheduler (belonging to the same BROADCAST AGENT) on the same or on another machine.

Failover

Failover ensures that BROADCAST AGENT automatically restarts tasks which fail, for example if the server machine crashes. BROADCAST AGENT provides two types of failover:

- When a task fails, a BROADCAST AGENT component called the Scheduler can be configured to automatically restart the task.
- If the Scheduler itself fails, the Scheduler is restarted automatically.

Managing Broadcast Agent activity and errors

Once BROADCAST AGENT is up and running, BUSINESSOBJECTS and WEBINTELLIGENCE users can begin processing documents. The administrator can trace the processing activity by viewing the list of documents that users have sent in an application called the BROADCAST AGENT Console.

Using the Console, administrators can modify the scheduling and distribution of documents. For example, they can add users to a distribution list. They can also view error messages that occur when tasks fail, as well as interrupt or resubmit tasks.

You can run the Console on any machine that has access to the repository in the CORBA network.

Access and security

BROADCAST AGENT benefits from the access and security provided by the BUSINESSOBJECTS repository, which the supervisor sets up.

The supervisor provides each group of BUSINESSOBJECTS and WEBINTELLIGENCE users with access to BROADCAST AGENT. The supervisor also defines BROADCAST AGENT properties, such as:

- BROADCAST AGENT name
- BROADCAST AGENT password
- The channels on which BUSINESSOBJECTS users can publish documents that BROADCAST AGENT has processed
- The document domain that will store documents before and after processing

The security features that the repository provides are also available for BROADCAST AGENT, namely:

- Pass-through database security

  BROADCAST AGENT can be configured to connect to the database using the user name and password of the user who scheduled the document. This allows the use of only one password, to simplify administration work.

- Multiple repositories

  The Scheduler must be connected to the BROADCAST AGENT through whicha document is being processed, using the correct security information. When a document is scheduled, it therefore asks for the correct security domain, then passes this information to the selected server at runtime.

Broadcast Agent Components

The BROADCAST AGENT system contains a number of independent software components, described in the following table. Each of these components can be configured and monitored from the Business Objects Services Administrator, an easy-to-use graphical tool.

Summary of Broadcast Agent 5.5 Server Components

Component	Purpose
Broadcast Agent Manager **	Starts and controls multiple Schedulers
Schedulers**	At correct times, signal the local CORBA daemon to signal a BOManager or WIGenerator on the first available machine to process scheduled documents
BOManager *	Starts and controls multiple BUSINESSOBJECTS processes (called BusObj.exe under Windows, `bolightsvr` under UNIX), which refresh BUSINESSOBJECTS documents. Batch processing can be enabled or disabled on each BOManager node to allow or prevent BROADCAST AGENT tasks being scheduled on it.
WIGenerator	Starts and controls multiple WIQT processes, which refresh WEBINTELLIGENCE documents. Batch processing can be enabled or disabled on each WIGenerator node to allow or prevent BROADCAST AGENT tasks being scheduled on it.
ORB Activation Daemon (OAD) and OSAgent	CORBA daemons which maintain communication between processes and servers to provide a distributed infrastructure with failover and scalability.

* BOManager under Windows can process documents containing VBA custom macros, and can also use the Publish to Channel feature to push documents to web subscribers via Microsoft Internet Information Server (IIS). These features are not currently available under UNIX.

** These components must be activated exclusively on Windows or on UNIX servers.

Heterogeneous cluster configuration

Each of the components described above can run on either Windows or UNIX servers. However, in a heterogeneous environment, the following limitations apply:

Component	Limitation in a Heterogeneous Cluster
BOManager	The BOManagers in a cluster which are used for interactive processing (i.e., with *Enable Interactive Processing* set to On) must be all on UNIX or all on Windows. When used for batch processing (i.e., used for BROADCAST AGENT, with *Enable Batch Processing* set to On) BOManager can run on both UNIX and Windows systems. BROADCAST AGENT will redirect tasks as needed.
Scheduler	Must all run on the same platform (i.e., all on UNIX or all on Windows).
WIGenerator	Must run on UNIX only in a heterogeneous deployment.

The Business Objects Services Administrator

The Business Objects Services Administrator allows the BROADCAST AGENT administrator to set parameters for processing documents. For example, the administrator can specify how many processes may run concurrently on the server. The Administrator provides settings for the Scheduler and the BOManager, which together are responsible for launching processes.

The Administrator runs both on Windows and in a web browser. In full client deployments, the Windows version can be used while the web version of the Administrator can be used for administering WEBINTELLIGENCE and BROADCAST AGENT. The web version requires an HTTP server.

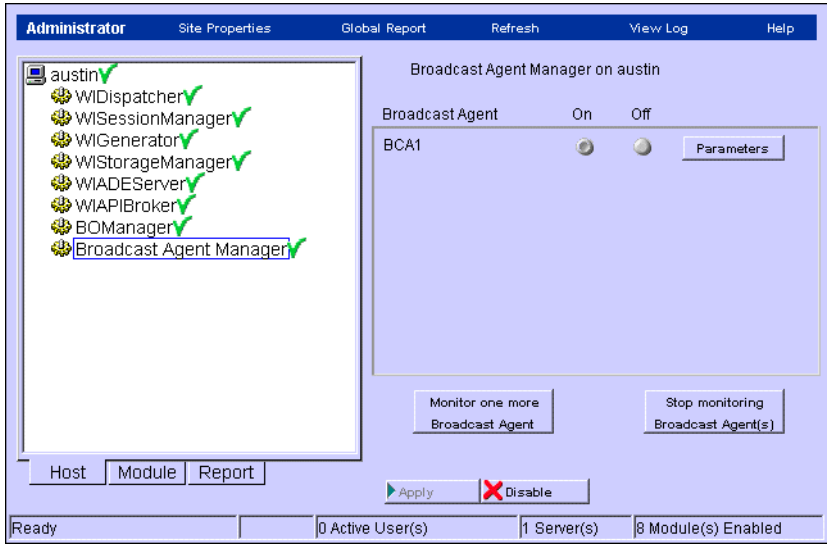

The Scheduler

The Scheduler periodically queries the repository to determine which documents are due for processing. You determine how frequently the Scheduler queries the repository using the Administrator.

When a scheduled task is due, the Scheduler instructs either the BOManager or WIGenerator to process it. Tasks can run concurrently on the server. In the event of one of these tasks failing, the Scheduler automatically starts another BUSINESSOBJECTS or WIQT process to continue processing the interrupted task.

What happens if the Scheduler shuts down or fails?

Using the Administrator, you can configure the Scheduler so that it automatically shuts down at specific intervals, in order to clear the memory on the Scheduler machine. Following a shutdown, the Scheduler is automatically restarted by the BCSKick utility, which is delivered with the Scheduler, and which kick-starts the Scheduler.

Another utility (BCSWatch) constantly monitors the Scheduler while it is running, and restarts it if there is an unexpected failure.

You can also avoid Scheduler failure by deploying additional Schedulers on different machines in your configuration. In this case, if one Scheduler machine fails, a Scheduler on another machine will take over the failed Scheduler machine's jobs.

The BOManager

The BOManager receives instructions from the Scheduler to process documents that are due. The documents are then extracted from the repository's document domain, and the BOManager launches one BUSINESSOBJECTS session on the server per document extracted.

Example **Scheduling a BusinessObjects document for automatic refresh**

For example, if a user has requested that BROADCAST SERVER refresh a document named Sales.rep at 8p.m. on Friday:

- As soon as the Scheduler queries the repository after 8p.m., it retrieves the information that Sales.rep is due for processing and passes this to the BOManager.
- The BOManager launches BUSINESSOBJECTS on the server.
- BUSINESSOBJECTS extracts Sales.rep from the repository.
- BUSINESSOBJECTS on the server runs the Refresh command, saves the updated document, then returns it to the repository.

The WIGenerator

The Scheduler scans the repository for scheduling jobs that are due. When a WEBINTELLIGENCE document is due, the Scheduler requests the WIGenerator to launch an instance of WIQT.exe to process the document.

WIQT fetches the document, which is stored in the repository document domain as a .wqy file, and processes the task (this generally means the document is refreshed, then either published as a corporate document or sent to another user via the repository).

The Broadcast Agent Console

The Broadcast Agent Console is a desktop program that allows administrators to monitor the status of all processed and pending tasks, and to execute actions such as killing a task, changing its distribution list, or running a task immediately.

The Console has direct access to both the Scheduler and the repository. This component comes with the full-client BUSINESSOBJECTS product, and with the 5.1/2.7 release can also be installed on any client machine on the subnet.

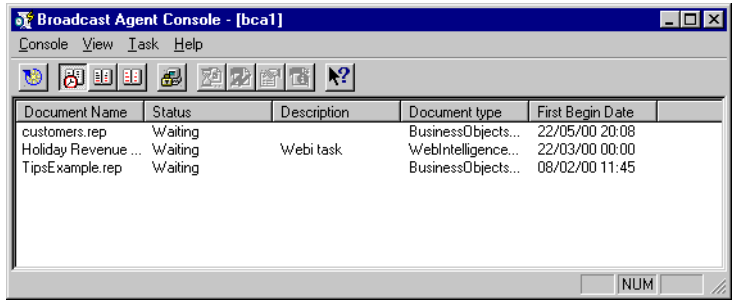

A read-only form of the Console is also available with WEBINTELLIGENCE. Users can therefore use it to monitor the documents they send to BROADCAST AGENT.

Broadcast Agent workflow

Once you have installed BROADCAST AGENT and set parameters using the Administrator, end users can begin submitting documents for processing and distribution. The illustration below shows the workflow.

