

PeopleSoft®

EnterpriseOne 8.93
Server and Workstation Administration
PeopleBook

May 2003

EnterpriseOne 8.93
Server and Workstation Administration PeopleBook
SKU TT893SWA0504

Copyright© 2003 PeopleSoft, Inc. All rights reserved.

All material contained in this documentation is proprietary and confidential to PeopleSoft, Inc. ("PeopleSoft"), protected by copyright laws and subject to the nondisclosure provisions of the applicable PeopleSoft agreement. No part of this documentation may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, including, but not limited to, electronic, graphic, mechanical, photocopying, recording, or otherwise without the prior written permission of PeopleSoft.

This documentation is subject to change without notice, and PeopleSoft does not warrant that the material contained in this documentation is free of errors. Any errors found in this document should be reported to PeopleSoft in writing.

The copyrighted software that accompanies this document is licensed for use only in strict accordance with the applicable license agreement which should be read carefully as it governs the terms of use of the software and this document, including the disclosure thereof.

PeopleSoft, PeopleTools, PS/nVision, PeopleCode, PeopleBooks, PeopleTalk, and Vantive are registered trademarks, and Pure Internet Architecture, Intelligent Context Manager, and The Real-Time Enterprise are trademarks of PeopleSoft, Inc. All other company and product names may be trademarks of their respective owners. The information contained herein is subject to change without notice.

Open Source Disclosure

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>). Copyright (c) 1999-2000 The Apache Software Foundation. All rights reserved. THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PeopleSoft takes no responsibility for its use or distribution of any open source or shareware software or documentation and disclaims any and all liability or damages resulting from use of said software or documentation.

Table of Contents

About These EnterpriseOne PeopleBooks Preface	1
EnterpriseOne Application Prerequisites	1
Obtaining Documentation Updates	1
Typographical Conventions and Visual Cues.....	2
Typographical Conventions.....	2
Visual Cues	2
Comments and Suggestions.....	3
Overview of Server and Workstation Administration	4
Understanding Roles	5
Custom Solution Consultants and Application Developers.....	6
Application Consultants and Application Project Leaders	6
CNC Consultant and CNC Administrator	6
Hardware, Network, and Third-Party Software Consultants and Administrators	6
Understanding the Server and Workstation Administration Guide.....	6
Server Administration - iSeries	8
Understanding the EnterpriseOne iSeries Architecture and Process Flow for iSeries.....	8
Understanding EnterpriseOne Initialization for iSeries	11
Starting the Enterprise Server for iSeries.....	12
Understanding the iSeries Library Structure for EnterpriseOne	12
Manually Starting the Enterprise Server for iSeries	14
Shutting Down the Enterprise Server for iSeries.....	16
Using iSeries Integrated File System Logging Support	17
Example: Easy Access to Log Files	17
Cleaning Up the Enterprise Server for iSeries	18
Setting Up a Printer for iSeries	19
Printing Multiple Copies to a Remote Printer	20
Administrating Batch Processes for iSeries	20
Reviewing Batch Output Files	21
Encoding the Passwords of Users Who Submit Batch Jobs	22
Running Batch Process Reports from the Server Command Line	22
Running Multiple Instances of EnterpriseOne on the iSeries	23
Identifying an EnterpriseOne Instance on the iSeries.....	24
Identifying an EnterpriseOne Instance for Applying Security on the iSeries	25
Creating an Instance of EnterpriseOne on the iSeries	25
iSeries EnterpriseOne Database Security.....	30

Setting Up iSeries EnterpriseOne Database Security	33
Explanations of iSeries Database Security Parameters	35
Adding Administrators	40
Removing Administrative Authority from User Profiles	41
Sample Results for SETOWAUT	42
Displaying User Profile Information.....	46
Sample Results for Authorization Lists	51
Server Administration for UNIX and Linux	54
EnterpriseOne Directory Structure for UNIX and Linux	54
EnterpriseOne Architecture and Process Flow for UNIX and Linux.....	56
jdenet_n Operation.....	57
jdenet_k Operation	57
Understanding EnterpriseOne Initialization for UNIX and Linux	58
Starting the Enterprise Server for UNIX or Linux	59
RunOneWorld.sh.....	59
Database Settings	62
Shutting Down the Enterprise Server for UNIX or Linux.....	62
EndOneWorld.sh.....	62
Setting Up a Printer for UNIX or Linux	63
Administrating Batch Processes for UNIX or Linux	63
Listing Batch Output Files	64
Running Reports from the Command Line for UNIX or Linux.....	66
Example: Running Reports from the Command Line for UNIX or Linux.....	66
Scheduling Reports from the Command Line for UNIX or Linux	67
Maintaining File Security for UNIX and Linux.....	68
Setting Specification File Security.....	68
Setting Business Function File Security.....	69
Setting Executables Security	69
Setting jde.ini File Security.....	70
Understanding HP-UX and Solaris Kernel Parameter Settings	71
Message Queues	73
Semaphores	74
Shared Memory.....	75
File Descriptors	76
Processes.....	76
Understanding Linux Kernel Parameter Settings	76
IPC Resources	77
File Limits	77
Example: /etc/sysctl.conf.....	78
Understanding AIX Kernel Parameter Settings for EnterpriseOne	78
System Parameters.....	78
Tune Parameters	79
Example: Disk Striping	80
Running Multiple Instances of the EnterpriseOne Enterprise Server.....	81
Server Administration - Windows	84
EnterpriseOne Directory Structure for Windows	84

EnterpriseOne Architecture and Process Flow for Windows	85
EnterpriseOne Initialization for Windows	87
Setting Up a Printer for Windows	89
Understanding Windows Services, Accounts, and Permissions	89
Adding a Printer	90
Determining or Changing Printer Ownership	91
Setting Up User Accounts on an Enterprise Server	91
Changing the Domain	91
Adding a Local Account	92
Adding a User to the Administrators Group	93
Working with Network Services	94
Setting Up the Network Service	95
Starting the Network Service	95
Stopping the Network Services	96
Cleaning Up the Enterprise Server for Windows	96
Uninstalling the Network Service	97
Manually Starting the Enterprise Server for Windows	97
Verifying the EnterpriseOne Installation	98
Administering Batch Processes for Windows	99
Monitoring Batch Processes	99
Reviewing Batch Output Files	99
Running Reports from the Command Line for Windows	100
Example: Running Reports from the Command Line for Windows	101
Scheduling Reports from the Command Line for Windows	101
Example: Scheduling Reports from the Command Line for Windows	102
Understanding JDE.INI Settings for Starting Batch Queues on Windows	102
Using Active Directory	103
SCP Object in Active Directory	104
Additions to the Server JDE.INI file	104
Additions to the Workstation JDE.INI File	105
Maintaining File Security for Windows	105
Specification File Security	106
Business Function File Security	106
EnterpriseOne Executables Security	106
JDE.INI File (Enterprise Server) Security	106
Running Multiple Instances of EnterpriseOne on Windows	107
Generating a Unique Identifier	107
Modifying the Server JDE.INI Files	108
Installing Services for a EnterpriseOne Instance	110
Modifying the Workstation JDE.INI File	110
Uninstalling EnterpriseOne Services	111
Moving or Changing an EnterpriseOne Directory Tree	111
Server Administration Workbench	113
Setting up Security Access for SAW	113
Monitoring EnterpriseOne with SAW on Windows	114

Adding to the JDE.INI File	114
Accessing SAW	114
Adding an EnterpriseOne Server for Querying	115
Checking EnterpriseOne Server Configurations	116
Checking EnterpriseOne Server Connectivity	116
Removing an EnterpriseOne Server from the Query List	117
Configuring EnterpriseOne Server Monitoring Settings	117
Monitoring EnterpriseOne Server Processes	120
Monitoring Enterprise Server Statistics	122
Monitoring EnterpriseOne Web Server Statistics	127
Querying EnterpriseOne Enterprise Servers	134
Configuring SAW and Monitoring EnterpriseOne Servers from the Web	139
Creating an SMC File	140
Modifying Advanced SMC Settings	142
Adding EnterpriseOne Servers	143
Adding Ports for Multiple Servers	145
Adding Events to Multiple Server Ports	145
Adding Statistics to Multiple Server Ports	148
Adding Profiles to Events	149
Customizing Server Ports	149
Running the SAW Agent	152
Performing SMC Actions	153
Working with SMC Views	154
Accessing the Server Monitoring Interface	157
Monitoring Enterprise Servers from the Web	158
Server Summary View	159
All Log Files View	164
Active Log Files View	165
PrintQueue Log Files View	165
Disk Usage View	165
INI Settings View	166
Package Build Files View	167
Environment Variables View	167
Troubleshooting Call-Object Processes	167
Monitoring the JAS Servers from the Web	167
Web Server Monitor Technical Considerations	168
Configuring the Web Server Monitor	168
Monitoring the Web Server System	169
Monitoring Web Server Users	171
Monitoring JDENET Connection Pools	173
Monitoring Business Functions Running on Connected Enterprise Servers	174
Monitoring Database Connection Pools	176
Monitoring Virtual Clients	178
Monitoring Web Server Log Files	181
Troubleshooting with the Web Server Monitor	186
Locations of Key Parameter Values for Web Server Monitor Settings	187
Monitoring System XPIe Servers from the Web	187
Broker Summary View	188
Event Types View	189
Client Groups View	190
Client States View	190
Broker Log View	191
Monitoring EnterpriseOne with SAW on the iSeries	191

Accessing SAW	192
Monitoring Remote Servers	192
Starting, Stopping, and Displaying ERP Server Processes	192
Displaying Server IPC and Disk Space Resources	194
Displaying, Printing, and Deleting Server Log Files	195
Modifying the Server jde.ini File	196
Monitoring EnterpriseOne with SAW on UNIX	196
Accessing SAW	197
Monitoring Remote Servers	197
Starting, Stopping, and Displaying Server Processes	198
Displaying, Printing, and Deleting Server Log Files	198
Displaying Server IPC and Disk Space Resources	200
Working with the Server jde.ini File	202
Monitoring EnterpriseOne with the Knowledge Module	203
Operating System-Specific Considerations	204
Adding OWKM Files to the PATROL Environment	204
Loading OWKM Files	205
Adding a Host	205
Adding an EnterpriseOne Server	206
Configuring OWKM	207
Enterprise Server Components Monitored by OWKM	208
Monitoring Enterprise Servers Using KM Commands	209
Monitoring Enterprise Server Components	210
Monitoring Web Server Components	214
EnterpriseOne on Windows Terminal Server Edition	219
Understanding Windows Terminal Server Edition	219
Incorporating Citrix MetaFrame with TSE	221
TSE Restrictions in Multi-user Mode	222
Network Considerations	223
Performance Considerations	223
Setting Up EnterpriseOne on the Terminal Server	223
Troubleshooting EnterpriseOne on Windows Terminal Server Edition	224
Troubleshooting: UBE Output Security on TSE	224
Troubleshooting: Submit UBE Locally to TSE	225
Troubleshooting: Import/Export with Microsoft Excel	226
Troubleshooting: Specification Files are Locked	226
Troubleshooting: User Cannot Restart EnterpriseOne	227
Troubleshooting: Logging Off Versus Disconnecting	227
Troubleshooting: Shortcuts Do Not Work in E-Mail Messages	228
Troubleshooting: Data Selection and Sequencing Criteria Lost	228
Troubleshooting: Run-Time Error Occurs During Server Connection Test	228
Troubleshooting: EnterpriseOne Development Tools Are Disabled	228
Troubleshooting: Users Experience Problems Accessing EnterpriseOne	228
Troubleshooting: Log Path is Incorrect	229
Troubleshooting: Shortcut Path is Incorrect	229
Troubleshooting: Only One User Can Log in to EnterpriseOne	229
EnterpriseOne on a Cluster	230
Hp-UX Clustering	230

Configuring Oracle Parallel Server (OPS)	231
Setting Up an Oracle Package for MC/ServiceGuard	231
Setting Up an EnterpriseOne Package	232
Maintaining Multiple Instances of EnterpriseOne	235
Troubleshooting: HP-UX Clustering	236
HACMP for AIX Clustering	237
Shared File Considerations	237
How HACMP Works	237
Installation Considerations	238
Creating Group and User Accounts	239
Setting Up EnterpriseOne for HACMP	240
Creating an Application Server	242
Maintaining Multiple Instances of EnterpriseOne	244
Troubleshooting: AIX Clustering	244
Sun Solaris Clustering	245
Requirements for Clustering on Sun	246
Clustering Scripts and How to Modify Them	246
Registering EnterpriseOne with SUNClustering	248
Windows Clustering	249
Setting Up EnterpriseOne on a Windows Cluster	250
Adding EnterpriseOne Network Resources to the Group	251
Starting and Stopping the EnterpriseOne Resources on the Cluster	252
Testing EnterpriseOne Connections on the Cluster	252
Setting Advanced Options for EnterpriseOne Resources	253
Setting Failover and Failback Parameters	253
Backing Up EnterpriseOne Tables	255
Understanding Backup Requirements for Servers	255
Deployment Server	255
Enterprise Server	256
EnterpriseOne Tables and Object Owner IDs	258
Backing Up EnterpriseOne Tables on Servers	264
SQL Server Parameters	269
SnapShot	271
Using SnapShot	271
Running Snapshot from the Command Line	274
Generating Serialized Objects for the EnterpriseOne Web Server	276
Installing eGenerator	277
Setting the Default Storage Parameter (Windows and Unix)	277
Setting Up the Win32 Client to Generate Serialized Objects	277
Configuring the Generation Machine	278
Preparing EnterpriseOne for Serialized Objects	279
Modifying gen.bat	280
Generating Serialized Objects	282
Logging In	283

jas.ini Settings	284
JDBj.ini Settings	284
Running the eGenerator Diagnostic Tool.....	285
Generate to the Web Server	286
Bypass the Web Server and Generate Directly to a Database.....	287
Configuring eGenerator	288
Setting eGenerator Options	288
Generation Modes.....	289
Choosing the HTMLOverrides File.....	290
Choosing Languages	290
Generating All Standard Serialized Objects.....	294
Running the TAM Analyzer	294
Generating All EnterpriseOne Objects.....	295
Executable Files on the Workstation	296
EnterpriseOne Linked Executable Files	296
Standalone Executable Files	299
Troubleshooting the Workstation	306
Understanding Error Messages	306
Troubleshooting: Looking At Error Message Details	306
Troubleshooting: Error Messages Generated By Applications	307
Troubleshooting: Frequent Generic Error Messages.....	307
Troubleshooting: Memory Violations.....	308
Troubleshooting: Form and Grid Add Failures.....	308
Troubleshooting: Communication Failure	308
Troubleshooting the Production Workstation	309
Performing Preliminary Troubleshooting.....	309
Troubleshooting Interactive Application Problems.....	309
Troubleshooting Batch Process Problems.....	310
Troubleshooting: Local Data Availability Problems.....	312
Troubleshooting: Standalone Installation of EnterpriseOne	313
Troubleshooting: Enterprise Server Data Availability Problems	313
Troubleshooting: Printing Problems	313
Troubleshooting: .DLL Problems on a Production Workstation	314
Troubleshooting: Data Source Setup Problems.....	314
Troubleshooting the Development Workstation	315
Troubleshooting: .DLL Problems on a Development Workstation	315
Troubleshooting: Event Rule Problems	316
Troubleshooting: Business Function Problems.....	316
Working with the Workstation Log Files	316

Global Tables	317
Logic Processing Logs	318
Application Development Logs.....	318
Using Log Files to Troubleshoot Strategies	319
Viewing Log Files	319
Setting Up the Workstation jde.log.....	320
Setting Up the Workstation jdedebug.log.....	322
Setting Up the Batch Process Log	325
Troubleshooting: Working with the Compile Error Log	326
Troubleshooting: Working with jdecpy.log	327
Troubleshooting: Working with the sql.log	328
Troubleshooting: Reading sql.log.....	328
Working with the jdeinst.log	330
Troubleshooting the Enterprise Server	331
General Information for Troubleshooting Enterprise Servers	331
Troubleshooting General Problems	344
Troubleshooting the iSeries Enterprise Server.....	348
Troubleshooting iSeries Enterprise Server Installation	349
Troubleshooting: Multiple Release Setup	356
Troubleshooting: JDBNET	357
Troubleshooting: Interprocess Communications.....	357
Troubleshooting: JDE.INI File	358
Troubleshooting the UNIX/Linux Enterprise Server	359
Troubleshooting the UNIX Enterprise Server Installation	360
Troubleshooting the JDE.INI File	360
Troubleshooting: Copying EnterpriseOne to a Server	361
Troubleshooting: Configuring Database Tables.....	361
Troubleshooting: Setting Up a Printer	362
Troubleshooting: E-Mail	362
Troubleshooting: Multiple Release Setup	362
Troubleshooting: Finding the Report Files	362
Troubleshooting: JDBNET Server Not Found	362
Troubleshooting: EnterpriseOne Testing	363
Troubleshooting the Windows Enterprise Server	363
Troubleshooting the Windows Enterprise Server Installation	364
Troubleshooting Web Servers	373
Troubleshooting: IIS and IBM HTTP Web Servers	374
Troubleshooting: JAS	374
Troubleshooting: Serialized Database and Generation Issues.....	374
Troubleshooting: SQL Server Issues	374
Troubleshooting: Problems Using Log Files	375
EnterpriseOne PeopleBooks Glossary	377
Index	411

About These EnterpriseOne PeopleBooks

Preface

EnterpriseOne PeopleBooks provide you with the information that you need to implement and use PeopleSoft EnterpriseOne applications.

This preface discusses:

- EnterpriseOne application prerequisites
- Obtaining documentation updates
- Typographical elements and visual cues
- Comments and suggestions

Note

EnterpriseOne PeopleBooks document only fields that require additional explanation. If a field is not documented with the process or task in which it is used, then either it requires no additional explanation or it is documented with common elements for the section, chapter, PeopleBook, or product line.

EnterpriseOne Application Prerequisites

To benefit fully from the information that is covered in these books, you should have a basic understanding of how to use EnterpriseOne applications.

See the *Foundation Guide*.

You might also want to complete at least one EnterpriseOne introductory training course.

You should be familiar with navigating the system and adding, updating, and deleting information by using EnterpriseOne menus and forms. You should also be comfortable using the World Wide Web and the Microsoft Windows or Windows NT graphical user interface.

These books do not review navigation and other basics. They present the information that you need to use the system and implement your EnterpriseOne applications most effectively.

Obtaining Documentation Updates

You can find updates and additional documentation for this release, as well as previous releases, on the PeopleSoft Customer Connection Website. Through the Documentation section of PeopleSoft Customer Connection, you can download files to add to your PeopleBook Library. You can find a variety of useful and timely materials, including updates to the full PeopleSoft documentation that is delivered on your PeopleBooks CD-ROM.

Note

Before you upgrade, you must check PeopleSoft Customer Connection for updates to the upgrade instructions. PeopleSoft continually posts updates as the upgrade process is refined.

See Also

PeopleSoft Customer Connection Website, <http://www.peoplesoft.com/corp/en/login.jsp>

Typographical Conventions and Visual Cues

This section discusses:

- Typographical conventions
- Visual cues

Typographical Conventions

The following table contains the typographical conventions that are used in EnterpriseOne PeopleBooks:

Typographical Convention or Visual Cue	Description
<i>Italics</i>	Indicates emphasis, topic titles, and titles of PeopleSoft or other book-length publications. Also used in code to indicate variable values.
Key+Key	A plus sign (+) between keys means that you must hold down the first key while you press the second key. For example, Alt+W means hold down the Alt key while you press W.
Monospace font	Indicates a PeopleCode program or other code example.
“ ” (quotation marks)	Indicates an adjective that is used in a way that might not be readily understood without the quotation marks, for example "as of" date, "as if" currency, "from" date, and "thru" date.
Cross-references	EnterpriseOne PeopleBooks provide cross-references either below the heading "See Also" or preceded by the word See. Cross-references lead to other documentation that is pertinent to the immediately preceding documentation.

Visual Cues

EnterpriseOne PeopleBooks contain the following visual cues:

- Notes

- Cautions

Notes

Notes indicate information that you should pay particular attention to as you work with the PeopleSoft system.

Note

Example of a note.

Cautions

Text that is preceded by *Caution* is crucial and includes information that concerns what you must do for the system to function properly.

Caution

Example of a caution.

Comments and Suggestions

Your comments are important to us. We encourage you to tell us what you like, or what you would like to see changed about PeopleBooks and other PeopleSoft reference and training materials. Please send your suggestions to:

PeopleSoft Product Documentation Manager, PeopleSoft Inc., 4460 Hacienda Drive, Pleasanton CA 94588

Or you can send e-mail comments to doc@peoplesoft.com.

While we cannot guarantee an answer to every e-mail message, we will pay careful attention to your comments and suggestions.

Overview of Server and Workstation Administration

Configurable Networking Computing (CNC) specialists, EnterpriseOne system administrators, and network and server administrators use the guides within the Configuration Planning and Setup suite. Throughout these guides, we assume that the initial PeopleSoft EnterpriseOne installation is complete, and that the standard data sources, path codes, and environments are defined. These guides tell you how to make changes or additions to the configuration setup after the initial installation.

The Configuration Planning and Setup suite consists of the following guides:

- *Configurable Network Computing Implementation Guide*
Intended primarily for CNC specialists, this guide contains information about the following topics:
 - Middleware
 - Data sources
 - Path codes and environments
 - Object Configuration Manager
 - Modes of processing
 - A typical EnterpriseOne customer configuration
- *System Administration Guide*
Intended for system administrators, this guide contains information about the following topics:
 - Data replication
 - Printer setup
 - The Work with Servers program
 - User profile setup
 - Security setup
 - Data dictionary administration
 - Vocabulary overrides
 - Transaction processing
 - Media objects and imaging
 - The universal table browser
 - Naming conventions
 - The jde.ini file
- *Package Management Guide*

Intended for EnterpriseOne system administrators and others who manage custom modifications to the environments, this guide contains information about the following topics:

- Package management planning and setup
- Modification rules
- Object management
- Building packages
- Deploying packages
- Server packages
- Multitier deployment
- *Server and Workstation Administration Guide*

Written primarily for EnterpriseOne network administrators, this guide contains information about the following topics:

- Snapshot multiclient installer
- Server administration
- Troubleshooting the workstation
- Troubleshooting the server

Although the guides are organized according to related tasks, a CNC specialist, EnterpriseOne administrator, or network administrator might need to review information in more than one guide. For example, the person responsible for setting up path codes, environments, and data sources (described in the *Configurable Network Computing Implementation Guide*) might also be responsible for building and deploying packages (described in the *Package Management Guide*).

The Configuration Planning and Setup suite is the central location for all configuration-related tasks except the following:

- Initial installation of EnterpriseOne
- EnterpriseOne upgrade and cumulative updates
- Network infrastructure and third-party software setup and maintenance. The applicable software or hardware vendor provides this information. PeopleSoft does not provide documentation but does provide the following documents on the PeopleSoft Web site:
 - *EnterpriseOne and Microsoft System Management Server Technology Considerations*
 - *EnterpriseOne and Tivoli TME 10 Technology Considerations*
 - *EnterpriseOne and BMC Software PATROL Technology Considerations*

Understanding Roles

The PeopleSoft EnterpriseOne implementation methodology defines the following specific roles:

- CNC consultant and CNC administrator

- Custom solution consultant and application developers
- Application consultants and application project leaders
- Hardware, network, and third-party software consultants and administrators

During the implementation, each of these roles is performed by both a consultant and a customer. After implementation, the role of the consultant is diminished. Therefore, customers must ensure that adequate training occurs for the person in each of these roles.

Custom Solution Consultants and Application Developers

PeopleSoft EnterpriseOne custom solution consultants resolve business issues by developing applications. Their primary responsibilities include designing the modifications with upgrades in mind and developing, testing, and introducing the customized software. While the CNC administrator performs the version control functions that build and deploy software, the customer solution consultant must help to develop the internal procedures for the application development cycle for your business.

Application Consultants and Application Project Leaders

After a release is installed, configured, and rolled out, the application consultants will continue in their role as product experts. Although application consultants do not implement the configurations, they must understand how the software handles distributed processing, data replication, environments, and so on, because these application issues influence the configuration decisions. In addition, application consultants must have good troubleshooting skills.

CNC Consultant and CNC Administrator

The CNC consultant and CNC administrator are involved with the installation of PeopleSoft EnterpriseOne and the setup of environments, users, security, distributed processing, and data replication. They are also responsible for the setup of version control and testing of various CNC configurations. The CNC consultant and CNC administrator control the deployment of PeopleSoft EnterpriseOne software throughout the company.

Hardware, Network, and Third-Party Software Consultants and Administrators

Implementing PeopleSoft EnterpriseOne includes many tasks that are outside the scope of PeopleSoft services. Third-party consultants provide these services, as well as supplementing our staff as CNC consultants, network architects, custom modification consultants, and so on.

Understanding the Server and Workstation Administration Guide

Although this guide targets EnterpriseOne network administrators, those with other job functions might find the information useful or essential to their positions as well.

The *Server and Workstation Administration Guide* focuses on how to do the following:

- Use the flexibility of the CNC architecture to optimize the PeopleSoft EnterpriseOne installation for your enterprise
- Extend an initial installation (CRP) to meet practical requirements
- Recognize, address, and solve daily issues that arise in a dynamic enterprise


Server Administration - iSeries

PeopleSoft supports EnterpriseOne enterprise servers on the iSeries platform. The iSeries enterprise server can operate in a logic server or database server environment. You need to perform certain administration procedures on your enterprise server to ensure that EnterpriseOne runs properly.

Understanding the EnterpriseOne iSeries Architecture and Process Flow for iSeries

The following graphic illustrates the actions that the host server processes perform:

Iseries Host Server Processes



All communications between the client and the host server occur using sockets. The communications between JDENET_N (network processes) and JDENET_K (kernel processes) occur with shared memory. JDENET_N and the queue kernel communicate using the Job Control Status Master database table (F986110).

The process flow is as follows:

1. An iSeries JDENET job starts.

- Executing the iSeries STRNET command results in a single JDENET_N job starting and a SENTINEL job starting in a newly started subsystem. The JDENET_N job handles the communications requests. The SENTINEL job monitors the number of JDB_x programs that are available and in use. The JDB_x programs ensure commitment control is correctly scoped. A SENTINEL job should always be running while JDENET is active.
 - Each time a processing request (such as a UBE or business function) is submitted to the server and all previously started NET processes are busy, EnterpriseOne starts another JDENET_N job in the same subsystem until the number of JDENET_N jobs equals the maxNetProcesses field value in the [JDENET] section of the .INI file.
 - Each time a request to the server is submitted, a JDENET_N job starts another JDENET_K job until the number of JDENET_K jobs equals the value in the maxKernelProcesses field in the [JDENET] section of the .INI file.
2. When a user submits a batch application or a server package installation, JDENET_N, as part of the host server, communicates with the client, and the following actions occur:
 - The host server programs become E810SYS/JDENET_N.
 - The client environment is initialized.
 - The client tells the host server (using a socket) to initialize its environment.
 - The host server (JDENET_N) initializes its environment and gets environment and user handles.
 - The host server passes the environment and user handles to the client (using a socket).
 - The client launches the batch application or deploys the server package installation. It then sends data to the host server (using a socket).
 - If the maximum number of kernel (JDENET_K, where K represents kernel) jobs is not met, JDENET_N can start a new JDENET_K job.
 - If the maximum number of JDENET_K jobs is met, JDENET_N sends a message to a queue for each waiting JDENET_K job.
 - The client frees the user environment.
 - The client tells the host server (using a socket) to free the server's user environment.
 - The host server frees its user environment.
 3. JDENET_K performs the requested action.
 4. Messages from JDENET_K are passed back to the request client workstation or server via JDENET_N.
 5. If a request to run a batch process is submitted, the following events occur:
 - JDENET_K adds a record to the F986110 database table with a status of W for waiting.
 - JDENET_K submits a job to the queue
 - If you are using native iSeries job queues, JDENET_K submits a job to the iSeries queue. This job calls the EnterpriseOne program PRINTUBE on the iSeries enterprise server.
 - If you are using the EnterpriseOne queue kernel, JDENET_K sends a message to the queue kernel, alerting it that a new job request was submitted. When the job is ready, the queue kernel executes the PRINTUBE program.

- The PRINTUBE process runs the batch application.
- PRINTUBE changes the status of the record in the F986110 table to P for processing.
- If the batch application runs successfully, the software changes the status of the record in the F986110 table to D for done.
- If the batch application fails, EnterpriseOne changes the status of the record in the F986110 table to E for error.

Understanding EnterpriseOne Initialization for iSeries

The following initialization occurs when you start PeopleSoft EnterpriseOne programs such as PRINTUBE:

- The EnterpriseOne environment name is passed as an argument to the program.
- This environment might be translated to a different environment, based on the settings in the [SERVER ENVIRONMENT MAP] section of the .INI file.
- The software verifies that the environment is a valid entry in the Library ListMaster File table (F0094) and that it has a valid corresponding path code in the Environment Detail - One World table (F00941).
- The Library .INI file setting in the [DB SYSTEM SETTINGS] section indicates where the EnterpriseOne server startup tables, such as Data Source Master (F98611), Object Configuration Master (F986101), and so on, are located.
- Using the above information, the software opens the F986101 (OCM) table in the specified database on the server.
- If an override for a given table, BSFN, and so on, or the current user exists, that data source (the OMDATP field in the F986101 table) is used for the given object or user and environment. Otherwise, the data source in which OMOBNM=DEFAULT for the given environment is used. Ignore any inactive records (that is, OMSTSO=NA).

Warning

PeopleSoft *highly* recommends that you do not have any default (OMOBNM=DEFAULT) records for reports (OMFUNO=UBE) or for BSFNs that are mapped to the server. These records might prevent report interconnections (one report calling another report) from starting correctly.

Each unique data source in the F986101 table should correspond to one entry in the F98611 table. The corresponding information in the F98611 table must be correct. In particular, the OMDLLNAME field must show the correct SRVPGM (.DLL) for the database to which the data source points:

- DBDR for files located on the iSeries enterprise server
- JDBNET for files not located on the iSeries enterprise server

Starting the Enterprise Server for iSeries

You can start the EnterpriseOne enterprise server for the iSeries either manually or automatically.

Prerequisites

- ❑ Install PeopleSoft EnterpriseOne as described in the *EnterpriseOne Installation Guide*. In that guide, you should have performed all steps up to the Installation Workbench.
- ❑ Run the `clear CLRIPC` command before you start the server to ensure that the server is clear. If you do not run this command prior to starting a server, the startup process will fail.

Understanding the iSeries Library Structure for EnterpriseOne

You can set up an initial program to create the library list. Also, you should add the following library to the top of your library list before you start EnterpriseOne on the enterprise server: *releaseSYS* (or your system library name). The variable *release* is the EnterpriseOne release level, such as E810SYS.

The *releaseSYS* library contains the following:

INI	The JDE.INI file used to initialize EnterpriseOne on the iSeries enterprise server.
*PGM and *SRVPGM	The various programs and service programs required to run the EnterpriseOne iSeries enterprise server.
*MODULE	The objects used to create EnterpriseOne database programs that the software uses to interface with DB2/400.
H (*FILE)	The file containing the EnterpriseOne header file members required to build new objects on the iSeries enterprise server.
JDE_HP4PS (*WSCST)	The workstation customization file used to transform the EBCDIC PostScript file to ASCII for autodetecting PostScript printers.
CHGLIBOWN (*CMD)	An EnterpriseOne utility command used to change ownership of all objects contained in a library.
SHOW (*CMD)	An EnterpriseOne utility command used to show runtime output.
UPDLF (*CMD)	An EnterpriseOne utility command used to modify the maintenance attribute of logical files.
DPSPSTMF (*CMD)	The display stream file, which displays iSeries Integrated File System (IFS) text stream files. The EnterpriseOne log files, JDE.LOG and JDEDEBUG.LOG, typically reside in a directory called <i>PSFTrelease</i> , where <i>release</i> represents the EnterpriseOne release, such as /PSFT810.
LINKBSFN (*CMD)	The command used to relink business functions to their respective service programs (*SRVPGM). Typically, the system uses this command during an upgrade of the EnterpriseOne system library.
PID2JOB	The Convert Process ID to Job command, which returns the job information when the system

(*CMD)	passes a process ID to the command. The system writes the process ID in the JDE.LOG files. This command returns job information only while the job is still active.
PORTTEST (*CMD)	The command that runs the EnterpriseOne test program PORTTEST.
RUNUBE (*CMD)	The command that interactively runs a batch program. If you need to run a batch program, use the SBMJOB command to submit the RUNUBE command to batch.
SAW (*CMD)	The command that starts the Server Administration Workbench.
QTXTSRV (*FILE)	The file that contains the source for the workstation customization object for autodetecting PostScript printers.
PRINTQUEUE (*FILE)	The file that contains the output from a batch program. This output is stored as ASCII PDF members.
*PGM and *SRVPGM	The programs and server programs required to run the EnterpriseOne network.
JDENET (*JOBQ)	The job queue used by the EnterpriseOne iSeries network jobs.
NETJOB (*JOB)	The job description used by EnterpriseOne iSeries network jobs.
DBMONJOB (*JOB)	The job description used by the database to monitor jobs.
JDENET (*CLS)	The class used to create the routing entry for the JDENET subsystem.
ENDNET (*CMD)	The command that ends the EnterpriseOne iSeries network jobs and cleans up the network runtime structures.
IPCS (*CMD)	The utility command that indicates the status of objects used by the EnterpriseOne iSeries network jobs and as a backup method for cleaning up the IPCS objects.
STRNET (*CMD)	The command that starts the EnterpriseOne iSeries network jobs.
CLRIPC (*CMD)	The command used to clear IPC structures.
DSPIPC (*CMD)	The command used to display IPC structures.
PSFT<i>release</i> (*SBSD)	The subsystem description under which the EnterpriseOne network jobs run. The variable <i>release</i> is the EnterpriseOne release level, such as PSFT810.