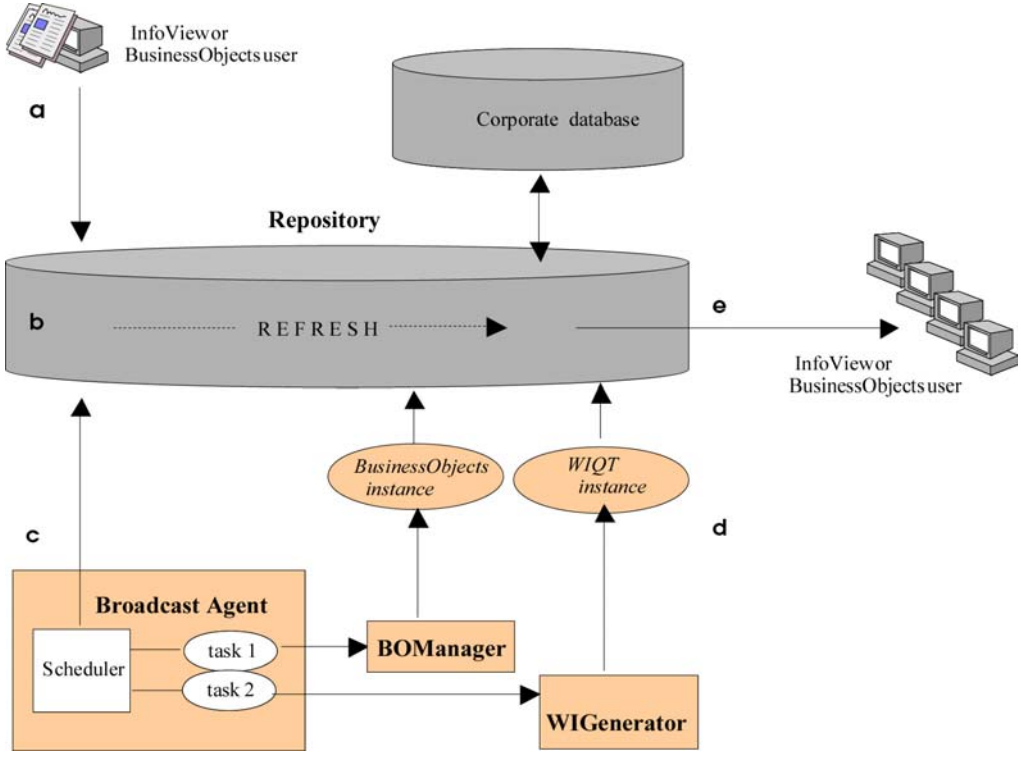

a. Using BusinessObjects or InfoView, end users send documents to Broadcast Agent. They set the actio they want Broadcast Agent to perform, plus options for scheduling, categories and distribution.

b. The documents that users send to Broadcast Agent are stored in the repository's document domain

c. The Scheduler periodically scans the domain and extracts documents that are due for processing.

d. The Scheduler passes due BusinessObjects documents such as Task 1 to the BOManager, which launches one instance of BusinessObjects on the server per document. Similarly, WebIntelligence documents such as Task 2 are passed to the WIGenerator, which launches one instance of WIQT to process each.

e. On the server, BusinessObjects or WIQT performs the actions the user requested on the document. Tl can mean querying the corporate database to refresh the document.

Broadcast Agent Sizing Guidelines

The type of machine that you need to run your BROADCAST AGENT depends very much on the size of the documents that are to be processed: A document that fetches 100 rows of data does not require the same amount of computer resources as another which fetches 100,000 rows of data.

However, the following general guidelines are true in all situations:

- Reasonably fast server processors are needed, such as a Pentium 200. Multiple and uniformly-configured processors can improve performance.
- You need at least 64 MB of RAM, but we recommend 256 MB or more.
- You need to reserve additional disk space to store temporary documents.

Remember also that you do not necessarily need a dedicated machine for your BROADCAST AGENT. You can run it together with a BOManager on the same machine.

Note: BROADCAST AGENT uses both the security and document domains. For better performance, the BROADCAST AGENT should be located as close as possible to the machine on which the repository has been installed.

For more sizing guidelines, see "System Configuration, Sizing, and Benchmarking" on page 155.

Troubleshooting and Optimization

There are a number of things you can do to avoid problems and to optimize the performance of the BROADCAST AGENT. These fall into the following categories:

- Connectivity
- Managing memory
- Using Scan and Repair

Connectivity

Any problems you may experience are often due to connectivity issues, since they tend to occur when exchanging data with the repository. This is the first area you should check if troubleshooting is required.

The BROADCAST AGENT is one of the few products in your configuration that establishes possibly hundreds of connections per day (to both the security and document domains). You therefore need to be very careful about how you define these connections.

When configuring your connections, you can choose from the following options:

- Keep the connection active during the whole session
- Keep the connection active for X minutes
- Disconnect after each transaction

Choose either *Keep the connection active for X minutes* or *Disconnect after each transaction*. Do not choose *Keep the connection active during the whole session.*

The reason for this is that the module called SQLBO handles a pool of connections to the different domains involved. As soon as a connection is closed, it can be physically closed (*Disconnect after each transaction*) or only logically closed (*Keep the connection active during the whole session*).

Imagine the BROADCAST AGENT has been running for 3 days and the connection is defined as *Keep the connection active during the whole session*. This means that the connection is established when the application is starting. If for some reason the connection is physically broken, the logical connection will not be aware of this and will hold the connection open.

At this point, you will get an error such as:

Cannot connect to the security domain

This is because the connection is logically open, although it is physically closed.

You can recover from this situation as follows:

- Perform a safe recovery on the security domain on the machine where BOManager is running and select the Advanced tab in the connection properties.
- Run SUPERVISOR as a General Supervisor and edit the connection property on the document domain (you cannot perform this operation on the security domain).

Note: You cannot modify the security domain connection from SUPERVISOR. The only way to modify the security domain connection is to perform a safe recovery.

Memory management

Not all problems are due directly to the BROADCAST AGENT. Problems are often due to another part of the product that does not affect BUSINESSOBJECTS in normal circumstances, but does so when the BROADCAST AGENT is running.

The BROADCAST AGENT passes all processing tasks to a BOManager or WIGenerator, which then runs one or more instances of the BUSINESSOBJECTS program or WIQT process "behind the scenes." As these processes are progressively launched, stopped, and re-launched, there may be consequences in terms of memory usage, and the availability of free memory at any given time.

How much memory should you allocate to your Broadcast Agent?

It is the machine running the BOManager or WIGenerator which must actually process the scheduled tasks.

This machine should be equipped with 256 MB of RAM or more, plus an additional:

- 16 MB of RAM for each **full-client job** that it is configured to run concurrently (as this depends on the size of the document being processed, however, this figure can go as high as 35-40 MB).
- 12 MB of RAM for each thin-client job configured to run concurrently (the precise RAM can vary between 6-39 MB).

You should also configure sufficient swapping space to allow for peak conditions. You need to bear in mind the following:

- If a job cannot be handled in the available RAM, swapping occurs and processing slows down.
- Virtual memory "swapping" is not performed by the BROADCAST AGENT - it is done by your operating system's memory management, as with all other programs that run on your machine. If swapping occurs and the swapping space is exceeded, reliability then becomes an issue.

Note: For large volumes and/or large numbers of jobs, it is not recommended to install the BROADCAST AGENT on the same machine as the database server.

Using Scan and Repair

Using Scan and Repair from SUPERVISOR, you can scan the repository tables used by the BROADCAST AGENT. This allows you to detect problems, such as a BROADCAST AGENT without any timestamps declared, or with invalid dates.

If you have a problem (not necessarily specific to the BROADCAST AGENT), you should run a Scan and Repair.

If a date is negative, the scan will detect it, and the repair will fix it. The repair operation deletes all invalid timestamps. If this leaves the BROADCAST AGENT with no timestamps defined, then the Repair operation will detect that no timestamp has been defined, and will recreate one.

Remember that the BROADCAST AGENT uses both the security and the document domains. You therefore also need to scan the document domain. If you don't know which document domain is used by the BROADCAST AGENT, run SUPERVISOR, and check the group properties.

Where to get more information

You can also find information on BROADCAST AGENT in the following documentation:

- The *Broadcast Agent Administrator's Guide* provides information on setting up BROADCAST AGENT, upgrading from DOCUMENT AGENT SERVER 4.1, the day-to-day running of BROADCAST AGENT, and troubleshooting.
- The *Installation and Update Guide* explains how to install server and desktop components step-by-step.
- The *WebIntelligence System Administrator's Guide* describes how to use the Business Objects Services Administrator to set Scheduler and BOManager parameters, how to audit Broadcast Agent activity, and more. It also contains a detailed BOManager troubleshooting guide.
- The *InfoView User's Guide* contains several chapters dedicated to automated document and processing from the end user's point of view.

Chapter 11 Windows, UNIX or Both?

In this chapter

- **Overview 360**
- **Why Use UNIX Servers? 360**
- **Version 5.1/2.6 Clusters 361**
  When do you need a Windows node in your cluster? 361
  What administrative functions require a Windows machine? 361
  OLE 2 Objects and UNIX 362
  Configuring version 2.6 heterogeneous clusters 363
- **Migrating Storage from UNIX to Windows servers 365**
- **UNIX Character Set Support 366**
- **Refreshing BusinessObjects documents in Heterogeneous Clusters 367**
  Should you deploy BOManager on Windows or UNIX? 367
  You prefer to refresh BusinessObjects documents on UNIX... 368

Overview

This chapter is designed to answer your questions about Business Objects support for UNIX, a popular operating system that is well suited to running large enterprise applications on very powerful computers. The section at the end of this chapter also provides specific information on configuring heterogeneous clusters, which include both UNIX and Windows machines.

Why Use UNIX Servers?

Organizations need robust, dependable servers -- servers that crash rarely, and that don't have to be rebooted every time new software is installed. But organizations also need powerful, user-friendly software.

WEBINTELLIGENCE for UNIX offers a high-performance combination of power and flexibility that benefits everyone. UNIX servers can be twice as fast as Windows servers if there are no bottlenecks elsewhere in the system, and they can handle substantially greater numbers of users. For this reason, UNIX servers are especially adapted to large-scale Business Objects deployments serving thousands of users.

Servers running WEBINTELLIGENCE on UNIX offer system administrators the advantages of administering robust, reliable UNIX software, while still enabling PC users to make use of the powerful functionality, yet easy-to-use interface that are the hallmarks of WEBINTELLIGENCE software.

UNIX machines running WEBINTELLIGENCE can be configured as part of a cluster that includes Windows machines. This type of cluster is called a **heterogeneous cluster**.

With the 5.1/2.7 release, you can perform almost all tasks on UNIX cluster nodes. There are, however, some restrictions. For information, see page 361.

Note: OLE 2 objects are specific Microsoft objects. Consequently, users may encounter problems when viewing, editing or sharing documents containing OLE 2 objects, using a WebIntelligence Server on UNIX.

Version 5.1/2.7 Clusters

One of the most important features in this release is that BOManager and Broadcast Agent Manager can now run on UNIX servers.

- You no longer need to install BROADCAST AGENT on a machine running Windows in order to schedule both BUSINESSOBJECTS and WEBINTELLIGENCE document refreshes.
- You no longer need to run Broadcast Agent Manager, the module which runs the scheduling process, on a Windows machine. This module can run under Windows or UNIX.

This means that you can run BROADCAST AGENT in a UNIX-only cluster if you want.

When do you need a Windows node in your cluster?

Some tasks or task aspects must still be run on Windows nodes in your cluster. These tasks include:

- All the OLAP sources
- Visual Basic Procedures Data providers
- Personal data files
- Custom macros for BCA task
- Publish to channels

What administrative functions require a Windows machine?

Administrative functions belong to two broad categories:

- Those pertaining to the overall Business Objects solution
- Those pertaining to the distributed architecture running BROADCAST AGENT and/or WEBINTELLIGENCE

Business Objects system administration

Business Objects does not support UNIX for SUPERVISOR or DESIGNER. Therefore, you need a Windows machine to:

- Create and manage the Business Objects repository
- Create and manage users and user groups
- Define database connections
- Define fundamental BROADCAST AGENT criteria, such as the BROADCAST AGENT's user name and password, the groups to which each BROADCAST AGENT is assigned, and the document domain and channels it uses
- Create, manage, import and export universes

The Windows machine you use for these critical functions does *not*, however, need to be part of a cluster in your distributed system. All of these definitions are stored in the repository, where the nodes in your cluster(s) can access them.

Monitoring Broadcast Agent and the WebIntelligence system

BROADCAST AGENT administrators and end users use the Broadcast Agent Console to track and modify the scheduled processing of documents.

If the BROADCAST AGENT they are using is running on an exclusively UNIX cluster, however, they must install the Console as a Desktop Product on their client Windows 95, 98, 2000 or NT4 workstations.

If they wish, they can also install a standalone copy of the Business Objects Services Administrator on their client machines.

For more information, see the *Installation and Update Guide*.

OLE 2 Objects and UNIX

OLE 2 objects are specific Microsoft objects. Consequently, users may encounter problems when viewing, editing or sharing documents containing OLE 2 objects, using a WebIntelligence Server on UNIX.

To insert a logo or other image in a report and share the report via Enterprise Server products installed on UNIX, user's can save the image as a bitmap (.bmp). You can do this in Microsoft Paint or a similar application.

Configuring version 2.7 heterogeneous clusters

Note that:

- The cluster manager must be a UNIX server;
- WISessionManager is not supported on a Windows or UNIX node.
- WIDispatcher WIStorageManager and WIGenerator are not supported on Windows nodes in a heterogeneous cluster.
- Schedulers must all reside on the same platform.
- All WIGenerators must run on UNIX servers.
- Because the Cluster Manager must be a UNIX Server, you cannot use NT Challenge/Response authentication.

The following table provides general configuration details.

Type of document processing	Action	Module or component	No. instances	UNIX cluster mgr	UNIX cluster node	Windows cluster node
	Serve web pages	Web server*	1*	x	x	x
	Control WEBINTELLIGENCE server traffic	WIDispatcher**	1	x	x	Should disable on Windows.
	Save personal documents, categories, manage user sessions	WIStorageManager*** and WISessionManager***	1	x	x	
All documents	Program-mability	WIAPIBroker via ASP (WICOM)	>1			x
		WIAPIBroker via JSP (WIBEAN)	>1	x	x	

Type of document processing	Action	Module or component	No. instances	UNIX cluster mgr	UNIX cluster node	Windows cluster node
Web-Intelligence documents	View, save, refresh	WIGenerator	>1	x	x	Should disable on Windows.
	Schedule with BROADCAST AGENT	Broadcast Agent Manager (Scheduler)	>1	x****	x****	Or x****
Business-Objects documents	View, save, refresh	BOManager	>1	x	x	x
	View, save, refresh using Zero Admin Business-Objects	WIGenerator *and* WIADE Server	>1	x	x	
	Schedule with BROADCAST AGENT	Broadcast Agent Manager (Scheduler)	>1	x****	x****	Or x****
	Program-mability ADE ActiveX (ADE Viewer)	BOManager *and* WIADE Server	>1	x	x	

* You may have more than one web server if the cluster previous had no Microsoft IIS and you have added it to support WEBINTELLIGENCE SDK's ASP programmability features.

** For optimal performance, deploy on the same node as the web server

*** Should not be moved after being initially enabled. WISessionManager must be on the cluster manager, and WIStorageManager is recommended on the cluster manager.

**** Deploy on either UNIX or Windows node(s), but not both.

1: only one instance of this item should be deployed in the cluster

>1: multiple instances of this item may be deployed in the cluster

"x": where this item may be deployed in the cluster

Migrating Storage from UNIX to Windows servers

The WIStorageManager stores users' personal documents, cached documents, and BUSINESSOBJECTS presentations on the server which hosts it.

If you decide to switch some or all of your system's Storage areas from a UNIX server to a Windows server, you'll need to make sure that none of the files in your Storage on the UNIX server have the same names, but with different capitalization (for example, `Document.rep` and `document.rep`). UNIX is case-sensitive, Windows is not. Once the files in Storage on UNIX are moved to Windows, these types of like-named files will be considered duplicates, and one will overwrite the other.

So before moving the Storage contents, verify that there are no duplicate names in UNIX Storage by typing the following from the Storage directory:

```
#!/bin/sh

os=`uname`
wdir=""

if test ! -z "$1" ; then
  wdir=$1
fi

if test ! -d "$wdir" ; then
   wdir=`ckpath -y -p "Which directory to check (.)?" -d "."`
fi

if test ! -d $wdir ; then
 echo "Not a valid directory" ; exit
fi

find $wdir | tr '[:upper:]' '[:lower:]' | sort | uniq -c | awk '
  $2 ~ /:/ {cdir=$2}
  $1 !=1 { print "==> DUPLICATED FILES ==> " cdir " " $2}'
```

This script lets you use the directory you want to check (by default the . directory) as an argument, or prompt for a directory if the specified directory doesn't exist.

This script returns the names of all duplicate files in your storage.

Example **Example results of duplicate file check**

```
<username>-<machinename>:~/bench-utility_WD/bench-utility/
bin>./wchkstorage  /home/users/webi/webi26/

==> DUPLICATED FILES ==> /home/users/webi/webi26/server/
webintelligence2.5/http-server/classes/images: b1icon.gif

==> DUPLICATED FILES ==> /home/users/webi/webi26/server/
webintelligence2.5/http-server/images: partialresults.gif
```

If the check does find duplicate files, you should rename them.

UNIX Character Set Support

Charset support for WEBINTELLIGENCE for UNIX is restricted to Western European languages compliant with ISO-Latin-1 Charset + Euro currency formatting.

The Windows-only environment for servers and BUSINESSOBJECTS and WEBINTELLIGENCE clients, extends support to the Windows 1252 Charset.

Refreshing BusinessObjects documents in Heterogeneous Clusters

How you configure your heterogeneous cluster for the refreshing of BUSINESSOBJECTS documents depends on whether the documents contain features that aren't supported on UNIX, and *how* you want the documents to be refreshed.

Should you deploy BOManager on Windows or UNIX?

In most cases, you can refresh BUSINESSOBJECTS documents using a BOManager on UNIX. If the document contains any of the following, however, you will need a Windows node on your cluster:

- Multidimensional databases (all OLAP sources)
- Visual Basic Procedures data providers
- Personal data files
- Custom macros for Broadcast Agent tasks

Note: If the document is being published to channels through BROADCAST AGENT, i.e. via a batch refresh, the Broadcast Agent Manager will need to be enabled on a Windows machine as well.

The Business Objects Services Administrator now provides two new parameters for BOManager:

- Enable interactive processing - all the BOManagers must be on the same operating system.
- Enable batch processing - the BOManagers can be on both Windows and UNIX.

If you have BUSINESSOBJECTS documents that need to be refreshed interactively, you need to put all the BOManagers on Windows and enable interactive processing.

You prefer to refresh BusinessObjects documents on UNIX...

You might prefer to use UNIX servers when possible as UNIX machines are more powerful, and they can process more BUSINESSOBJECTS processes faster than Windows servers.

However, some of your documents may need to be refreshed on Windows. For example, 99% of your BUSINESSOBJECTS documents use an Oracle database, which can be refreshed on UNIX or Windows but 1% of your documents need to be interactively refreshed on Windows, because they use an OLAP database source or contain macros, and this functionality isn't supported on UNIX.

You have two options:

- Deploy all your BOManagers on Windows and enable interactive processing. Your documents will all be refreshed interactively with no errors, but performance decreases.
- Put all BOManagers on UNIX with interactive processing enabled. Schedule the BUSINESSOBJECTS documents that need to be refreshed on Windows, so that BROADCAST AGENT can process them.

Appendix A The Repository Tables

In this appendix

- **Overview 370**
  New tables 370
  Modified tables 370
- **Broadcast Agent 371**
- **Document Properties 372**
- **Categories 374**
- **Channels 375**
- **Generic table 376**

Overview

This appendix lists the new and modified repository tables in BUSINESSOBJECTS 5.x. These additions and modifications only affect the security domain tables.

For information about other tables in the repository, see the repository.htm file in the Freeware directory on your installation CD.

New tables

The new tables are:

- OBJ_M_DOCAT - Stores information on documents and their associated properties.
- OBJ_M_DOCATVAR - Stores information on data providers.
- OBJ_M_CATEG - Stores information on categories.
- OBJ_M_DOCCATEG - Stores information on documents and their associated categories.
- OBJ_M_CHANNEL - Stores the main channel information.
- OBJ_M_OBJSLICE - Stores additional information on attributes and channels.

Modified tables

The modified tables are:

- DS_PENDING_JOB - New columns have been added to take into account the new distributed architecture and BROADCAST AGENT.
- OBJ_M_MAGICID - Values have also been added to this table for the new channels and categories which need unique IDs. During the update phase, at least two new rows are inserted in this table (one for the channel IDs and one for the categories IDs). These values are automatically added when you create a 5.x repository from scratch.
- OBJ_M_GENPAR - The Repository Identifier: Each repository has a version identifier. This is used by the Repair function, which checks that the repository version is the same as the SUPERVISOR version.

Broadcast Agent

The DS_PENDING_JOB table already existed in 4.1 for handling BROADCAST AGENT tasks. In version 5.x, new columns have been added to take into account the new distributed architecture and BROADCAST AGENT.

DS_PENDING_JOB	Main table for BROADCAST AGENT
VERSION	The version of the submitting application.
IN_NAME_OF	Boolean value which indicates whether or not the task is refreshed with the profile of the recipients.
IP_ADDRESS	The Scheduler machine name.
ERROR_TEXT	Full error information is exported to the repository. These errors can be seen on all consoles.
JOB_DATA	Password for RDBMS security. This can be updated using the *Resolve* button in the Reporter Console.
JOB_PLATFORM	Indicates if the task can be run on a Unix platform.

Additionally, the JOB_SCRIPT column has been extended from 35 to 254 characters.

Document Properties

The following tables are new in version 5.x, and are used to store document properties.

Note: The “Comments”, “Keywords”, “Title” and “Subject” entries are stored in the OBJ_M_OBJSLICE table. This is because these entries can exceed 254 characters.

OBJ_M_DOCAT	Main table for document properties
M_DOC_N_ID	Document ID.
M_DOCAT_N_LPRINT	Date when the document was last printed.
M_DOCAT_N_LAD	Date when the properties were last updated.

OBJ_M_DOCATVAR	Data Provider properties
M_DOC_N_ID	Document ID.
M_DOCATV_C_DPNAME	Name of the Data Provider.
M_SRC_C_NAME	Name of the universe and secured connection used by the Data Provider.
M_DOCATV_N_REF	Date when the Data Provider was last refreshed.
M_DOCATV_N_NBROWS	Number of rows for the current Data Provider.
M_DOCATV_N_REFDUR	Number of seconds taken for the last Data Provider refresh.
M_DOCATV_N_LAD	Date when the Data Provider properties were last updated.
M_DOCATV_N_PARTIAL	Indicates if the current Data Provider returned a Partial result or not.

Categories

A category is a descriptive label that can be used as a search key in the repository. This lets users filter their document searches to see only the documents or available documents for selected categories, no categories or all categories. This is a simple and effective way to search the repository more powerfully, especially when there are many documents.