Manually Starting the Enterprise Server for iSeries

You manually start the enterprise server for iSeries by starting JDENet from the command line, and then starting the PORTTEST program, which verifies that the enterprise sever software was installed correctly. If it was, PORTTEST initializes an environment and user.

► To manually start the enterprise server for iSeries

1. Sign on to the iSeries as ONEWORLD.
2. Start JDENet using the following command:
`STRNET`
3. Start the PORTTEST program using the following command to verify that the basic enterprise server software was correctly installed:
`PORTTEST userID password environment`

Where *userID* represents the EnterpriseOne iSeries user ID, *password* represents the password, and *environment* represents the environment that you want to test.

The PORTTEST program initializes an environment and user if EnterpriseOne was correctly installed and configured. The program opens a table and displays up to 99 rows of data. You should see results similar to those shown in the following example:

```
Running porttest for JDESVR on M9ASD2 with password JDESVR
Initializing Environment M9ASD2,...
Environment M9ASD2 was initialized successfully.
Initializing JDESVR/JDESVR (User/Password),...
JDESVR/JDESVR (User/Password) Initialized successfully.
Opening table F986110,...
Opened table F986110 successfully.
Closing table F986110,...
Closed table F986110 successfully.
Opening table F0902,...
Opened table F0902 successfully.
Performing select all on table F0902,...
Select all on F0902 succeeded.
Printing up to 99 records in the table F0902,...
f0902.gbaid f0902.gbawtd
-----
[98] 00009697 24060973
[97] 00009806 13540877
[96] 00010102 3140380
...
[1] 00068798 10000
```

```

[0] 00058798 250000
Total number of rows printed = 99
Calling DataDictionary Validation function,...
Data Dictionary Validation Succeed for "CO 00001".
Closing table F0902,...
Closed table F0902.
Freeing user JDESVR,...
Freed user JDESVR successfully.
Cleaning up environment M9ASD2,...
Cleaned up environment M9ASD2 successfully.
Congratulations! Porttest completed successfully.
All Done!
BYE!

```

If the table in the environment that you specified is empty, the total number of records that the program prints will equal zero.

4. Enter the following command:

```
WRKACTJOB SBS(PSFTrelease)
```

The variable *release* is the EnterpriseOne release level that your site is using, such as PSFT810.

5. Verify that the entry NETWORK with function PGM-JDENET_N and status of SELW is running (until a net request is performed, the CPU will be 0).

The STRNET command also starts a second job called SENTINEL. SENTINEL monitors the number of JDB_*x* programs that exist and are available, where *x* is an integer. The JDB_*x* programs connect to the iSeries database (DB2) and monitor commitment control. Information regarding the number of these programs that is available and in use resides in a user space called JDE_PGMCTL (*USRSPC). When the software starts JDENET for the first time, the system creates these programs based on .INI file settings.

► To automatically start the enterprise server for iSeries

1. Create a CL program.

You will use this program to establish the appropriate EnterpriseOne library list and execute the command to start the iSeries server job (JDENet).

```

The CL program should be similar to:
PGM
CHGLIBL LIBL(E*)E810SYS QTEMP QGPL)
STRNET
ENDPGM

```

2. Identify and modify the program called during the iSeries IPL to submit a job to call the previous program.

The program name and location are set in the iSeries system value, QSTRUPPGM.

3. Determine the QSTRUPPGM value by entering the following command:
`DSPSYSVAL SYSVAL(QSTRUPPGM)`
4. Determine where the source of the program is located by executing the following command against the library and program (as set in the system value):
`DSPPGM LIBRARY/PROGRAM NAME`
5. Modify the source of the startup library and program by inserting a SBMJOB command that calls the program created in Step 1.
6. Verify that the startup program is created correctly by recreating it and ensuring that it is created in the library specified by the system value.

Use CRTCLPGRM and prompt (using F4) for the appropriate parameters.

PeopleSoft recommends that you separate the EnterpriseOne add library list entry (ADDLIBL) and startup (STRNET) commands from the iSeries startup program. You should create a separate EnterpriseOne startup program and call that program from the iSeries startup program. This action ensures that commands subsequent to the EnterpriseOne add library list entry and startup are not associated with the modified library list. Following this recommendation also ensures that your EnterpriseOne library list is set correctly before issuing the STRNET command. In addition, the separately-called program provides you with a single location in which to locate and maintain EnterpriseOne startup commands on the iSeries.

Shutting Down the Enterprise Server for iSeries

You can manually shut down the enterprise server for the iSeries.

Prerequisite

- ☐ Ensure that your library is set correctly before performing this command.

Note

For more information, see *Understanding the iSeries Library Structure for EnterpriseOne* in the *Server and Workstation Administration Guide*.

► To shut down the enterprise server for iSeries

From an iSeries console, end JDENet using the following command:

`ENDNET`

This command is in the system library. For example, E810SYS causes EnterpriseOne to end the JDENet jobs and clean up all JDENet runtime structures.

Using iSeries Integrated File System Logging Support

To achieve better performance and to allow easier access to log files from the workstation, EnterpriseOne generates log files for the iSeries in the Integrated File System (IFS) rather than the traditional file system on the iSeries.

With IFS, EnterpriseOne generates log files as stream files (STMF) in an IFS directory, based on the iSeries jde.ini file settings.

Example: Easy Access to Log Files

The following examples illustrate how to modify the JDE.INI file to allow easier access to log files from the workstation.

<pre>[DEBUG] DebugFile=jddebug JobFile=jde.log</pre>

EnterpriseOne generates log files in the IFS root directory.

<pre>[DEBUG] DebugFile=/psft810_a/jddebug JobFile=/psft810_a/jde.log EnterpriseOne generates log files in the IFS directory called "/psft810_a."</pre>

Note

The directory must exist with proper authority granted to the logging job.

► To view log files on the iSeries

From an iSeries console, enter either of the following commands:

- WRKLNK, to show directory and stream files in the IFS. Page up and down to locate the log files.
- DSPSTMF (shipped with EnterpriseOne in the SYSTEM library), to view the IFS stream file.

For example, enter the command DSPSTMF
STMF(`/psft810_a/jddebug_248135.log').

Cleaning Up the Enterprise Server for iSeries

When PeopleSoft EnterpriseOne ends abnormally, you might need to manually perform cleanup tasks on the iSeries enterprise server. Interprocess Communication (IPC) structures might not be cleaned up following an execution of ENDNET, which might cause further problems when you start JDENet. If the IPC structures are not properly removed by ENDNET, you can manually remove them. IPC structures might become locked by an interactive job. For example, you might have to sign off and sign back on in order to perform a successful cleanup.

The EnterpriseOne iSeries server is shipped with the DSPIPC and CLRIPC commands, which allow you to display the IPC-related information and to remove IPC structures.

If tracing is turned on in addition to IPC, you should clear the jde.log and jddebug files. This action keeps the files from becoming too large and removes old messages from it.

Caution

Clear IPC structures only when you are ready to restart the JDENet process.

Prerequisite

- ❑ Ensure that your library list is correct before executing the IPC commands. Each of the commands calls the IPCS command for all of the IPC types. Each command has two parameters: *from* and *to*. Use these parameters to specify the starting and ending IPC addresses on which you want to operate. The default value for the *from* parameter is *INI; this is the address specified in the .INI file. The default value for the *to* parameter is *CALC; this means that the value is calculated based on the value of the *from* parameter. For example, you could specify 999 more than the *from* parameter.

Note

IBM Opti-Connect and Opti-Mover products use the IPC shared memory address 9999. Avoid setting the jde.ini file setting IPCStartKey to a starting value that uses the range of 9000 to 9999.

► To clean up the Enterprise server for iSeries

From an iSeries command line, enter the following IPCS commands:

- DSPIPC
- CLRIPC

► To clear the jde.log and jddebug files for iSeries

1. To clear the jde.log stream files, enter the following command:

```
DEL `/PSFTrelease/jde_*
```

Where *release* is the EnterpriseOne release, such as psft810.

2. To clear the Jdedebug log, enter the following command:

```
DEL `/PSFTrelease/jdedebug_*
```

Where *release* is the EnterpriseOne release, such as psft810.

Setting Up a Printer for iSeries

For printing, EnterpriseOne iSeries servers generate PostScript, PCL, or line printer reports. The line printer OUTQ configuration is similar to most typical iSeries OUTQ configurations. This section provides the steps necessary to set up the Postscript and PCL OUTQ configurations.

Unless otherwise specified in the printer definition, the default OUTQ used for printing batch process reports is the same as the default OUTQ of the user submitting the job.

See Also

- ❑ *Working with the Printers Application* in the *System Administration Guide* for more information about setting up printers

► To create the OUTQ

Enter the following command:

```
CRTOUTQ OUTQ(QGPL/outqname) RMTSYS(*INTNETADR) RMTprtQ(` `)  
CNNTYPE(*IP) DESTTYPE(*OTHER) TRANSFORM(*NO) INTNETADR(`IP Address of  
your printer')
```

Note

Some printers require that you set the parameter RMTprtQ to something other than ` `. See the instruction manual for your printer for additional information. For example, you must set this parameter to "PASS" for the IBM Network Printer 4317.

► To start the OUTQ

1. Enter the following command:

```
STRMTWTR outqname
```

For example:

```
STRMTWTR QGPL/JDE_HP4PSB
```

2. If you need to release the outqueue before using it, enter the following command:

```
RLSOUTQ outqname
```

For example, enter `DEL ' /PSFTrelease`, where *release* is the EnterpriseOne release, as in PSFT810.

Printing Multiple Copies to a Remote Printer

This task is necessary only when the output queue does not support printing multiple copies, and it applies to remote output queues only. Only system administrators can print multiple copies to a remote printer.

► To print multiple copies to a remote printer

1. End the remote writer to which the output queue is connected.
2. Use the Change Output Queue (CHGOUTQ) command to change the Display Options (DSPOPT) parameter so that it contains the value XAIX.
3. Restart the remote writer.

Your output queue should now be able to send multiple copies of your documents to the remote printer.

Administering Batch Processes for iSeries

Administering batch processes involves knowing what processes run when EnterpriseOne starts, where files are placed before and after printing, and how to watch those processes.

Depending on how the software is installed, jobs run under several subsystems on the iSeries. The first subsystem, PSFT810, is created during the installation process and is responsible for running the EnterpriseOne net and kernel processes. QBATCH is the default subsystem in which jobs run, but you can use other subsystems to distribute the workload.

► To monitor batch processes

1. Sign on to the iSeries enterprise server using an administrative account.
2. Enter the following command, substituting *Subsystem* with the appropriate subsystem name:
`WRKACTJOB SBS(Subsystem)`

When you send a batch process report to an iSeries server for processing, the network jobs are responsible for accepting and queuing the request, while the QBATCH subsystem is responsible for executing the report. To monitor the batch requests, use the WRKACTJOB command, specifying QBATCH as the subsystem.

A job appears indented underneath the subsystem. A job such as the R0006P job is the actual report that is running at this time. The program PRINTUBE is the job that is responsible for running and printing the request. When the job is finished, it leaves the queue, and the print job is either printed and deleted, or saved in the E810SYS/PRINTQUEUE file.

Reviewing Batch Output Files

When users submit batch reports to run on the iSeries, a corresponding output file is created on the enterprise server. When the report completes, it can either appear on the user's workstation or print on the server.

The default location for Portable Document Format (PDF) files is under the system directory in IFS. This directory is normally E810SYS, but the installation process can override the default setting.

The system directory in IFS contains a PRINTQUEUE file. This file is the location for PDF report requests. Members are added to this file each time a PDF report is completed.

► To review batch output files

1. From Windows Explorer, use the following command to map a drive to the root directory of IFS on the iSeries machine:

```
//machinename/root
```

2. Navigate to the PrintQueue folder in the System directory (for example, the directory might be called /E810SYS/PrintQueue), and view the PDF files.

The naming convention for each member is based on the EnterpriseOne job number, which is a unique number that the system assigns when the report is submitted. This number is a unique print request ID and is incremented each time a report is submitted to the enterprise server, regardless of whether the job is successful or fails. It is not related to the process ID or job number that the iSeries assigns to the batch job.

If you submit a batch process report to a specific server, the OUTQ for printing depends on the jde.ini file settings for the workstation. You must make the following changes to the workstation and enterprise server jde.ini files:

- Change the default OUTQ specified in the jde.ini file of the enterprise server. This setting is in the [Network Queue Settings] section and is called DefaultPrinterOUTQ. This OUTQ is used when no OUTQ is passed to the enterprise server from the workstation, or when the OUTQ name that is passed to the enterprise server is Default.

Two other settings, based on the jde.ini file on the workstation, tell the server whether to print the report immediately upon completion and whether to save the output from the report or delete it. Both of these settings are set in the following manner:

<pre>[NETWORK QUEUE SETTINGS] SaveOutput=TRUE PrintImmediate=TRUE</pre>

Setting SaveOutput to TRUE causes the enterprise server to save the PDF files in E810SYS/PRINTQUEUE until you explicitly delete them. Setting PrintImmediate to TRUE tells the enterprise server to print the job immediately after completing the report.

You should encourage workstation users to use the SaveOutput=FALSE entry in their jde.ini files. If users at workstations decide to save their output, they should periodically delete the

entries using the correct EnterpriseOne Job Master Search in the Job Control Master program (P986110B).

Note

To display job numbers, end-users can use the Job Control Master program (P986110B). Similarly, system administrators can use the Work With Servers application (P986116). While both applications perform similar functions, most sites generally use security to restrict access to the Work With Servers application to system administrators. Both programs use the Job Master Search form to display job numbers that correspond to member names. You can use either program to delete .PDF files by deleting appropriate entries.

Encoding the Passwords of Users Who Submit Batch Jobs

On the iSeries, when you want to encode user passwords for batch jobs, you need to change settings in the [SECURITY] section of the JDE.INI file.

► **To encode passwords of users who submit batch jobs**

Change the following setting in the JDE.INI file to False to deactivate encoding:

```
[SECURITY]
ServerPswdFile=TRUE
```

Running Batch Process Reports from the Server Command Line

If you have the proper authority, you can run batch process reports from the server command line.

► **To run batch process reports from the server command line**

On the command line, enter the following command:

<pre>RUNUBE USER(USER) PASSWORD(PASSWORD) ENVIRON(ENVIRONMENT) REPORT(REPORTNAME) VERSION(VERSION)</pre>

Example: Running Reports from the Command Line for iSeries

This example shows a command for executing the Business Unit Report (R0006P):

```
RUNUBE USER(SF5488324) PASSWORD(PASSWORD) ENVIRON(PD810)
REPORT(R0006P) VERSION(XJDE0001)
```

This command begins processing version XJDE0001 of the report in the PD810 environment. After completion, the PostScript spool file resides on the printer_1 OUTQ. The spool file leaves printer_1, and the .PDF file is not deleted.

Example: Scheduling Reports from the Command Line for iSeries

You can schedule a report from the command line for processing on a future date. You do this with the SBMJOB (submit job) command. Many options are available for this command, but the general form will be similar to the following example:

```
SBMJOB CMD(RUNUBE USER(SF5488324) PASSWORD(PASSWORD) ENVIRON(PD810)  
REPORT(R0006P) VERSION(XJDE0001)) SCDDATE(*FRI) SCDTIME(0600)
```

This command schedules the XJDE0001 version of the Business Unit Report (R0006P) to run on the next Friday at 06:00am. This job is submitted in the default job queue for the user who submitted the job. You can specify overrides on the command line or by prompting (F4) for more information.

You can review reports that have been submitted in this method by using the WRKSBMJOB command. This command displays all jobs submitted by the current user for batch processing. Information that this command displays includes the job name, the user who submitted the job, the type of job (BATCH), and the status. Using F11 also displays scheduling information for jobs that have been submitted but not yet run.

Running Multiple Instances of EnterpriseOne on the iSeries

You might want to run multiple instances of EnterpriseOne on an iSeries server for the following reasons:

- To test a new service pack
- To upgrade to a new version of EnterpriseOne

Note

You cannot use EnterpriseOne Planner to help you set up data for multiple instances of EnterpriseOne. Be prepared to manually copy data and to set up new Object Configuration Manager (OCM) mappings for each new instance.

Identifying an EnterpriseOne Instance on the iSeries

An EnterpriseOne instance on the iSeries server is uniquely identified by the following:

- EnterpriseOne system directory (integrated file system, or IFS) and library (QSYS file system)
- Path codes (IFS and QSYS file systems)
- Use of selected .ini file settings

The JDE.INI settings that you use to uniquely define an EnterpriseOne instance are summarized in the following table:

Section in server JDE.INI file	Parameter	Purpose
[INSTALL]	DefaultSystem=	The name of the EnterpriseOne System library. This value must be unique for each EnterpriseOne instance.
[JDEIPC]	StartIPCKeyValue=	The value of the first interprocess communication (IPC) ID of a range of keys, which JDEIPC uses for shared memory. This value, plus the value of the maxNumberOfResources parameter, defines the range of IPC IDs that the software uses for an instance of EnterpriseOne.
[JDENET]	ServiceNameListen=	The TCP/IP port number that the server uses for receiving communications packets from workstations and other EnterpriseOne servers.
[JDENET]	ServiceNameConnect=	The TCP/IP port number that the server uses for sending communications packets to workstations or other EnterpriseOne servers.
[DBSYSTEM SETTINGS]	Default Env=	The default environment for an instance of EnterpriseOne.
[DB SYSTEM SETTINGS]	Default PathCode=	The data source for an instance of EnterpriseOne.
[DB SYSTEM SETTINGS]	Library=	The database library that stores the system tables used by EnterpriseOne at startup.

Identifying an EnterpriseOne Instance for Applying Security on the iSeries

Additionally, to apply EnterpriseOne security throughout multiple instances, you use the following items to uniquely identify an instance:

- OCM mappings
- Database
- EnterpriseOne User Profile (the owner and default user ID under which EnterpriseOne jobs start)
- Selected settings in the JDE.INI file

The JDE.INI settings that you use to uniquely define an EnterpriseOne instance when you are applying security throughout multiple instances are summarized in the following table:

Section in server JDE.INI file	Parameter	Purpose
[DEBUG]	DebugFile	Specifies the location of the jdedebug.log file.
[DEBUG]	JobFile	Specifies the location of the jde.log file.
[DEBUG]	JDTSFile	Specifies the location of the lock manager trace file on the iSeries.
[DB SYSTEM SETTINGS]	Database	Specifies the name of the database in which the system tables reside.
[SECURITY]	DataSource	Specifies the name of the EnterpriseOne data source that contains the security tables and is used for user validation.

See Also

- ❑ *Setting Up iSeries EnterpriseOne Database Security in the Server and Workstation Administration Guide* for more information about securing multiple instances of EnterpriseOne

Creating an Instance of EnterpriseOne on the iSeries

To create an instance of EnterpriseOne on the iSeries, complete the following:

- Copy needed libraries and directories and modify the values of selected parameters in the .ini library
- Apply security to multiple instances of EnterpriseOne, if you desire to do so
- Create a new EnterpriseOne subsystem identification
- Start EnterpriseOne services

Prerequisites

- ❑ Verify that enough space exists on the direct access storage device (DASD) to create a new instance of EnterpriseOne.
- ❑ Assess data storage and backup requirements.
- ❑ Consider the procedure that you will follow for updating the EnterpriseOne server with new versions of EnterpriseOne.
- ❑ Determine the strategy for performing server package builds and updates. This might include, for example, setting up a second deployment server.
- ❑ Create a new environment for use with each new EnterpriseOne instance. For information about creating a new environment, see *Working with an Environment* in the *Configurable Network Computing Implementation Guide*.

Copying Libraries and Directories

To create an instance of EnterpriseOne on the iSeries, you copy the following:

- System library
- System directory
- Path code library
- Path code directory

► To copy libraries and directories

1. End EnterpriseOne services, if necessary.
2. Remove EnterpriseOne security, if necessary.
3. From the iSeries main menu, copy the EnterpriseOne system library in the QSYS file system by typing the following command:

```
CPYLIB 810SYS 810CST
```

Where 810CST is the name for the system library in the new instance of EnterpriseOne.

4. From the iSeries main menu, copy the EnterpriseOne system directory in the IFS by first the using the following command to create a temporary library:

```
CRTLIB TEMPLIB
```

5. Create a save file in the temporary library for the system directory by typing the following command:

```
CRTSAVF FILE (TEMPLIB/810SYS)
```

6. Save the system directory into the save file by typing the following command:

```
SAV DEV ( '/QSYS.LIB/TEMPLIB/810SYS.FILE' ) OBJ ( ( '/E810' ) )  
USEOPTBLK( *NO ) DTACPR( *YES )
```

7. Restore the save file for the system directory to a directory with a new name by typing the following command:

```
RST DEV('QSYS.LIB/TEMPLIB/810SYS.FILE') OBJ((' /810sys/* '
*INCLUDE/810cst'))
```

Where B810cst is the name of the new system directory.

Note

Throughout the entire copying procedure, the name for the new directories and libraries must match.

8. From the iSeries main menu, copy the path code library in the QSYS file system by typing the following command:

```
CPYLIB PRD810 CST810
```

Where CST810 is the name for the path code library in the new instance of EnterpriseOne. The name of the library for the new instance cannot exceed eight characters in length.

Note

The path code directory for any environment that you intend to use for a new instance of EnterpriseOne must be copied to the new directory. You cannot share path code directories between two or more instances of EnterpriseOne because this might corrupt the specification file.

9. From the iSeries main menu, copy the path code directory in the IFS by first using the following command to create a save file in the temporary library:

```
CRTSAVF FILE(TEMPLIB/PRD810)
```

Note

You must follow the procedure for copying the path code directory for each path code that you copy.

10. Save the path code directory into the save file by typing the following command:

```
SAV DEV('QSYS.LIB/TEMPLIB/PRD810.FILE') OBJ((' /prd810/* '))
USEOPTBLK(*NO) DTACPR(*YES)
```

11. Restore the save file for the path code directory to a directory with a new name by typing the following command:

```
RST DEV('QSYS.LIB/TEMPLIB/PRD810.FILE') OBJ((' /prd810/* ' INCLUDE
'/cst810'))
```

Where cst810 is the name of the new path code directory.

12. From the iSeries main menu, create an EnterpriseOne subsystem from the system library by typing the following command:

```
CRTOWSBS <subsystem name> <system library>
```

Where *<subsystem name>* is the name you give to the EnterpriseOne subsystem for the new instance of EnterpriseOne, and *<system library>* is the name of the system library in the QSYS file system for the new instance of EnterpriseOne.

Note

You can use the same subsystem for multiple instances of EnterpriseOne.

13. Modify the following parameters in the INI library:

```
[INSTALL]
DefaultSystem=<System Library>
[JDEIPC]
startIPCKeyValue=<Unused start key not within another instance's IPC range>
[JDENET]
serviceNameListen=<Available port>
serviceNameConnect=<Available port>
[DB SYSTEM SETTINGS]
Default Env=<New environment>
Default PathCode=<New path code>
```

Applying Security to Multiple Instances of EnterpriseOne on the iSeries

If you want to apply security to multiple instances of EnterpriseOne, complete the steps in the following task. If you do not want to apply security to multiple instances, proceed to the steps for creating an EnterpriseOne subsystem and starting an EnterpriseOne service.

Prerequisite

- ❑ Set up security for multiple instances of EnterpriseOne. See *Setting Up iSeries Database Security for Multiple EnterpriseOne Instances* in the *Server and Workstation Administration Guide*.

► To apply security to multiple instances of EnterpriseOne on the iSeries

1. Copy the OCM library.
2. Copy the database libraries, such as SYS810, 810MAP, and so on.
3. Create a new iSeries user profile for each new instance of EnterpriseOne.
4. From the iSeries main menu, create a new log path in the IFS by typing the following command:

```
CRTDIR DIR( '/810CSTLOG' )
```

Where CSTLOG is the name of the new IFS log directory.

5. Modify the following parameters in the INI library:

```
[DEBUG]
DebugFile=<new log path>/JDEDEBUG.LOG
JobFile=<new log path?>/JDE.LOG
JDETSFile=<new log path>/JDETS.LOG
[DB SYSTEM SETTINGS]
Database=<new OCM library>
[SECURITY]
DataSource=<Location of new F980WSEC library>
```

Note

The parameter values in the [DEBUG] section must be uppercase.

Creating an EnterpriseOne Subsystem on the iSeries

On the iSeries platform, a subsystem is a logical process that is used to run system jobs, whether they are EnterpriseOne or other application jobs. EnterpriseOne network and kernel jobs run under the iSeries subsystem, which PeopleSoft ships with a default description. You can use this description without modification when you are running a single instance of EnterpriseOne on the iSeries server.

If you decide to run multiple instances of EnterpriseOne, you need to create a new subsystem with a unique description for each instance of EnterpriseOne that you create. To create a new EnterpriseOne subsystem description, you use the CRTOWSBS command.

Prerequisites

- ❑ Stop EnterpriseOne services before you run the CRTOWSBS command.
- ❑ Review information about subsystems. See *PeopleSoft EnterpriseOne Subsystems* in the *System Administration Guide*.

► To create an EnterpriseOne subsystem on the iSeries

1. From the iSeries main menu, type the following command, and then press Enter or press the F4 key:

CRTOWSBS

The CREATE New EnterpriseOne Subsystem (CRTOWSBS) form appears.

2. On the CREATE New EnterpriseOne Subsystem form, enter character values for the following parameters, and then press Enter:
 - SUBSYSTEM
 - SYSLIB

Caution

The maximum number of characters allowed for the description of each parameter is 10. Verify that the name of the system library matches the name that you created when you copied the EnterpriseOne system library in the QSYS file system.

The CRTOWSBS command creates a new subsystem description in the EnterpriseOne system library and updates the STRNET and ENDNET programs with the new subsystem name as the default parameter.

3. To delete the old subsystem description from the system library, type the following command, and then press Enter or press the F4 key:

```
WRKOBJ OBJ <SUBSYSTEM NAME>/<SYSTEM LIBRARY NAME>) OBJTYPE(*SBSD)
```

Where *SUBSYSTEM NAME* is the subsystem description that you want to delete and *SYSTEM LIBRARY NAME* is the system library where the subsystem description is located.

4. In the Work with Objects form, type 4 for Delete, and then press Enter.

Starting EnterpriseOne Services on the iSeries

You start EnterpriseOne services by clearing IPC memory, using the CLRIPC command, and then running the STRNET command.

► To start EnterpriseOne services on the iSeries

1. From the iSeries main menu, clear IPC memory by typing the following command:
CLRIPC
2. From the iSeries main menu, start EnterpriseOne iSeries services by typing the following command:
STRNET

iSeries EnterpriseOne Database Security

You can secure profiles and objects for EnterpriseOne on the iSeries with the Set Up OneWorld Authority (SETOWAUT) command. When you enter this command, a form appears that allows you to enter specific security information for your system. The authority is implemented only on the iSeries machine on which you execute the command.

Note

If you upgraded to EnterpriseOne from an existing ERP installation and do not intend to rerun SETOWAUT, then you must manually add the PSFT user profile to the existing security profile authorization list. (The default name for authorization list is OneWorld.)

The SETOWAUT command allows you to set up security for a single instance of EnterpriseOne or for multiple instances of EnterpriseOne. If you run multiple instances of EnterpriseOne, you can set up separate user profiles for each instance. The SETOWAUT command sets up the authorities for each EnterpriseOne instance, adds profile names to an authorization list, and sets object ownership for each EnterpriseOne instance.

Two separate authorization lists exist for maintaining security. Values in two parameters of the SETOWAUT program specify the authorization lists.

The USRPRF parameter value specifies the EnterpriseOne user profile. When you run the SETOWAUT program, the program automatically creates a user profile authorization list with the same name. This list secures all EnterpriseOne objects.

Note

The ALLOBJECTS parameter determines how SETOWAUT secures EnterpriseOne objects. The recommended setting for this parameter is *NONCOEXIST. Using this setting, the resulting authorization list secures only the root directories and the libraries. This is true for all libraries except the System library; SETOWAUT secures all of the objects in the system library. The value ALLOBJ secures every object in all EnterpriseOne libraries and directories. This parameter is not recommended because it negatively affects EnterpriseOne performance.

The COEXIST option can be used for OneWorld Xe, but never for EnterpriseOne. COEXIST is not valid with EnterpriseOne.

This release of EnterpriseOne introduces the PSFT user profile. To use EnterpriseOne software, this user profile must have access to objects that are owned by this instance of the software, regardless of whether SETOWAUT is used (that is, the default profile is the ONEWORLD user profile). To provide the PSFT user profile access to objects, you must do the following:

- Change PSFT user profile attribute GRPPRF to the name of the EnterpriseOne or multiple instance USRPRF (the default value is ONEWORLD).
- Verify that the PSFT user profile attribute OWNER is set to *GRPPRF. If it is not, manually set this value to *GRPPRF.

The USRAUTL parameter value specifies the administrative authorization list. When you run the SETOWAUT program, the program automatically creates an administrative authorization list that gives specified users administrative access to EnterpriseOne. Any user who will perform basic EnterpriseOne administration tasks, such as Start, End, Clear IPC, and so on, on the iSeries must be added to this list. CRTOWADPRF is a supplied command that adds administrative users to this list; RMVOWADPRF is a supplied command that removes such users from the list.

Use PROFTYPE(*USER) to perform basic EnterpriseOne administrative tasks. Use PROFTYPE(*ADMIN) for users who need access to all EnterpriseOne objects. (*ADMIN is similar to security officer but can only be used for EnterpriseOne.

Whether you want to set up security for one instance of EnterpriseOne or for multiple instances, the Setup OneWorld Authority (SETOWAUT) form appears when you run the SETOWAUT command. However, the parameter values that you enter and the parameter fields that appear on the form differ, depending on whether you set up security for one instance or for multiple instances. These parameter differences are explained in the following three tables:

Parameters Present in SETOWAUT Form for Both Single and Multiple Instances of EnterpriseOne	Meaning	Value to be Entered for a Single Instance of EnterpriseOne	Value to be Entered for Multiple Instances of EnterpriseOne
USRPRF	EnterpriseOne User Profile	EnterpriseOne	Configurable. Enter a new value for each instance of EnterpriseOne.
USRAUTL	Admin. Authorization List	OWADMINL	Configurable. Enter a new value for each instance of EnterpriseOne.

Parameters Present in SETOWAUT Form for Single Instance of EnterpriseOne Only	Meaning	Value to be Entered for a Single Instance of EnterpriseOne	Value to be Entered for Multiple Instances of EnterpriseOne
OWPRF	Modify ONEWORLD Profile	Y is the default value.	Parameter is not present if you enter a value other than ONEWORLD for the USRPRF parameter.
JDEPRF	Modify JDE Profile	Y is the default value.	Parameter is not present if you enter a value other than ONEWORLD for the USRPRF parameter.

Parameter Present in SETOWAUT Form for Multiple Instances of EnterpriseOne Only	Meaning	Value to be Entered for Multiple Instances of EnterpriseOne	Value to be Entered for Single Instance of EnterpriseOne
OBJOPT	Secure All Objects	N is the default value. Enter Y if you want to secure all objects that appear in one or more directories. Because it can degrade system performance, entering Y is not recommended.	Parameter is not present if you enter OneWorld as the value for the USRPRF parameter.