The following tables are new in version 5.x, and are used to store the categories.

OBJ_M_CATEG	Main category table
M_CATEG_N_ID	Category ID.
M_CATEG_C_NAME	Category name.
M_ACTOR_N_ID	Creator of the category.
M_CATEG_N_LAD	Date when the category was created.
M_CATEG_N_PARENTID	Reserved for future use.

OBJ_M_DOCCATEG	Link between documents and categories
M_CATEG_N_ID	Category ID
M_DOC_N_ID	Document ID

Channels

The following table is new in version 5.x, and is used to store information on channels The OBJ_M_OBJSLICE table is also used.

OBJ_M_CHANNEL	Main channel table
M_CHNL_N_ID	Channel ID.
M_CHNL_C_NAME	Channel name.
M_CHNL_C_ COMMENT	Comment (tooltip) of the channel.
M_ACTOR_N_ID	Creator ID.
M_CHNL_N_LAD	Date when the channel was created or last modified.
M_DASACTOR_N_ID	ID of the BROADCAST AGENT associated with the channel.
M_CHNL_N_CDF_ VERSION	CDF version number.

Generic table

This table is used by the Push, Report Properties, and BROADCAST AGENT tables. It is used to slice information exceeding 254 characters, such as folder names, document comments, and so on.

OBJ_M_OBJSLICE	Used by Push, Report Properties and BCS
M_OBJS_N_ID	Object ID (can be a DocID, a ChannelID, or a BatchID).
M_OBJS_N_TYPE	Type of information stored in the table.
M_OBJS_N_BLOCK_ID	Number of the block, used to sort the slices.
M_OBJS_C_VALUE	A slice of data.

Appendix B Databases and Database Connectivity

In this appendix

- **Overview 378**
- **Database Connectivity Basic Concepts 378**
  - What is a database? 378
  - What is database middleware? 380
  - What is a Business Objects driver? 380
  - How are they all integrated? 381
- **Where To Get More Information 381**

Overview

When you install Business Objects products for the first time, or when you update to the most recent version, you will need to specify how the products will communicate with your database. This is done by setting up a series of Business Objects connections between the client PC or BROADCAST AGENT and the database server.

These connections allow you to access:

- The repository
  The database account(s) created in SUPERVISOR for the secured deployment of Business Objects resources in a client/server architecture.
- The corporate database
  The database accounts which store data for query and analysis.

For a complete list of the connectivities supported for your version of Business Objects products, see the ReadMe.

Database Connectivity Basic Concepts

This section is divided into four subsections that present the working environment for database connectivity:

- What is a database?
- What is a database middleware?
- What is a BUSINESSOBJECTS driver?
- How are they all integrated?

What is a database?

A database is a structure to store and provide data on demand for various users. For example, you could create a database based on your personal address book. This would then allow you to print out selective listings of friends who live in particular cities or countries.

Inside a database, you find:

- Tables
  These are groupings of data. For example, you may have a `CUSTOMER` table in your database.

- Columns
  These are single attributes of a table. For example, the `CUSTOMER` table may contain the following columns: `name`, `address` and `telephone number.`
- Rows
  These are individual records of information. For example, you may want to retrieve a particular customer from the `CUSTOMER` table.

Relational database

A relational database has the following attributes:

- Relationships may be established between its tables based on common information. For example, your database may have a `CUSTOMER` table and an `INVOICE` table. Both may contain a `CUSTOMER NUMBER` column from which you create a relationship between the two tables.
- The data is arranged into tables of rows and columns.
- The database is accessible via Structured Query Language (SQL) - a database access language.

In a client/server system, relational databases are the most common type of database.

Database server

A database server is a program that runs on a server computer which lets you set up, access, and maintain a database. In a client/server environment, database servers are almost always SQL servers, i.e. relational databases.

A database server has the following properties:

- Client programs and users can create and manipulate relational databases.
- A locking mechanism is provided which prevents users from accessing the same data at the same time.
- Performance optimization is provided to process SQL requests in the most efficient manner.
- A security system is provided to control access to data.
- Backup and recovery tools exist to restore data if the system should fail.

Database server programs are commonly called Relational Database Management System (RDBMS).

What is database middleware?

Client machines and database servers must be able to communicate with each other. This functionality is provided by the database middleware.

The role of middleware is to establish and maintain a connection between the client application and the server, and to exchange messages between them. A standard middleware provides both client-to-server and server-to-server communications across a network. When a connection is made, the middleware enables client tools to access, modify, share, and store data on database servers over a network.

What is a Business Objects driver?

A Business Objects driver lets you access your middleware, which in turn gives you access to your RDBMS. To enable the Business Objects driver to access your middleware, you must first configure the middleware and then apply the configuration to the driver.

For more information, see your Business Objects user's guides for the appropriate database.

How are they all integrated?

As explained above, you need the following components to create a connection between Business Objects products and your database server:

- Database server
- Database middleware
- Business Objects driver

The following figure shows how these elements are linked together:

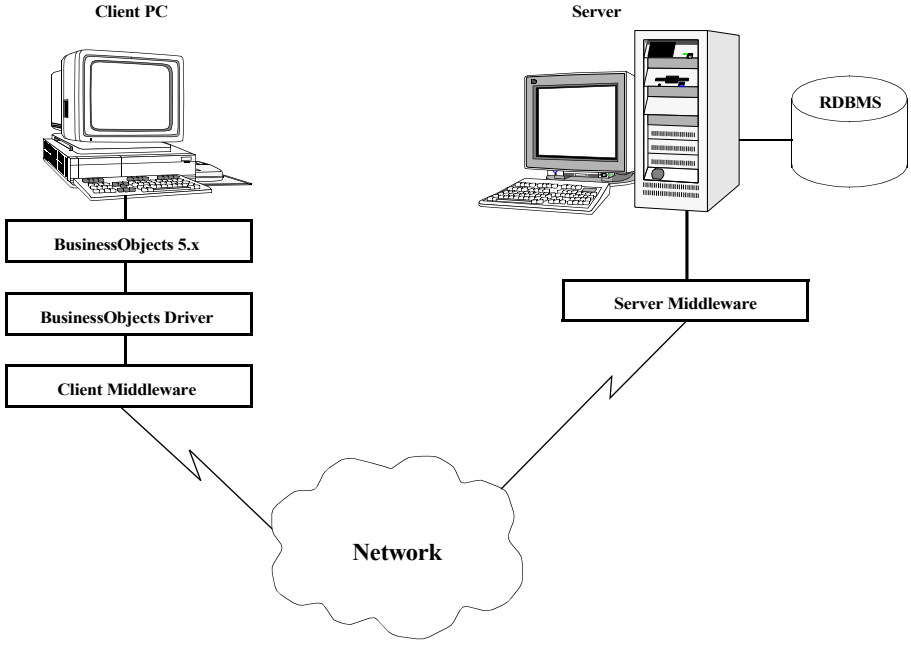

Figure B.1: Connections between BusinessObjects and your database server

Where To Get More Information

For complete information on creating connections between your Business Objects application and your particular RDBMS, refer to your BUSINESSOBJECTS database User's Guide. A list of guides for supported RDBMSs is provided in the Preface.

Appendix C Understanding Secure HTTP Transactions

In this appendix

- Introduction to Secure Transactions 384
- Public-Key Encryption 385
- Client and Server Certification 387
- Making a Secure Connection 389
- Configuring Your Web Server to Use SSL 390

Introduction to Secure Transactions

This appendix provides a brief overview of secure HTTP transactions using Secure Socket Layers (SSL).

The rapid growth of networks and networking technology, both within and between organizations, has greatly increased people's access to information. However, putting information within the reach of those who you want to have it also puts that information within the reach of those who you *don't* want to have it. Open network communication is vulnerable to:

- Electronic wiretapping, where anyone can monitor a data exchange and decipher the conversation, including stealing financial and other data
- Sabotage, which can be as simple as disrupting the data flow or as complex as modifying purchase orders or database queries
- False identification, where a site claims to be a company or entity that it's not or a client claims to be another person

The vulnerability of network data has produced a need to make sure that information transactions between computers are secure. There are a number of ways to do this:

- Isolate the network from all means of access, such as modem lines and outside network connections. This is the most secure way to protect a network, but also prevents the network's users from easily contacting other networks and external resources.
- Place a *firewall* around the network. A firewall limits or prevents outside access to the network, while allowing network users access to networks outside. Some firewalls prevent any access to the network from outside, others require user authentication before allowing access inside, and some even require authentication before allowing users inside to go outside the firewall.
- Secure transactions to and from the local network by encrypting or encoding the data between the sending and receiving computers.

For distributing information over public networks like the Internet, the only viable option is encrypting or encoding information. The other security measures are effective at keeping intruders out of a private network, but have the unfortunate side effect of also keeping out the people you *want* to have access.

SSL is a means of encrypting or encoding data as it passes between a Web server and a browser. Most popular free and commercial Web servers provide SSL encryption. SSL provides three major benefits over most other encryption or data protection schemes:

- Public-key encryption provides a means of exchanging data without transmitting unsecured data or data that would permit snoopers to decrypt the transmissions.
- Authentication certificates permit positive identification of the server and optionally of the client through the exchange of certificates that have been validated by a certificate authority.
- Transaction validation prevents data from being modified between the sender and the receiver. Changes in the data invalidate its authentication code and alerts the receiver that the data has been modified.

Public-Key Encryption

There are two primary methods of encrypting data:

- Symmetric-key cryptography
- Public-key cryptography

In this context, the "key" is a random number, usually quite large and complicated. Encryption engines take the key and the data, use the key to seed the encryption algorithm, and process the data accordingly.

Symmetric-key cryptography is the simpler of the two methods. It uses the same key both to encrypt and decrypt the data. So the sender gives the data and the key to the encryption engine and passes the resulting encrypted data to the receiver. The receiver passes the encrypted data and key to the decryption engine, which runs the process in reverse. The problem with doing this over the network is that, at some point, the key used to encrypt and decrypt the data has to be exchanged. Since the person receiving the key doesn't already have the key, transmission of the key can't be encrypted. Someone else could then intercept the unencrypted key and use it to decrypt (and listen in on or disrupt) any later communications that use that key.

Public-key cryptography solves this problem through the use of two keys, one public and one private:

- The public key is available to anyone who wants to send encrypted messages or data to you.
- The private key stays private and is used to decrypt data encrypted with the public key.

Data encrypted with the public key can *not* be decrypted with it and can only be decrypted through the use of your private key. The process of public-key encryption works like this:

1. To begin a secure session with a server, a browser requests the servers' public key.
2. The server sends its public key to the browser.
3. The browser receives the servers' public key, encrypts a message using the key, and sends the message to the server.
4. The server receives the encrypted message and decrypts it using its private key.

Note that, in Step 2, the communication is not yet secure, just like in symmetric-key cryptography. Someone *could* intercept B's unencrypted public key as it makes its way to A. The difference is that the symmetric key can be used to decrypt the encoded data. But a public key can only be used to *encrypt* data. Someone who intercepted B's public key would only be able to send secure messages to B, *not* decrypt A's subsequent messages to B.