The following information provides a summary of the security model when you run a single instance of EnterpriseOne:

Library

Description of Security

EnterpriseOne
System Library

SETOWAUT secures all of the objects in the system library. Administrative programs, such as CLRIPC, STRNET, ENDNET, and PORTTEST, are set to adopt the authority of the owner.

Prerequisite

- ❑ Review information about each of the parameters that appears on the SETOWAUT form. See *Explanations of iSeries Database Security Parameters* in the *Server and Workstation Administration Guide*.

Setting Up iSeries EnterpriseOne Database Security

You can set up security for a single instance of EnterpriseOne, or you can set up security for separate EnterpriseOne instances. In the latter case, the SETOWAUT program creates a user profile and individual authorization lists for each instance, which establishes object ownership.

Setting Up iSeries Database Security for a Single EnterpriseOne Instance

You can set up security for a single instance of EnterpriseOne. To do so, you enter ONEWORLD as the USRPRF (user profile) parameter value and OWADMINL (admin. authorization list) parameter value.

► To set up iSeries database security for a single EnterpriseOne instance

1. In the SETOWAUT library, on the command line, type the following command, press F4, and then press F11:

SETOWAUT

Note

Verify that the SETOWAUT library is in the library list. If it is not, run the ADDLIBL command.

The Set Up OneWorld Authority (SETOWAUT) form appears.

2. On Set Up OneWorld Authority (SETOWAUT), complete the following field, and then press Enter:
 - USRPRF
Enter OneWorld
3. The form displays additional security parameters. You can specify various security settings, including library access.
4. Complete the following required fields, and then press Enter:
 - USRAUTL
Enter OWADMINL.
 - TYPE
 - INILIB

5. Complete any additional fields, if necessary.
6. Press Enter.

Caution

Use caution when you use EnterpriseOne security to lock a library that contains third-party software. PeopleSoft does not support iSeries EnterpriseOne database security with third-party software.

Setting up iSeries Database Security for Multiple EnterpriseOne Instances

You can set up security for separate instances of EnterpriseOne. To do so, you enter a value other than ONEWORLD for the User Profile parameter and a value other than OWADMINL for the Admin. Authorization List parameter. You enter different values for these parameters for each instance of EnterpriseOne that you run.

Prerequisite

- ❑ Before you enter a value for the USRPRF and USRAUTL parameters, verify that the value is not being used for an authorization list for any other instance of EnterpriseOne. To do so, run the DSPAUTL command. On the Display Authorization form, you can enter the value that you intend to use to make sure that it is unique.

► To set up iSeries database security for multiple EnterpriseOne instances

1. In the SETOWAUT library, on the command line, type the following command and press F4:
SETOWAUT
2. On Setup OneWorld Authority (SETOWAUT), complete the following field, and then press Enter:
 - USRPRF
The SETOWAUT program uses this name when it creates a user profile authorization list.
3. The form expands to reveal an additional security parameter. The Modify OneWorld Profile (OWPRF) and Modify JDE Profile (JDEPRF) parameters, which appear when you enter OneWorld as the User Profile parameter value, do not appear when you enter a value other than OneWorld.
4. Complete the following required fields and press Enter:
 - USRAUTL
Enter a name that identifies the administrative authorization list.
 - TYPE
 - INILIB
5. Complete any additional fields, if necessary.
6. Press Enter.

See Also

- *Explanations of iSeries Database Security Parameters* in the *Server and Workstation Administration Guide* for details about the available options for each field

Explanations of iSeries Database Security Parameters

You use the iSeries database security parameters to modify user and administrator profiles, to secure objects, and so on. These parameters appear on the Setup OneWorld Authority (SETOWAUT) form.

Type

Depending on the value that you enter in this field, you can implement a full security setup, modify only the security profiles, or modify only the datapaths authority. A full security setup includes the system library, datapath, pathcode, and user profiles.

- Use ***FULL** when you initially implement SETOWAUT. This value directs SETOWAUT to perform all of the security routines.
- Use ***DTAPATH** only when you need to secure one or more datapaths.
- Use ***PROF** to perform only the user profile routines. SETOWAUT uses the user profile settings in the command to direct the process.
- Use ***SYSTEM** to perform the System library authority functions. If you are running a single instance of EnterpriseOne, ***SYSTEM** secures the System library and all of the objects within it with the AUTL OWADMINL. If you are running multiple instances of EnterpriseOne, ***SYSTEM** secures the library and all the objects contained within it with the administrative authorization list created by the SETOWAUT program for each individual instance of EnterpriseOne. Note that SETOWAUT must be run separately for each instance of EnterpriseOne.

Additionally, all the ***PGM** objects with attributes of ***CLP**, ***CLLE**, or ***CLE** will have the program attributes modified for adopt authority. The system library is treated differently to allow the administration of EnterpriseOne.

You can use this parameter to lock other nonsystem libraries that contain objects that you can use to administer EnterpriseOne.

Additional Profile Work That SETOWAUT Performs When You Use Types ***FULL** or ***PROF**

When you enter Type ***FULL** or ***PROF**, SETOWAUT does the following:

- Creates the ONEWORLD and OWADMINL authorization lists (if they do not already exist) if you are running a single instance of EnterpriseOne. If you are running multiple instances of EnterpriseOne, SETOWAUT creates both authorization lists and uses the names that you specified for each instance of EnterpriseOne.
- Changes the owner of both lists to ONEWORLD if you are running a single instance of EnterpriseOne. If you are running multiple instances of EnterpriseOne, SETOWAUT changes the owner of both lists to the user profile name that you specified for each instance of EnterpriseOne.
- Adds JDE to both lists if you are running a single instance of EnterpriseOne.

- Adds PSFT to both lists if you are running a single instance of EnterpriseOne.
- Changes *PUBLIC entry to *EXCLUDE in both lists.

INILIB (INI Library)

This field identifies the library in which the JDE.INI file resides for the security application. The *NONE value allows you to specify that the JDE.INI file is either not needed or not available.

Caution

You cannot use the parameter value *NONE if the Type parameter value is *FULL or *SYSTEM.

Use a library name if all of the following requirements are true:

- An EnterpriseOne INI library is located on the host system.
- The control files (OCM) are located on the host system.
- The JDE.INI file references the OCM library.

When the Type field contains the value *FULL or *SYSTEM, the library and all of the objects will be secured with SYSTEM attributes. SETOWAUT uses the JDE.INI file to perform all of the INI retrievals.

When any of the previous requirements are false, use *NONE. This setting requires you to enter actual values in any parameter that allows the value *INI.

DTAPATH Datapath (library)

Note

If you set the INI library field to *NONE, you must manually set datapaths in this field.

Type *INI in this field to use the datapaths that are set in the JDE.INI file. You can also type specific datapaths in this field. You can type up to 10 datapaths at a time.

Use *INI when the following are true:

- SETOWAUT will modify each library based upon the ALLOBJECTS parameter.
- The INILIB parameter contains the library name in which the JDE.INI file is located (the INILIB value is not *NONE). This parameter tells SETOWAUT to use the JDE.INI file to retrieve the datapath libraries. SETOWAUT retrieves the library name from the JDE.INI value in the [DB SYSTEM SETTINGS] Library and uses this setting to access the Object Configuration Master (F986101) and Data Source Master (F98611) tables. SETOWAUT selects all of the library names (F98611.OMLIB) that meet the following criteria:
 - F986101.OMDATP = F98611.OMDATP
 - OMUGRP = *PUBLIC, OMSTSO = 'AV'
 - OMSRVR = *the host name*

Modify System Profile

Valid values for this field are Y and N.

Note

This field does not appear when you set up authorization for multiple instances of EnterpriseOne and you enter a value other than ONEWORLD in the USRPRF field.

Enter Y when you want to do the following:

- Modify or create the system profile that has not yet been modified. For example, you might enter the following information to modify a system profile:
 - *NONE in the GRPPRF field
 - *NONE in the SUPGRPPRF field
 - *USER in the USRCLS field
 - *SIGNOFF in the INLMNU field
 - *NONE in the INLPGM field
 - *JOBCTL in the SPCAUT field
- Grant authority to change the profile ONEWORLD to *USE profile QSECOFR
- Revoke *ALL authority from *PUBLIC

Enter N only when the system profile has the correct attributes.

Modify JDE Profile

Valid values for this field are Y and N.

Note

This field does not appear when you set up authorization for multiple instances of EnterpriseOne and you enter a value other than ONEWORLD in the USRPRF field.

Enter Y when you want to do the following:

- Modify or create the JDE profile that has not been modified. For example, you might enter the following to modify a JDE profile:
 - *NONE in the GRPPRF field
 - *NONE in the SUPGRPPRF field
 - *USER in the USRCLS field
 - *NONE in the INLPGM field
 - *JOBCTL *SAVSYS in the SPCAUT field
- Revoke *ALL authority from *PUBLIC

Enter N only when the profile JDE has the correct attributes.

Modify Security Profile

Note

It is recommended that you delete existing security profiles before running SETOWAUT. After running SETOWAUT and creating security profiles, the passwords must be changed to correspond with passwords that were set up using EnterpriseOne User Security. The Security user is used as the System user in EnterpriseOne User Security.

You can enter up to 10 security profiles at a time in this field to modify using the SETOWAUT program. SETOWAUT must be run with the PSFT user profile specified as a security profile when using EnterpriseOne. If you enter a security profile that does not already exist, SETOWAUT creates the profile and modifies the profile accordingly. You can do any of the following:

- Create or modify the profile by entering information such as the following:
 - *USER in the USRCLS field
 - *SIGNOFF in the INLMNU field
 - *NONE in the INLPGM field
 - *NONE in the SPCAUT field
 - ONEWORLD in the GRPPRF field, if you are running a single instance of EnterpriseOne. If you are running multiple instances of EnterpriseOne, enter in the GRPPRF field the EnterpriseOne User Profile name that you entered in the USRPRF field.
 - JDE in the SUPGRPPRF field, if you are running a single instance of EnterpriseOne. If you are running multiple instances of EnterpriseOne, enter in the SUPGRPPRF field the EnterpriseOne User Profile name that you entered in the USRPRF field.
- Revoke *ALL authority from *PUBLIC
- Grant profile ONEWORLD *CHANGE authority to the security profile
- Grant security profile *CHANGE authority to ONEWORLD

See Also

- *Sample Results for SETOWAUT in the Server and Workstation Administration Guide* for more information about specific profile properties

EnterpriseOne DB Admin Profile

When you type *INI in this field, SETOWAUT retrieves the user and password values from the [SECURITY] section of the JDE.INI file. If you type a value that does not exist, SETOWAUT creates a profile with a password that is the same as the profile name. If the profile exists, SETOWAUT modifies the profile to be an EnterpriseOne database administrator.

Enter a profile to be used as a database administrator. This profile will have all rights to all EnterpriseOne objects. These database administrator profiles are allowed to perform certain

EnterpriseOne processes (RUNUBE and PORTTEST) that an administrator with normal privileges cannot perform.

If the profile does not exist, the system creates the profile with a password that is the same name as the profile. If the profile does not exist, you should set the password to expire (PWDEXP = *YES). For example, the following occurs:

- If BV3C is in library list, SETOWAUT will place this program as the initial program. (This program lists all of the EnterpriseOne occurrences to allow the user to select one occurrence at signon).
- USRCLS is set to *PGMR
- SPCAUT is set to *NONE
- GRPPRF is set to ONEWORLD if you are running a single instance of EnterpriseOne. If you are running multiple instances of EnterpriseOne, GRPPRF is set to the EnterpriseOne User Profile name that you entered in the USRPRF parameter field.

This profile revokes *ALL authority from *PUBLIC and grants ONEWORLD *USE rights to the DB ADMIN profile.

BSFNLIB (Libs or *INI (Default PathCode))

Note

If you enter *NONE in the INI library field, you must set pathcodes in this field.

Type *INI in this field to use the pathcode library and the associated specification file directory that is set in the JDE.INI file. You can also type specific pathcode libraries in this field. You can type up to 10 pathcodes at a time.

Use *INI when the INILIB parameter contains the library name in which the JDE.INI file is located (INILIB does not contain *NONE). This parameter tells SETOWAUT to use the JDE.INI file to retrieve the application pathcode libraries. SETOWAUT retrieves the library name from the JDE.INI value in "[DB SYSTEM SETTINGS] Library" and uses this setting to access the Object Configuration Master (F986101) and Data Source Master (F98611) tables. SETOWAUT selects all of the library names (F98611.OMLIB) that meet the following criteria:

- F986101.OMDATP = F98611.OMDATP
- OMUGRP = *PUBLIC
- OMSTSO = 'AV'
- OMDBNM = F00942

SETOWAUT retrieves EMPATHCD (pathcode) from each record in the Object Path Master File table (F00942) for each library (F98611.OMLIB).

For each pathcode, SETOWAUT modifies the library and associated IFS directory (specifies path) accordingly.

Secure Log Path

Y and N are valid values for this field. The recommended value is N.

Enter N when you do not want to secure JDE log paths.

Enter Y only when you need to secure the log paths. One situation in which you might secure JDE log paths is when logs are being deleted without permission.

Only DB administrators have permission to access the logs in the log path.

Secure All Objects

Use this field to secure objects when you are running multiple instances of EnterpriseOne. The parameter appears on the SETOWAUT form only when you configure an instance of EnterpriseOne by entering a value other than ONEWORLD in the USRPRF field.

*NONCOEXIST is the default value for the Secure All Objects parameter, and PeopleSoft recommends that you use this value. This value secures all directories, but not the files in the directories.

Entering COEXIST secures the files as well as the directories. Entering COEXIST can degrade performance because the system must verify authority for every object that the user wants to access. This value is the equivalent of entering *ALLOBJECTS when you run a single instance of EnterpriseOne. The value *COEXIST can only be used for OneWorld Xe, and must never be used for EnterpriseOne.

Adding Administrators

You can add administrators to the administrative authorization list by running the CRTOWADPRF command. The command also allows you to designate levels of authority to the administrators whom you are adding to the list.

► To add administrators

1. On the command line, enter the following command and press F4:
`CRTOWADPRF USRPRF`

The Setup OW User Profile (CRTOWADPRF) form appears.
2. On Setup OW User Profile (CRTOWADPRF), complete the following fields and press Enter:
 - **ADMIN USER Profile**
Type the name of an administrator whom you want to add to the administrative authorization list. You can add up to 10 administrators at a time.
 - **EnterpriseOne USER Profile**
Your entry must match the EnterpriseOne user profile name that you entered in the USRPRF field during setup.
 - **ADMIN Authorization List**
Your entry must match the Admin. Authorization List name that you entered in the USRAUTL field during setup.
 - **Profile Type**

Type *USER to give the profiles basic administration capabilities, such as STRNET, ENDNET, CLRIPC, SAW, CLRLCK, DSPIPC, DSPSTMF, IPCS, LINKBSFN, and PID2JOB.

Type *ADMIN if the profiles need rights to PORTTEST and RUNUBE, as well as the basic administration capabilities.

- Initial program to call

Type BV3C if you want the system to display a list of environments when the administrators sign on to EnterpriseOne, *SAME to use the current initial program setting, or *NONE to remove the initial program setting.

Note

For EnterpriseOne, the initial program to call by default is BV3C. This program sets the iSeries to provide a choice of environments at signon. A user with an administrator profile who signs on to an environment can then perform EnterpriseOne commands on the iSeries server.

Removing Administrative Authority from User Profiles

To remove a user's administrative authority, you run the RMVOWADPRF command and complete the Remove OW Profile Authority form.

► **To remove administrative authority from user profiles**

Note

Submit this command to a batch subsystem.

1. On the command line, enter the following command and press F4:

RMVOWADPRF

The Remove OW Profile Authority (RMVOWADPRF) form appears.

2. On Remove OW Profile Authority (RMVOWADPRF), complete the following fields and press Enter:

- User Profile

Enter the name of the user from whom you want to remove authority.

- Admin. Authorization List

Your entry must match the Admin. Authorization List name that you entered in the USRAUTL field during setup.

- EnterpriseOne User Profile

Your entry must match the EnterpriseOne user profile name that you entered in the USRPRF field during setup.

Sample Results for SETOWAUT

You can expect the following examples for each of the various commands. Using Client Access, sign onto the iSeries, type each command on the command line, and press F4. For libraries (data sources and pathcodes), the required parameters are object type (*LIB) and the name of the library.

If you set up multiple instances of EnterpriseOne, the owner of each instance is the user profile that you entered in the EnterpriseOne User Profile parameter during the authority setup. If you set up a single instance of EnterpriseOne, the owner is ONEWORLD.

Similarly, if you set up multiple instances of EnterpriseOne and you display object authority, the value that appears is the name of the user profile for all objects except the SYSTEM library. The object authority for the SYSTEM library when you run multiple instances of EnterpriseOne is the name of the Admin. Authorization List. If you set up a single instance of EnterpriseOne, all objects are secured by the authorization list, except the SYSTEM library, which is secured by the OWADMINL authorization list.

Data source DSPOBJAUT:

```
JDEDEV - [24 x 80]
File Edit View Communication Actions Window Help

Display Object Authority

Object . . . . . : PRODDTA      Owner . . . . . : ONEWORLD
Library . . . . . : QSYS        Primary group . . . : *NONE
Object type . . . . : *LIB

Object secured by authorization list . . . . . : ONEWORLD

User      Group      Object
*PUBLIC   Group      Authority
*PUBLIC   Group      *AUTL

Press Enter to continue.

F3=Exit  F11=Display detail object authorities  F12=Cancel
F14=Display authorization list  F17=Top      F18=Bottom

MA a MW 01/001
I902 - Session successfully started
```

Pathcode DSPOBJAUT:

```
JDEDEV - [24 x 80]
File Edit View Communication Actions Window Help

Display Object Authority

Object . . . . . : PY9FA      Owner . . . . . : ONEWORLD
Library . . . . . : QSYS      Primary group . . . : *NONE
Object type . . . . : *LIB

Object secured by authorization list . . . . . : ONEWORLD

User      Group      Object
*PUBLIC   Authority
*AUTL

Press Enter to continue.

F3=Exit  F11=Display detail object authorities  F12=Cancel
F14=Display authorization list  F17=Top         F18=Bottom

MA a MW 01/001
I902 - Session successfully started
```

Data source DSPLIBD:

```
JDEDEV - [24 x 80]
File Edit View Communication Actions Window Help

Display Library Description


Library . . . . . : PRODDTA
Type . . . . . : PROD
ASP of library . . . . . : 1
Create authority . . . . . : ONEWORLD
Create object auditing . . . . . : *SYSVAL
Text description . . . . . : Production Data Tables

Press Enter to continue.

F3=Exit  F12=Cancel
(C) COPYRIGHT IBM CORP. 1980, 2000.

MA a MW 01/001
I902 - Session successfully started
```

Pathcode DISLIBD:



Note

Authority for objects in data sources and pathcodes should remain the same after you run SETOWAUT. You can see this by displaying the authority for an object in each library before and after you run SETOWAUT. The forms should be identical. The required parameters are object name, object type (*FILE or *PGM), and the library name in which the object resides.

Owner, object security, and authority creation differ depending on whether you are running a single instance of EnterpriseOne or multiple instances.

SETOWAUT changes the authority on system libraries. You can view this for both DSPOBJAUT and DSPLIBD on system libraries. The shaded information in the following illustrations should correspond to the information that appears on your form. The required parameters are the object name, object type (*PGM), and the name of the library in which these objects reside.

System library DSOBJAUT:

The screenshot shows a terminal window titled "JDEDEV - [24 x 80]" with a menu bar (File, Edit, View, Communication, Actions, Window, Help) and a toolbar. The main display area is titled "Display Object Authority" and shows the following information:

```
Object . . . . . : B9SYS      Owner . . . . . : ONEWORLD
Library . . . . . : QSYS      Primary group . . . : *NONE
Object type . . . . : *LIB

Object secured by authorization list . . . . . : OWADMINL

User      Group      Object
*PUBLIC   *AUTL      Authority
```

Below the information, it says "Press Enter to continue." and "Bottom". At the bottom, there are function key shortcuts: F3=Exit, F11=Display detail object authorities, F12=Cancel, F14=Display authorization list, F17=Top, and F18=Bottom. The status bar at the bottom shows "MA a MW 01/001" and a message "I902 - Session successfully started".

System library DSPLIBD:

The screenshot shows a terminal window titled "JDEDEV - [24 x 80]" with a menu bar (File, Edit, View, Communication, Actions, Window, Help) and a toolbar. The main display area is titled "Display Library Description" and shows the following information:

```
Library . . . . . : B9SYS
Type . . . . . : PROD
ASP of library . . . . . : 1
Create authority . . . . . : ONEWORLD
Create object auditing . . . . . : *SYSVAL
Text description . . . . . :
```

Below the information, it says "Press Enter to continue." and "Bottom". At the bottom, there are function key shortcuts: F3=Exit, F12=Cancel, and the copyright notice "(C) COPYRIGHT IBM CORP. 1980, 2000.". The status bar at the bottom shows "MA a MW 01/001" and a message "I902 - Session successfully started".

The authority changes for objects within system libraries that either contain the attributes CLLE or CLP or that share the same name. You can use commands to review the authority on these objects. The required parameters are object name, object type (*PGM or *CMD), and the name of the library in which these objects reside.

Displaying User Profile Information

After you run SETOWAUT, you can review user profiles and authorization lists to verify that the information is correct.

► To display user profile information

1. On the command line, enter the following command:

DSPUSRPRF

The Display User Profile (DSPUSRPRF) form appears.

2. On Display User Profile (DSPUSRPRF), type the name of a user profile in the User Profile field, and then press Enter.

Information similar to the following appears; compare the shaded information for accuracy:

User profile	: ONEWORLD
Previous sign-on	: 03/23/04 15:16:53
Sign-on attempts not valid	: 0
Status	: *ENABLED
Date password last changed	: 02/27/03
Password expiration interval	: *NOMAX
Set password to expired	: *NO
User class	: *USER
Special authority	: *JOBCTL
Group profile	: *NONE
Owner	: *USRPRF
Group authority	: *NONE
Group authority type	: *PRIVATE
Supplemental groups	: *NONE
Assistance level	: *SYSVAL
Current library	: *CRTDFT
Initial program	: *NONE
Library	:
Initial menu	: *SIGNOFF
Library	:
Limit capabilities	: *NO
Text	:
Display sign-on information	: *SYSVAL
Limit device sessions	: *SYSVAL

Keyboard buffering	: *SYSVAL
Storage information:	
Maximum storage allowed	: *NOMAX
Storage used	: 286236536
Storage used on independent ASP	: *NO
Highest scheduling priority	: 3
Job description	: ONEWORLD
Library	: QGPL
Accounting code	:
Message queue	: ONEWORLD
Library	: QUSRSYS
Message queue delivery	: *NOTIFY
Message queue severity	: 00
Output queue	: *WRKSTN
Library	:
Printer device	: *WRKSTN
Special environment	: *SYSVAL
Attention program	: *SYSVAL
Library	:
Sort sequence	: *SYSVAL
Library	:
Language identifier	: *SYSVAL
Country identifier	: *SYSVAL
Coded character set identifier	: *SYSVAL
Character identifier control	: *SYSVAL
Locale job attributes	: *SYSVAL

User profile	: JDE
Previous sign-on	: 03/23/04 15:25:53
Sign-on attempts not valid	: 0
Status	: *ENABLED
Date password last changed	: 02/27/03
Password expiration interval	: *NOMAX
Set password to expired	: *NO
User class	: *USER
Special authority	: *JOBCTL
	: *SAVSYS
Group profile	: *NONE
Owner	: *USRPRF
Group authority	: *NONE
Group authority type	: *PRIVATE
Supplemental groups	: *NONE

Assistance level	: *SYSVAL
Current library	: *CRTDFT
Initial program	: J98INIT
Library	: JDF0BJ7R2
Initial menu	: *MAIN
Library	: *LIBL
Limit capabilities	: *NO
Text	:
Display sign-on information	: *SYSVAL
Limit device sessions	: *SYSVAL
Keyboard buffering	: *SYSVAL
Storage information:	
Maximum storage allowed	: *NOMAX
Storage used	: 11243168
Storage used on independent ASP	: *NO
Highest scheduling priority	: 3
Job description	: JDE
Library	: QGPL
Accounting code	:
Message queue	: JDE
Library	: QUSRSYS
Message queue delivery	: *NOTIFY
Message queue severity	: 00
Output queue	: *DEV
Library	:
Printer device	: *WRKSTN
Special environment	: *SYSVAL
Attention program	: *SYSVAL
Library	:
Sort sequence	: *SYSVAL
Library	:
Language identifier	: *SYSVAL
Country identifier	: *SYSVAL
Coded character set identifier	: *SYSVAL
Character identifier control	: *SYSVAL
Locale job attributes	: *SYSVAL

User profile	:	JDEOW
Previous sign-on	:	03/23/04 15:28:02
Sign-on attempts not valid	:	0
Status	:	*ENABLED
Date password last changed	:	02/27/03
Password expiration interval	:	*NOMAX
Set password to expired	:	*NO
User class	:	*USER
Special authority	:	*NONE
Group profile	:	ONEWORLD
Owner	:	*GRPPRF
Group authority	:	*NONE
Group authority type	:	*PRIVATE
Supplemental groups	:	JDE
Assistance level	:	*SYSVAL
Current library	:	*CRTDFT
Initial program	:	*NONE
Library	:	
Initial menu	:	*SIGNOFF
Library	:	
Limit capabilities	:	*NO
Text	:	
Display sign-on information	:	*SYSVAL
Limit device sessions	:	*SYSVAL
Keyboard buffering	:	*SYSVAL
Storage information:		
Maximum storage allowed	:	*NOMAX
Storage used	:	147904
Storage used on independent ASP	:	*NO
Highest scheduling priority	:	3
Job description	:	QDFTJOB
Library	:	QGPL
Accounting code	:	
Message queue	:	JDEOW
Library	:	QUSRSYS
Message queue delivery	:	*NOTIFY
Message queue severity	:	00
Output queue	:	*WRKSTN
Library	:	
Printer device	:	*WRKSTN
Special environment	:	*SYSVAL

Attention program	: *SYSVAL
Library	:
Sort sequence	: *SYSVAL
Library	:
Language identifier	: *SYSVAL
Country identifier	: *SYSVAL
Coded character set identifier	: *SYSVAL
Character identifier control	: *SYSVAL
Locale job attributes	: *SYSVAL

User profile	: OWDBADMIN
Previous sign-on	: 03/23/04 15:30:12
Sign-on attempts not valid	: 0
Status	: *ENABLED
Date password last changed	: 02/27/03
Password expiration interval	: *NOMAX
Set password to expired	: *NO
User class	: *PGMR
Special authority	: *NONE
Group profile	: ONEWORLD
Owner	: *GRPPRF
Group authority	: *NONE
Group authority type	: *PRIVATE
Supplemental groups	: JDE
Assistance level	: *SYSVAL
Current library	: *CRTDFT
Initial program	: *NONE
Library	:
Initial menu	: MAIN
Library	: *LIBL
Limit capabilities	: *NO
Text	:
Display sign-on information	: *SYSVAL
Limit device sessions	: *SYSVAL
Keyboard buffering	: *SYSVAL
Storage information:	
Maximum storage allowed	: *NOMAX
Storage used	: 0
Storage used on independent ASP	: *NO
Highest scheduling priority	: 3
Job description	: QDFTJOB
Library	: QGPL


Accounting code	:	
Message queue	:	JDEOW
Library	:	QUSRSYS
Message queue delivery	:	*NOTIFY
Message queue severity	:	00
Output queue	:	*WRKSTN
Library	:	
Printer device	:	*WRKSTN
Special environment	:	*SYSVAL
Attention program	:	*SYSVAL
Library	:	
Sort sequence	:	*SYSVAL
Library	:	
Language identifier	:	*SYSVAL
Country identifier	:	*SYSVAL
Coded character set identifier	:	*SYSVAL
Character identifier control	:	*SYSVAL
Locale job attributes	:	*SYSVAL


Sample Results for Authorization Lists

Use the following commands to view the authorization list authorities. The name of the list is the only necessary parameter:

- IFS directories (specification files)
- WRKLNK – option 9 Work with authority

DSPAUTL





Important

This folder should be inaccessible from Windows Explorer.

Server Administration for UNIX and Linux

PeopleSoft supports EnterpriseOne enterprise servers for UNIX operating systems on the Hewlett-Packard HP 9000 (HP-UX), the IBM RS/6000 (AIX), and the Sun Microsystems SPARC (Solaris) platforms. In addition, beginning with PeopleSoft 8.93, we also support RedHat Enterprise Linux AS on the Intel Architecture. To operate your UNIX or Linux enterprise server, you need to perform administrative procedures on your server to ensure that EnterpriseOne will run properly.

Note

Some information in this and other guides refers to UNIX generically and includes the supported Linux platforms unless otherwise noted.

EnterpriseOne Directory Structure for UNIX and Linux


The following is a list of directories that are shipped on the UNIX and Linux EnterpriseOne Server Installation CD. They should be installed under the EnterpriseOne base directory; for example, install them in /u01/PeopleSoft/E810. Indented names indicate subdirectories of the directories, which are not indented.

pathcode	The main directory for the business function shared libraries, C header files, object files, source files, and specification (spec or TAM) files. Upon installation, this directory is copied to the correct path codes, such as PD810 and DV810. The following subdirectories are included: • bin32 – Business function shared libraries • Spec files – Binary data files in a Peoplesoft proprietary format
system	The main directory for the system-level executables, shared libraries, C header files, libraries, and localization files. The following subdirectories are included: • bin32 – System-level executables and scripts • include – System-level C header files • includev – System-level C header files provided by third-party vendors, such as Vertex • lib – System-level shared libraries and export files • libv32 – System-level shared libraries provided by third-party vendors
ini	The location of the JDE.INI file.
PrintQueue	The location to which all .PDF file outputs for reports are written.
log	The location to which the jde_xxx.log and jdedbug_xxx.log files are written.
packages	The server package installation base directory. Directories exist here only when a package has been installed. Under the package directory are subdirectories named for each package that has been installed. Located under each package are the following directories: • bin32 – Business function shared libraries • include – Business function header files • obj – Business function object files These are divided among lower-level subdirectories that correspond to each shared library in the bin32 directory. • source – Business function source files These are divided among lower-level subdirectories that correspond to each shared library in the bin32 directory. • spec – Specification files These binary data files are in a PeopleSoft proprietary format.

EnterpriseOne Architecture and Process Flow for UNIX and Linux

The following host server processes perform the indicated actions.

UNIX® Host Server Process



The following information explains the process flow:

1. The jdenet_n Master process spawns jdenet_n Slave and jdenet_k processes (also called kernels) at startup or as they are needed. EnterpriseOne uses a number of different types of kernels to handle different types of processing, even though all of these have the same

process name in the operating system (jdenet_k). The definitions for the number of processes to start and what types to start are stored in the JDE.INI file.

2. The queue kernel process spawns the runbatch process whenever a relevant batch process request is placed in the Job Control Status Master table (F986110). The runbatch process completes the job, updates the F986110 table, and then quits. In EnterpriseOne, you use the Job Queue Maintenance program (P986130) to set up and manage the job queues.

Nearly all jdenet_k processes access various other database tables as needed. The runbatch process, for instance, accesses and modifies any database table that is relevant to the particular business logic it is running.

3. Message queues are a type of interprocess communication (IPC) resource. They are allocated by the jdenet_n processes by calls to the operating system. While the software is running, operating system information about the message queues can be obtained by using the command `ipcs`.

When message packets are routed to the jdenet_n process from a client or another server, the jdenet_n process places them in the appropriate message queue according to the type of message. For example, when a client submits a batch process, a message is routed to the batch process kernel; when business logic needs to be run on the server, a request is routed to the CallObject kernel; when a user signs on to the system, a request is routed to the security kernel, and so on.

Each message queue has an identifier (IPC key) so that multiple processes can access them. EnterpriseOne uses a configurable IPC key range, which is controlled by the `startIPCKeyValue` in the JDE.INI file, in case a conflict occurs with other software that is using IPC resources.

jdenet_n Operation

The jdenet_n process usually starts when you run the supplied EnterpriseOne startup script: `RunOneWorld.sh`, which then starts all other processes as needed.

The jdenet_n process listens to the socket (port) as specified in the JDE.INI file by the keywords `ServiceNameListen` and `ServiceNameConnect`. These two keywords should be set to the same number; this number must be the same for every client who wishes to connect to the EnterpriseOne server.

The definitions for the particular jdenet_k processes to start are also given in the JDE.INI file. They are listed in the sections headed by `[JDENET_KERNEL_DEFx]`. Each of these entries lists the type of jdenet_k processes to start and the maximum number of jdenet_k processes of this type to start.