Most encryption schemes, including SSL, are actually a combination of symmetric- and public-key encryption. While more secure, public-key encryption uses a complex encryption algorithm. Symmetric-key encryption is much faster and cheaper. Public-key encryption is used to actually establish the secure connection where a randomly generated symmetric key can be exchanged, with the rest of the transaction encrypted by the quicker symmetric-key encryption.

Client and Server Certification

Another benefit of SSL is positive identification of servers and optionally clients through the use of digital certificates. Digital certificates provide a means of positive identification over the Internet. The certificates used by SSL are defined in the X.509 standard issued by the International Telecommunications Union (ITU). Among the items contained in an X.509-compliant digital certificate are:

- Certificate owner's name
- Certificate issuer's name
- Owner's public key
- Issuer's signature
- Valid dates

If your company has its own internal certificate authority (CA), you should request your certificate from them. To purchase a certificate from a commercial CA, you must submit a Certificate Signing Request (CSR) to the one of the various CAs. With most servers, the SSL configuration tool asks you for the required information and automatically sends a CSR.

A CSR is an encrypted file containing information about your company. The CA will use the CSR to verify that your organization is legitimate and to issue you a certificate. You must then install the certificate on your Web server.

Verisign and Thawte are examples of CAs, for more information, see http://www.verisign.com or http://www.thawte.com.

Users who want to access your secure Web server must have a root certificate from the same authority. A Verisign root certificate is installed with most browsers .

SSL provides three different levels of security based on certificates.

Level	Requirements
No security	No certificate is required and transactions are not encrypted.
Server certification	Server passes a certificate to client. The client checks to make sure that the certificate is still valid and that it was issued by a trusted certificate authority. If the certificate is valid and from a trusted issuer, the client uses the public key contained in the server certificate to begin the encrypted session.
Client certification	Both server and client exchange certificates. Both must check the validity issuing authority of the certificate they receive. If both server and client accept the other's certificates, the session begins just like with server certification. **Note:** You cannot use client certification to manage which users you want to grant access to your site. Although client certificates do contain the user's distinguished name, this is verified but not used to decide whether the user gets access to the site. Access is based solely on whether the certificate is valid and the issuer of the certificate is trusted by the server.

Making a Secure Connection

Creating and using a secure SSL connection is a simple process, especially from the standpoint of the client. The process works much like the process described for beginning a public-key encrypted session, with a couple of extra details:

1. The browser sends an HTTP request to the server. The difference between this request and a standard non-secure request is that it uses the secure HTTP protocol. A secure HTTP URL differs from a standard in that it begins with **https** instead of the normal **http**.
2. The server passes its certificate to the browser.
3. The browser checks the certificate for validity. If there's a problem with the certificate (for example, it has expired or the authority is not trusted), the browser prompts the user to decide whether to accept the certificate anyway.
4. If there is no problem, the browser generates a set of random numbers, encrypts them using the server's public key, and passes them to the server. These numbers are used to encrypt all further communications.
5. If the server requires browser certificates, it then requests one. If the browser can't produce a trusted certificate, the connection is terminated. If it can, the transaction continues normally.

SSL is convenient and easy to use, since the entire transaction is handled by the software. Assuming that the client is attempting a valid transaction, with a valid certificate recognized by the server, the connection hardly appears different from an ordinary one.

Besides the https URL, there is generally an indication in the browser that a secure transaction is taking place. For example, in Netscape Navigator, the key in the lower-left-hand corner is whole, while for non-secure transactions, the key appears "broken."

With https, applet communications are also secure since they use the same path as the HTML pages.

Configuring Your Web Server to Use SSL

To enable SSL for WEBINTELLIGENCE transactions, you must configure your Web server to use SSL and request and install a certificate. For more information on certificates, refer to "Client and Server Certification," on page 387. Then you must enable SSL for each path in your WEBINTELLIGENCE system mapping on your Web server.

For information on how to configure your Web server and apply for a certificate, refer to your Web server administration manual.

You can also find online information about configuring your Web server by searching for SSL or security at:

- http://msdn.microsoft.com for Microsoft Internet Information Server
- http://docs.iplanet.com for Netscape and iPlanet servers.

For Apache you must compile your Web server with an SSL module.

For more information, refer to the *Installation and Update guide for UNIX* or to http://www.modssl.org.

Index

*.lsi files 301
*.ssi files 301
.profile files (UNIX) 237
.rkey file 55
.unv file 166
/etc/system file 238

A

access logs 239
    Netscape 239
accessing
    documents from WebIntelligence InfoView 106
    the repository 280
Act as part of the operating system rights (Windows) 207
Active Server Pages (ASP) 40
ActiveX
    document catalog and viewer control 124
    including in Web pages 32
    reader 108
Admin Tool *see Business Objects Services Administrator*
administering
    the WebIntelligence system 96
administration
    and Windows machines 361
    managing default options for WebIntelligence/InfoView users 59
    WebIntelligence system administrative components 99
Administration files 300
Adobe Acrobat PDF files
    how they are cached 63
    processing in WebIntelligence 60
    saving BusinessObjects documents as 61
Adobe Acrobat PDF format 163
Adobe Acrobat Reader 61
Alerter service 222
Audit directory
    optimizing (UNIX) 241
Audit Kit 56, 259
auditing
    how it is started 102
    role of WISiteLog 103
    WebIntelligence system activity 56
authentication
    in distributed architecture 118
    LDAP 315
automated document scheduling and distribution 344

B

bandwidth 168
    increasing 227
Basic authentication
    and LDAP server authentication 315
batch mode 65
BCSKick utility 351
BCSWatch utility 351
benchmarking 198
BLOB (Binary Large Object) 164
    management 285
boconfig.cfg file 247
bolightsvr processes 93
BOMain.key file 288
BOManager
    and InfoView 110
    and presentation cache 114
    and UNIX 66, 361
    as Broadcast Agent component 348, 351
    assigning number of processes 210
    deploying 181
    Enable Batch Mode parameter 65
    Enable Interactive Mode parameter 65
    enablement in cluster 206
    how many should you enable? 160
    its role in the WebIntelligence system 93
    limitations in heterogeneous cluster 349
    measuring HTML overhead 212
    on what platform should it be enabled? 367
    on Windows or UNIX machines 65
    relieving overloaded BOManager machines 212
    running interactively 269
    setting BusObj.exe automation server identity 268
    setting system heap size for server 266
    tracing 271
    tracing activity 269
    tuning 264
    where to find troubleshooting information 262
BOPASS variable
    and NT authentication driver 310
Broadcast Agent
    access and security 347
    and document domain 166
    and PDF files 61
    and the Business Objects Services Administrator 96
    and UNIX-only clusters 66, 361
    BOManager's role 94
    components 348–352
    connectivity 355
    date and time handling 66
    deploying 181
    failover 346
    failure recovery 119
    how many Schedulers? 181

installed server components 197
load balancing 346
managing activity and errors 346
memory management 356
memory requirements 356
monitoring activity 170
OSAgent port number 120
providing failure recovery 196
RAM required for BusinessObjects processes 182
scan and repair 357
scheduling WebIntelligence documents with prompts 65
security domain selection 347
server sizing 182
sizing guidelines 354
tracing Scheduler activity 271
troubleshooting & optimization 355
updating from Broadcast Server 151
where to find troubleshooting information 262
workflow 353
Broadcast Agent Console 352
and all-UNIX systems 362
as Windows application 59
selecting server platform for processing 65
vs. Business Objects Services Administrator 96
Broadcast Agent Manager 348
and UNIX 101, 361
and WIClusterNode 101
tracing Scheduler activity 271
Broadcast Agent Publisher 68
Broadcast Agents
getting list of 102
Browser Check page 130
browser configuration check 273
Bus Mastering Cards 169
Business Objects
access and security 295—342
Consulting 81
Customer Support 80
desktop products 34—43
ongoing documentation service 262
Online Customer Support web site 262
secure configuration 26
Business Objects desktop products
Business Objects Services Administrator 97
Business Objects enterprise server products 29—33
Business Objects installer 136
Business Objects product line 28
is it server only? 115
UNIX support 361
Business Objects security
and WIAdminBOTools 102
Business Objects Semantic Layer 88
Business Objects Services Administrator 42
accessing administrator's manual in PDF format 70
and the PDF file cache 114
and the WebIntelligence system 96
and WINotify 104
as Windows application 58
enabling modules 140
forms 97
in UNIX-only clusters 66
tracing system activity 56, 259—272
using with Broadcast Agent 349
vs. Broadcast Agent Console 96
BusinessMiner 37
with BusinessObjects 37
BusinessObjects
and Microsoft Office 124
and ReportScript 124
and VBA 124
automatically updating versions 54
BLOB support 166
customizing 39
documents vs WebIntelligence documents 161
drivers 380
enhancing system performance from 218
Enterprise mode 291
extending capabilities 124
extensible architecture 125
Faster Document Viewing Over the Web option 162, 219
importance of user groups 145
installing via web browser 53
logging in in offline mode 219
object model 125
online/offline modes 291
opening WebIntelligence documents in 162
optimizing HTML generation from 215
Save for all users option 291
scripting language 151
security 299
starting automatically 141
using distributed system middleware from full-client 48
using the presentation cache 218
vs. Zero Admin BusinessObjects 50
Workgroup mode 291
zero administration deployment 48—55
BusinessObjects 4.1
BLOB management 285
BusinessObjects 5.0
BLOB management 286
disabling login from previous versions 309
how has repository structure changed? 292
BusinessObjects documents
adding hyperlinks to 61
and system resources 77
and the PDF file cache 114
how many users are processing 161
how they are cached 114
refreshing in heterogeneous clusters 367
saving and distributing as PDF files 61
scheduling from InfoView 57
speeding up processing of 211
viewing from web browser 107
viewing with web server in CGI mode 361
vs. WebIntelligence documents 162
BusinessObjects InfoView 36

BusinessObjects OLAP Connect 62
BusinessObjects processes
and CPU 171
assigning number to BOManager 210
how many? 160
RAM required 182
BusinessObjects Reader 104
BusinessQuery 36
BusObj 348
BusObj processes 93, 264
load balancing 208
BusObj.exe automation server
setting identity 268

C

caches 113
for Adobe Acrobat PDF files 63
IIS cache flushing 233
keeping data in web server cache 235
login 113
optimizing IIS cache 233
PDF file 114
pre-loading with document presentations 114
presentation 114
purging inbox documents 67, 293
using presentation cache 218
categories 109
repository tables 374
channels 345
repository tables 375
character set support
under UNIX 366
Cisco LocalDirector 189
client
automatically updating versions of BusinessObjects 54
in an enterprise server deployment 115
installing BusinessObjects through web browser 53
client-server deployment
advantages 87
and associated products 89
performance factors 158
ClipBook Server service 222
cluster manager 116, 205
deploying redundant cluster manager 186
failover capacities 119
if it fails 118, 196
installing 139
overview 117
providing fault tolerance 184
WISiteLog 103
cluster managers
main and backup 184
cluster nodes 116
and networks 169
enabling modules on 140
installing 139
module enablement 205
overview 118
WIClusterNode 101
clusters 116
and time zones 122
deploying multiple 195
enabling/disabling server components 121
geographically distributed 119
how many machines 196
installing 121, 139
module enablement advice 141
multiple 119
OSAgent port number 120
setting up 121
starting 141
types of servers 157
using multiple 119
version 5.1/2.7 361
when do you need a Windows node? 361
where should you locate Storage? 183
columns 379
communication between processes 348
Computer Browser service 222
configurations
partitioned 206
supported 175
configuring
BusinessObjects services to start automatically 141
heterogeneous clusters 363
connectivity
and Broadcast Agent 355
Consulting 81
Content Index service 224
CORBA
backbone 92
daemons 348
modifying kernel components 56, 247–256
modifying timeout for object implementation 247
core dump
setting maximum size 238
corporate documents
and the PDF file cache 114
CPU
monitoring usage (UNIX) 243
requirements 171
creating
a database reporting environment 144
a repository 144
a universe 145
CSV files 67
Custom installation profile 143
Customer Support 80
customization
Developer Suite 39
document catalog and viewer control 124
integrating applications 124
WebIntelligence SDK 32
customizing
BusinessObjects and Designer 39
kernel components 56, 247
the WebIntelligence client 32
WebIntelligence 32
WebIntelligence InfoView's start page 64
WebIntelligence system's kernel components 56
your e-BI solution 123

D

Data Access modules 33
data analysis
with Set Analyzer 40
data management 76
data mining 37
data sets 41
data warehouses
and the repository 281

database files
  location under UNIX 243
databases
  and WIQT processes 94
  as performance factors 164
  Binary Large Objects (BLOBs) 164
  choosing database platform for repository 284
  connectivity overview 378
  creating a reporting environment 144
  definition 378
  featuring row-level locking 284
  middleware 380
  optimizing performance 217
  relational 379
  row-level locking 164
  server 379
  speed of 164
  that support BLOBs 285
  type being used 164
datawarehouses *see databases*
date and time handling 66
debug log 270
Deferred Procedure Call (DPC) 229
defining
  secured connection to a universe 146
  users and user groups 145
defragmentation utility 220
deploying
  BOManagers 181
  Broadcast Agent 181
  Broadcast Agents 181
  multiple clusters 195
  Schedulers 181
  WebIntelligence 183
  WIStorageManager 183
deployment
  and the repository 276
  definition 24
  disaster recovery 79
  environments 85—125
  geographical distribution 76
  geographically distributed 119
  limitations 195
  performance factors 157—174
  pilot project 80
  planning 73
  pre-deployment checklist 74
  products associated with each type of deployment 88
  recommended server machines 170
  troubleshooting 260—273
  UNIX 359
deployment architectures 87
deployment *see also distributed architecture*
deployment strategy 157
Designer 37
  and UNIX 362
  creating universes 145
  creating universes manually 146
  customizing 39
  enhancing system performance from 217
  importing/exporting universes 146
  object model 125
  using external strategies 217
desktop products 34
  BusinessMiner 37
  BusinessObjects 34
  BusinessQuery 36
  Designer 37
  Developer Suite 39
  installing (Master) 143
  installing service pack updates 153
  Set Analyzer 40
  Supervisor 38
desktop variable 173
Developer Suite 39
  Report Viewer control 94
DHCP Client service 222
digital certificates 387—390
Directory Replicator service 222
Disable login from 4.x release or earlier option 309
disabling
  login from previous releases 309
disabling NT services 222
disaster recovery 79
disk access 173
distributed architecture
  administration on Windows machines 361
  administrative layer components 99
  advantages 88
  and DMZs 319—342
  and e-Business Intelligence 123
  and user activity ratio 159
  associated products 89
  clusters 116
  component distribution 174
  components 174
  enhancing performance from BusinessObjects 218
  enhancing performance from Designer 217
  enhancing performance from Supervisor 216
  hard disk space 171
  illustrated 97
  in double-firewall configuration 249
  installed server components 197
  memory requirements 172
  modifying kernel components 247—256
  modules 204
  on LANs 168
  overview 91—115
  performance factors 157—174
  performance optimization using caches 113
  portal 104
  refreshing full-client documents 162
  server components 92, 113
  supported 175
  tracing activity 170
  troubleshooting 260—273
  typical features 91
  use of its middleware by full-client BusinessObjects 48
  WebIntelligence backbone 92
  when do you need a Windows node in your cluster? 361
  without an HTTP server 97
distributed architecture *see also deployment*

distributed deployment
example implementation 176—180
DMZ 319—342
and HTTP servers 249
definition 249
setting server ports to use 249
Document Agent Server
updating to Broadcast Agent 151
document agnosticism 60, 163
accessing third-party files 106
document domain 278
how much space do you need? 166
reducing size 167, 293
document lists
how they are cached 113
document presentations 114
document properties
repository tables 372
documentation xi
deployment and administration xi
ongoing documentation web site 262
what's new with this release 68
documents 106
adding hyperlinks to 61
choosing server platform for processing 65
distributing 110
exporting 4.1 documents to 5.0 repository 150
full-client 77
generating 147
how they are processed 93
improving computation time 211
products needed for creating, modifying and drilling 112
purging inbox documents from Supervisor 67, 293
script execution disablement 313
thin-client 77
thin-client vs. full-client 161
updating to 5.x format 150
uploading from WebIntelligence InfoView 61
using categories 109
viewing in enhanced document format 108
what WebIntelligence InfoView can users do with 111
domains
and Windows NT security 307
creating multiple 278
document 278
repository 277
security 277, 299
universe 278
using multiple repository domains 167
drivers
for database access 380
DS_PENDING_JOB table 371
Duplex operation 169
duplicate file check 366
Dynamic Host Configuration Protocol (DHCP) 222

E

e-Business Intelligence 27, 122
customizing the e-Business Intelligence solution 123
e-commerce 122
Enable Batch Mode parameter (BOManager) 65
Enable Interactive Mode parameter (BOManager) 65
Enable WinNT authentication option 308
enabling modules on cluster nodes 140
enabling/disabling server components 121
enhanced document format 108
Enhanced Document Format option 216
Enterprise mode 291
enterprise server products
installing 139
environments
for deployment 85—125
errors
managing Broadcast Agent errors 346
Ethernet 168
extending using hubs 168
Event Log service 222
Explorer
BusinessObjects 36
WebIntelligence 31
exporting
universes to the repository 146
exporting documents
version conflicts 150
extensibility 124
external strategies 217
extranets 122
and firewalls 319
customizing WebIntelligence for 32
sharing information over 110

F

failover 348
and Schedulers 346
for cluster manager 119
failure
of cluster manager 118, 196
fast file server 206
Faster Document Viewing Over the Web option 162, 219
fault tolerance
for cluster managers 184
fiber channels 168
files
internal.log 260
mm.log 260
oad.log 261
osagent.log 261
sm.log 260
user.log 260
userdet.log 260
viserr.log 261
vislog.log 261
visout.log 261
WebiServer.log 260, 261
filtering
document lists 109
firewalls
and the Gatekeeper 251
and the HSAL 250
definition 319
setting ports to use 249
fragmentation 220
Freeware directory 153, 259

FTP Publishing Service 224
full client
load balancing 210
full-client documents *see BusinessObjects documents*

G

Gatekeeper 56, 102
and MachineManager 101
in firewall configurations 251
modifying settings 251
VBjavavm parameter 256
VMs 256
Gatekeeper Configuration Manager 252
gatekeeper.properties file 252
GatekeeperCmdLine key 254
generic table 376
Global Heap setting (Windows NT) 174

H

hard disk space 171
hard limit 237
determining 237
heap settings 173
heap size 266
recommended settings 264
heterogeneous clusters 360
configuring 363
deployment and associated products 90
deployment restrictions 363
full-client refresh 367
HSAL 141, 249
in DMZ configurations 250
on cluster nodes 205
HTML
reducing HTML overhead 212
HTML generation
optimizing from BusinessObjects 215
HTML overhead
performance factors 214
HTTP server layer 92
HTTP servers
and DMZs 249
on cluster nodes 205
when you don't need one 97
HTTP servers *see also web servers*
hubs 168
placing RDBMS on 169
hyperlinks
adding to BusinessObjects documents 61

I

IIOP protocol 102
IIS Admin Service 224
importing/exporting
user and group lists 290
inbox documents
purging 67, 293
Inetinfo.exe 230, 231
InfoView
accessing third-party files 106
accessing user manuals in PDF format 70
and the PDF file cache 114
and third-party files 163
BOManager's role 94
Browser Check page 130
BusinessObjects Reader 104
caching PDF files 63
categories 109
Corporate Documents page 106
customizing by programming 32
distributing documents with 110
document agnosticism 60
Enhanced Document Format option 216
enhanced document reader 108
enhancing system performance from 216
how BusinessObjects documents are cached 114
how system processes requests 93
Inbox Documents page 106
installing BusinessObjects via InfoView session 53
interactive refresh of BusinessObjects documents from 57
interface 105
managing default user options 59
My InfoView 64
overview 99
Personal Documents page 106
processing third-party documents 60
scheduling BusinessObjects documents 57
uploading documents 61
using enhanced document viewers with web servers in CGI mode 361
viewing BusinessObjects documents 107
WIQT processes 110
InfoView Quick Tour 69
InfoView *see also WebIntelligence InfoView*
Inprise Smart Finder 263
installation
Custom profile 143
fresh installation or update? 129
installed server components 197
required disk space 171
shared 143
system requirements 130
installation profile 139
installing
BusinessObjects via web browser 53
clusters 121, 139
desktop products 143
enterprise server products 139
service pack updates 153
software 136
Zero Admin BusinessObjects on server 52
interactive mode 65
BOManager 269
Interactive setting (Windows NT) 174
interactive users 65
internal trace 270
internal.log 260
internal.log file 270
Internet browsers *see web browsers*
Internet Engineering Task Force (IETF) 314

Internet Information Server (IIS)
Cache flushing 233
desktop variable 173
Object Cache 232
optimizing cache 233
tuning memory usage 232
tuning queue size 234
intranets 122
customizing WebIntelligence for 32
IPlanet
optimizing web servers 239

J

Java applets
use in distributed architecture 88
JavaServer Pages (JSP) 40

K

kboui50.dll 305
kbwui50.dll 306
kernel
configuring (UNIX) 238
kernel components 56, 247
modifying 56
key files 299
used by WebIntelligence system 288, 289
kglui50.dll 305
kntui50.dll 305
kntuidrv.dll 306

L

LANs 168
trunking LAN connections 227
latency 230
LDAP (Lightweight Directory Access Protocol) 314
License Logging Service service 222
limit command 237
linked universes 217
ListenBacklog key 234
load average (UNIX) 243
load balancing 204
full-client 210
load ratio 209
Node Load Factor 209
round robin 208
thin-client 209
load ratio 209
local.lsi file 49
LocData folder 300
log files
generating debug log 270
internal.log 270
UNIX-specific 271
VisiBroker 272
logging in
and security domain selection 278
in offline mode 219
preventing users from logging in from previous releases 309