The number of jdenet_n slave processes to start is listed in the JDE.INI file under the keyword `maxNetProcesses`. The purposes of these slave processes are to provide parallel processing for the job of listening to the socket and to put the associated messages on the message queues for the jdenet_k processes to finish.

jdenet_k Operation

jdenet_k processes are referred to as kernel processes. They do the actual work on the enterprise server. When a jdenet_k process starts, it can be any type of kernel process. The jdenet_n process instructs each kernel process to be of a certain type.

The jdenet_k process that becomes a CallObject kernel has the job of calling business function logic on the server. Business function logic is written in C code and compiled into UNIX-shared libraries. The shared libraries are loaded onto the jdenet_k processes and then called directly through a C function call.

The jdenet_k process that becomes a batch process kernel waits for requests to run batch processes from the client. These batch processes are then placed in the Job Control Status Master table (F986110). The processes are then picked up by the queue kernel processes that launch runbatch processes, as required.

Many other types of jdenet_k processes exist. Review your JDE.INI file for a complete list.

Understanding EnterpriseOne Initialization for UNIX and Linux

The following initialization occurs when you start EnterpriseOne programs, such as the queue kernel, runbatch, and so on:

- The environment name is passed as a command line argument to the program (such as porttest, runbatch, and so on).
- This environment can be translated to a different environment, based on the settings in the [SERVER ENVIRONMENT MAP] section of the JDE.INI file.
- The environment that is used must be a valid entry in the Library List Master File table (F0094). Likewise, it must have a valid corresponding pathcode in the Environment Detail table (F00941).
- The following JDE.INI settings in the [DB SYSTEM SETTINGS] section are used to determine where the EnterpriseOne server startup tables, such as the Data Source Master (F98611) and the Object Configuration Master (F986101), are located:
 - Base Datasource
 - Object Owner
 - Server
 - Database
 - Load Library
 - Type
- Using the above information, the F986101 table in the specified database on the server is opened.
- When an override for a given table or the current user exists, that data source (the OMDATP field in the F986101 table) is used for the given object or user and environment. Otherwise, the data source in which OMOBNM=DEFAULT for the given environment is used. Ignore any inactive records (that is, OMSTSO=NA).

PeopleSoft strongly recommends that you do not have any default records for reports (OMOBNM=DEFAULT and OMFUNO=UBE) on the server. These records might prevent report interconnections (that is, one report calling another report) from starting correctly.

- Each unique data source in the F986101 table should correspond to one entry in the F98611 table.
- The corresponding information in the F98611 table must be correct. In particular, the OMDLLNAME field must show the correct library for the database to which the data source points.
- For an Oracle database, the OMDATB field from the F98611 table maps to an entry in the tnsnames.ora file. This tnsnames.ora file must be set up correctly. (Ask an Oracle database administrator to verify the setup).

Starting the Enterprise Server for UNIX or Linux

You can start the enterprise servers either manually at the command line or automatically when the server boots. The manual process is the same for all supported platforms, but the automatic process varies slightly by platform.

Note

If you are running EnterpriseOne on the same server as your Oracle database, you must make sure that Oracle is running before you start EnterpriseOne. In particular, if you are starting EnterpriseOne at system boot time, you must make sure the Oracle startup processes are completed first.

RunOneWorld.sh

RunOneWorld.sh is the script that starts the EnterpriseOne system on the enterprise server. This script does the following:

- Checks for existing EnterpriseOne processes. The script returns an error if it detects that EnterpriseOne is already running.
- Runs the rmics.sh script to clear IPC resources. This script ensures no IPC resources conflict with other software.
- Starts jdenet_n, which is the EnterpriseOne network listener that receives requests from EnterpriseOne workstations.
- Runs a program called cleanup that checks for unfinished batch processes from a previous shutdown.

► To manually start the enterprise server for UNIX or Linux

Note

This procedure is the same for all supported UNIX or Linux operating systems.

1. Sign on to the machine using the appropriate user ID, as set up during the installation process. If you used the PeopleSoft-recommended user ID, the user ID is jde.

2. Enter the following commands:

- `cd log_directory`

This command moves the user's current directory to the log directory. The administrator determines the name of this directory.

- `rm -f jde*log*`

This command deletes the log files in the directory.

Caution

Use extreme care when you enter this command. A syntax error in this command can cause severe problems on your system.

- `RunOneWorld.sh`

This script starts the EnterpriseOne system.

3. Sign off the system.

► **To automatically start the enterprise server for HP-UX**

1. Create a script named `psft` in `/sbin/init.d` with all necessary permissions for execution.

The script should contain only the following:

```
#!/sbin/sh
```

```
/bin/su - psft_user -c ` $SYSTEM/bin32/RunOneWorld.sh `
```

The value *psft_user* is the name of the user who owns the shell script `$SYSTEM/bin32/RunOneWorld.sh`. Make sure that no interactive commands appear in the *psft_user* profile, and that `RunOneWorld.sh` has all necessary permissions for execution.

2. Using the following command, create a soft link named `S995psft` to the `psft` script in the directory named `/sbin/rc2.d`.

```
ln -s /sbin/init.d/psft /sbin/rc2.d/S995psft
```

3. Verify that the following line is present in the profile of the user who owns `RunOneWorld.sh`:

```
/usr/local/bin/oraenv
```

Before you execute `oraenv`, ensure that the Oracle environment variables of `ORACLE_BASE`, `ORACLE_HOME`, `ORACLE_SID`, `ORACLE_TERM`, and `ORAENV_ASK` are properly assigned and exported. Also, you must add `$ORACLE_HOME/bin` to the `PATH` environment variable.

4. Set `ORACLE_TERM` to `hp`.

5. Set `ORAENV_ASK` to `NO`.

6. If the following command is in the profile, delete it:

```
unset ORAENV_ASK
```

► To automatically start the enterprise server for AIX and Solaris

1. Create a script named `rc.psft` in `/etc` with all necessary permissions for execution.

The script should contain only the following:

```
#!/bin/sh
```

```
/bin/su - psft_user -c ` $SYSTEM/bin32/RunOneWorld.sh '
```

The value `psft_user` is the name of the user who owns the shell script `$SYSTEM/bin32/RunOneWorld.sh`. Make sure there are no interactive commands in the `psft_user` profile, and that `RunOneWorld.sh` has all the necessary permissions for execution.

2. Add the following line at the end of the text file named `inittab` in `/etc`:

```
psft:2:wait:/etc/rc.psft
```

3. Verify that the following line is present in the profile of the user who owns `RunOneWorld.sh`.

```
. /usr/bin/oraenv
```

Before you execute `oraenv`, ensure that the Oracle environment variables of `ORACLE_BASE`, `ORACLE_HOME`, `ORACLE_SID`, `ORACLE_TERM`, and `ORAENV_ASK` are properly assigned and exported. Also, you must add `$ORACLE_HOME/bin` to the `PATH` environment variable.

4. Set `ORACLE_TERM` to `hp`.
5. Set `ORAENV_ASK` to `NO`.

To see a list of valid values for `ORACLE_SID`, look at the `oratab` text file in `/etc`.

6. If the following command is in the profile, you must delete it:

```
unset ORAENV_ASK
```

► To automatically start the enterprise server for Linux

1. Add the following line to the `rc.local` file in the `/etc` directory:

```
/bin/su - psft_user -c ` $SYSTEM/bin32/RunOneWorld.sh '
```

The value `psft_user` is the name of the user who owns the shell script `$SYSTEM/bin32/RunOneWorld.sh`. Make sure there are no interactive commands in the `psft_user` profile (or `.bash_profile`), and that `RunOneWorld.sh` has all the necessary permissions for execution.

► To verify the EnterpriseOne installation

After you start EnterpriseOne, execute the following commands:

```
cd $SYSTEM/bin32
```

`porttest userID password environment`

The porttest program initializes an environment, initializes a user, opens the Account Balances table (F0902), and displays up to 99 rows of data.

Note

The parameters for userID, password, and environment should be any valid EnterpriseOne user ID, password, or environment.

Database Settings

The default database parameters for UNIX might not fully support multiple users. You might reach the maxprocess limit for the database. The initial settings are for a small database, so you should change these parameters to a medium setting to avoid database problems. These settings reside in the init.ora file. The following path is an example of where you might typically find this file:

`/u01/app/oracle/product/8.0.5/dbs/init.ora`

Shutting Down the Enterprise Server for UNIX or Linux

The shutdown process is identical for all supported UNIX or Linux operating systems.

EndOneWorld.sh

EndOneWorld.sh is the script that stops the EnterpriseOne system on the enterprise server. This script completes the following functions:

- Checks for existing runbatch processes. If any runbatch (batch process) is running, the user is prompted to make sure that he or she wants to shut down the enterprise server.
- Checks for and ends EnterpriseOne processes other than jdenet_n and jdenet_k.
- Shuts down jdenet_n and jdenet_k processes by running endnet.
- Runs the rmics.sh script to clean up any remaining IPC resources.

► **To shut down the enterprise server for UNIX or Linux**

1. Sign on using the appropriate user ID that you set up during the installation process.
2. Execute the following commands:

`cd $SYSTEM/bin32`

`EndOneWorld.sh`

Setting Up a Printer for UNIX or Linux

Each supported UNIX system use different processes for setting up printers. HP-UX uses a tool called SAM to help in setting up a printer; AIX uses a tool called SMIT; Solaris uses a tool called Admintool; and RedHat Enterprise Linux AS uses a tool called printgui-conf. Each of these processes requires a privileged account to access the specific setup tasks. Normally, you will need to use the root account of the system. For more information about printer setup, see the appropriate HP-UX, AIX, or Solaris documentation.

See Also

- ❑ *Printing Reports* in the *System Administration Guide* for information specific to EnterpriseOne printer setup

Administering Batch Processes for UNIX or Linux

Administering batch processes involves knowing what processes run when EnterpriseOne starts, where files are placed before and after printing, and how to watch those processes.

Processes running for EnterpriseOne are owned by the user who started the EnterpriseOne software. The user ID for this user is set up during the installation of the software, and is site dependent. When EnterpriseOne starts, the following processes start and run under the environment and security of the user who started them:

jdenet_n The network listener that listens for connection requests.

jdenet_k The jdenet_n process starts the jdenet_k processes, which control EnterpriseOne components, such as the security server, the transaction monitor, and data replication.

► To monitor batch processes

From the operating system prompt, enter the following command, replacing *userid* with the user ID of the user who started EnterpriseOne:

```
jdejobs <userid>
```

If you omit the user ID, the current user is assumed.

jdejobs is a script in the EnterpriseOne \$SYSTEM/bin32 directory that uses the UNIX ps command to display job information.

Output from this command looks similar to the following:

```
pfst810 (EnterpriseOne Admin,,):  
Semaphores: 1   Shmem Segs: 5   Msg.Queues: 13  
Jobs on ent-1:  
  6137 ttyp6      0:43 jdenet_n  
  6163 ttyp6      0:44 jdenet_k  
  6188 ttyp6      0:44 jdenet_k  
  7213 ttyp6      2:12 jdenet_n  
  7241 ttyp6      0:47 jdenet_k  
  9008 ttyp6      1:36 jdenet_n  
  9009 ttyp6      0:45 jdenet_k  
 11042 ttyp6      0:09 runbatch
```

In the output, jdenet_n jobs are listening for requests, and four jdenet_k jobs are handling various EnterpriseOne kernel functions. A runbatch job is processing a report.

The first column of the output displays the UNIX process ID that is associated with each process. For more information about a particular process, look for the files in the log directory that have the same process ID as part of the file name.

Listing Batch Output Files

All output from each report, regardless of whether it is a preview, is placed in the PrintQueue directory under the installation directory of EnterpriseOne before printing. Depending on the JDE.INI settings for the workstation, the job might not be deleted after printing.

Jobs are printed to the location specified in the JDE.INI file unless an EnterpriseOne program overwrites them. Use the Printers program to specify default printers.

Two settings in the JDE.INI file for the workstation tell the server whether to print the report immediately upon completion, and whether to save the output from the report or delete it. These settings are as follows:

```
[NETWORK QUEUE SETTINGS]  
SaveOutput=TRUE  
PrintImmediate=TRUE
```

Setting SaveOutput to TRUE causes the JDE.INI to hold the jobs within the PrintQueue directory until the user explicitly deletes them. Setting PrintImmediate to TRUE tells the JDE.INI file to print the job immediately after completion of the report.

► To list batch output files

1. From the operating system prompt, enter the following command:
`cd $EVRHOME/PrintQueue`

This command changes the directory to the PrintQueue directory. The environment variable EVRHOME should be set to the EnterpriseOne installation directory.

2. Enter the following command to list the files:

```
ls
```

The output looks similar to the following:

```
R014021_XJDE0001_4554_PDF
R014021_XJDE0001_4554_PDF.jde.log
R014021_XJDE0001_4554_PDF.jdedebug.log
R31515_XJDE0001_4566_PDF
R31515_XJDE0001_4566_PDF.jde.log
R31515_XJDE0001_4566_PDF.jdedebug.log
R94NM08_XJDE0008_4568_PDF
R94NM08_XJDE0008_4568_PDF.jde.log
R94NM08_XJDE0008_4568_PDF.jdedebug.log
R94NM10_XJDE0016_4526_PDF
R94NM10_XJDE0016_4526_PDF.jde.log
R94NM10_XJDE0016_4526_PDF.jdedebug.log
R94NM10_XJDE0016_4526_PDF.ps
R94NM10_XJDE0016_4527_PDF
R94NM10_XJDE0016_4527_PDF.jde.log
R94NM10_XJDE0016_4527_PDF.jdedebug.log
R94NM10_XJDE0016_4527_PDF.pcl
```

The file names in this example are the actual reports that were generated when the job was executed. The file naming conventions are as follows:

R31515	The report name
XJDE00001	The report version executed
1914	The request number assigned by EnterpriseOne
PDF	A PDF file, meant for viewing on the workstation
.jde.log	The file extension that indicates the log file for the report
.jdedebug.log	The file extension that indicates the debug log for the report
.ps	The file extension that indicates a file formatted for

postscript printing

.pcl The file extension that indicates a file formatted for pcl printing

You should encourage workstation users to use the SaveOutput=FALSE entry in their jde.ini file. If users at workstations decide to save their output, they should periodically delete the entries through EnterpriseOne. When you delete .PDF files from the operating system, the corresponding EnterpriseOne print job entries in the Job Control Status Master table (F986110) are not deleted. You must manually delete these entries from EnterpriseOne using the Work with Servers program (P986116).

Running Reports from the Command Line for UNIX or Linux

You can initiate batch process reports from the server command line by issuing the following command (you must have the proper authority and the path equal to the description in the installation instructions):

```
runube UserName Password Environment Role ReportName VersionName JobQueue  
Interactive|Batch Print|Hold Save|Delete [PrintQueue]
```

For the command parameters, only the first character of the parameter name is required. The vertical bar symbol (|) indicates that you must specify one of the parameters on either side of the vertical bar. The brackets indicate an optional parameter. The following options apply to the runube command:

- Interactive** The system holds the current terminal session until the entire report is processed.
- Batch** The runube command starts a runbatch job and returns control of the terminal to the user.
- Print** The batch process report is spooled to the PrintQueue directory and then printed on the specified printer (OutQ). If you do not specify a printer, the system uses the default printer that you specified for the enterprise server in the Printers program.
- Hold** The system places the spool file in the PrintQueue directory for later printing at the user's request.
- Save** The system saves the file after printing.
- Delete** The system removes the file from the PrintQueue directory after the report prints.

Example: Running Reports from the Command Line for UNIX or Linux

The following example shows a command for executing a batch process report:

```
runube KL5952 mypass PROD *ALL R0006P XJDE0001 QBATCH I P D printer_1
```

Scheduling Reports from the Command Line for UNIX or Linux

You can schedule a report from the command line for processing on a future date, daily, or even a recurring day of the week. This task can be accomplished by using the operating system utilities called `at`, `batch`, and `cron`. The `batch` and `at` utilities are used to schedule single occurrence jobs; `cron` can be used to schedule recurring jobs. Use the `at` command or the `batch` command to schedule a job at a later time. The command line structure of these commands is identical, but you use them differently.

The `batch` command is intended to run a job immediately in the background, providing that the system load is low enough to handle the request. If the system load is not low enough, the job is held until system activity is low enough to handle the new request load.

The `at` command also runs jobs in the background, but allows you to schedule the job to run at a future time. You can use this utility to run the batch job during off-peak hours.

The command format for the `batch` command is as follows:

`batch command`

The command format for the `at` command is as follows:

`at -t CCYYMMDDHHMMSS command`

The `-t` switch is used to schedule the time. The following table describes the `CCYYMMDDHHMMSS` variable:

CC	Century (first two digits of the year)
YY	Year (last two digits of the year)
MM	Two-digit value of the month (such as 02 for February)
DD	The day of the month (01 - 31)
HH	The hour to start the job (00 - 23)
MM	The minute to start the job (00 - 59)
SS	The second to start the job (00 - 59)

command The command to run at the specified time. To schedule a report, use the `runube` command.

You can use the `cron` UNIX utility to run jobs at a scheduled time. You can specify variable times, such as once a year or once every hour. The operation of this utility is controlled by a table of events based upon each user.

Enter the following command to modify the `cron` schedule and edit the `cron` table for the current user:

`crontab -e`

The format of the `cron` table is as follows:

mm HH DD MM W command

The following table describes the variables for this command:

mm	The minute to run the job (00 - 59, or * for any minute)
HH	The hour to run the job (00 - 23, or * for any hour)
DD	The day of the month to run the job (0 - 31, or * for any day)
MM	The month to run the job (1 - 12, or * for any month)
W	The day of the week to run the job (0 - 6, with 0 being Sunday)

command The command to run at the specified time

After exiting the editor, the operating system should respond with a message stating that the crontab has been modified.

Example: Scheduling Single-Occurrence Reports from the UNIX or Linux Command Line

The following example shows a command line used to schedule a report to run at 06:00 on February 26, 2004:

```
at -t20040226060000 runube KL595218 mypass PROD *ALL R0006P XJDE0001
QBATCH Interactive Print Delete printer_1
```

Example: Scheduling Recurring Reports from the UNIX or Linux Command Line

This example shows a command line used to schedule a report to run at 06:00, any Sunday in the month of February (by the use of * for the day of the month and 0 for the day of the week).

```
00 06 * 02 0 runube KL5952 mypass PROD *ALL R0006P XJDE0001 QBATCH
Interactive Print Delete printer_1
```

Maintaining File Security for UNIX and Linux

Overall, only two accounts ever need operating system access to the EnterpriseOne environment files and version executables: the account that starts and stops EnterpriseOne, and the account that builds the environment SPEC and BSFN files. Normally, these accounts are the same.

Setting Specification File Security

Specification (SPEC) files are the first part of the environment files. You access these files by the EnterpriseOne kernel processes. These files should never be accessed directly by an operating system

user. Because of this, security on these files should be read/write for the user and role. They are not executables, so no reason exists for setting the executable option for any user, or role.

► To set specification file security

Add the following line to your .profile:

```
umask 022
```

This command sets the default file security for files that get created on the server. When a package build completes, SPEC files and business functions should be created with read permission for everyone, and with write permission for only the file owner. In general, both business function changes and SPEC file changes are controlled by the deployment server.

The security for your SPEC files should look similar to the following example:

```
-rw-r--r-- psft psft jdeblc.xdb
```

```
-rw-r--r-- psft psft jdeblc.ddb
```

Setting Business Function File Security

Business function security should be similar to SPEC file security. This allows the business function code to be viewed, but not modified directly on the server. In general, both business function changes and SPEC file changes are controlled by the deployment server.

► To set business function file security

1. Enter the following command in the BSFN Source directory:

```
chmod 644 *.c
```

2. Enter the following command in the BSFN Include directory:

```
chmod 644 *.h
```

The security for your BSFN files should look similar to the following example:

```
-rw-r--r-- psft psft b4200100.c
```

```
-rw-r--r-- psft psft b4200100.h
```

Setting Executables Security

You should prevent access to the EnterpriseOne executable files to prevent other users from attempting to start EnterpriseOne. Running the same version of EnterpriseOne on the same system and using the same JDE.INI settings can cause unpredictable results. In most cases, the second startup will fail, but giving users access to the shutdown procedures can allow them to shut down EnterpriseOne.

► **To set executables security**

Enter the following command:

```
cd $SYSTEM/bin32
```

```
chmod 540 *.sh
```

The access granted by this command allows all users in the EnterpriseOne role read-only permission to the files, but does not grant them execute privilege. You can omit read access if desired.

The security for your EnterpriseOne executables should look similar to the following example:

```
-r-xr----- psft psft RunOneWorld.sh
```

```
-r-xr----- psft psft EndOneWorld.sh
```

Setting jde.ini File Security

You must keep the jde.ini file as secure as possible. This file contains a database user name and password that enables EnterpriseOne security to function. This database account is given read authority to the EnterpriseOne Security table (F98OWSEC), which controls EnterpriseOne access.

Access to the F98OWSEC table, which contains privileged database user names and passwords, could give a user the ability to manipulate any data in the database, regardless of its sensitivity or security. Because of this, you should restrict access to the jde.ini file as much as possible.

► **To set JDE.INI file security**

1. Enter the following command:

```
cd $JDE_BASE
```

```
chmod 600 JDE.INI
```

This command sets maximum security for the JDE.INI file. The JDE_BASE environment variable is set to the directory that contains the JDE.INI file.

Note

The file name is case-sensitive.

The security for your JDE.INI file should look similar to the following:

```
-rw----- psft psft JDE.INI
```

Denying write access to the user psft is not strictly necessary, but can prevent accidental modification of JDE.INI settings, which could adversely affect the operation of EnterpriseOne.

2. If you want to deny the user write access, enter the following command:

```
chmod 400 JDE.INI
```

Because it is so important to keep access to the JDE.INI file as secure as possible, you should also limit the amount of access to the user psft (or the user account that starts and stops EnterpriseOne) to a minimum. Users with access to this account might obtain the user names and passwords in the F98OWSEC table, and, thus, gain privileged access to the database.

Understanding HP-UX and Solaris Kernel Parameter Settings

The kernels for HP-UX and Solaris include a long list of configurable parameters. These parameters control the quantity of various resources available within the HP-UX and Solaris kernels. Also, the EnterpriseOne server software, specifically the IPC facilities, is sensitive to numerous kernel parameters for operation. These parameters differ among the various vendor implementations of UNIX. To change the values of kernel parameters for HP-UX, you must use the System Administration Management (SAM) tool to modify the parameters, which might require rebooting the system. For Solaris, you must reboot the system after you modify kernel parameters in the /etc/system file. The proper values of these parameters depend on various criteria, such as number of users on the system, active applications, and the resource requirements for the active applications.

For HP-UX, you set kernel parameters with the SAM tool. To modify these parameters for Solaris, open the /etc/system file with the editor of your choice. You can set any given parameter to either a simple numerical value or an expression, based on the values of other parameters. The system administrator must set the kernel parameters. UNIX security refers to users with access to administrative functions as superusers.

When you first set up an HP-UX or a Solaris machine for EnterpriseOne, you should run SAM for HP-UX or an editor for Solaris, and change the kernel parameters. On an HP-UX system, you can see the current values of kernel parameters by running the kmtune command, or by running SAM. On a Solaris system, type the command sysdef -i to see the current kernel settings.

Note for Solaris

Some machines might require an additional parameter in the /etc/system file to activate messaging and semaphore processing. After you enter the sysdef -i command, if some parameters appear with 0 (zero) values, you might need to add one or more of the following parameters to the /etc/system file:

-
- forcload: sys/msgsys
 - forcload: sys/semsys
 - forcload: sys/shmsys

EnterpriseOne is not the only software to use the resources that the kernel parameters control. Therefore, for each parameter, the requirements for EnterpriseOne are either the minimum defaults provided with HP-UX and Solaris, in addition to the defaults provided with HP-UX and Solaris, or the requirements of other software installed on the system.

The following list provides the definitions of terms essential to the understanding of HP-UX and Solaris kernel parameters:

jdenet_n The maximum number of jdenet_n (net) processes that can be created for an instance EnterpriseOne server software running on the system. This is controlled by the maxNetProcesses parameter in the JDENET section of the JDE.INI file for each instance of EnterpriseOne.


jdenet_k The maximum number of jdenet_k (kernel) processes that can be created for an instance of EnterpriseOne server software running on the system. This is controlled by the maxKernelProcesses parameter in the JDENET section of the JDE.INI file for each instance of EnterpriseOne. Note that the maxNumberOfProcesses parameters in the JDENET_Kernel_Def sections do not matter here.

Note

The number of EnterpriseOne users that a machine serves, the number of instances of EnterpriseOne server software running on a machine, and the size of any databases on the machine are primary factors that affect the settings for HP-UX and Solaris kernel parameters. The number of jdenet_n, jdenet_k, and runbatch or runube processes running should reflect this information.

The following illustration provides an example of a Solaris editor that displays information for shared memory segments. The parameter name appears at the end of each line in the editor, such as shmmax at the end of the following line:

```
set shmsys:shminfo_shmmax=4294967295
```



Message Queues

Generally, the system clears queues quickly. If a problem arises, you can revise values for the following parameters to rectify the situation:

mesg This value must be 1. System-V style message queues are valid.

msgmni The value of msgmni represents the number of message queue identifiers. These identifiers determine the number of message queues that can exist throughout the system. In addition to the system default value and the requirements of other software, calculate what is needed for your EnterpriseOne installation (per EnterpriseOne instance). You can use the following equation to estimate the number of message queues necessary for EnterpriseOne:

$$1 + \text{jdenet\_n} + 2 \times \text{jdenet\_k} + (\text{max number of concurrent runbatch, runube, and runprint processes})$$

msgtql The value of msgtql represents the number of message headers. This number determines the total number of messages that can be in all the message queues at the same time. In addition to the requirements of other software, allow a value equal to $10 \times \text{msgmni}$ for the requirements of EnterpriseOne.

msgmap The value for msgmap represents the number of entries in the map of free message segments. The default value of $\text{msgtql} + 2$ should be used. If the value of msgmap is less than the value of $\text{msgtql} + 2$, attempts to create a message queue or to send a message might fail.

Note

This parameter is no longer used in Solaris 8.

msgmnb The value of msgmnb represents the maximum number of bytes that can reside on a single message queue at the same time. You should set the value for msgmnb at only a fraction of $\text{msgseg} \times \text{msgssz}$. For EnterpriseOne, a value of 32768 is reasonable. You can set a larger value as long as the product of $\text{msgseg} \times \text{msgssz}$ is large enough. The minimum value is 8192. Additional requirements of this parameter might increase the value of msgmnb.

msgmax The value of msgmax represents the maximum size, in bytes, of a single message. Do not set msgmax with a larger value than the value of msgmnb. The recommended setting is $\text{msgmax} = \text{msgmnb}$. The minimum value is 1024. Additional requirements of this parameter might increase the value of msgmax.

Inside the HP-UX and Solaris kernels (prior to Solaris 8), messages in message queues reside in message segments. The following parameters, which do not apply to Solaris 8, determine the size and number of segments available throughout the system:

msgssz The value of msgssz represents the size of each message segment in bytes. For EnterpriseOne, a value of 64 is adequate for most situations.

msgseg The value of msgseg represents the number of message segments throughout the system. In addition to the requirements of other software, allow a value equal to $50 \times$ the msgmni requirement for EnterpriseOne, or approximately 4096 per instance.

Semaphores

The following definitions apply to semaphores:

sema This value must be 1. System-V style message queues are valid.

semgni The value of semgni represents the maximum number of semaphore identifiers that can exist throughout the system.

For EnterpriseOne, two identifiers exist for each instance of EnterpriseOne, so the default value supplied with the HP-UX and Solaris systems should suffice.

semmap The value of semmap represents the number of entries in the map of free semaphores. The default value of semgni + 2 should suffice. If you decrease the value of semmap, attempts to create a semaphore set, which occurs during JDEIPC initialization, might fail.

Note

This parameter is not used in Solaris 8.

semnms The value of semnms represents the maximum number of semaphores that can exist throughout the system. Each instance of EnterpriseOne allocates 200 semaphores by default. However, you can customize this value in the JDE.INI file. In the [JDEIPC] section, modify the parameter `maxNumberOfSemaphores` to customize the number of semaphores that an instance of EnterpriseOne allocates.

For all releases of EnterpriseOne, the EnterpriseOne requirement is in addition to the requirements of other software. A good starting point for a typical EnterpriseOne installation (single instance) with Oracle should be 500.

semnmu The value of semnmu represents the maximum number of semaphore undo structures for the entire system. Effectively, this value is the maximum number of semaphores that the system can lock at the same time. For EnterpriseOne, allow one for each EnterpriseOne process that can exist for all installations of EnterpriseOne on the system. Use the following equation to determine this value:

$$1 + \text{jdenet\_n} + \text{jdenet\_k} + \text{maximum number of runbatch processes} + \text{maximum number of runprint processes} + \text{maximum number of runube processes}$$

Note

This equation is similar to the equation used to calculate the value for msgmni. If you will be running a large number of batch queues or print jobs, you might need to increase the value of this parameter.

The number of outstanding print requests at a given time, whether printing or waiting for a printer, determines the number of jdeprint processes. A reasonable estimate for the upper limit of this value is 10. However, this estimate is application-dependent. For example, a large warehouse that constantly prints pick slips might have more requests.

The number of batch processes that run directly on the server, not from a client, determine the number of runube processes. This value depends on the use of the system. Theoretically, this value has no limit.

- run** The value of semume represents the maximum number of semaphore undo structures per process. Effectively, this value is the maximum number of semaphores that a given process can lock at the same time. EnterpriseOne requires a minimum value of 4 for semume. This minimum value is not in addition to the system default and the requirements of other software. This value is a simple minimum. The default value provided with the system should suffice.
- semmsl** The value for semmsl, which applies to Solaris only, represents the maximum number of semaphores per unique identifier. For EnterpriseOne, this must be set equal to or higher than the maxNumberOfSemaphores setting in the JDE.INI file. For the default installation, you should set this parameter to 200.

Shared Memory

The following definitions apply to shared memory:

- shmem** The shmem value must be 1 to enable shared memory.
- shmmax** The value of shmmax represents the maximum size, in bytes, of a single shared memory segment. The default value provided with the system should suffice. Other software packages, such as Oracle, might require an increase in this value.
- shmmni** The value of shmmni represents the maximum number of shared memory segments throughout the system. For EnterpriseOne, allow 20 per instance of the EnterpriseOne server software running on the system. This requirement is in addition to the system default value and the requirements of other software.
- shmseg** The value of shmseg represents the maximum number of shared memory segments to which any one process can attach at a given moment. The default value provided with the system should suffice.

File Descriptors

The following definitions apply to file descriptors:

nfile The value of nfile represents the maximum number of open files, or sockets, throughout the system. The default value should be enough to handle most EnterpriseOne needs. However, you must make explicit allowance for the maximum number of sockets that jdenet_n processes can create to communicate with clients. This number is the sum of all sockets for all instances of EnterpriseOne server software that runs on the system. The maxNetConnections parameter in the [JDENET] section of each JDE.INI file indicates this sum. This requirement is in addition to the system default value and the requirements of other software.

maxfiles (rlim_fd_cur in the Solaris /etc/system file) The value of maxfiles represents the default soft limit on the number of file descriptors that any given process can have. A system call can raise the soft limit of a process as high as maxfiles_lim. For EnterpriseOne, the minimum value for maxfiles should equal at least the largest of all the maxNetConnections values in all the JDE.INI files in use + 10. This requirement is a minimum value, not a value in addition to the system default value and the requirements of other software.

Note

If this parameter is too small, EnterpriseOne might not open the log file to generate an error message.

maxfiles_lim (rlim_fd_max in the Solaris /etc/system file) The value of maxfiles_lim represents the hard limit of file descriptors that any given process can have. For EnterpriseOne, the minimum value for maxfiles should equal at least the largest of all the maxNetConnections values in all of the JDE.INI files in use + 10. This requirement is a minimum value, not a value in addition to the system default value and the requirements of other software.

Processes

The following definition applies to processes:

maxuprc The value of maxuprc represents the maximum number of processes that can run under a single user ID. This number is of particular concern on systems with either a very large EnterpriseOne installation or multiple instances running under the same user ID. You must allow for the total number of EnterpriseOne processes that might run at one time, plus other system processes that the EnterpriseOne user might be running.

Understanding Linux Kernel Parameter Settings

The Linux operating system uses many of the same kernel parameters as Solaris, but they are managed in a slightly different way. In the Linux 2.4 kernel, IPC parameters are defined and maintained in the /proc file system, in the directory /proc/sys/kernel. They can be modified dynamically by editing the appropriate file, but for enterprise applications, you should override the default parameters at boot time. In RedHat Enterprise Linux, the default parameters can be overridden

at boot time by adding entries to the `/etc/sysctl.conf` file. Use the command `ipcs -l` to view the current values for IPC resource limits.

IPC Resources

The following five entries in the `/etc/sysctl.conf` file affect EnterpriseOne IPC resources:

kernel.sem	This setting controls the following four different semaphore limits: • Maximum number of semaphores per array (semmsl on Solaris) • Maximum number of semaphores in the system (semmns) • Maximum operations per semop call (semopm) • Maximum number of semaphore arrays (semmni) For EnterpriseOne, you might need to increase the first value, semaphores per array, particularly if you increase the value of <code>maxNumberOfSemaphores</code> in the <code>jde.ini</code> file. Some database products also require that the fourth value, number of semaphore arrays, be increased from the default value.
kernel.shmmax	The default value for this parameter might be sufficient for Enterprise One, but some database products recommend that this be set to 256 Mb, or 90 percent of total memory, whichever is greater.
kernel.msgmax	This parameter defines the maximum size of a message. The recommendation for EnterpriseOne is 65535.
kernel.msgmnb	This parameter defines the maximum number of bytes on a message queue. The recommendation for EnterpriseOne is 65535.
kernel.msgmni	This parameter defines the maximum number of message queues (identifiers) in the system. You can use the following equation to estimate the number of message queues that are necessary for EnterpriseOne: $1 + jdenet\_n + 2 \times jdenet\_k + (\text{max number of concurrent runbatch, runube, and runprint processes})$

File Limits

In addition to the IPC resource limits, WebSphere and the EnterpriseOne HTML Server can require a large number of open files. To see your current values, review the file `/proc/sys/fs/file-nr`. This read-only file contains the following three values:

- Total allocated file handles
- Currently used file handles
- Maximum file handles

The first value represents a peak, so when this value approaches the maximum value, consider raising the limit. If the peak value reaches the limit, you will get unpredictable results because processes will not be able to open files. To change the maximum file handle limit, use the following setting:

fs.file-max

This controls the maximum number of files that can be simultaneously open throughout the entire system. The recommendation for EnterpriseOne is 32768, and this number might need to be increased to 65536 for larger installations.

Example: /etc/sysctl.conf

Below are lines from a typical sysctl.conf file that are used to set kernel parameters based on the previous information:

```
fs.file-max = 32768
kernel.shmmax = 268435456
kernel.sem = 500 32000 32 1024
kernel.msgmax = 65535
kernel.msgmnb = 65535
kernel.msgmni = 1024
```

Understanding AIX Kernel Parameter Settings for EnterpriseOne

AIX contains a set of kernel parameters (system parameters) that determine functionality and a separate set of performance parameters (tune parameters) that determine performance.

See Also

- ❑ The appropriate Oracle and IBM documentation for more information about AIX performance tuning

System Parameters

Setting the kernel parameters requires you to run the system management tool (SMIT). AIX has few configurable parameters that influence EnterpriseOne software; of those that influence EnterpriseOne, just one can cause the software to become inoperable. This parameter is maxuproc. The maxuproc parameter controls the number of processes that a single user can run simultaneously.

► To set the value of maxuproc

1. Sign on as the root user.
2. On the command line, enter the following command:

```
smit
```
3. In SMIT, choose the System Environments item and then choose the Change/Show Characteristics of Operating System item.
4. Change the value of Maximum number of processes to allow for all EnterpriseOne processes that might run at one time, plus any other system processes the EnterpriseOne user might be running.

Accept the default values for all other system parameters. The following lists these system parameters for general reference:

maxbuf	Max pages in block I/O buffer cache
maxmbu	Max real memory for Mbufs
autorestart	Automatically reboot after crash
iostat	Continuously maintain disk I/O history
maxpout	High water mark for pending write I/O per file
minpout	Low water mark for pending write I/O per file
keylock	State of system keylock at boot time
fullcore	Enable full core dump
pre43core	Use pre-430 style core dump (AIX 4.3 only)
logfilesize	Error log file size
memscrub	Enable memory scrubbing
dcache	Size of data cache in bytes
icache	Size of instruction cache in bytes
realmem	Size of usable physical memory
primary	Primary dump device
conslogin	System console login

► To view the system parameters

Enter the following command:

```
lsattr -E -lsys0
```

To change a system parameter, you must navigate to the correct SMIT menu option.

Tune Parameters

Setting the tune parameters requires you to run the following commands:

- For network parameters:

no

- For device parameters:
chdev
- For nfs parameters:
chnfs
- For general tuning parameters:
vmtune

Tune parameters can also be kept at their default values. Changes to tune parameters are generally needed only for performance reasons. Proper settings for optimal performance might vary with changes in the underlying database, hardware configuration, and EnterpriseOne configuration.

Performance tuning for AIX running EnterpriseOne or Oracle involves setting parameters that control virtual memory for paging, Raid, disk system types, and CPU scheduling.

Example: Disk Striping

Disk striping is the technique of spreading sequential data across multiple disk drives so data can be accessed in parallel from several drives at once. If striping is used, then the following tune parameters are set:

stripe size	64KB
max_coalesce	64KB
minpgahead	2
maxpgahead	16 x number of disk drives
maxfree	minfree + maxpgahead

Running Multiple Instances of the EnterpriseOne Enterprise Server

Common reasons for running multiple instances of the EnterpriseOne enterprise server are to test a new service pack or to upgrade to a new version of EnterpriseOne. You can run multiple instances of the EnterpriseOne server on the same machine by following a few simple guidelines.

Note

The following steps do not create a new database or any new database tables. Therefore, you will be using the same data tables that are used by the original instance of EnterpriseOne that was installed. If you want to create a completely separate set of database tables, follow the instructions for setting up a new environment. See *Adding an Environment* in the *Configurable Network Computing Implementation Guide*.

Prerequisite

- ❑ Verify that you have enough disk space to create copies of your current EnterpriseOne system directory and at least one path code directory.

► To run multiple instances of the EnterpriseOne enterprise server

1. The system administrator should create a new user ID that owns the new EnterpriseOne instance. Create the user ID using the appropriate administration tool, such as smit, SAM, admintool, or useradd.

Note

Although you can run multiple instances of the EnterpriseOne server using the same UNIX or Linux user ID, it is not recommended. The software depends on certain environment variables to function correctly, and these variables are easier to manage under different user IDs.

2. Sign on using the new user ID.
3. Copy the .profile and .psft files from the home directory of the original user ID to the home directory of the new user ID.
4. Change the .profile file for the new user ID, if appropriate.
5. Change the .psft file for the new user ID to reference the new directory path in which you will create the new EnterpriseOne instance.

For example:

Original .psft file:

```
export EVRHOME=/u01/PeopleSoft/E810
```

New .psft file:

```
export EVRHOME=/u02/PeopleSoft/E810
```

6. Create the directory in which the new EnterpriseOne instance will reside.

For example, type the following:

```
mkdir -p /u02/PeopleSoft/E810
```

7. Copy the system directory, the ini directory, and at least one path code directory from the original instance of EnterpriseOne to the new directory path.

The following sample commands accomplish this:

```
cp -R /u01/PeopleSoft/E810/system /u02/PeopleSoft/E810
cp -R /u01/PeopleSoft/E810/ini /u02/PeopleSoft/E810
cp -R /u01/PeopleSoft/E810/DV810 /u02/PeopleSoft/E810
```

Note

The path code directories for any environments that you intend to use for this second instance of EnterpriseOne must be copied to the new directory. You cannot share path code directories between two or more instances of EnterpriseOne, as this sharing might corrupt specification files.

8. Create an empty log directory under the new path using a command such as the following:

```
mkdir -p /u02/PeopleSoft/E810/log
```

9. In the new JDE.INI file, change all references to the original directory name to the new directory name, including the [INSTALL], [DEBUG], and [BSFN BUILD] sections.

For example:

```
[DEBUG]
DebugFile=/u02/PeopleSoft/E810/log/ jdedebug.log
JobFile=/u02/PeopleSoft/E810/log/jde.log
[INSTALL]
B9=/u02/PeopleSoft/E810
[BSFN BUILD]
BuildArea=/u02/PeopleSoft/E810/packages
```

10. Change the new JDE.INI file to reference a port number and starting IPC key that are different from the original EnterpriseOne instance.

These are defined by the following parameters; the numbers shown are examples only:

```
[JDENET]
serviceNameListen=6009
serviceNameConnect=6009
[JDEIPC]
startIPCKeyValue=9000
```

11. From the client workstation JDE.INI file, change the serviceName parameters to match those of the server JDE.INI file.

After you make all of the changes described in this chapter, you can start and stop the new EnterpriseOne instance independently of the original instance.

All existing EnterpriseOne environments will be valid for the new instance, provided that you have copied the corresponding path code directory for a given environment. All current logical data sources and OCM mappings will be recognized by the new instance.

Server Administration - Windows

PeopleSoft supports EnterpriseOne enterprise servers that run the Microsoft Windows Server. You can operate the enterprise server for Microsoft Windows in a logic or database server environment. You need to perform certain administration procedures on your enterprise server to ensure that the software runs properly.

EnterpriseOne Directory Structure for Windows


The following table lists the directories that are copied to the Windows enterprise server when EnterpriseOne is installed. They should be installed under the EnterpriseOne base directory (such as z:\PeopleSoft\E810\ddp). Indented names indicate subdirectories of the directories.

pathcode	The main directory for the business function shared libraries, C header files, object files, source files, and specification (spec or TAM) files. Upon installation, this directory will be copied to the correct path codes, such as PD810 and DV810. The following subdirectories are included: • bin32, which includes business function shared libraries. • spec, which includes specification files. These binary data files are in a PeopleSoft proprietary format.
system	The main directory for the system-level executables, shared libraries, C header files, libraries, and localization files. The following subdirectories are included: • bin32, which includes system-level executables and shared libraries. • include, which includes system-level C header files. • includev, which includes system-level C header files provided by third-party vendors such as Vertex. • lib, which includes system-level shared libraries and export files. • libv32, which includes system-level shared libraries provided by third-party vendors.
PrintQueue	The directory to which all .PDF file output for reports is written.
log	The directory to which jde_xxx.log and jdedbug_xxx.log files are written.
packages	The server package installation base directory. Directories exist here only if a package has been installed. Under the package directory are subdirectories named for each package that has been installed. Located under each package are the following subdirectories: • bin32, which includes business function shared libraries. • include, which includes business function header files. • obj, which includes business function object files. These are divided among lower-level subdirectories that correspond to each DLL in the bin32 directory. • source, which includes business function source files. These are divided among lower-level subdirectories that correspond to each DLL in the bin32 directory. • spec, which includes specification files. These binary data files are in a PeopleSoft proprietary format.

EnterpriseOne Architecture and Process Flow for Windows

The following host server processes perform the indicated actions.

Host Server Processes for Windows



All communications between the client and the host server occur using sockets. The communications between jdenet_n and jdenet_k occur with shared memory. jdenet_n and queue kernel communicate using the Job Control Status Master database table (F986110).

The following text explains the process flow:

- During Windows system startup, jdesnet runs automatically, provided that it is installed to start automatically. Otherwise, it must be started manually.
- The following information applies to the PeopleSoft network service:

- The program is system\bin32\jdesnet.exe.
 - Each time that a new server or workstation connects to this server, jdesnet might start another jdenet_n until the number of jdesnet and jdenet_n jobs equals the value in the maxNetProcesses field in the [JDENET] section of the JDE.INI file.
 - Each time that a new request, such as a batch application or CallObj is submitted, jdesnet (and any jdenet_n processes) might start another jdenet_k process until the number of jdenet_k jobs equals value in the maxKernelProcesses field in the [JDENET] section of the JDE.INI file.
 - Jdenet_n can be run manually by running system\bin32\jdenet_n.
4. The following information applies to the PeopleSoft queue service:
- The program is system\bin32\jdesque.exe.
 - The service runs the number of instances of queue kernels specified in the UBEQueues, PackageQueues, and SpecInstallQueues fields in the [NETWORK QUEUE SETTINGS] section of the JDE.INI.
- When a user submits a batch application, jdesnet or jdenet_n (as part of the host server) communicates with the client as follows:
 - The host server programs are system\bin32\jdesnet.exe and system\bin32\jdenet_n.exe.
 - The client environment is initialized.
 - The client tells the host server (using a socket) to initialize its environment.
 - The host server (for example, jdenet_n) initializes its environment and gets environment and user handles.
 - The host server passes the environment and user handles to the client (using a socket).
 - The client launches the batch application and then sends data to the host server (using a socket).
 - If the maximum number of kernel (for example, jdenet_k; the k stands for kernel) processes has not been met, jdesnet or jdenet_n might start a new jdenet_k process.
 - If the maximum number of jdenet_k processes has been met, jdesnet or jdenet_n puts the message in a queue for a jdenet_k process.
 - The client frees the user environment.
 - The client tells the host server (using a socket) to free the user environment for the server.
 - The host server frees its user environment.
 - The client tells the host server (using a socket) to free the environment for the server.
 - The host server frees its environment.
5. When the UBE Jdenet_k (the kernel) writes to the database (batch application only), the following occurs:
- The program is system\bin32\jdenet_k.exe.
 - Jdenet_k adds a record in the F986110 database table. The record has a status of W (Waiting).

- The Queue Kernel periodically checks the contents of table F986110 and launches a runbatch process
- When runbatch processes the batch application, the following occurs:
 - The program is system\bin32\runbatch.exe.
 - The system changes the status stored in table F986110 to P (Processing).
 - The system starts the batch application.
 - If the batch application completes successfully, it changes the status in table F986110 to D (Done).
 - If the batch application does not complete successfully, it changes the status in table F986110 to E (Error).
- Unlike the many processes that execute when a batch application is submitted, jdenet_k performs the processing when a user submits a CallObject and the following actions occur:
 - Cannot start the service name service on the enterprise server
 - Error 1069: The service did not start due to a logon failure

EnterpriseOne Initialization for Windows

The following initialization occurs when you start EnterpriseOne programs such as queue kernel, runbatch, and so on:

- The environment is passed as a command line argument to the program (such as porttest, queue kernel) or retrieved by jdenet_k from the QEnv key in the [NETWORK QUEUE SETTINGS] section of the JDE.INI file.
- This environment might be translated to a different environment, based on the settings in the [SERVER ENVIRONMENT MAP] section of the JDE.INI file.
- The environment that is used must be a valid entry in the Library ListMaster File table (F0094) and must have a valid corresponding path code in the Environment Detail - One World table (F00941).
- The following JDE.INI settings in the [DB SYSTEM SETTINGS] section specify where the EnterpriseOne server startup tables, such as the Data Source Master (F98611) and Object Configuration Master (F986101) tables, are located:
 - Base Datasource
 - Object Owner
 - Server
 - Database
 - Load Library
 - Type
- Using this information, the F986101 table opens in the specified database on the server.
- When an override exists for a given table, BSFN, or the current user, that data source (OMDATP field in the F986101 table) is used for the given object or user and environment.

Otherwise, the data source in which OMOBNM=DEFAULT for the given environment is used. Ignore any inactive records (that is, OMSTSO=NA). PeopleSoft *strongly* recommends that you do not have any default records (OMOBNM=DEFAULT) for batch applications (OMFUNO=UBE). These records might prevent report interconnections (such as one report calling another report) from starting correctly.

- Each unique data source in the F986101 table should correspond to one entry in the F98611 table.
- The corresponding information in the F98611 table must be correct. In particular, the OMDLLNAME field must display the correct DLL for the database to which the data source points.
- For an Oracle database, the OMDATB field from the F98611 table maps to an entry in the tnsnames.ora file. This tnsnames.ora file must be set up correctly (check with an Oracle database administrator).
- For a DB2 UDB database, the OMDATB field from the F98611 table maps to an entry in the ODBC data source. This datasource must be set up correctly (check with a DB2 UDB database administrator).
- For a Microsoft SQL Server, Microsoft Access, or Client Access database, the OMDATB field from the F98611 table maps to a data source specified in the ODBC Data Source Administrator applet in the Windows Control Panel. This data source must be set up correctly. If multiple users plan to sign on to this Windows platform and run EnterpriseOne or PORTTEST, the data sources must be defined on the System DSN tab. Otherwise, User Data Sources can be used.

If you are using Microsoft Windows 2000 to open the ODBC Data Source Administrator, from the Start menu, choose Programs, then Administrative Tools, and then Data Sources (ODBC).

- The following information pertains to the setup of different types of ODBC drivers, using the ODBC Data Source Administrator applet:

SQL Server:

- The data source name must match the name in the F98611 table.
- The description can be anything that you want.
- The server is the name of the database server.
- The network address includes the database server name, a comma, and a port in which the database user listens.
- Network Library should be set to Default.
- Click the Options button for more settings.
- The database name is usually set to JDE. You can set it to Default.
- The language name should be set to Default.
- The Generate Stored Procedure for Prepared Statement option should be turned off.
- The Use ANSI Quoted Identifiers option should be turned on.
- The Use ANSI Nulls, Padding and Warnings option should be turned on.
- The Convert OEM to ANSI characters option should be turned off.

Client Access:

- On the General tab the data source name must match the name in the F98611 table. The system is the name of the database server.
- On the Server tab, the default libraries should be the iSeries library, and the commit mode should be Commit immediate (*NONE).
- On the Format tab, the naming convention should be System naming convention (*SYS).
- On the Other tab, if the data that you are transferring using this data source contains a Binary Large Object (BLOB), translation should be set to Do not translate CCSID 65535. If the data that you are transferring using this data source does not contain a BLOB, translation should be set to Translate CCSID 65535.

Setting Up a Printer for Windows

Setting up a printer for a Microsoft Windows enterprise server involves setting up accounts under which EnterpriseOne runs, establishing printer ownership, and defining the printer. The default printer used for printing reports will be the system default printer.

Understanding Windows Services, Accounts, and Permissions

Before you can successfully set up a printer for Windows, you should understand the relationship of EnterpriseOne to Windows services, accounts, and permissions, which involves the following:

- Assigning permissions to the accounts under which EnterpriseOne services run
- Making printers accessible from the service programs
- Assigning ownership for accounts to enable access to printers

As shown in the following table, you must define two types of service accounts and printer ownerships for the two types of printers:

Local The service account type can be local or network. The printer owner account can be local or network.

Network The service account type must be network. The printer owner account must be network.

Windows services allow programs to run on a Windows platform even when no user is signed on to the machine. For the EnterpriseOne enterprise server, you must run the following two service programs:

- Network. This program provides the network connection between the EnterpriseOne workstation and the EnterpriseOne enterprise server.
- Queue. This program starts jobs (either batch reports or server package installations) on the enterprise server.

The accounts under which Windows services run must have permissions to start and stop services on the local machine. You must specify permissions for one of the following:

- Individual users, such as administrator and guest accounts
- Groups of users, such as administrators (note the plural; administrators are different than an individual administrator)

The accounts that automatically have permissions to start and stop services include:

- The Administrator user
- Users specifically designated by the Administrator user
- Users who belong to the Administrators group (which is different from an individual administrator)
- Users that belong to the Power Users group

Caution

PeopleSoft *strongly* recommends that you use an account for a user who belongs to the local Administrators group.

Adding a Printer

You must add a printer in Microsoft Windows before you can use it in EnterpriseOne.

► To add a printer

1. Click the Windows Start button.
2. Choose Settings, and then choose Printers.
3. Choose Add Printer.
4. On Add Printer Wizard, follow the system-guided steps.

For a local printer, these steps include selecting the port to which the printer is attached, specifying the type of printer that you are installing, specifying a name for the printer, and indicating where the drivers are located, if needed.

For a network printer, these steps involve selecting a print server and printer and indicating whether the printer is the default printer for the enterprise server.

Note

When you are defining a printer, do not use a space character in the name. If you do, EnterpriseOne will not be able to correctly read or access the physical printer.

Determining or Changing Printer Ownership

Every Windows printer is associated with one network account called the printer's owner. When EnterpriseOne runs a batch report, service programs must be able to access a printer. You can define this printer to be locally accessible only by the enterprise server or remotely accessible by other network resources (for example, it might be attached to a print server). You can specify a printer that is connected directly to an enterprise server as a local or network printer, depending on how you added the printer from the Control Panel.

The type of account that you use to run EnterpriseOne services depends on how you set up the printer, as follows:

- Local printer. Only local accounts can access the printer.
- Network printer. Only network accounts (that is, accounts that do not belong to the local domain) can access the printer. For example, the printer name is in the form `\\print_server\printer_name`.

► To determine or change printer ownership

1. From Control Panel, choose Printers.
2. Choose a printer, right-click, and choose Properties.
3. Click the Security tab.
4. Click the Ownership button.

The Owner dialog box displays the current owner of the printer.

5. On Owner, to make the account that you are currently signed onto the owner of the printer, choose Take Ownership, and then click OK.

Setting Up User Accounts on an Enterprise Server

You can set up local users to add local and network accounts to groups.

► To set up user accounts on an enterprise server

1. On the enterprise server, under Windows, choose Start, Settings, Control Panel, Administrative Tools, and then Computer Management.
2. On the Tree tab, choose Local Users and Groups, and then click the Users folder.

Changing the Domain

When you create a Windows user account, you must associate that account with one of the following two domains:

- Local. This domain is associated with a particular Windows machine. For example, each Windows machine has a local *administrator* account. Local accounts cannot access network resources, such as network printers. Any account names that do not begin with a domain name are considered to belong to the local domain.

- Network. This domain is spread across a Windows network. Users in the network domain can access network resources, such as printers and disk drives, on other servers. Account names that are assigned to the network domain must begin with a domain name, such as domain1\john_doe.

► To change the domain

1. From the main menu of User Manager, choose User.
2. Choose User Domain.

The Select Domain form displays all domains. The local domain is named the same as the enterprise server and does not appear in the list. However, you can still type the name of the enterprise server in the Domain field.

In this example, the name of the local machine is the same as the domain: DEVS5. That name is appears in the title bar as \\DEVS5. Although that syntax might typically indicate a network machine, in this case it represents a local machine name because the name of the machine and the domain are the same.

3. Click OK.

The User Manager form displays all of the accounts for the domain that you chose. If you choose a network domain, all listed names represent network accounts. Likewise, if you choose the local domain, all listed names represent local accounts.

Adding a Local Account

If you are using a local printer, you can use either a local or network account to run the EnterpriseOne services.

► To add a local account

1. Sign onto Windows as a user with administrative privileges in the local domain.
2. From Computer Management, choose System Tools, and then choose Local User and Groups.
3. From the Action menu, choose New User.
4. On New User, complete the following fields:
 - User name
 - Full Name
 - Description
 - Password
 - Confirm Password
5. Complete the following options, as appropriate for your installation:
 - User must change password at next logon
 - User cannot change password
 - Password never expires
 - Account disabled

6. Click Create.
7. Click Cancel.


Adding a User to the Administrators Group

To add an existing account (either local or network), you must use the local domain.


► To add a user to the Administrators group

1. From the User Manager main window, double-click on the Administrators group.

The user Administrator belongs to the Administrators group. Local accounts are not preceded by a domain name, and network accounts are preceded by a domain name. For example, the domain member with a name JDE is a local account, and a member with the name JDEMD1\AY5600427 is a network account.



2. On Administrators Properties, click Add.
A list displays all users in the selected domain.



3. On Select Users or Groups, choose the domain of the user whom you want to add to the Administrators group.
4. Choose the user whom you want to add to the Administrators group.
5. Click Add to add the user to the group, and then click OK.

Working with Network Services

EnterpriseOne uses the Network service on the enterprise server. This service is installed during the installation process using the `jdesnet -i` service from the `system\bin32` directory.

When you install this service, the system adds the following entries to the Windows registry:

- The name of the service that appears on the Services form (used when controlling the services)
- The location of the EnterpriseOne executable files

After the initial installation, you will need to reinstall the Network service only when it has been uninstalled. You will need to uninstall this service only when the EnterpriseOne directory tree is renamed, moved, or deleted.

After the Network service is installed, you must set up the service under a network account, if you are using a network printer, or a local account, if you are using a local printer. If you are using a network account, it must be in either the Administrators or Power Users group.

Caution

We *strongly* recommend that you use a user who belongs to the local Administrators group.

Setting Up the Network Service

Set up the Network service under a network account, if you are using a network printer, or a local account, if you are using a local printer. Whether you use a network or a local account, it must be in either the Administrators or Power Users group.

► To set up the network service

1. From the Start menu, choose Programs, Administrative Tools, and then Services.
2. Choose the EnterpriseOne Network service.
The name of the service is in the form JDE *release* Network, where *release* is the current EnterpriseOne release. For example, the Network service's name for Release E810 is JDE 810 Network.
3. Click Action, and then click Properties.
4. On the General tab, if you want EnterpriseOne to start automatically when the enterprise server boots, click the Automatic option under Startup Type.
5. On the Log On tab, click the This Account option.
6. Enter the account name under which the EnterpriseOne Network service will run.
7. Enter the password for the account and a confirmation of the password.
8. Click OK.

Starting the Network Service

After you have installed and set up the Network service, you must start the service before EnterpriseOne can use it.

► To start the Network service

1. From the Services window, choose the EnterpriseOne Network service.
The name of the service is in the form JDE *release* Network, where *release* is the current EnterpriseOne release. For example, the Network services name for E810 is JDE 810 Network.
2. From the Action menu, click Start.
3. Use the Windows Task Manager to ensure that the following processes are running:
 - jdesnet.exe.
 - jdenet_k.exe processes. (None, one, or more might exist.)

Stopping the Network Services

When you stop the Network service, follow the steps in the proper sequence.

► To stop the Network service

1. From the Services window, choose the Network service.
The name of the EnterpriseOne Network service is in the form JDE *release* Network. For example, the Network services name for EnterpriseOne 8.10 is JDE 810 Network.
2. Use the Windows Task Manager to ensure that all EnterpriseOne processes are terminated.
This might take several minutes. The following processes should be terminated and, therefore, should not appear in the list of processes in Task Manager:
 - jdesnet.exe
 - jdenet_n.exe
 - jdenet_k.exe
 - runbatch.exe
 - ipcsrv.exe

See Also

- *Cleaning Up the Enterprise Server for Windows in the Server and Workstation Administration Guide* for more information about terminating EnterpriseOne processes

Cleaning Up the Enterprise Server for Windows

After EnterpriseOne is shut down, you can determine whether any processes completed abnormally. If so, you need to clean up the enterprise server. Unforeseen circumstances can cause EnterpriseOne processes to terminate abnormally. Processes that terminate abnormally are called *runaway processes*. After shutting down EnterpriseOne, look for any runaway processes and, if any exist, manually terminate them.

► To clean up the enterprise server for Windows

1. In the Processes tab of Task Manager, search for any EnterpriseOne Host Server processes, such as jdesnet, jdenet_n, jdenet_k, and runbatch.
Wait until all the EnterpriseOne Host Server processes are terminated. If all processes terminate, you do not need to perform the remaining steps in this task. Otherwise, continue with the next step.
2. Choose a process in Task Manager.
3. Click End Process.
4. If the runaway process does not terminate, continue with the next step.
5. In Task Manager, right-click on the process and select debug.

6. When the Visual C++ main window appears, choose the Stop debugging option from the Debug menu.
7. Exit from Visual C++, and then repeat these steps for each runaway process.
8. If none of the previous steps stops the runaway process, reboot the enterprise server.

Uninstalling the Network Service

As an administrator, you must uninstall the Network service for EnterpriseOne before you rename, move, or delete the EnterpriseOne directory tree. The process to uninstall this service removes the following entries from the Windows registry:

- The names that appear for the service on the Services form
- The location of the EnterpriseOne executable files

During a new installation, or after you have renamed or moved the directory tree for an existing installation, you should reinstall the services.

See Also

- ❑ *Starting the Network Service* in the *Server and Workstation Administration Guide* for more information about the Services form

► To uninstall the Network services

Run the following program from the \system\bin32 directory:

- jdesnet -u

Manually Starting the Enterprise Server for Windows

If EnterpriseOne does not run through the Control Panel Services applet, you can run Network manually.

Note

If you start EnterpriseOne manually, you must stop the EnterpriseOne processes using the Windows Task Manager.

See Also

- ❑ *Setting Up the Network Service* in the *Server and Workstation Administration Guide* for instructions on automatically starting EnterpriseOne

► To manually start the enterprise server for Windows

1. On the enterprise server for Windows, sign on with administrator privileges.
If you used the user ID that PeopleSoft recommends, the value is PSFT.

2. On the Windows toolbar, from the Start menu, choose Run, and then enter the following commands:

drive: installpath\system\bin32\jdenet_n

Where *installpath* is the path to the EnterpriseOne installation.

This command launches an executable program that starts the EnterpriseOne network (JDENet) internal processes.

If you run *jdenet_n* from a command prompt, ensure that the working directory is the subdirectory *\system\bin32*.

Verifying the EnterpriseOne Installation

You can verify your EnterpriseOne installation with the PORTTEST program.

Note

When you run PORTTEST, make sure that one of the following is true:

- If the network service, such as *jdesnet.exe*, is running, make sure that you are signed on to Windows under the same user account as the net service is running. You can then run PORTTEST from a command prompt.
 - If the network process, such as *jdenet_n.exe*, is run from the command prompt, you can run PORTTEST from the command prompt.
-

► To verify the EnterpriseOne installation

In the command line, enter the following commands:

```
cd \PeopleSoft\E810\ddp\system\bin32
```

```
porttest <userid> <password> <environment>
```

The program initializes an environment, initializes a user, opens the Account Balances table (F0902), and displays up to 99 rows of data. The number of rows of data that the program displays depends on the data in the table. If you run the program before anyone enters data into the table, you will not see any data on the screen. In this case, the lack of data does not indicate an error. Review the messages on the form and the corresponding *jde.log* file to determine the results of the program.

Administering Batch Processes for Windows

Administering batch processes involves knowing the processes that run when EnterpriseOne starts, where files are placed before and after printing, and how to watch those processes.

The user who started the EnterpriseOne software owns the processes that are running for EnterpriseOne; Windows Task Manager cannot track this information. When the software starts, a number of processes start and run under the environment and security of the user who started them. These processes are as follows:

- jdesnet.exe** The network listener that listens for connection requests.
- jdenet_n.exe** A network listener that listens for connection requests. Depending on the jde.ini setting, zero, one, or more of these processes can run simultaneously.
- jdenet_k.exe** The job responsible for coordination between the net and queues. It is not started until the first batch job is submitted to the server.
- runbatch.exe** The job responsible for executing the submitted reports.
- ipcsrv.exe** The process responsible for passing Binary Large Objects (BLOBs) between other processes.

Monitoring Batch Processes

You can use the Task Manager to continuously monitor the performance of each job, the amount of CPU time it is consuming, and the amount of memory it is using. By default, the display refreshes every second.

Reviewing Batch Output Files

All output from each report, regardless of whether it is a preview, is placed in the PrintQueue directory under the EnterpriseOne installation directory before it is printed. Depending on the JDE.INI settings of the workstation that submitted the job, the job might or might not be deleted after being printed. Unless the submitter identified a printer, jobs are printed to the default printer that you specified for the enterprise server.

Two settings, based upon the workstation's JDE.INI file, tell the server whether to print the report immediately upon completion and whether to save the output from the report or delete it. Both of these workstation settings are shown below:

<pre>[NETWORK QUEUE SETTINGS] SaveOutput=TRUE PrintImmediate=TRUE</pre>

Setting SaveOutput to TRUE causes the enterprise server to hold the jobs within the PrintQueue directory until the user explicitly deletes them. Setting PrintImmediate to TRUE tells the enterprise server to print the job immediately after completion of the report.

Users should be strongly encouraged to use the SaveOutput=FALSE entry in their JDE.INI file. When users decide to save their output, they should periodically delete the entries through EnterpriseOne. Deleting the output files from the operating system will not delete the corresponding EnterpriseOne print job entries (for example, entries might still exist in the database). These print job entries still have to be deleted manually.

To list all files in the PrintQueue directory, use Windows Explorer to change your working directory to the PrintQueue directory.

These file names are the actual reports that were generated when the job was executed. The file names follow these conventions:

S_	Identifies the first part of a file name. Indicates that a specification installation was performed by the workstation. The system omits this prefix when no specification installation was performed.
R0006P	Identifies the report name.
XJDE0001	Identifies the report version.
UBE	Identifies the type of request.
216	Identifies the request number assigned by EnterpriseOne.
PS	Indicates a PostScript file.
PDF	Indicates a PDF (Portable Document Format) file. This file can be viewed on the workstation using Adobe Acrobat.