M

MachineManager 92
and WIClusterNode 101
Master installation
desktop products 143
MaxFreeTcbs key 228
MaxHashTableSize key 228
MaxPoolThreads key 235
MaxProcs parameter 239
memory
bottlenecks 224
shortage 224
tuning IIS memory usage 232
usage per user 172
memory management
Broadcast Agent 356
MemoryCacheSize key
Messenger service 222
Microsoft Excel
and BusinessQuery 36
Microsoft Internet Explorer
and Windows NT authentication 303
Microsoft Office
and BusinessObjects 124
Microsoft SMTP Service 224
Microsoft SQL Server Services
improved connectivity for 62
middleware 380
checking 263
migration 129
mkfile command 238
mm.log 260
modules 204
enablement in clusters 141
enablement on cluster nodes 140, 205
enabling/disabling 121
setting internal trace 270
multihomed configuration 72
multiple clusters 119
deploying 195
rules 196
sharing Storage between 207
multiple repositories 76, 120, 287
synchronizing data between 288
mutexes 57
My InfoView 64
Mystic River SBM 151

N

Net Logon service 223
Net Mon command 170
Netscape
optimizing web servers 239
plug-in 108
netstat command 170
Network DDE DSDM service 223
Network DDE service 223
network disk mapping
and shared Storage 207
network interface cards 169
using multiple under UNIX 72
networks
as performance factors 168
bandwidth 168
checking configuration 263
Ethernet 168
in web solutions 169
latency 230
monitoring 170
Token Ring 168
NIC card
optimizing usage 229
Node Load Factor 209
Non-interactive setting (Windows NT) 174

NT authentication driver 311
and the WebIntelligence system 310
benefits 302
BOPASS variable 310
Broadcast Agent compatibility 310
frequently asked questions 310—311
offline mode 310
NTFS 220

O

OAD (ORB Activation Daemon) 56, 100, 116, 348
and WIKill 103
what happens when it crashes 100
oad.log 261
oad.log file 272
OadCmdLine key 248
OBJ_M_CATEG table 374
OBJ_M_CHANNEL table 375
OBJ_M_DOCAT table 372
OBJ_M_DOCATVAR table 373
OBJ_M_DOCCATEG table 374
OBJ_M_OBJSLICE table 376
Object Activation Daemon (OAD)
setting server port 249
object model
BusinessObjects 125
Designer 125
new for WebIntelligence 68
Reporter 125
ObjectCacheTTL key 233
offline mode 219
using Zero Admin BusinessObjects 49
OLAP
BusinessObjects OLAP Connect 62
Olap Access Packs 33
OLAP modules
with WebIntelligence SDK 33
OLAP Panel 62
on-demand paging 109
online/offline modes 291
OSAgent
and Broadcast Agent 348
as administrative component 100
as kernel component 56
on cluster manager 116
OSAgent Port number 100
default value 140
what is it? 120
osagent.log 261
osagent.log file 272
OSFind 263

P

paging file 224
bottlenecks 225
size 172
pass-through database security 347
pdac.lsi file 300
pdac.ssi file 300
PDF file cache 114
performance enhancement
enhancing from BusinessObjects 218
from Designer 217
from Supervisor 216
performance factors 157—174
and HTML overhead 214
databases and data warehouses 164
networks 168
number of users 159
server machines 170
types of processed documents 161
Performance Monitor 170, 212
Memory object / Available Bytes 231
Memory object/Available Bytes 224
Network Segment object / %Network Utilization 227
Paging File object 225
Process object 231
Processor object / % Processor Time 225
System object / Total Processor Time 225
TCP object / Connections Established 227
Personal Data Account file (pdac.lsi or pdac.ssi) 300
personal documents
storage space for 183
pilot project 80
platforms
choosing server platform for document processing 65
supported 130
Plug and Play security driver mechanism 223, 306
portal 104
power users 75
presentation cache 114, 218
Print to Table Format 67
printer driver
installing on BOManager machines 211
processes
checking WebIntelligence server processes 272
product line
Business Objects 28
profile
installation 139
user 77
prompts
embedding in web pages 32
scheduling WebIntelligence documents with 65
public key encryption 385—386

Q

queue
tuning IIS queue size 234
Quick Design
description of 145

R

RAID drive 183, 206
RDBMS modules
with WebIntelligence SDK 33
RDBMS servers 379
RDBMSs
placement on switched hub 169
security 311
Refresh according to the profile of each recipient option 67, 167, 293

refreshing
full-client documents in heterogeneous clusters 367
interactive refresh of BusinessObjects documents from InfoView 57
regedit.exe 247
registry
backing up 226
emulation on UNIX 247
GatekeeperCmdLine key 254, 256
ListenBacklog key 234
MaxFreeTcbs key 228
MaxHashTableSize key 228
MaxPoolThreads key 235
MemoryCacheSize key 232, 233
modifying CORBA kernel components 247
OadCmdLine key 248
ObjectCacheTTL key 233
ProcessorAffin key 229
relational databases 379
remote file systems
minimizing (UNIX) 237
remote key file 55
Remote Procedure Call 223
report bursting 345
Report Viewer control 94
Reporter
BusinessObjects 36
object model 125
WebIntelligence 31
reporting environment
creating 144
ReportScript
and VBA 124
converting into VBA 151
repositories
setting up multiple 290
using multiple 120
repository
access 280
and login cache 113
and WIQT processes 94
and your data warehouse 281
Broadcast Agent and 347
choosing database platform 284
compatibility between versions 293
creating 144
database as performance factor 164
definition 26, 276
domains 277
exporting universe to 146
how much space do you need? 164
modified tables 370
multiple 287
new 5.x features 287
new tables 370
setting up 144
space required for 293
structure 292
synchronizing data between multiple repositories 288
updating to 5.x 149
updating to version 5.x 292
using multiple repository domains 167
using Scan and Repair 357
working with or without 291
repository cache 113
repository database
choosing 76, 284
row-level locking 284
Restrict SQL to 'Select' Only option 313
round robin load balancing 208
row-level locking 164
databases that feature 284
rows 379
RqThrottle parameter 239
run queue (UNIX) 243

S

SAP Business Information Warehouse
improved connectivity for 62
Save for all users option 291
scalability 348
Scan and Repair
utility 167, 293, 357
Schedule service (Windows NT) 223
Scheduler
in heterogeneous cluster 349
Schedulers 348, 350
and failover 346
deploying 181
how many? 181
tracing 271
what if it fails? 351
scheduling
BusinessObjects documents from InfoView 57
Faster Document Viewing Over the Web option 162
WebIntelligence documents with prompts 65
scripts
execution disablement for scheduled documents 313
sdac.lsi file 301
sdac.ssi file 301
secure HTTP transactions 383—390
Secure Socket Layers (SSL) 384—390
Secure Socket Layers protocol (SSL) 314
security 295—342
BusinessObjects 299
complexity as performance factor 165
DMZ 319—342
extended security commands 313
firewalls 319
for Broadcast Agent 347
Lightweight Directory Access Protocol (LDAP) 314
RDBMS 311
Secure Socket Layers protocol (SSL) 314
SSL 314
Windows NT authentication 302
security and administration files
default locations 302
security commands 165
at group level 316
security domain 165, 277
how much space do you need? 165
using multiple 165
security domain selection 165, 278
and Broadcast Agent 347
and Zero Admin BusinessObjects 55

Servconf utility 264
for running BOManager interactively 269
tuning your BusinessObjects server 264
Server Manager 139
server products
Broadcast Agent 31
disk access 173
stopping on servers 103
WebIntelligence 29
WebIntelligence SDK 32
Zero Admin BusinessObjects 32
Server service (Windows NT) 223
server sizing 157—174
Broadcast Agent 182
servers
and time zones 122
checking WebIntelligence processes on 272
common NT services 222
components in a distributed system 113
database 379
determining WIQT transaction loads 209
distributed architecture server components 92
enabling modules on 140
enabling multiple WIGenerators on a single machine 57
enabling/disabling components 121
how many? 196
in clusters 157
installation profiles 139
installing Zero Admin BusinessObjects on 52
logging out unused 221
measuring HTML overhead 212
memory usage per user 172
migrating Storage from UNIX to Windows 365
minimum memory requirements 172
multiple disks on UNIX server 171
optimal configuration 170
recommended hardware and CPU 170
relieving overloaded BOManager servers 212
setting number of BusinessObjects processes 210
stopping products on 103
tuning BOManager server configuration 264
tuning NT and web 226—236
tuning UNIX servers 237—243
when do you need a Windows node? 361
why use UNIX servers? 360
Windows NT tuning tips 220—225
service packs
installing 153
services
common Windows NT services 222
NT services to disable 222
Set Analyzer 40
integration with BusinessObjects and WebIntelligence 41
set-based analysis 41
setting up a cluster 121
Shared Data Account file (sdac.lsi or sdac.ssi) 301
shared installation 143
shared storage 207
ShData folder 300
shortcut joins 217
single-server environment 139
SiteManager 102
sizing *see server sizing*
Slice and Dice functionality 112
sm.log 260
soft limit 237
Softbridge Basic Language 151
Solaris *see UNIX*
Spooler service (Windows NT) 223
SQL
standard reporting configuration 25
starting
BusinessObjects automatically 141
clusters 141
Storage
migrating from UNIX to Windows servers 365
where should you locate? 183
Storage directory
optimizing (UNIX) 241
Summary of Broadcast Agent 348
Supervisor 38
and UNIX 362
CSV files 67
defining users and user groups 145
Disable login from 4.x release or earlier option 309
Enable WinNT authentication option 308
enhancing system performance from 216
importing/exporting universes 67, 146
purging inbox documents 67, 293
Restrict SQL to Select' Only option 313
Scan and Repair utility 167
timestamp 313
updating to 5.x 151
Use BusinessObject's Username and Password option 312
using security commands 165
using to set up repository 144
swap command 237
swap file management (UNIX) 238
synchronizing
data between multiple repositories 288
storage 185
system
heap size 266
improving performance 198
system requirements 130

T

tables 378
tasking
foreground vs background 221
TCB hash table size 228
TCB timewait table size 228
TCP/IP
and intranets 122

Temp directory (UNIX) 240
thin-client documents *see WebIntelligence documents*
third-party files
accessing 106
and WIQT processes 94
time zones
and clusters 119, 122
timeout
modifying for CORBA object implementation 247
TMPDIR environment variable 242
Token Ring systems 168
using switching hubs 168
total supported users ratio 160
trace
setting internal 270
tracing
BOManager activity 269
WebIntelligence system activity 259
tracing methods 259—272
troubleshooting your deployment 260—273
checking middleware 263
checking network configuration 263
checking WebIntelligence client software 273
checking WebIntelligence server processes 272
generating debug log 270
running BOManager interactively 269
tuning BOManager server configuration 264
using Servconf 264
where to find detailed information 262
trusted NT domains 307