Running Reports from the Command Line for Windows

If you are a user with the proper authority and path (equal to that described in the installation instructions), you can run batch report processes from the server command line by first changing to the EnterpriseOne system directory (system\bin32) and then entering the following commands:

```
runube UserName Password Environment ReportName VersionName JobQueue Interactive|Batch  
Print|Hold Save|Delete [OutQ]
```

For the command parameters Interactive, Batch, Print, Hold, Save, and Delete, only the first character of the parameter name is required. The vertical bar symbol (|) indicates that you must specify one of the parameters on either side of the vertical bar. The bracket symbols ([and]) indicate an optional parameter. The following options apply to the runube command:

Interactive	Runube processes the report and outputs a .pdf file. The queuing mechanism is skipped altogether.
Batch	The runube command starts a runbatch job and returns control of the terminal to the user. Runube writes a record to the Job Control Status Master table (F986110) and sends a message to the Queue Kernel, which launched Runbatch. Runbatch processes the report and outputs a .pdf file.
Print	The batch process report spools to the PrintQueue directory and then prints on the specified printer (OutQ). If you do not specify a printer, the system uses the default printer that you have specified for the enterprise server.
Hold	The system places the spool file in the PrintQueue directory for later printing at the user's request.
Save	The system saves the file after printing. If you specify Delete, the system removes the file from the PrintQueue directory after it is printed.
Delete	The system removes the file from the PrintQueue directory after the report prints.
OutQ	Optional. This is the printer name on which the given report is printed. If this option is not specified, the report will be printed on the enterprise server's default printer.

Example: Running Reports from the Command Line for Windows

This example shows commands for executing a batch process report:

```
cd \PeopleSoft\E810\ddp\system\bin32
runube KL5595218 KL5595218 PROD R0006P XJDE0001 QBATCH Interactive
Print Delete printer_1
```

Scheduling Reports from the Command Line for Windows

You can schedule a report from the command line for processing on a future date, daily, or even on a recurring day of the week. To schedule one-time only reports, use the at command.

When you issue jobs with the at command, they run in the background. However, the at command is designed to allow you to schedule a future time of execution. You can use this command to run a batch job during off-peak hours.

Note

Use of the at command depends on how security is configured on the Windows enterprise server. You should limit the amount of access that users have to submit jobs on the server. If possible, only an administrator should do this type of scheduling.

The command format for the at command is as follows:

```
at [\computername\ time [/INTERACTIVE] [/EVERY:date[,...]] | /NEXT:date[,...]] command
```

Where the following options apply:

\\computername	Identifies the computer on which to run the program. If you do not specify a value, the default is the local machine.
time	Specifies the time to run the job, such as 08:00.
/Windows INTERACTIVE	Allows the program to interact with the Windows operating system desktop.
/EVERY:date	Specifies the days on which to run the job. Valid values are M, T, W, Th, F, S, and Su.
/NEXT:date	Specifies the next date for the first execution. If you do not specify a value, the default value is today's date.
command	Specifies the command to run. To run batch jobs here, use the runube command with any of its parameters.

Example: Scheduling Reports from the Command Line for Windows

This example shows a sample at command that you can use to schedule a EnterpriseOne batch report to run on the DEPLOY machine at 06:00 every Sunday:

```
at \\DEPLOY 06:00 /EVERY:Su z:\b731\system\bin32\runube KL5595218
KL5595218 PROD R0006P XJDE0001 QBATCH Interactive Print Delete
printer_1
```

Understanding JDE.INI Settings for Starting Batch Queues on Windows

The following JDE.INI settings are used to start batch queues on the Windows enterprise server:

```
[NETWORK QUEUE SETTINGS]
UBEQueues=number of batch queues
UBEQueue1=batch queue name
UBEQueue2=batch queue name
PackageQueues=number of package queues
PkgQueue1=package queue name
PkgQueue2=package queue name
SpecInstallQueues=number of spec install queues
SpcQueue1=spec install queue name
QEnv=queue environment
QUser=queue user
QPassword=queue user password
```

The following table describes each setting:

number of batch queues	Identifies the number of batch queues available. If you do not specify a number of batch queues that matches the number specified here, EnterpriseOne uses QBATCH when a missing queue is called.
batch queue name	Identifies the name of the batch queue. For example, for UBEQueue2, you might specify the queue as QBATCH2. You should specify a number of batch queue names that is equal to the value that you specify for the number of batch queues.
number of package queues	Identifies the number of package queues that are available. If you do not specify a number of package queues that matches the number specified here, EnterpriseOne uses QBATCH when a missing queue is called.
package queue name	Identifies the name of the package queue. For example, for PkgQueue2, you might specify the queue as XBATCH2. You should specify a number of package queue names that is equal to the value that you specify for the number of package queues.
number of spec install queues	Identifies the number of specification install queues available. If you do not specify a number of specification install queues that matches the number specified here, EnterpriseOne uses QBATCH when a missing queue is called.
spec install queue name	Identifies the name of the specification install queue. For example, for PkgQueue2, you might specify the queue as XBATCH2. You should specify a number of specification install queue names equal to the value that you specify for the number of specification install queues.
queue environment	Identifies the EnterpriseOne environment under which the Windows operating system starts the queues.
queue user	Identifies a valid EnterpriseOne user.
queue user password	Identifies the password for the queue user.

Using Active Directory

Windows Active Directory is Microsoft's implementation of a hierarchical, object-based directory service for managing system resources, including developers, end users, and groups. If you publish EnterpriseOne server information in Active Directory, client workstations use this information to locate and connect to the server dynamically. If EnterpriseOne service changes from one server to another, workstations can still connect to the server by referencing published server information in Active Directory.

Note

Active Directory is a Windows feature, and its use with EnterpriseOne is platform-specific and optional. If you are running EnterpriseOne enterprise servers on Unix or iSeries platforms, client workstations still reference their jde.ini files to connect to the server.

SCP Object in Active Directory

EnterpriseOne NT service installation creates a Service Connection Point (SCP) object in Active Directory. The SCP object specifies the following information:

- The server name
- The port number

Starting EnterpriseOne service on a server automatically updates the SCP object with the server name and port number, and establishes the SCP object status as Running. When service stops, the status of the SCP object automatically changes to Stopped.

Note

EnterpriseOne Windows service installation creates the SCP object in Active Directory only if you have added an [Active Directory] section to the jde.ini file on the server *before* installation.

When a user signs on to EnterpriseOne, EnterpriseOne searches Active Directory for an SCP object with a service name that matches the parameter value in the [Active Directory] section of the workstation jde.ini file. EnterpriseOne chooses an SCP object that has a status of Running and retrieves the server name and port number, which enables the workstation to make a connection to the server.

See Also

- ❑ *Windows-Based Systems* in the *Windows EnterpriseOne Install Guide* for details about Windows installation
- ❑ *Additions to the Server jde.ini file* in the *Server and Workstation Administration Guide* for details about adding an [Active Directory] parameter value to the server jde.ini file

Additions to the Server JDE.INI file

For each server that you publish in Active Directory, you must add an [Active Directory] section in the JDE.INI file on the server. In the [Active Directory] section, you include the SCPToPublish entry, which identifies the SCP object in the Active Directory.

The value of the SCPToPublish parameter should be unique for each object, and you should consistently adhere to a naming convention for ease of administration. For example, the value of each SCPToPublish parameter might represent a version of EnterpriseOne.

The following is a sample entry in the [Active Directory] section of the server JDE.INI file.

SCPToPublish PEOPLESFT_ENTERPRISEONE_810_SP1

If you move EnterpriseOne service from one server to another or change the service port number, no changes to the workstation JDE.INI file are needed, so long as the name of the SCP object in Active Directory and the parameter values of the [Active Directory] section of the workstation JDE.INI file match.

Note

Although users can automatically connect to a new server when a change in service is made, batch processes and business functions are not automatically mapped to the new server. Therefore, you typically need to change OCM mappings for the users so that they use the new data source.

Additions to the Workstation JDE.INI File

You also add an [Active Directory] section to the workstation JDE.INI file that specifies the name of the SCP object that contains port number and server name information.

The following parameters are included in the [ActiveDirectory] section of the workstation JDE.INI file:

- JdenetSCP (the connection port)
- SecurityServerSCP (the security server)
- LockManagerSCP (the Lock Manager)
- UnifiedLogonServerSCP (unified logon server)

For each of these parameters, you assign as the value the name of the SCP object in the Active Directory file. For example, enter PEOPLESOFT_ENTERPRISEONE_810_SP1.

The following table presents an example of the parameters that you add to the [Active Directory] section of the workstation JDE.INI file. The value of each parameter is the SCP object name in Active Directory.

Parameter of [Active Directory] Section of Workstation JDE.INI File	Meaning	Parameter Value: name of SCP Object in Active Directory
JdenetSCP	Connection port	PEOPLESOFT_ENTERPRISEONE_810_SP1
SecurityServerSCP	Security server	PEOPLESOFT_ENTERPRISEONE_810_SP1
LockManagerSCP	Lock manager	PEOPLESOFT_ENTERPRISEONE_810_SP1
UnifiedLogonServerSCP	Unified logon server	PEOPLESOFT_ENTERPRISEONE_810_SP1

Maintaining File Security for Windows

You should be aware of the security that is set up for the files on an EnterpriseOne enterprise server. System-wide, only the following two accounts will ever need operating system access to the EnterpriseOne environment files and version executables:

- The account that starts and stops EnterpriseOne
- The account that builds the environment specification (SPEC) and business function (BSFN) files (if this account is separate from the startup and shutdown account)

Specification File Security

Specification files are the first part of the environment files. You access these files using the EnterpriseOne kernel processes. These files should never be accessed directly by an operating system user; therefore, security for these files should be read/write for the user and group. These files are not executables, so you do not need to set the executable option for any user, group, or other.

Business Function File Security

You should keep business functions secure. In an environment in which development takes place, you must have a strict form of version control on source and object files. If the business function files change without the knowledge of the EnterpriseOne administrators, rebuilding them might produce unknown or undesired results. Most likely, a developer is working to correct a problem, but the problem could become worse.

You should set a high level of security on the source, include, and object files.

EnterpriseOne Executables Security

You should prevent access to EnterpriseOne executable files to prevent other users from attempting to start up EnterpriseOne. Running the same version of EnterpriseOne on the same system, using the same JDE.INI settings, can cause unpredictable results. In most cases, the second startup will fail, but giving users access to the shutdown procedures allows them to shut down EnterpriseOne.

JDE.INI File (Enterprise Server) Security

You must keep the JDE.INI file on the Windows enterprise server as secure as possible. This file contains a database user name and password that enables EnterpriseOne security to function. This database account is given read authority to the OneWorld Security table (F98OWSEC), which controls EnterpriseOne access.

Note

The F98OWSEC table contains privileged database user names and passwords, which could give a user the ability to manipulate any data in the database, regardless of its sensitivity or security. Therefore, access to the enterprise server JDE.INI file should be minimized.

Denying written access to EnterpriseOne is not necessary, but prevents accidental modification of JDE.INI settings that could adversely affect the operation of EnterpriseOne.

Because of the importance of limiting access to the JDE.INI file for security reasons, you also should limit access to the EnterpriseOne account (or the user account that starts and stops EnterpriseOne). Users with access to this account can easily obtain the F98OWSEC user names and passwords, and gain privileged access to the database.

Running Multiple Instances of EnterpriseOne on Windows

You can run multiple instances of EnterpriseOne on a Windows 2000 server. You might do so to test a new service or to upgrade to a new version of EnterpriseOne. You do not need to install a separate machine to run multiple instances of EnterpriseOne, so long as you follow a series of recommended steps.

Prerequisites

- ❑ Verify that you have enough disk space to create copies of your current EnterpriseOne system directory and at least one path code directory.
- ❑ Verify that you install each new instance of EnterpriseOne in a separate directory tree and that the version-level directories are different. For example, EnterpriseOne version 1 might be installed in the `z:\PeopleSoft\b9` directory tree, while EnterpriseOne version 2 might be installed in the `z:\PeopleSoft\E810` directory tree.

Generating a Unique Identifier

Each instance of EnterpriseOne must have a unique identifier. You set the value of this identifier in the CLSID parameter of the server JDE.INI file. To generate the identifier, you run the uuidgen program.

► To generate a unique identifier

1. From the Start menu on the Windows taskbar, choose Run, and then enter the following command:

```
uuidgen-oFILENAME
```

Where *FILENAME* is the name of the file that will contain the new identifier.

Note

For help about the options for the uuidgen program, run the following command:

```
uuidgen-?
```

The uuidgen program creates a unique identifier and stores it in the file that you specified.

2. Copy the identifier.
3. Open the server JDE.INI file and paste the identifier into the CSLID parameter under the [JDEIPC] section of the file.

Modifying the Server JDE.INI Files

For each new instance of EnterpriseOne, you modify the values of parameters in the JDE.INI file on the server. Each value for each EnterpriseOne instance must be unique.

The following table presents the server jde.ini file parameters that require modification, the purpose of each, and example values for each:

Section of server JDE.INI file	Parameter	Purpose	Example Value
[DEBUG]	DebugFile=	Name of the log file that contains debugging data.	z:\PeopleSoft\E810_2\log\jdedebug.log
[DEBUG]	JobFile=	Name of the log file that contains log data.	z:\PeopleSoft\E810_2\log\jde.log
[INSTALL]	StartServicePrefix=	Prefix that is used for names of the EnterpriseOne network and queue services.	Instance 2
[INSTALL]	B9=	Base directory of the EnterpriseOne installation.	z:\PeopleSoft\E810_2
[JDEIPC]	StartIPCKeyValue=	Integer that indicates an arbitrary starting point in memory for interprocess communications. For multiple instances of EnterpriseOne, differences between the values of the parameter must be at least 1000.	6000
[JDEIPC]	CLSID=	Unique string generated by the NT guidgen program. The string identifies each instance of EnterpriseOne.	1E0CF350-AF81-11D0-BD7B-0000F6540786
[JDENET]	serviceNameListen=	The TCP/IP port number used by the server to receive communication packets from	6005

		workstations.	
[JDENET]	serviceNameConnect=	The TCP/IP port number used by the server to send communications packets to servers.	6005

► To modify the server JDE.INI file

1. In the system\bin32 subdirectory for each new EnterpriseOne instance, open the server JDE.INI file.
2. In the [DEBUG] section of the JDE.INI file, modify the values of the following parameters:
 - DebugFile=
This is the name of the log file that will contain debugging information.
 - JobFile=
This is the name of the file that will contain log information.
3. In the [INSTALL] section of the JDE.INI file, modify the values of the following parameters:
 - StartServicePrefix=
The value of this parameter is used for the names of the EnterpriseOne network and queue services. The names are listed in the Services window under Control Panel.

The default value is JDE followed by the current version number, such as 810. The default value produces the service names JDE 810 Network and JDE 810 Queue.
 - B9=
The value of this parameter represents the base directory of the EnterpriseOne installation. The EnterpriseOne server uses this value to determine the location of the executables and DLLs used to run EnterpriseOne programs.
4. In the [JDEIPC] section of the JDE.INI file, modify the values of the following parameters:
 - StartIPCKeyValue
The value of this parameter indicates a starting point in memory for interprocess communications. For multiple instances of EnterpriseOne, verify that the difference between starting point values for each instance is at least 1000. The default value is 5000.

Caution

To ensure that the difference between starting point values is at least 1000, review the maxNumberofResources parameter in the [JDEIPC] section of the JDE.INI file. If the parameter value is less than 1000, change the value.

- CLSID=

The value of this parameter is a unique string that is generated by the NT guidgen program.

5. In the [JDENET] section of the JDE.INI file, modify the values of the following parameters:

- serviceNameListen=

The value of this parameter specifies the TCP/IP port used by the server to receive communications packets from the workstations. Each instance of EnterpriseOne must communicate with workstations through a different port.

The default value is jde_server.

- serviceNameConnect=

The value of this parameter specifies the TCP/IP port used by the server to send communications packets to the workstations. Each instance of EnterpriseOne must communicate with workstations through a different port.

The default value is jde_server.

See Also

- ❑ *Generating a Unique Identifier in the Server and Workstation Administration Guide* for directions for running the uuidgen program

Installing Services for a EnterpriseOne Instance

You are not required to install network and queue services for an existing EnterpriseOne instance unless you change the location of the system\bin32 directory for the new instance. For example, you might decide to put the directory on a new disk.

Modifying the Workstation JDE.INI File

After you have installed services for each EnterpriseOne installation, you must modify the workstation JDE.INI file so that the values of the following parameters match those that you set up in the server JDE.INI file:

- serviceNameListen=
- serviceNameConnect=

► To modify the workstation JDE.INI file

1. In the Windows directory on the workstation, locate and open the jde.ini file.
Examples of the windows directory include c:\winnt and c:\windows.
2. Modify the values of the following parameters to match the values in the server jde.ini file:
 - serviceNameListen=
 - serviceNameConnect=

Uninstalling EnterpriseOne Services

To delete an instance of EnterpriseOne after you install its services, you must uninstall the services for that instance before you delete the EnterpriseOne directory tree.

► To uninstall EnterpriseOne services

1. From a command line prompt, change directories to the system\bin32 directory of the EnterpriseOne instance.
2. For example, enter the following command:

```
C:\> d:\E810\system\bin32
```

3. To uninstall network services, enter the following command:

```
jdesnet -u
```

This command removes some settings in the Windows registry that were created when you installed EnterpriseOne services.

Moving or Changing an EnterpriseOne Directory Tree

To move or rename a directory for an EnterpriseOne instance after you install its services, you must uninstall the network service and uninstall the IPC Automation Server (ipcsrv.exe). You can then move or rename the EnterpriseOne directory and reinstall the network service. The IPC Automation Server automatically reinstalls itself when it is first used.

► To move or change an EnterpriseOne directory tree

1. From a command line prompt, change directories to the system\bin32 directory of the EnterpriseOne instance.

For example, enter the following command:

```
C: \> d:\E810\system\bin32
```

2. To uninstall network services, enter the following command:

```
jdesnet -u
```

Note

You do not need to reregister ipcsrv.exe in the new directory because the executable is automatically registered when a binary large object is first transferred via interprocess communications.

3. Move or change the directory tree.
4. Reinstall EnterpriseOne Services.

See Also

- ❑ *Setting up the Network Service* in the *Server and Workstation Administration Guide* for more information about installing network and queue services

Server Administration Workbench

The server administrator needs to be aware of JDENET to ensure that EnterpriseOne functions properly. JDENET starts child processes that control additional communication functions and kernel processes that determine the actual EnterpriseOne client requests. For example, these client requests might include Security Server or JDE CallObject.

The EnterpriseOne Server Administration Workbench (SAW) provides a unified interface for administrators to use while monitoring the processes of JDENET.

Setting up Security Access for SAW

You set up security access for Server Administration Workbench (SAW) to ensure that only select individuals have access to privileged operations, such as deleting server log files and updating the jde.ini file on the server, from within SAW. Any user can run SAW to monitor servers, but only authorized users can affect the operation of the server.

You use the Security Workbench program (P00950) to restrict permission to privileged operations when a user runs either the Web or fat-client versions of SAW. (The fat-client version of SAW is called SERVERADMINISTRATIONWORKBENCH.EXE and is found in the system\bin32 directory on the fat client.)

The SAW program (P9861100) is not supported in EnterpriseOne and later releases.

See Also

- ❑ *Security* in the *System Administration Guide* for more information about setting up role and user security and using the Security Workbench program (P00950)

► To set up security access for SAW

From the System Administration Tools menu (GH9011), choose Security Workbench (P00950).

1. On Security Workbench, then choose Setup Security and External Calls from the Form menu.
2. In Security Workbench [External Calls Security], enter the user's EnterpriseOne ID in the User/System Role control.

Do not enter a system role into this control; you must enter the user ID.

3. Enter "SERVERADMINISTRATIONWORKBENCH.EXE" in the following field, and then click Find:
 - Display Secured Item Executable
4. Expand the Unsecured node in the tree at the bottom of the form.
5. Choose the Run Security option.
6. Drag the "SERVERADMINISTRATIONWORKBENCH." item to the Secured node in the tree and click Find.

The "SERVERADMINISTRATIONWORKBENCH." item should appear under the Secured node.

7. Click Close.

After you have completed these steps, you have secured unauthorized individuals from performing privileged actions in the Web and fat-client versions of SAW. To grant a user access to privileged actions in SAW, drag the “SERVERADMINISTRATIONWORKBENCH.” item from the Secured node to the Unsecured node.

Monitoring EnterpriseOne with SAW on Windows

The Server Administration Workbench (SAW) provides a unified interface in which administrators can review information about EnterpriseOne server processes. Administrators can monitor and modify the processes that they access using SAW. For example, you can change the maximum number of kernels and net processes that SAW monitors. From SAW, administrators can also access the jde.ini file for the local EnterpriseOne server and access log files for any server on the network.

The interface includes the following two tabs:

- Query. This tab lets you add servers and review statistics about their components at a particular point in time.
- Monitor. This tab allows you to periodically monitor the EnterpriseOne servers that you have configured for monitoring. When you have configured a server for monitoring, you can quickly see whether processes on the server are running. Under the monitor tab, SAW also maintains an ongoing record of each test of the server processes.

Adding to the JDE.INI File

Before monitoring EnterpriseOne on Windows with Server Administration Workbench (SAW), you should add a variable to the [JDENET] section of the jde.ini file on the Windows enterprise server. This variable allows SAW to create, transfer, and remove temporary files that are larger than 5 MB. The following is an example of the variable:

```
[JDENET]
netTemporaryDir=<temp_dir>
```

In the new jde.ini entry, <temp_dir> is a temporary directory that JDENET uses to create temporary files, transfer the log files, and remove them after the transfer.

Accessing SAW

Server Administration Workbench (SAW) is a separate EnterpriseOne executable that resides by default in the \E810\SYSTEM\Bin32 directory.

When you sign on to SAW, you must enter your EnterpriseOne user name and password as a validation for all servers. If, after signing on, you add a new server for monitoring, SAW uses the cached signon information to validate the server that you added.

To perform the following operations when you work with SAW, you must have system administration privileges:

- Change the maximum number of net processes
- Change the maximum number of kernel processes
- Delete server log files

► **To access SAW**

1. Open the \E810\SYSTEM\Bin32 directory and then run SERVERADMINISTRATIONWORKBENCH.exe.
2. On Sign On, complete the following fields, and then click OK:
 - User ID
 - Password

See Also

- ❑ *Setting Up Security Access for SAW* in the *Server and Workstation Administration Guide* for a discussion about granting SAW administrative rights

Adding an EnterpriseOne Server for Querying

From Server Administration Workbench (SAW), you can add servers for querying. Doing so enables you to administer EnterpriseOne processes and resources on the server and review statistics at a specific point in time. To add a server for querying, you must know the server name and port number.

Note

When you add a server for querying, SAW displays it under the Query tab. SAW does not display the server under the Monitor tab until you have configured it for monitoring.

► **To add a server for querying**

1. On Server Administration Workbench, click the Query tab.
2. Click the Add button on the toolbar.

Note

You can identify the purpose of any button in the toolbar of the SAW form by passing the cursor over the button.

3. On Server Information, in the EnterpriseOne enterprise server type field, choose the Windows option to monitor an EnterpriseOne server on Windows.
4. In the Server Name field, enter the name of the server that you want to monitor or click the ellipsis button and choose the name of an EnterpriseOne server that is running on the Windows platform.

The ellipsis button is disabled if you choose UNIX or iSeries.

5. In the Port Number field, enter the port number for the server that you are adding, and then click OK.

Checking EnterpriseOne Server Configurations

Server Administration Workbench (SAW) stores the server name and port connection number each time that you add an EnterpriseOne server for querying. You can check each server name and port connection and change a port connection, if necessary.

► To check server configurations

1. On Server Administration Workbench, click the Query tab.
2. Click the Settings button.
3. On Server Configuration, choose a server name.
4. If you added more than one server for monitoring, click the scroll button in the Server Name control and choose a name.

SAW shows the connection port number when you choose a server name.

5. To change the connection port number, enter a new valid number in the Connection Port control, and then click OK.

Checking EnterpriseOne Server Connectivity

Server Administration Workbench (SAW) allows you to run a test to ensure that you are connecting to the EnterpriseOne server that you want to monitor. SAW returns an error message if you fail to make a connection. When SAW fails to connect to the server, you should verify that the name of the server that you want to monitor is correct. You also might need to verify where the connection to the server exists.

► To check EnterpriseOne server connectivity

1. On Server Administration Workbench, click the Query tab.
2. From the File menu, choose Servers, and then choose Connectivity.
3. On Server Connectivity, choose a server name, and then click Start.

SAW displays in the Attempts and Successes controls the number of attempted connections to the server, as well as the number of successful connections made. If these numbers increment simultaneously and consistently, the connection is good.

4. If SAW connects to the server several times successfully, click the Stop button.
5. If SAW displays a form with the message Unable to Resolve <Server Name>, click OK and investigate the problem.

Removing an EnterpriseOne Server from the Query List

At any time, you can use Server Administration Workbench (SAW) to remove a server from your list of monitored EnterpriseOne servers.

► To remove a server from the query list

1. On Server Administration Workbench, click the Query tab.
2. If necessary, expand the Query Servers node on the SAW form.
3. In the expanded tree, click the icon of the server that you want to remove.

Note

You must click the server icon, not a server component, such as Processes.

4. Click the Remove button.

Configuring EnterpriseOne Server Monitoring Settings

You use the Settings button from the Monitor tab to configure each EnterpriseOne server that you want to monitor. Server monitoring configuration allows you to do the following:

- Choose the servers that you want to monitor
- Choose the processes that Server Administration Workbench (SAW) monitors
- Choose the server events that trigger notification messages for you
- Set up the way in which you are notified
- Set up the frequency of notification
- Set up thresholds that determine when server events trigger notification

SAW displays information about servers that you configure for monitoring each time it retrieves server data.

You configure your server monitor settings in SAW using the EnterpriseOne Server Monitor Settings form. You make choices from the following three tabs to configure your monitoring settings:

- Port/Server Selection, which allows you to set up servers, port connection numbers, and server processes to monitor.
- Notification Configuration, which allows you to choose the server events that trigger administrator notification, the ways in which you are notified, and the frequency of notification.
- General, which allows you to configure how frequently SAW tests the server and the thresholds that must be met for server processes in order for you to receive notification.

After you apply the server configuration parameters, SAW displays a server icon for each server that you configured. SAW runs periodic tests on the processes that you specified and displays a message in a window that indicates whether the processes are running or not.

Selecting the Server, Port, and Processes to Monitor

Using the EnterpriseOne Server Monitor Settings form, you can set up EnterpriseOne servers, ports, and processes for monitoring. Then you can use Server Administration Workbench (SAW) to quickly see whether the processes that you chose for monitoring are running on the server.

► To select the server, port, and processes to monitor

1. On SAW, click the Monitor tab.
2. Click the Settings button on the toolbar.
3. On Server Monitor Settings, click the Port/Server Selection tab.
4. In the Port field, click the + button and enter a port number.
5. In the Server field, click the + button and enter a server name.
6. If the server is a Web server, choose the following option:
 - This is a WebServer
7. To set up server processes to monitor, choose any of the following options, if applicable:
 - Off line
 - UBE
 - Lock Manager
 - Replication
 - Security

Note

If you choose Off line, SAW does not monitor EnterpriseOne server processes for the server that you specified.

8. Click Apply to save the settings.

Setting up Event Notification

After you have chosen the EnterpriseOne server and port for monitoring and chosen the processes that you want to monitor, you can specify the type of events that will trigger notification to you, as well as the methods of notification.

► To set up event notification

1. On the EnterpriseOne Server Monitor Settings form, click the Notification Configuration tab.
2. Under Notification Events, choose from the following available options to specify the server events that will trigger administrator notification:
 - Dead processes
 - UBEs exceed run-time threshold
 - Disk usage exceeds threshold

- Number of jobs in queue exceeds threshold
 - System errors in web CallObjects
3. To set up e-mail notification, click the + button in the E-mail address field, enter an e-mail address, and then click OK.
 4. To set up pager notification, click the + button in the Pager address field, enter a pager address, and then click OK.
 5. Under Notification Settings, complete the following options to indicate how you want to be notified:
 - Choose the Repeat Notification option if you want SAW to notify you again after the first notification
 - Enter the interval (frequency) between notifications, in seconds
 - Enter the pager message size in number of lines
 6. Click Apply to save the settings.

Setting up SAW Run-time Parameters

You can use the EnterpriseOne Server Monitor Settings form to configure Server Administration Workbench (SAW) run-time parameters. For example, after you establish the server events for which you want to receive notification, you can set up the thresholds that trigger notification. You also can set up how often you want SAW APIs to retrieve and display server information, and you can specify the directory to which you want SAW to write server log files.

From the General tab, you specify the following run-time parameters:

- Log file, which specifies the directory to which you want SAW to write log files
- Test frequency, which specifies in seconds how frequently you want SAW APIs to retrieve and display server information
- Outstanding request threshold, which specifies the maximum number of requests for a kernel process that SAW will allow to back up on the server before administrator notification is triggered
- Disk usage percent threshold, which specifies the maximum percentage of used space on a server disk that SAW allows before administrator notification is triggered
- UBE run-time threshold, which specifies the number of minutes for which SAW allows a batch process to run before administrator notification is triggered
- Queue threshold, which specifies the number of jobs in a queue that SAW allows before administrator notification is triggered

The following threshold settings are recommended:

Test frequency (seconds)	300
Outstanding request threshold	10
Disk usage threshold (percentage)	80
UBE run-time threshold (minutes)	90
Queue threshold (number of jobs)	20

► **To set up SAW run-time parameters**

1. On the EnterpriseOne Server Monitor Settings form, click the General tab.
The program displays a list of monitoring parameters to which you can assign values.
2. Set values for the following parameters:
 - Log file
 - Test frequency
 - Outstanding request threshold
 - Disk usage percentage threshold
 - UBE run-time threshold
 - Queue threshold
3. Click Apply to save the settings.


Note

If you enter 0 for a parameter value, SAW does not test that parameter.

Monitoring EnterpriseOne Server Processes

After you have configured EnterpriseOne servers for monitoring, Server Administration Workbench (SAW) runs tests at regular intervals on the processes that you chose. You can quickly determine whether the servers that you are monitoring are up or down and whether the server processes that you are monitoring are running.

SAW displays this information in the Server Monitor window.



The information in the Server Monitor window allows you to see, at a glance, whether a SAW-monitored server process is up or down. A stop sign icon indicates that a server process is down, while a spotlight icon indicates that a process is running.

The following table summarizes the server process parameters that appear in the Server Monitor window:

Date & Time	The date and time that SAW ran a server check. The parameter value that you set during server-monitoring configuration determines the frequency of the checks.
Server	The name of the EnterpriseOne server that you are monitoring. Sleep indicates that the Server Monitor has entered sleep mode.
Service Name	The port connection number.
Server Type	The server process that you are monitoring, such as UBE or Lock Manager.
Status	The status of the server process, which is either Up or Down.
Details	Additional information about a server process.

Prerequisite

- ❑ Verify that you have configured the monitoring settings for each EnterpriseOne server for which you want to monitor processes. You must configure the monitoring settings before you can monitor the server processes in the Server Monitor window. Adding an EnterpriseOne server for querying is not sufficient.

► To monitor EnterpriseOne server processes

1. Sign onto SAW.

During your initial signon, a brief delay occurs as SAW retrieves information about server processes.


2. After SAW updates the Server Monitor window, review the Date & Time and Status parameters for Stop values, which indicate whether a server is down or a process on the server terminated abnormally.

Note

If you change the view in the SAW form and need to return to the Server Monitor window, click the Restore icon on the menu bar.

Monitoring Enterprise Server Statistics

The Server Administration Workbench (SAW) form displays performance statistics for each enterprise server that you configure for monitoring. A clock icon designates a parent node in the Statistics window. The parent node displays the date and time of each data retrieval.



You expand the node to review performance statistics on each server. You can review detailed statistics about the following:

- Kernel processes
- Network processes
- batch processes
- Server disks

The Statistics window offers you another quick way to identify server problems. When a process terminates abnormally or some other problem occurs, SAW displays the monitoring time, the port connection node, and the server node in red.

The statistics for the process or server component, such as a disk, that is causing the problem are highlighted in yellow in a separate window.

Finally, SAW allows you to view and print log files for kernel and network processes that run on servers that you have configured for monitoring. You can also turn logging on or off for any server process.

Monitoring Process Statistics for Enterprise Servers

The Statistics window allows you to review information about network and kernel processing for each monitored enterprise server. To display the processing statistics, you click a server icon. Statistics about each process appear in the Process Statistics window.

The following table describes the parameters that appear in the Process Statistics window:

Process Name	The name of the kernel or network process.
PID	The process identification number.
Start Time	The date and time that the process began.
Total Process Time	The total time, in milliseconds, that the server required to run the process.
Total Requests	The total number of requests for a particular process handled by the server, up to the time that Server Administration Workbench (SAW) retrieved server process data.
In Connections	The number of workstations that are connected to the server.
Outstanding Requests	The number of requests in a queue for a process.
Avg. Processing Time – Overall	The average amount of time the server required to handle a kernel process request (Total Process Time divided by Total Requests).
Avg. Processing Time – Last Period	The average amount of time the server required to handle a kernel process request during the last server session. Compare this parameter value with the value for the Avg. Processing Time – Overall parameter to assess server performance.
Lock Conditions	The number of processes that are trying to access the same resource.
Kernel Max Count	The maximum number of kernel processes that can run.
Kernel Current Count	The number of kernel processes that are currently running.
Version	The current version of SAW. The value is for internal use only

You can right-click any kernel or network process to review additional information and complete the following additional tasks:

- View the debug log
- View the JDE log
- Print the debug log
- Print the JDE log

- Process users
- Turn on server logging
- Turn off server logging
- Change the maximum number of a process type
- Disable kernels

Note

You can change the maximum number of a process type and disable kernels only if you have SAW administrative privileges.

► To monitor process statistics for enterprise servers

1. On SAW, click the Monitor tab.
2. Click the Refresh button on the toolbar to get the latest server information.
3. On Statistics, expand the top clock icon node.
4. Expand the port connection icon node for the enterprise server that you want to monitor.
5. Click the server icon.

Statistics for kernel and network processes that are running on the enterprise server appear in the Process Statistics window.

6. To sort the items in a column, click the column heading.

Changing Process Settings

If you have administrative rights, you can use Server Administration Workbench (SAW) to change the number of net (jde_n) and kernel (jde_k) processes. You might do so when the server is experiencing performance problems. You can also disable processes that are causing problems.

Your ability to use SAW to increase the number of jde_k processes running on the server is constrained by the server's JDE.INI file, specifically the `maxNumberOfProcesses` parameter in the `[JDENET_KERNEL_DEF_x]` section and the `maxKernelProcesses` parameter in the `[JDENET]` section. The maximum number of processes that you can increase for a kernel of a specific type, such as `CallObject`, or for a combination of kernel types, must not exceed the product of the following formula:

$$\text{maxKernelProcesses} \text{ minus the sum of individual kernel type } \text{maxNumberOfProcesses}.$$

For example, if the `maxKernelProcesses` parameter value is 50, and the sum of the values for the `maxNumberOfProcesses` parameter is 34, you can use SAW to increase the number of processes for an individual kernel type or a combination of kernel types by no more than 16. If the `maxNumberOfProcesses` for the `CallObject` kernel is 10, you can use SAW to increase the number of processes to no more than 26 ($50 - 34 = 16$, $10 + 16 = 26$).

► **To change process settings**

1. On SAW, click the Monitor tab.
2. Click the server icon for the server that you want to monitor
SAW updates the Process Statistics window with data on each net and kernel process that is running on the server.
3. Choose a kernel or net process and right-click.
4. To change the number of net or kernel processes, choose Change max number of this process type, enter a new value, and then click OK.
5. To disable a kernel process, choose Disable Kernels.

Viewing Log and Debug Log Files for Enterprise Server Processes

Server Administration Workbench (SAW) allows you to review server log and debug log files for individual kernel and network processes that run on enterprise servers that you have configured for monitoring.

Note

If you attempt to view a file that has been deleted from the server, SAW displays a message saying that the file is no longer available.

► **To view log and debug log files for enterprise server processes**

1. On Statistics, click an enterprise server icon.
2. On Process Statistics, right-click a network or kernel process and choose View Debug Log or View JDE Log.
The log or debug log file appears in a text editor.

Note

You can also choose a network or kernel process on the Process Statistics form and click the Debug Log or JDE Log buttons on the toolbar.

Printing Log and Debug Log Files for Enterprise Server Processes

You can print server log or debug log files from your workstation. Server Administration Workbench (SAW) prints a text editor document to your local printer. Remember that debug log files can be very large, which might make printing them impractical.

► To print log and debug log files for an enterprise server process

1. On Statistics, click an enterprise server icon.
2. On Process Statistics, right-click a network or kernel process and choose Print Debug Log or Print JDE Log.

SAW prints the file to your local printer. For server logs, you can also click the Print JDE Log button on the toolbar.

Turning Logging on and off for Enterprise Server Processes

You can control logging for processes that run on enterprise servers. If you observe that a problem has occurred with a process, you might turn on logging so that Server Administration Workbench (SAW) adds information to the log file for that process. When you have enough information, you can turn off logging so that you can more easily isolate in the log file in which an error occurred. Even with logging turned off, you can view the log file for a process, although the entries stop at the point that you turned off logging.

► To turn logging on and off for enterprise server processes

1. On Statistics, click an enterprise server icon.
2. On Process Statistics, right-click a network or kernel process and choose Turn on Server Logging or Turn off Server Logging.

Viewing Process User Statistics for Enterprise Servers

You can view statistics about users who are running business functions on the enterprise server. A separate window in the Server Administration Workbench (SAW) form displays the user's ID and machine name for each CallObject kernel process that you choose.

► To view process user statistics for enterprise servers

1. On Statistics, click an enterprise server icon.
2. On Process Statistics, right-click a CallObject kernel process and choose Process Users.

SAW displays user and machine information for the CallObject process that you chose.

Viewing Statistics for Batch Processes

The UBE icon appears as a child of the enterprise server icon in the Statistics window only if a batch process is running on the enterprise server. If you click the UBE icon, you can review the following additional information about the batch process:

- UBE name
- Process ID number
- Date and time the process began
- Running time for the process, in minutes

► To view statistics for batch processes

1. On Statistics, expand an enterprise server icon.
2. Click a UBE icon, if one appears.

Statistics for batch processes that run on the enterprise server appear on the Monitor UBE form.

Viewing Statistics for Server Disks

The disk icon appears at all times as a child of the enterprise server icon. If you click the disk icon, you can review the following additional information about server disks:

- Disk name
- Total space
- Used space
- Available space
- Percent used
- Disk type

► To view statistics for server disks

1. On Statistics, expand an enterprise server icon.
2. Click the Disk icon.

Statistics for each server disk appear on the Monitor Disk form.

Monitoring EnterpriseOne Web Server Statistics

You can also use Server Administration Workbench (SAW) to monitor statistics about EnterpriseOne Web servers from your fat-client workstation. The Statistics window on the SAW form displays no information about kernel and network processes for Web servers because these processes run on enterprise servers. However, you can monitor information about Web server users, connections to enterprise servers, and business functions that run on enterprise servers.

If you are working on a thin- or zero-client workstation, you can monitor Web servers from your workstation using the Web Server Monitor utility.

When you configure a Web server for monitoring in SAW, you can set up notification triggers to let you know when system errors occur in business functions that run on enterprise servers that are connected to the Web server. The Web Server Monitor does not allow you to set up notification triggers.

See Also

- ❑ *Monitoring the Web Server System* in the *Server and Workstation Administration Guide* for more information about monitoring Web server statistics

Monitoring Data on Enterprise Servers Connected to the Web Server

When you click a Web server icon in the Statistics window, Server Administration Workbench (SAW) displays a separate window with information about each enterprise server that is connected to the Web server.

This window contains important connection information, as well as information about business functions running on the enterprise servers.

The following table summarizes the data SAW displays on each enterprise server that is connected to the Web server:

Host Name	The enterprise server name.
Port	The enterprise server port connection number.
Max Conn Pool	The maximum number of enterprise servers that can be connected to the Web server, as defined in the server INI file.
Max Pool Wait	The number of users waiting for a connection when the Max Conn Pool value has been exceeded. For example, if the Max Conn Pool value is 30, and five users are waiting to connect, the Max Pool Wait value is 5. As soon as the Max Conn Pool value is exceeded, the value is incremented with each new user who attempts to connect.
Current Pool Wait	The number of users who are waiting for a server connection. This parameter has a value only if the Max Conn Pool value is exceeded.
Current Connections	The number of users currently connected to the enterprise server from the Web server.
No. of CallObjects	The number of business functions that have run on the enterprise server during a session.

► To monitor data on enterprise servers connected to the Web server

1. On SAW, click the Monitor tab.
2. In the toolbar, click the Refresh button to get the latest server information
3. On Statistics, expand the top clock icon node.
4. Expand the port connection icon node for the Web server that you want to monitor.
5. Click the Web server icon.

SAW updates the Monitor WebServer window with the name of and information about each enterprise server that is connected to the Web server.

Monitoring Web Server CallObjects

Server Administration Workbench (SAW) allows you to view additional data about business functions running on an enterprise server connected to the Web server that you are monitoring. You can troubleshoot the data on business functions from SAW. For example, SAW displays system errors and application errors that occur during the running of business functions.

The following table summarizes the data that SAW displays about business functions that run on an enterprise server that is connected to the Web server:

CallObject Name	The name of the business function.
Min Time	The minimum time, in milliseconds, required to run the business function.
Max Time	The maximum time, in milliseconds, required to run the business function.
Avg. Time	The average time, in milliseconds, required to run the business function.
Timeouts	The number of times a business function timed out before completion.
System Errors	The number of system errors that occurred during the running of a business function.
Application Errors	The number of application errors (usually related to business logic) that occurred during the running of a business function.
Times Called	The number of times a business function was called to run on the enterprise server.

► To monitor Web server CallObjects

1. On Monitor WebServer, right-click a line containing the name of an enterprise server and then choose Display CallObjects and Connections.

SAW opens a new window with the following two tabs:


- Web Server CallObjects
- Web Server Connections Statistics

2. Click the Web Server CallObjects tab.

SAW updates the window with information about each business function that ran on the enterprise server during a Web server session.

Monitoring Web Server Connection Statistics

You click the Web Server Connections Statistics tab to review data on the connections made from the Web server to the enterprise server.



The following table summarizes the data that Server Administration Workbench (SAW) displays on connections between the Web server and the enterprise server:

Connection ID	A number that identifies the socket connection between the web server and the enterprise server.
In Use	The status of the connection, either Active or Inactive.
Bytes Sent	The number of bytes transmitted from the connection socket.
Bytes Received	The number of bytes received from the connection socket.
Messages Sent	The number of JDENET messages transmitted from the connection socket.
Messages Received	The number of JDENET messages received from the connection socket.

► To monitor Web server connection statistics


1. On Monitor WebServer, right-click a line containing the name of an enterprise server, and then choose Display CallObjects and Connections.
2. Click the Web Server CallObjects tab.

SAW updates the window with information about each business function that ran on the enterprise server during a Web server session.

Monitoring Web Server User Statistics

You can also use the Statistics window on the Server Administration Workbench (SAW) form to view statistics about Web server users.

When you expand the node for a Web server, only a User icon appears as a child of the server icon.



The Web Users window displays information about each user who is connected to the Web server. The following table summarizes the data that SAW displays for each Web user:

Session ID	Unique numeric identifier for each user session.
User Name	User ID.
Environment	The EnterpriseOne environment in which the user is working, such as PDEVNIS2. Environment = Path Code + Object Configuration mappings.
Computer Name	Name of user's local workstation.
Login Time	Date and time that the user signed on to the Web server.
Idle Time	The amount of time, in minutes, that the user has been inactive.
Browser	The Web browser name.
User Type	The way information appears on the workstation: JavaApplet, HTML, or Portal.
Maximum Users (Allowed)	The maximum number of users who can be connected to the Web, as defined in the server INI file.
Maximum Users (Used)	The number of users currently connected to the server.
Session Timeout (Minutes)	The amount of time, in minutes, that a user can be idle before the session times out and the user is disconnected from the server.

► To monitor Web server user statistics


1. On SAW, click the Monitor tab.
2. In the toolbar, click the Refresh button to get the latest server information
3. On Statistics, expand a clock icon node.
4. Expand the port connection icon node for the Web server that you want to monitor.
5. Expand the server icon.
6. Click the Users icon.

SAW updates the Web Users window with data about Web server users for the current session.

Monitoring Web Server User Connection Statistics

You can gain additional data about connections that Web server users make to enterprise servers by right-clicking a record in the Web Users window.

Choosing Show Web User Connections opens a window that displays each Web server user's connection to the enterprise server.



Note

After you open the window that displays the connection data, you can review each user's connection data by passing the cursor over a record in the Web Users window.

The following table summarizes the connection data that appears for each Web user:

Host	The name of the enterprise server to which the web user connects.
Port Number	The number of the port used for the connection.
Number of Kernels	The number of kernel processes the user is running on the enterprise server.
Kernels Being Used (PID)	The process ID number for each kernel the user is running on the enterprise server.

► To monitor Web server user connection statistics

1. On Web Users, right-click a record that contains a user ID, and then choose Show Web User Connections.
SAW opens a window that contains connection data for the user you chose.
2. To view connection data for another user, pass the cursor over another record in the Web Users window.
SAW changes the connection data in the window to reflect the record that you selected.

Monitoring Web Server Logs

When you monitor a Web server, Server Administration Workbench (SAW) presents additional buttons on the toolbar. You can use these buttons to do any of the following:

- Get details about the Java virtual machine, which functions as an interface between the Web server and the enterprise server.
- Get the JDE log from the Web server.
- Get the debug log from the Web server.
- Get the JAS log, which enables you to turn on or off JAS logging. JAS log gathers information about processes running on the Java Application Server.
- Get the NET log, which enables you to turn on or off Net logging. Net log contains information about JDENET processes.
- Delete the JDE log from the Web server, provided that you have SAW administrative privileges.
- Delete the debug log from the Web server, provided that you have SAW administrative privileges.

You can accomplish any of these tasks by clicking the appropriate button on the toolbar. If you want to get the JDE log or the debug log from the Web server, SAW displays a form that states the length of the file in kilobytes.

You can shorten the file by entering the number of bytes that you want. You must enter an integer between 0 and 20,000.

Querying EnterpriseOne Enterprise Servers

Even if you have not configured an enterprise server for monitoring, you can get detailed statistics about server components, resources, and log files using the Query tab on the Server Administration Workbench (SAW) form. To query a server, you must add the server for monitoring.

Querying an enterprise server offers the following types of enterprise server data:

- Processes
- Disks
- Interprocess communications (IPCs)
- Logs
- Users
- Job queues

SAW updates the server component information each time it retrieves data from the enterprise server. Unlike when you use the Statistics window of the Monitor tab, however, you cannot review information from previous data retrievals when you query a server.

SAW arranges server component data in a tree with Query Servers as the parent node. Servers, server components, and server subcomponents are child nodes.

When you click a component or subcomponent icon, SAW displays data about the component in a separate window.

For example, if you click the Kernel component under the Processes node, SAW displays data for all kernel processes for the time of the SAW latest data retrieval.

You can also display data on any component by clicking a server icon, and then clicking File in the toolbar and choosing Open and the name of a component or subcomponent of the server node.

Note

SAW displays in the Query Servers tree each Web server that you added for monitoring. However, you cannot monitor Web servers from the Query Servers tree. Web server monitoring must be done from the Statistics window on the Monitor tab.

See Also

- ❑ *Adding an EnterpriseOne Server for Querying* in the *Server and Workstation Administration Guide* for more information about adding a server for monitoring

Querying Enterprise Server Process Components

Querying a server allows you to monitor kernel, network, and queue processes. When you click a process component, Server Administration Workbench (SAW) displays—in a separate window—the most recent data for each process of the type you chose. You can also view and print JDE log and debug logs, as well as turn logging on or off for a process.

SAW displays the following data for each server process type:

- Process Type, such as kernel, network, or queue
- Process ID
- Parent Process ID
- User Name
- Start Date/Time
- Last Date/Time
- OW - Status
- OS - Status

From the Query tab, you can gain additional data on each kernel and network process by right-clicking a line in the window and choosing Information.

If you choose Information after right-clicking a kernel process, the Detailed Info - Kernel Process form appears.

ProcessType :	Kernel	Total Kernel Process Time :	23
Server System Directory :		Total Kernel Requests :	4516
INI File Location On Server :		Kernel Process Range Index :	5
Process User Group ID :	0	Kernel Process Out Requests :	0
Process Last Active Time :	Mon Mar 22 15:54:45 2004		

OK

If you choose Information after right-clicking a network process, the Detailed Info - Network Process form appears.

Detailed Information : Net Process

ProcessType : Total Network Process Time :

Server System Directory : Network Stream Port Number:

INI File Location On Server : Network Data Port Number:

Process User Group ID : Number of In Connections :

Process Last Active Time :

Connection information:

No.	Type of Conn	IP Address	Long Address
1	0	10.0.11.140	.
10	0	10.0.11.140	.
11	0	10.0.11.140	.
12	0	10.0.11.140	.
13	0	10.0.11.140	.
14	0	10.0.11.140	.
15	0	10.0.11.140	.
16	0	10.0.11.140	.
17	0	10.0.11.140	.
18	0	10.0.11.140	.

OK

A separate window on the Detailed Info - Network Process form displays the following connection information:

- Connection number
- Connection type
- Internet Protocol (IP) address
- Long address

See Also

- ❑ *Monitoring Process Statistics for Enterprise Servers* in the *Server and Workstation Administration Guide* for more information on using SAW to monitor logs for individual processes

► To query enterprise server process components

1. On SAW, click the Query tab.
2. Expand the Processes node for an enterprise server you want to monitor and click one of the following icons:
 - Kernel
 - Network
 - Queue
 - All

SAW populates a form window.

3. Click a column heading to sort the data in a column.
4. For additional information on a process, right-click and choose Information.
5. To turn logging on or off for the process or to view and print logs, right-click and choose the appropriate option.

Querying Enterprise Server Disk Components

Querying the server disk components produces the same information that is available from the Monitor tab.

See Also

- ❑ *Viewing Statistics for Server Disks* in the *Server and Workstation Administration Guide* for more information about the disk data that SAW gathers

► To query enterprise server disk components

1. On SAW, click the Query tab.
2. Expand the Disks node for an enterprise server you want to monitor and click one of the following icons:
 - EnterpriseOne Disks
 - All Disks

Querying Enterprise Server IPC Resources

Server Administration Workbench (SAW) allows you to query interprocess communications, which allow applications to share data and memory and allow one application to control another.

SAW monitors the following IPC types:

- Mutexes, which are locks that allow only one process to hold the lock. This is the most common type of lock.
- SWMR locks, which allow only one writer to change data but allow multiple readers to access the data. The system grants reader locks if no outstanding writer locks exist; the system grants writer locks if no outstanding reader locks exist.
- Message queues, which allow one process to pass packets to another process.
- Shared memory, which is a segment of memory that appears in the address space of more than one process. For example, network processes might use a shared memory segment to monitor existing processes and their status.
- All

SAW displays the following data for each server process type:

- Resource Name
- IPC ID
- Key, which is an ID number the operating system uses to identify one or more groups of IPC resources

- IPC Type
- User ID
- Role

► To query enterprise server IPC resources

1. On SAW, click the Query tab.
2. Expand the IPC node for an enterprise server you want to monitor and click one of the following icons:
 - Mutexes
 - SWMR Locks
 - Message Queues
 - Shared Memory
 - All

Querying Enterprise Server Log Files

Server Administration Workbench (SAW) allows you to query for summary server log and debug log information, as well as summary information for print queue logs. You can also choose and display any log or debug log from your SAW sessions. If you have SAW administrator privileges, you can delete log and debug server log files.

► To query enterprise server log files

1. On SAW, click the Query tab.
2. Expand the Logs node for an enterprise server you want to monitor and click one of the following icons:
 - Server Logs
 - Print Queue Logs

A list of log files on the server appears in the workspace with the names of the files, the dates and times when EnterpriseOne created the files, and the sizes of the files.

3. To view a file, right-click a line containing the name of a log and choose View Selected. SAW displays the log file or debug log file.

4. To delete the file, choose Delete Selected from File menu.

If you have administrator privileges, SAW deletes the file and then refreshes the workspace. If you do not, SAW displays a form advising you that you must have administrative privileges to perform the deletion.

► To query the enterprise server users

1. On SAW, click the Query tab.
2. Expand the node for an enterprise server you want to monitor and click View Users.

SAW populates a separate window with information on current users on the enterprise server.

3. To view the EnterpriseOne processes being run for an individual user, double-click a line containing an asterisk in the Using Processes column.

Querying the Enterprise Server Job Queue

Server Administration Workbench (SAW) allows you to query for data about job queues on the enterprise server. When you click the Queue icon, SAW displays—in a separate window—information about each job in the queue. The information is displayed in the following columns:

- Queue Name
- Queue Type
- Queue Status
- Queue Max Jobs

► To query the enterprise server job queue

1. On SAW, click the Query tab.
2. Expand the node for an enterprise server you want to monitor and click Queues.

SAW populates a separate window with information on jobs in the queue on the enterprise server.

Querying the Enterprise Server Users

Server Administration Workbench (SAW) allows you to query for data about users connected to the enterprise server. When you click the View Users icon, SAW displays—in a separate window—information about each user. The information is displayed in the following columns:

- User Name
- Machine Name
- Using Process

An asterisk in the Using Process column indicates that a user is currently running EnterpriseOne processes on the server.

You can view the list of active processes for a user by double-clicking a line that contains an asterisk in the Using Process column. SAW displays identifying data for the active process.

Configuring SAW and Monitoring EnterpriseOne Servers from the Web

The Server Administration Workbench (SAW) provides a unified Web interface for administrators to monitor EnterpriseOne Servers (Enterprise, Java, and XPI) remotely through a browser. From the Web, administrators can set up a Server Monitoring Configuration (SMC) to monitor different instances of EnterpriseOne servers and notify assigned users via e-mail or pager when an assigned agent detects a faulty event with any of the instances.

SAW for the Web can be installed as a standard part of the EnterpriseOne Web server. In future releases, SAW for the Web will be installed as standalone. If you install a Web server, SAW is automatically installed and can be accessed like any of the other Java server components. However, if you want to use SAW as a standalone product, you must purchase the SAW package and install the included jar file on a Web server connected to the EnterpriseOne network. The SAW package includes all the components needed to run SAW and monitor the EnterpriseOne servers.

The process of setting up an SMC consists of the following:

- Create an SCM file
- Set up a group of system servers to monitor
- Add port numbers to each server
- Add events and statistics to each server port
- Add profiles (e-mail and pager addresses) to each event
- Add, change, or delete ports, processes, and e-mail addresses for selected server ports

Creating an SMC File

Before you can monitor server ports via the Web, you must first access an administration Web server and create a Server Monitoring Configuration (SMC) file. This file stores the information that is required for monitoring EnterpriseOne servers, such as server names, port numbers, monitored processes, and so on. The SMC file stores all the vital configuration data that is accessed by an agent that runs on the Web server and performs the monitoring duties.

► To create an SMC file

1. From a web browser, enter the URL of the Server Administration Workbench (SAW) web page located on the web server where the SAW package is installed. The path to the html page is `saw/smc/smc.html`.

For example, use the following format for the URL:

```
http://host_name:port_number/web_alias/saw/smc/smc.html
```

Where *host_name* is the name of the web server where the SAW package is installed, *port_number* is the number of the port that the server is running on, and *web_alias* is the alias of the configuration for accessing the web pages on the web server.

The *web_alias* can point to a directory on the server (for example, `c:\internet\distribution`) that is configured to access Web Server documents.

2. Add the URL to your list of favorites.
3. On the Server Administration Workbench (SAW) WEB Version page, complete the following fields:
 - Agent Web Server Name
Host name of the web server.
 - Agent WEB Server port number

Port number on which the Web Server runs.

- Agent SMC File Name

Type a path and file name for the configuration file, for example, c:\saw\saw.smc.

Note

The Agent Web Server Name and the Agent WEB Server Port number must be entered the first time you access SAW. These will be values used when you access this page.

4. In the Action field, verify that Work with SMC is selected, and click the action button.
5. If the file name you entered does not exist, the program prompts you to create the new file. Click OK to create the file.
6. On the SAW Home Page, complete the following fields:

- JAS Install Directory

Install directory for the JAS server or the SAW package.

To enter the correct value for this field, search for the file “saw.css” on your web server. For example, if saw.css is found in the directory “/dir1/dir2/saw”, then the correct value for this field is “/dir1/dir2”.

- SAWDB Root directory

Enter a path for the directory to store data collected by the SAW agent.

Note that for AS400 and UNIX®, the userid that is running the web server must have full access rights to this directory.

- HTML Status Page

The name of the HTML Page that SAW agent will create as an interface to SAW. (Use the default value if you run only one SAW Agent, or enter new names if you run additional SAW Agents.)

- Java Install directory

The directory path to the java.exe file, or java for UNIX® and AS400, on the SAW Agent host.

To enter the correct value for this field, look for the file java.exe (or java for UNIX® and AS400). If the file is found in the directory “/dir1/dir2/bin”, then the correct value for this field is “/dir1/dir2/bin”.

- Mail Server

Enter your company mail server, for example, mail.company_name.com. The SAW agent uses this server to send e-mail/pager notifications when it detects faulty processes.

To access more information about any of the fields, click the Help button at the top of the screen.

7. Click Submit.

► **To regenerate an SMC file**

If you update SAW to a new version for which the SMC file is no longer compatible, or if your SMC file is corrupted for any reason, you can regenerate your SMC file from a text file.

Each time you change or save your SMC file, SAW creates a text file that contains the same information as your SMC file, but in a text file format. If your SMC file is called `saw.smc`, then each time you change or save your SMC file, the system will save a text file called `saw.smc.txt`.

1. To regenerate an SMC file from the text file, enter the name of the text file (for example, `saw.smc.txt`) in the Agent SMC File Name field.
2. Choose Importing SMC from text file from the action.

SAW regenerates the SMC file with the same name as the name of the text file. The new file does not have the `.txt` extension.

Modifying Advanced SMC Settings

The Server Monitoring Configuration (SMC) file includes a list of global parameters that govern the general monitoring process. All of these settings have default values that should work fine under most SMC configurations. However, you can easily change these settings to suite your requirements.

► **To modify the advanced SMC settings**

1. On the SAW Home Page, select Advanced Settings from the Action field.
2. On SMC – Advanced Settings, modify any of the following fields to suite your SMC configuration:
 - Agent cycle
Enter time (in seconds) between the start of one agent monitoring cycle and the next.

Note

See *Performing SMC Actions* in the *Server and Workstation Administration Guide* for more information about Agent cycles.

- SAWDB used
Change this value to false if you do not want to save statistics or event data in the SAW database.
- Auto Monitoring SAW Agent
Change this value to false if you do not want the SAW agent to automatically restart when the agent fails.
- HTML page refresh
Enter time (in seconds) to refresh all html pages (including collected data) used in SAW. This value is not related to the agent cycle.

- **Agent read SMC cycle**
Enter a time interval (in seconds) between times when the SAW Agent checks if the user forced the Agent Cycle. If the user forced the Agent cycle since the last read, the Agent automatically starts a new monitoring cycle.
- **Check e-mail address**
Change this value to false if you do not want the system to verify that the e-mail addresses you enter have the correct format.
- **Server timeout**
Enter time (in seconds) to wait for the response from a server before it assumes the server is down.
- **HTML page secondary path**
Enter a valid network path to store an additional status page. This path is useful if you want to access SAW data from another web server that cannot access the primary web page. This path is required to access the Port View.

Note

The Port view is defined in *Working with SMC Views* in the *Server and Workstation Administration Guide*.

- **HTML first favorite link**
Enter the web addresses for up to three web links that you want to include in the result pages of SAW.
- **Name of first favorite link**
Enter the names for each of your favorite links.

3. Verify that Submit Changes is selected in the Action field and click the Action button.

Adding EnterpriseOne Servers

After creating the configuration file, you can create a list of the EnterpriseOne Web, Enterprise, and XPIe servers that you want to monitor.

► To add EnterpriseOne servers

1. On SAW Home Page, verify that Work with Servers is selected in the Action field, and click the Action button.

Note

By default, the agent server name is automatically added to the list. This is to make sure that the agent will be restarted if it needs to be.

2. In the Action field, verify that Add new Server is selected, and click the Action button.
3. Complete the following fields:
 - Server Name
Type the name of the EnterpriseOne Server (Enterprise, Web, or XPI).
 - Server OS
Select the type of operating system for the server.
4. Click the Action button.
The program adds a new row to the Work with Servers screen listing the server you entered.
5. Repeat the steps above to add additional servers to the SMC configuration.
6. Verify each of the server names and operating systems you entered.
In the upper-right corner of SMC, click on the Global View link to view the current configuration. This screen shows at-a-glance the global settings and all the server/port information that you enter as you configure SAW. You can return to this screen at any time during the configuration process to view all the information you enter.
7. If you discover an error, return to the Work with Server screen, and click Modify Server next to the server you need to change, then correct the error.
8. When you have entered all the servers you want to monitor, select one of the following options:
 - To add a port number for several servers at the same time, complete the task *Adding Ports for Multiple Servers* in the *Server and Workstation Administration Guide*.
 - To add ports to a single server, complete the task *Customizing Server Ports* in the *Server and Workstation Administration Guide*.

Note

If you want to monitor several servers that have similar port configurations, it is quicker and easier to assign a port number to multiple servers by completing the first task listed above.

Adding Ports for Multiple Servers

You can add port numbers to the SMC file and assign these ports to several servers at once.

► To add ports to multiple servers

1. On SAW Home Page, select Work with Ports from the Action field.
2. On SMC – Work with Ports, select Add new Port from the Action field and click the Action button.
3. Complete the following sub steps to configure a port number:
 - a. On SMC – Adding new Port, complete the following fields:
 - Port Number
Enter the port number you want to monitor (for example, 6004).
 - Description
Enter a description of the port (for example, Xe_SP20).
 - Server Type
Click the down arrow and select the type of server from the list.
 - b. Check the box to the left of each server that you want to assign this port number.
 - c. Click the Action button to accept this configuration.
The program adds the port number to SMC – Work with Ports screen.
4. Repeat these steps for each port number that you want to add to multiple servers.
The screen displays the complete list of ports that you entered.
5. To view the complete configuration, click on the Global View link.

Adding Events to Multiple Server Ports

Once you have assigned port numbers to servers in Server Monitoring Configuration (SMC), you can assign a predefined event to each of the server ports and alert administrators if the event indicates a malfunction. SMC includes two types of events:

- Events that are either true or false (such as Server down)
SMC will alert assigned administrators if the event occurs.
- Events that require threshold values (such as Disk Percent Used)
You assign Warning and Alarm values to these types of events. For all threshold events, Warning = write in the event file, Alarm = write in the event file and send notification.

You can set up the following events and assign each to particular port numbers:

- Server Down

Monitor alerts administrators if it cannot communicate with a server port.

- Dead Process

Monitor alerts administrators if it detects a dead process on an enterprise server.

- Disk Percent Used

Administrator enters a warning and alarm threshold for the percent of disk space used (for example, 80%), and the monitor alerts administrators if the used disk space exceeds the threshold.

- Kernel Outstanding Request

Administrator enters a warning and alarm threshold for the number of unprocessed kernel requests, and the monitor alerts administrators if either threshold is exceeded.

- Web CallObject Errors

Administrator enters a warning and alarm threshold for the number of CallObject kernel process errors for a business function activated through a Web server, and the monitor alerts administrators if the threshold is exceeded.

- UBE Runtime Limit

Administrator enters a warning and alarm threshold for the maximum time (in minutes) that a UBE can run, and the monitor alerts administrators if either threshold is exceeded.

- XPIe Disk Low

Monitor alerts administrators if an XPI server port indicates it is low on disk space.

This task assigns an event to a port number; that is, it assigns the event to *all servers* that are associated with the port number. If you want to customize a port number on one server to monitor an event differently than other servers, you must complete customize your server ports.

► To add events to multiple server ports

1. On the SAW Home Page, select Work with Events from the Action field and click the Action button.

SAW - SMC - Work with Events



c:\saw\saw_doc.smc

Fri Mar 19 15:37:45 MST 2004

Select Task SMC Global View Cancel:

[Home](#) [Frame View](#) [Port View](#) [Global View](#)

#	Event Name	Action
1	Server Down	Activate/Deactivate
2	Email Dead Process	Activate/Deactivate
3	Disk Percent Used	Activate/Deactivate
4	Kernel Outstanding Request	Activate/Deactivate
5	Number of Connected Users	Activate/Deactivate
6	Web CallObject Errors	Activate/Deactivate
7	UBE Runtime limit	Activate/Deactivate
8	XPlc Server Disk Low	Activate/Deactivate
9	Agent Down	Activate/Deactivate
10	Clear-Email Dead Process	Activate/Deactivate
11	Clear-No Email Dead Process	Activate/Deactivate
12	Server Bounce	Activate/Deactivate

- On the SMC – Work with Events screen, click Activate next to the event you want to configure (in this example, Disk Percent Used).