U

ulimit command 237
UNC paths
and shared Storage 207
universe domain
how much space do you need? 166
universes 37
and size of universe domain 166
choosing connection with which to store 67
creating 145
creating manually 146
defining secured connections 146
domain 278
exporting to the repository 146
how big are they? 163
importing/exporting from Supervisor 67
linked 217
size of average .unv file 166
updating 149
UNIX
.profile files 237
/etc/system file 238
and BOManager 65, 66
and Broadcast Agent 66
and Broadcast Agent Manager 101
and Business Objects products 361
and remote file systems 237
and Supervisor and Designer 362
and the BOManager 368
and the Broadcast Agent Console 362
and the Wiorb process 101
character set support 366
configuring the kernel 238
deploying 359
determining hard limits 237
distribution over multiple disks on single server 171
duplicate file check 366
enabling access to shared Storage 207
GatekeeperCmdLine key 254
hard limit 237
limit command 237
load average 243
location of internal trace files 270
migrating Storage to Windows server 365
mkfile command 238
monitoring CPU use 243
monitoring systems 170
most resource-intensive directories 173
multiple network interface cards 72
netstat command 170
oad 100
optimizing the Audit directory 241
optimizing the Storage directory 241
optimizing the Temp directory 240
optimizing the var/tmp directory 242
optimizing web server performance 239
osagent 100
placement of database files 243
registry emulation 247
run queue 243
server tuning tips 237—243
setting maximum core dump size 238
setting number of file descriptors per process 237
soft limit 237
swap command 237
swap file management 238
thread per process constraints 174
TMPDIR environment variable 242
ulimit command 237
UNIX-specific log files 271
using the bocongif.cfg file 247
vmstat command 170, 237, 243
why use UNIX servers? 360
WICpw 103
WIOrb 100
Unix installation
audit log files 260
internal.log 260
mm.log 260
oad.log 261
osagent.log 261
sm.log 260
user.log 260
userdet.log 260
viserr.log 261
Visibroker log files 261

vislog.log 261
visout.log 261
WebIntelligence log files 260
WebiServer.log 260, 261
UNIX-only clusters
and Broadcast Agent 66
and the Business Objects Services Administrator 66
updating
BusinessObjects versions via web browser 54
desktop products 153
documents from 4.x format 150
from a previous version 129
from previous releases 148–152
from version 2.0 to 2.5/2.7 152
from version 3.1 to 5.x 148
from version 4.1 to 5.x 148
repository 149, 292
server products 143
Supervisor to 5.x 151
to Broadcast Agent 151
universes to 5.0 format 149
uploading documents from InfoView 61
UPS (Uninterruptible Power Supply) utility 223
Use BusinessObject's Username and Password option 312
user activity
peak 78
user activity ratio 159
user groups
defining 145
importance of 145
User Login Information files 301
user profile 77
defining with Supervisor 38
user.log 260
userdet.log 260
users
active 77
coping with too many 119
defining 145
importing user information from LDAP directory 315
LDAP authentication of 315
licensed 77
logged 159
logged active 159
memory usage per user 172
number using the system 159
power 75
registered users and security domain size 165
total supported users 160
user activity ratio 159
validating in distributed architecture 118

V

var/tmp directory (UNIX) 242
VBA (Visual Basic for Applications)
and BusinessObjects 124
installing after BusinessObjects via web browser 54
VBjavavm parameter
for Gatekeeper 256
Virtual Machine (VM)
and the Gatekeeper 256
virtual memory
determining system's total virtual memory 224
monitoring (UNIX) 237
size 172
viserr.log 261
viserr.log file 272
VisiBroker
log files 272
vislog.log 261
vislog.log file 272
visout.log 261
visout.log file 272
Visual Basic for Applications (VBA)
converting from ReportScript 151
vmstat command 170, 237, 243

W

web browsers
automatically updating versions of BusinessObjects 54
checking configuration for InfoView 130
installing BusinessObjects via 53
supported 130
viewing BusinessObjects documents 107
Web Panel
illustration 29
supported web browsers 183
web servers
access logs 239
interaction with LDAP server and WebIntelligence system 318
keeping data in cache and RAM 235
MaxProcs parameter 239
optimizing engine 229
optimizing IPlanet servers 239
optimizing Netscape servers 239
optimizing under UNIX 239
RqThrottle parameter 239
tuning to work optimally with NT servers 226–236
web servers *see also HTTP servers*
WebIntelligence
accessing user manuals in PDF format 70
and Lightweight Directory Access Protocol (LDAP) 314
checking client software 273
creating documents 30
customizing ASP pages 40
customizing by programming 32
deploying 183
documents vs BusinessObjects documents 161
for UNIX 360
Getting Started with 71
installed server components 197
key file used 289
managing default user options 59
monitoring activity 170
overview 29
updating from 2.0 to 2.5 152
updating from 2.0 to 2.5/2.7 152
WebIntelligence administrative components 100
WebIntelligence backbone 92
WebIntelligence documents
and system resources 77
and WIQT processes 94

opening in BusinessObjects 162
scheduling with prompts 65
vs. BusinessObjects documents 162
WebIntelligence Explorer 31
WebIntelligence InfoView 31, 104
accessing documents 106
processing third-party files 163
standard HTML viewer 107
what users can do with documents 111
WebIntelligence InfoView *see also InfoView*
WebIntelligence Manager service 142, 224
WebIntelligence Orb service 142, 224
WebIntelligence Reporter 31
WebIntelligence SDK 32
new object models 68
WebIntelligence system
administrative layer 99
and caches 113
and SSL 314
and the Business Objects Services Administrator 96
and the OSAgent port number 120
choosing server platform for document processing 65
enhancing performance from BusinessObjects 215
enhancing performance from InfoView 216
example deployment 176—180
how auditing is started 102
how it caches PDF files 63
how it interacts with web and LDAP servers 318
how it processes InfoView requests 93
HTTP server component 92
kernel components 56, 247
key file used 288
login mechanism 183
minimum memory requirements 172
modifying kernel components 56
monitoring activity 170
NT authentication driver support 310
optimizing processing of BusinessObjects documents 211
optimizing the Audit directory (UNIX) 241
optimizing the Storage directory (UNIX) 241
optimizing the Temp directory (UNIX) 240
optimizing the var/tmp directory (UNIX) 242
performance factors 158
performance factors and HTML overhead 214
portal 104
processing power 171
reducing HTML overhead 212
relieving overloaded BOManager machines 212
role of the BOManager 93
role of the WIGenerator 94
Secure Socket Layers protocol (SSL) 314
server components 113
server processes 272
setting internal traces 270
SQL engine 163
tracing methods 259—272
tracing system activity 259
where to find troubleshooting information 262
WebiServer.log 260, 261
WebiServer.log file 271
WebiSetup.log file 271
Where clause tables 211
WIADE Server
installing Zero Admin BusinessObjects 52
WIAdminBOTools 102
tracing 271
WIAdmToolStop 103
WIAPI Broker
setting server port 249
WIBean 68
WIClusterManager 102, 117
tracing 271
WIClusterNode 101
tracing 271
WICOM 68
WICpw 103
WIDispatcher
setting server port 249
tracing 271
WIGenerator 93, 94, 348
and WIQT processes 94
enablement in cluster 206
enabling multiple 57
increasing threads per process under UNIX 174
limitations in heterogeneous cluster 349
multiple instances on a single server 57
Node Load Factor 209
tracing 271
WIKill 103
Wimanager 101
Windows
Act as part of the operating system rights 207
Alerter service 222
and Business Objects administration 361
and the BOManager 368
Business Objects Services Administrator and Broadcast Agent Console 57
Business Objects Services Administrator as Windows application 97
character set support 366
Computer Browser service 222
configuring Business Objects services to start automatically 142
Content Index service 224
DHCP Client service 222
Directory Replicator service 222
duplicate file check 366
Event Log service 222
FTP Publishing Service 224
GatekeeperCmdLine key 254
heap settings 173
IIS Admin Service 224
implementing shared Storage 207
License Logging Service service 222

ListenBacklog key 234
location of internal trace files 270
MaxFreeTcbs key 228
MaxHashTableSize key 228
MaxPoolThreads key 235
MemoryCacheSize key
Messenger service 222
Microsoft SMTP Service 224
migrating Storage from UNIX 365
migrating Storage from UNIX to Windows servers 365
modifying the OadCmdLine key 248
Net Logon service 223
Network DDE DSDM service 223
Network DDE service 223
NTFS 220
OAD 100
ObjectCacheTTL key 233
optimizing processing of BusinessObjects documents 211
OSAgent 100
paging file 224
Plug and Play security driver mechanism 223
regedit.exe 247
registry 226
Remote Procedure Call 223
Schedule service 223
Server service 223
setting BusObj.exe automation server identity 268
Spooler service 223
stopping services with Wimanager 103
UPS (Uninterruptible Power Supply) utility 223
WebIntelligence Manager service 224
WebIntelligence Orb service 224
when do you need a Windows node in your cluster? 361
Wimanager 101
WINotify 104
Wiorb 100
Workstation service 223
World Wide Web Publishing Service 223
Windows 2000 130—135
"privatization" 135
Business Objects Enterprise Server products 132
compared with Windows NT 130
installing Business Objects desktop products 131
web servers 133
WebIntelligence SDK 133
Windows NT
authentication security 302
common services 222
foreground/background tasking option 221
Maximize Throughput for Network Application option 231
memory allocation limitations 264
Server Manager 139
server tuning tips 220—225
services to disable 222
trusted domains 307
tuning NT servers to work optimally with web servers 226—236
Windows NT LM Security Support Provider service 223
Windows NT authentication
and Internet Explorer 303
Windows NT LM Security Support Provider service 223
Windows NT registry *see registry*
Windows NT *see also Windows*
Windows *see also Windows NT*
Windows Task Manager
Performance tab 212
WINotify 104
Wiorb 100
and UNIX platforms 101
tracing 271
WIQT processes 94, 206
and Broadcast Agent 348
and CPU 171
and InfoView 110
at login 159
how many? 160
load balancing 208, 209
tracing 271
WISessionManager 117
enablement advice 141
sending transactions to WIGenerators 209
tracing 271
validating users 118
WISiteLog 103
tracing 271
WISTORAGE environment variable 207
WIStorageManager 141
and partitioned configurations 206
and repository cache 113
and time zones 122
deployment advice 183, 206
enablement advice 141
implementing shared storage 207
migrating from UNIX to Windows server 365
on cluster manager 205
synchronizing between 185
tracing 271
where should you locate Storage 183
Workgroup mode 291
working set 230
Workstation service (Windows NT) 223
World Wide Web Publishing Service 223

Z

Zero Admin BusinessObjects 48—55
and multiple security domain selection 55
and web security 55
and WIADE Server 52
automatically updating versions on the client 54
installing client software 53
installing server software 52
monitoring activity 170
offline mode 49
three-tier architecture 55
tracing activity 56
vs. BusinessObjects (full-client) 50