SAW - SMC - Activate Ports for Event



c:\saw\saw_doc.smc

Fri Mar 19 15:49:03 MST 2004

Working SMC file c:\saw\saw_doc.smc

Working event Disk Percent Used

Event type Threshold

Event Server type Ent

Warning value:

Alarm value:

[Home](#) [Frame View](#) [Port View](#)
[Global View](#)

Activate this Event for these Ports:

☒ 8004



Cancel:



- For each of the following events, enter a Warning Value and Alarm Value, or accept the default values:

Event	Parameter
Disk Percent Used	Number (in %) of disk space used.
Kernel Outstanding Request	Number of unprocessed kernel requests.
Web CallObject Errors	Number of CallObject kernel process errors for a business function activated through a web server.
UBE Runtime limit	Time (in minutes) that a UBE should normally take to complete.

4. Check the box next to each port number you want to monitor for this event.

Note

The event will be monitored on every server you assigned to the selected port numbers.

5. Click the Action button.
6. Repeat these steps to assign other events to port numbers.
7. Click on the Global View to see the current configuration.

PeopleSoft. back

SAW - SMC - SMC Global View
 c:\saw\saw_doc.smc Fri Mar 19 15:52:26 MST 2004

Select Task: Work with Servers Cancel: ☒ [Home](#) [Frame View](#) [Port View](#) [Global View](#)

Monitored Instances 3 Not monitored Instances 0 Total Instances 3
 Monitored Servers 3 Not monitored Servers 0 Total Servers 3

Saw Agent is not running!

Sort Server list

Server Name	OS	Ports	Server Type	Description	Monitored Events												
HPdev1 Stop Monitoring	HP-UX	6003 Stop Monitoring	JAS	E810 JAS	<table border="1"> <tr> <td>Server Down</td> <td>-</td> <td>-</td> </tr> <tr> <td>Number of Connected Users</td> <td>50</td> <td>50</td> </tr> <tr> <td>Web Call Object Errors</td> <td>15</td> <td>10</td> </tr> </table>	Server Down	-	-	Number of Connected Users	50	50	Web Call Object Errors	15	10			
Server Down	-	-															
Number of Connected Users	50	50															
Web Call Object Errors	15	10															
sundev1 Stop Monitoring	SUN-OS	6004 Stop Monitoring	Ent	E810_port	<table border="1"> <tr> <td>Server Down</td> <td>-</td> <td>-</td> </tr> <tr> <td>Email Dead Process</td> <td>-</td> <td>-</td> </tr> <tr> <td>Kernel Outstanding Request</td> <td>15</td> <td>10</td> </tr> <tr> <td>Number of Connected Users</td> <td>50</td> <td>50</td> </tr> </table>	Server Down	-	-	Email Dead Process	-	-	Kernel Outstanding Request	15	10	Number of Connected Users	50	50
Server Down	-	-															
Email Dead Process	-	-															
Kernel Outstanding Request	15	10															
Number of Connected Users	50	50															
Windev1 Stop Monitoring	Win-2000	6010 Stop Monitoring	XPlx	XPlx Server port													

Adding Statistics to Multiple Server Ports

Currently, the only choice of monitoring statistics is the User Connection Graph. The User Connection Graph is a bar chart showing the number of users connected to a server port as a function of time. You can configure multiple server ports to view this graph.

► To add statistics to multiple server ports

1. On SAW Home Page, select Work with Statistics from the Action field and click the Action button.
2. On SMC – Work with Statistics, click Activate next to User Connection Graph.
3. On SMC – Activate Ports for Statistics, check the box next to each port number for which you want to include statistics.
4. Click the Action button.

5. Click on the SMC Global View icon to see the current configuration.

Adding Profiles to Events

Once you have assigned events to server ports, you can assign e-mail and pager addresses of administrators to each event so that they are notified when an error is detected or a threshold exceeded with any of the server ports assigned to that event.

The e-mail and pager addresses of administrators are assigned to events as individual profiles. Each profile includes the administrator's name and up to three addresses. When you assign a profile to a particular event, this profile becomes attached to all server ports to which that event is assigned. If you want to assign different profiles to different server ports for the same event, you must customize your server ports.

► To add profiles to events

1. On SAW Home Page, select Work with Profiles from the Action field and click the Action button.
2. On SMC – Work with Profiles, verify that Add New Profile is selected, and click the Action button.
3. Enter the name and e-mail addresses of a person you want to include on the notification list.

You must enter an e-mail address in the First E-mail field. The 2nd and 3rd E-mail fields are optional. However, if you do not enter an e-mail address in the 2nd and 3rd E-mail fields, make sure to delete the default values from these fields.

4. Check each of the events for which you want the person to be notified.

Note

The person will be notified for every selected event on every server port to which the event is assigned.

5. Click the Action button.
6. Repeat these steps to configure other profiles.

You can view the current configuration in the SMC Global view.

Customizing Server Ports

In some cases, you may want to modify individual server ports and their events, statistics, and profiles without affecting other ports. You must configure each port for each server separately. If you were to choose Work with Ports from the Server Monitoring Configuration (SMC) action field, you would affect the configuration of every server that uses the same port number. Instead, you must go to a particular server configuration and modify each port number and its associated parameters individually.

► To change servers

1. On SAW Home Page, select Work with Servers in the Action field, and click the Action button.
2. To add a server, verify that Add new Server is selected in the Action field, and click the Action button.
3. On SMC – Adding new Server, complete the following fields:
 - Server Name
 - Description
 - Server OS
4. Click the Action button to add the new server to the SMC configuration.
5. Continue with other tasks in this section to add ports, events, statistics, and profiles.
6. To remove a server, click Remove Server next to the server you want to delete.
7. To modify a server configuration, click Modify Server next to the server you want to modify.
8. Change the fields as needed and click the Action button.

► To change server ports

1. On SAW Home Page, select Work with Servers in the Action field, and click the Action button.
2. In the row that lists the server you want to configure, click Work with Ports for this Server.
3. To add a port, verify that Add new Port is selected in the Action field, and click the Action button.
4. On SMC – Adding new Port, complete the following fields:
 - Port Number
 - Description
 - Server Type
5. Click the Action button to add the new port to the SMC configuration.
6. To remove a port, click Remove Port next to the port number you want to delete.
7. To modify a port, click Modify Port next to the port number you want to modify.
8. Change the fields as needed and click the Action button.

► To change statistics on a server port

1. On the SAW Home Page, select Work with Servers in the Action field, and click the action button.
2. In the row that lists the server you want to configure, click Work with Ports for this Server.
3. Click Work with statistics next to the port number you want to modify.
4. To add a statistic, verify that Add new Statistic is selected in the Action field, and click the Action button.
5. On Selecting SMC Statistic, click the Action button.

Note

SMC currently includes only one statistic you can add.

6. To delete a statistic, click Remove Statistic.

► To change events on a server port

1. On the SAW Home Page, select Work with Servers in the Action field, and click the action button.
2. In the row that lists the server you want to configure, click Work with Ports for this Server.
3. Click Work with Events next to the port number you want to modify.
4. To add an event, verify that Add new Event is selected in the Action field, and click the Action button.
5. On the SMC – Selecting new SMC Event screen, select an event and click the action button.
6. If the event includes threshold values, confirm the default values or modify them to suite your configuration.
7. Click the Action button.
8. To delete an event, click Remove Event next to the event you want to delete.
9. To modify an event, click Modify Event next to the event you want to modify.
10. Change the fields as needed and click the Action button.

► To change the profiles of a single event

1. On the SAW Home Page, select Work with Servers in the Action field, and click the action button.
2. In the row that lists the server you want to configure, click Work with Ports for this Server.
3. Click Work with Events next to the port number you want to modify.
4. Click Work with Profiles next to the event you want to modify.
5. To add a profile, verify that Add new Profile is selected in the Action field, and click the Action button.
6. On Adding new SMC Profile, complete the name and e-mail address of the person you want to notify.
7. Click the Action button.
8. To delete a profile, click Remove Profile next to the person you want to delete from the notification list.
9. To modify a profile, click Modify Profile next to the person whose profile you want to change.
10. Change the name and e-mail address as needed, then click the Action button.

Note

You can add new e-mail addresses, modify or delete an existing address, or change the person's name.

Running the SAW Agent

Once you configure Server Monitoring Configuration (SMC), you can run the SAW agent and begin monitoring the assigned server ports. When the agent is started, it cycles through each of the server ports, monitors the assigned events, and notifies the assigned users when a malfunction or warning threshold is crossed. The agent also collects and stores the statistical and event data in the SAW database.

Before you run the agent, it is important that your monitoring interval is longer than the time it takes the agent to cycle through all the server ports. The cycle time depends on the number of assigned server ports, the monitoring events, the statistics, and the number of e-mails it sends out. You can measure the cycle time by running the agent manually from the command prompt and noting the start times for each subsequent cycle. You can change or verify the cycle time on the SMC - Advanced Settings form.

The agent sends an e-mail message to assigned users when it first detects a malfunction or event threshold, but it does not send additional messages for the same event on subsequent cycles unless a new threshold is exceeded or a new event occurs.

► To run the SAW agent

The script file for the SMC file that you create is created under the SAW database directory.

Windows

The name of the file for Windows is:

`SawAgent_SMCFILE.bat`

Where SMCFILE is the name of the current SMC file (for example, saw.smc).

1. From Windows Explorer, go to the SAW database directory and double-click the SAW Agent script file.

UNIX

The name of the file for UNIX is:

`SawAgent_SMCFILE.sh`

Where SMCFILE is the name of the current SMC file (for example, saw.smc).

1. Go to the SAW database directory and run the script by entering the name of the script.
Make sure that the SAW database directory and the user running the agent have full access rights.

► To start the SAW Agent on the AS/400

Log in to the AS/400.

1. Enter the following command:

```
strqsh
```

This will start a Unix-like shell session.

2. Change directory to your SAW database directory.
3. Run the SAW agent by entering the name of the SAW Agent script. If your SMC file is `saw_xe.smc`. There will be a SAW Agent script created in the SAW database directory called:

```
SawAgent_saw_xe.smc.sh
```

4. In order to start the SAW Agent, enter the following command:

```
SawAgent_saw_xe.smc.sh
```

Performing SMC Actions

The Server Monitoring Configuration (SMC) Action field is a pull-down menu of functions that help administrators configure and monitor EnterpriseOne servers. The Action field is available at the top of most SMC forms. The Action fields are discussed in detail in the sections that describe how to configure the server ports. The following table briefly describes all of the available functions:

Action Field Function	Description
Work with Servers	Displays the Work with Servers configuration screen.
Work with Ports	Displays the Work with Ports configuration screen.
Work with Events	Displays the Work with Events configuration screen.
Work with Statistics	Displays the Work with Statistics configuration screen.
Work with Profiles	Displays the Work with Profiles configuration screen.
Work with SMC	Displays the SMC – Global Settings screen – same as home page.
Stop Monitoring All Servers	Deactivates all servers/ports in the SMC file so the agent will not monitor them. Servers/ports are not removed from the SMC file.

Action Field Function	Description
Start Monitoring All Servers	Activates all servers/ports in the SMC file so the agent will monitor them.
Force Agent Cycle	Manually restarts an agent cycle.
Submit	Validates the SMC settings you have entered.
Advanced Settings	Displays the advanced SMC settings.
Stop SAW Agent	If the SAW agent is running, stops the agent and cancels the monitoring cycle.
SMC Global View	Displays the SMC Global View.
View Agent log files	Displays agent logs.
Delete Agent log files	Deletes agent logs.
Perform SMC Actions	Access SMC configuration screens.

Working with SMC Views

Server Monitoring Configuration (SMC) includes the following views that allow you to view the status of each port and identify problems with the servers and processes:

- Work with SMC
- Port View
- Table View
- Frame View
- SMC Global View

You can access these views from most SMC screens by clicking the appropriate link on the right side of the screen.

Each view gives administrators a different perspective on the data and allows them to perform unique functions:

View	Description	Available Functions
Work with SMC	Displays the SMC configuration screens.	• Access configuration screens
Frame View*	Lists all the servers/ports by the type of server (enterprise, Web, XPI). This text-based frame view allows users to view the result of their requests in a separate frame in the same window.	• Force Agent Cycle • Delete Agent logs • Perform SMC Actions • Display status of server ports (up or down) • Link to a monitoring interface that is

		• Web, XPI) • Display agent logs • Display event logs • Display User Connection Chart • Display statistics file
Table View*	Lists all the server ports by the type of server (enterprise, Web, XPI). Graphical interface illustrates the status of a server port. The result of the requests will be displayed in a new window.	• Force Agent Cycle • Delete Agent logs • Perform SMC Actions • Display status of server ports (up or down) • Link to a monitoring interface that is specific to the server type (enterprise, Web, XPI) • Display agent logs • Display event logs • Display User Connection Chart • Display statistics file
Port View*	Displays an array of server ports by server name and port number. Graphical interface shows the status of each port. This view is created by the agent only if you enter a path for the HTML page secondary path on the Advanced Settings page.	• Display status of each port (up or down) • Link to a monitoring interface that is specific to the server type (enterprise, Web, XPI)
SMC Global View	Displays the contents of the entire SMC file and allows users to modify some of the settings.	• Force agent cycle • Modify events and e-mail addresses • Start or stop monitoring all servers • Start or stop monitoring all ports on a single server. • Display events and profiles for each server port. • Display global SMC settings

* The Frame View and Table View show the same information in different formats.

Note

The Work with SMC and SMC Global View forms are shown in *Adding EnterpriseOne servers* in the *Server and Workstation Administration Guide*. Examples of the Port View, Table View, and Frame View are shown below.

Frame View

The Frame View is organized by server type in left frame. The result of the request is displayed in the right frame.

PeopleSoft SAW - Agent Frame View - c:\saw\saw_doc.smc

View Events/Stats for last 1 day(s) and 0 hour(s)

SMC Global View

Monitored Servers (c:\saw\saw_doc.smc)

--- Enterprise Servers ---

+ (Ent) hp9000b/6012 : B9 - SP0.0

+ up for: 39 d. 1 h. 35 m. 37 s.

User chart Events Statistics file

--- JAS Servers ---

+ (JAS) den-sh5706272/80 : 0 user(s)

+ up for: 3 m. 11 s.

User chart Events Statistics file owhtml owportal

Server Administration Workbench (SAW) for Enterprise Servers

Host: hp9000b Port: 6012 Monitor

View: Server Summary sawEnt

1 zombie process(es) on the Server; Please refer to the process table

Server Information:

Server Status: Up for: 39 days, 1 hour, 46 minutes and 2 seconds

Version: B9 - SP0.0 Disk Used: 65% Log Size: 120 k Log Files: 9 Running Processes Log Size: 120 K

#	Process Name	Log	Debug	Process ID	Status	Requests	Outstand. Rqst	Users	Zombie	Max. #	Trace
1	ur15/shahin/b9_spt/system/bin32/idenet_n			2726	✓	6427	-	-	-	2	-
2	REPLICATION KERNEL			2728	✓	224999	0	-	-	1	On - Off
3	CALL OBJECT KERNEL			2729	ZOMBIE	0	0	-	1	-	-
4	SAW KERNEL			2747	✓	228533	0	-	-	1	On - Off
5	SECURITY KERNEL			2748	✓	224999	0	1	-	2	On - Off
6	idenet_n			6098	✓	1541	-	-	-	2	-

Table View

The Table View displays the result of the request is shown in a separate window.

PeopleSoft SAW - Agent Frame View - c:\saw\saw_doc.smc

View Events/Stats for last 1 day(s) and 0 hour(s)

Table View

Monitored Servers (c:\saw\saw_doc.smc)

--- Enterprise Servers ---

+ (Ent) hp9000b/6012 : B9 - SP0.0

+ up for: 39 d. 1 h. 35 m. 37 s.

User chart Events Statistics file

--- JAS Servers ---

+ (JAS) den-sh5706272/80 : 0 user(s)

+ up for: 3 m. 11 s.

User chart Events Statistics file owhtml owportal

SAW Agent Monitored Servers - c:\saw\saw_doc.smc

View Events/Stats for last 1 day(s) and 0 hour(s)

SMC Global View

Enterprise Servers	JAS Servers	XPIe Servers	XPIx Servers
 hp9000b/6012 : B9 - SP0.0 up for: 39 d. 1 h. 35 m. 37 s. User chart - Events - Statistics file	 den-sh5706272/80 : 0 user(s) up for: 3 m. 11 s. User chart - Events - Statistics file - owhtml owportal		
Enterprise Servers	JAS Servers	XPIe Servers	XPIx Servers

Port View

In the Port View, one server name is shown per column, and one port number is shown per row. The result of the request is shown in a separate window.

PeopleSoft® SAW - Agent Frame View - c:\saw\saw_doc.smc

View Events/Stats for last 1 day(s) and 0 hour(s) Port View

Monitored Servers (c:\saw\saw_doc.smc)

--- Enterprise Servers ---
 + (Ent) hp9000b/6012 : B9 - SP0.0
 + up for: 39 d. 1 h. 35 m. 37 s.
[User chart](#) [Events](#) [Statistics file](#)

--- JAS Servers ---
 + (JAS) den-sh5706272/80 : 0 user(s)
 + up for: 3 m. 11 s.
[User chart](#) [Events](#) [Statistics file](#) [owhtml](#) [owportal](#)

PeopleSoft®

PeopleSoft Mid-Enterprise Servers (Enterprise, JAS, XPI)- Effective Fri Mar 19 16:07:17 MST 2004

SMC File: c:\saw\saw_doc.smc

[Monitoring Enterprise Servers](#)
[Monitoring JAS Servers](#)
[Monitoring XPI Servers](#)

	den-sh5706272 Win-2000	hp9000b HP-UX	
80 B9-SP3	UP owhtml owportal	NA	80 B9-SP3
6012 B9	NA	UP Ent	6012 B9
	den-sh5706272 Win-2000	hp9000b HP-UX	

Accessing the Server Monitoring Interface

From the Port View, Table View, and Frame View, you can access a monitoring interface for each server by clicking the link for that server. Each interface allows you to monitor specific data and parameters for the selected server, and the interface is the same regardless of the view from which you start. Although the information on the form depends on the type of server being monitored (Enterprise, JAS, or XPI), the interface is the same.

The forms below show the monitoring interface for each of the three types of servers.

Enterprise Server Monitor

PeopleSoft®

Server Administration Workbench (SAW) for Enterprise Servers

Host: hp9000b Port: 6012 Monitor

View: Server Summary sawEnt

1 zombie process(es) on the Server; Please refer to the process table

Server Information:
 Server Status: Up for: 39 days, 1 hour, 48 minutes and 24 seconds
 Version: B9 - SP0.0 Disk Used: 65% Log Size: 120 k Log Files: 9 Running Processes Log Size: 120 K

#	Process Name	Log	Debug	Process ID	Status	Requests	Outstand. Rqst	Users	Zombie	Max. #	Trace
1	/u15/shahin/b9_spf/system/bin32/idenet_n			2726	✓	6427	-	-	-	2	-
2	REPLICATION KERNEL			2728	✓	225008	0	-	-	1	On - Off
3	CALL OBJECT KERNEL			2729	ZOMBIE	0	0	-		-	-
4	SAW KERNEL			2747	✓	228565	0	-	-	1	On - Off
5	SECURITY KERNEL			2748	✓	225008	0		-	2	On - Off
6	idenet_n			6098	✓	1589	-	-	-	2	-

Web Server Monitor

PeopleSoft® back

Server Administration Workbench (SAW) for JAS Servers

Host: Port:

Views: sawJas

Web Server: den-sh5706272:80

Server Status: Up for 19 minutes and 9 seconds

No users currently logged in.

Host	MaxWaited	Waiting	Current Size	Busy	CallObjs	SysErrs	Timeouts
HP9000B:6012	0	0	1	0	0	0	0

Heap memory allocated by VM: 5,935,096 bytes

Heap used memory in VM: 5,709,832 bytes

Available (free) memory in VM: 225,264 bytes

Total threads in VM: 60

XPI Server Monitor

PeopleSoft®

Server Administration Workbench (SAW) for XPI Servers

Host: Port: sawXPIe

Views:

Broker Summary Mon Mar 22 14:41:11 MST 2004 [New Window](#)

XPIe SERVER den-ur6864969/6849

Server Status Up for: -57 seconds

NON SSL CONNECTIONS 17 SSL CONNECTIONS 0

HIGHEST NON SSL CONNECTIONS 17 HIGHEST SSL CONNECTIONS 0K

DISK SPACE > 1 Meg [View Server Statistics](#)

#	broker name	broker host	territory	description	Event Types	Client States	Clnt Grps
1	Broker #1	den-ur6864969	null	-			

Monitoring Enterprise Servers from the Web

The Enterprise Server Monitor allows you to use the Web to monitor an EnterpriseOne enterprise server. The Enterprise Server Monitor uses SAW APIs that have been converted to Java code. To create the Web interface for the Enterprise Server Monitor, the JAS server calls newly developed servlets, which in turn call SAW APIs to retrieve information from the enterprise server.

To monitor an enterprise server from the OW-XPI workspace or from the sawent.html page, you must enter a valid host name and port number in the Enterprise Server Monitor workspace. Choose a view from the dialog box, and then click the Monitor button to view data. Each view contains a header with the following information on the enterprise server:

- Enterprise server name and port number, such as sundev/6012
- EnterpriseOne release, such as 810 - Service Pack 1.0
- Server status, either up or down
- Percentage of disk space used
- Server log file size (in kilobytes)
- Number of server log files
- Size of log files (in kilobytes) of process currently running on the server

Server Summary View

The Server Summary view allows you to conveniently monitor the status of the Enterprise server and the processes that are running on it.

1 zombie process(es) on the Server; Please refer to the process table

Server Information:
 Server Status: Up for: 39 days, 1 hour, 55 minutes and 22 seconds
 Version: B9 - SP0.0 Disk Used: 65% Log Size: 120 k Log Files: 9 Running Processes Log Size: 120 K

#	Process Name	Log	Debug	Process ID	Status	Requests	Outstand. Rqst	Users	Zombie	Max. #	Trace
1	/u15/shahin/b9_spt/system/bin32/idenet_n			2726	✓	6427	-	-	-	2	-
2	REPLICATION KERNEL			2728	✓	225036	0	-	-	1	On - Off
3	CALL OBJECT KERNEL			2729	ZOMBIE	0	0	-		-	-
4	SAW KERNEL			2747	✓	228608	0	-	-	1	On - Off
5	SECURITY KERNEL			2748	✓	225036	0		-	2	On - Off
6	idenet_n			6098	✓	1621	-	-	-	2	-

The Server Summary view displays in a table the following essential information about processes running on the server:

- Process name.
- Log file information. Click the magnifying glass to view the process log file.
- Debug log file information. Click the magnifying glass to view the process debug file.
- Process ID.
- Process status. A check mark indicates that the process is active. If the entire row displaying information on the process is red and the status is ZOMBIE, the server process is dead.
- Total number of requests processed.

- Number of outstanding requests for the process. An increase in this number indicates a potential problem with the process.
- Users of server processes. Click the magnifying glass to view the list of users connected to a kernel process.
- Zombie. Click the icon to remove the zombie process from the table.

The Server Summary view displays the following essential information only if the server you are monitoring supports it:

- Max #. Maximum number of processes to be created on the server. Click the number to change the value.
- Trace. Click on_off to activate or deactivate the debugging for the process.

Note

You can view further details on server processes by clicking on the process name in a row of the table.

See Also

- *Monitoring EnterpriseOne Server Processes* in the *Server and Workstation Administration Guide* for a detailed discussion of enterprise server processes that you monitor using SAW

Process Name

You can click the URL that identifies an enterprise server process name to view detailed information about network, kernel, and queue processes.

When you click the URL, the Process information view appears.

Server Administration Workbench (SAW) for Enterprise Servers

Host: hp9000b

Port: 6012

Monitor

View: Server Summary

sawEnt

 **1 zombie process(es) on the Server; Please refer to the process table**
Server Information:

Server Status: Up for: 39 days, 1 hour, 56 minutes and 31 seconds

Version: B9 - SP0.0 Disk Used: 65% Log Size: 120 k Log Files: 9 Running Processes Log Size: 120 K

Process debug size: 30750 Process log size: 365 Running processes: 6 IPC resources: 22

BSFN Tracing Information

No BSFN information for this process.

Process General Information

process name	process ID	Start time	Last active time	status	Rqst	Outstand. Rqst	Log	Debug
CALL OBJECT KERNEL	2729	Mon Feb 09 14:31:40 MST 2004	Mon Feb 09 15:43:10 MST 2004	ZOMBIE	0	0		

Parent ID	Job ID	User ID	Group ID	User name
2726	0	258	20	h5706272

process Argument list

6012

Kernel Specific Information

last message	total process time	total request	average time	index range	Outstand. Rqst
Mon Feb 09 14:31:40 MST 2004	0	0	0	0	0

In addition to general server information, this view displays the size of the log and debug log files for the process you chose, as well as the total number of processes running on the server and the total number of running interprocess communication (IPC) resources.

A table displays the parameters for the process, with values for each. Many of these parameters also appear in the Server Summary view. However, the view for a specific process contains the following additional information:

- Time and date the process started
- Time and date the process was last active or, for zombie processes, the time and date the process died
- Parent ID of the process
- Job ID of the process
- User ID of the process
- Group ID of the process
- User name

If the process is a security kernel or a CallObject kernel and a current user exists for the process, additional information appears.

You can view the log file and debug log file for a process by clicking the magnifying glass icon under the appropriate column in the table.

Note

You can also view the log file or the debug log file for a process by clicking the magnifying glass in the appropriate row and column of the table in the Server Summary view.

Depending on the kind of process that you choose, the Process information view displays additional information. For example, if you choose a network process, the Process information view displays the values for the following network messaging parameters:

- Time of the last message
- Total processing time (in milliseconds)
- Total number of messages
- Average processing time
- Stream port, which is the connection port for the movement of data
- Data port, which is the transfer port for the movement of data
- Number of in connections
- Number of out connections
- Connection type (incoming or outgoing)
- IP address

If you choose a kernel process, the Process information view displays the following kernel-specific information:

- Index range, which defines the kind of kernel process that is running. For example, the value of the index range parameter for a security kernel is 3. These values are defined in the server's jde.ini file.
- Number of outstanding requests for the process.

If you choose a queue process, the Process information view displays the following queue-specific information:

- Server on which the UBE is running
- EnterpriseOne environment on the server on which the UBE is running
- Name of the UBE batch queue and its number

Log

From the Server Summary view, you can view the log file for a server process by clicking the magnifying glass in the Log column of the table containing the process information.

When you click the magnifying glass icon, a web page containing the log file appears.

Enterprise Server Log file

Server	hp9000b
Port	6012
Log type	Log
Process name	/u15/shahin/b9_spf/system/bin32/jdenet_n
Process ID	2726
File size	15684

```

2726    Mon Feb  9 14:31:39.239807  ipcmisc.c234
        process 2726 (/u15/shahin/b9_spf/system/bin32/jdenet_n) registered in entry 0

2726    Mon Feb  9 14:31:39.457646  netconn.c1386
        Bind for socket 6 on port 6012 failed. Port in use, now forcing the bind. Please verify that no other

2726    Mon Feb  9 15:43:10.726354  netsig.c153
        net process: addip process 2729 in Unregister list.

2726    Mon Feb  9 15:43:10.726705  netsig.c333
        Kernel Process 2729 has died

2726    Mon Feb  9 15:43:10.727113  netsig.c161
        net process: process 2729 set to Zombie

2726    Mon Feb  9 15:43:10.735774  ipcmisc.c234
        API ipcSawZombieProcV1 : process 2729 set to Zombie in entry 2

2726    Mon Feb  9 17:51:27.998756  netrcv.c351
        232-Error from Host=(10.3.161.91) in rcv returned Connection reset by peer

```

The web page displays the log file with a header that contains the process name, the process ID, and the file size, in kilobytes.

Debug

From the Server Summary view, you can view the debug log file for a server process by clicking the magnifying glass in the Debug column of the table containing the process information.

When you click the magnifying glass, a web page containing the debug log file appears.

Enterprise Server Log file

Server	hp9000b
Port	6012
Log type	Debug
Process name	CALL OBJECT KERNEL
Process ID	2729
File size	30750

```

Feb  9 14:31:40.799407  DEBUG INIT0          - 2729  ****  jdeDebugInit -- output to file.
Feb  9 14:31:40.808350  ipcmisc.c233         - 2729  process 2729 (jdenet_k) registered in entry 2
Feb  9 14:31:40.808657  ipcmisc.c234         - 2729  process 2729 (jdenet_k) registered in entry 2
Feb  9 14:31:40.809026  jdenet_k.c174       - 2729  setrlimit not needed, open file soft limit (2048) was )= hard limit (2048)
Feb  9 14:31:40.841487  netcfg.c250         - 2729  Starting Kernel of Type:CALL OBJECT KERNEL
Feb  9 14:31:40.841746  netcfg.c250         - 2729  Starting Kernel of Type:CALL OBJECT KERNEL
Feb  9 14:31:41.493810  netflow.c977        - 2729  About to call dispatch function, flags=0x0, Type=11
Feb  9 14:31:41.494348  winansi.c1529       - 2729  Adding library (jdenet_k) to list
Feb  9 14:31:41.494515  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libjdelib.sl) to list
Feb  9 14:31:41.494592  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libjdeipc.sl) to list
Feb  9 14:31:41.494664  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libjdenet.sl) to list
Feb  9 14:31:41.494736  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libjdethread.sl) to list
Feb  9 14:31:41.494807  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libnettest.sl) to list
Feb  9 14:31:41.494878  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libcallobj.sl) to list
Feb  9 14:31:41.494948  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/liberror.sl) to list
Feb  9 14:31:41.495019  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libgentext.sl) to list
Feb  9 14:31:41.495091  winansi.c1529       - 2729  Adding library (/u15/shahin/b9_spf/system/lib/libjdb.sl) to list

```

Users

If you click a magnifying glass in the Users column of the table in the Server Summary view, the Enterprise Server Monitor displays a User List. The User List contains the user name, the name of the user's machine, and, if available, the time the user connected to the process.

Zombie

Zombie is the status of a process that dies on the server. When a process dies on the Enterprise server, a mechanism on the server keeps the dead process in Server Administration Workbench (SAW) registration shared memory and changes its status from Active to Zombie. The process appears in the process list with a status of Zombie even if a process is created and dies between SAW data-retrieval cycles. Although the process is not present in the operating system on the Server, the process appears in the process table until the Administrator removes it.

In the Server Summary view, the row containing the summary information for a Zombie process appears in red, and the Zombie parameter column contains a trashcan icon. If you have SAW administrative rights, you can click the icon to delete the dead process from the server.

All Log Files View

If you click All Log Files in the dialog box in the Enterprise Server Monitor workspace, the Log File Summary view appears.

PeopleSoft®

Server Administration Workbench (SAW) for Enterprise Servers

Host: hp9000b Port: 6012 Monitor

View: All Log Files sawEnt

1 zombie process(es) on the Server; Please refer to the process table

Server Information:

Server Status: Up for: 39 days, 2 hours, 1 minute and 38 seconds

Version: B9 - SP0.0 Disk Used: 65% Log Size: 120 k Log Files: 9 Running Processes Log Size: 120 K

Deleting files bigger than 2 M Delete

#	file name	process name	last modification time	file size	Delete
1	jdedebug_2726.log	/u15/shahin/b9_spt/system/bin32/jdenet_n	Wed Mar 10 10:18:47 MST 2004	30935	
2	jdedebug_2729.log	CALL OBJECT KERNEL	Mon Feb 09 15:42:59 MST 2004	30750	
3	jde_2728.log	REPLICATION KERNEL	Mon Feb 09 14:31:42 MST 2004	22536	
4	jde_2726.log	/u15/shahin/b9_spt/system/bin32/jdenet_n	Wed Mar 10 10:18:47 MST 2004	15684	
5	jdedebug_6098.log	jdenet_n	Tue Mar 16 16:13:54 MST 2004	9205	
6	jde_2748.log	SECURITY KERNEL	Tue Mar 09 09:59:40 MST 2004	4873	
7	jde_6098.log	jdenet_n	Tue Mar 16 16:13:54 MST 2004	4637	
8	jde_2747.log	SAW KERNEL	Mon Feb 09 14:31:45 MST 2004	1400	
9	jde_2729.log	CALL OBJECT KERNEL	Mon Feb 09 14:31:41 MST 2004	365	

The Log File Summary view displays the names of all the log files and debug log files on the server, as well as the time each file was created and the size of each file in kilobytes. You can set the threshold file size for the files to be deleted. If you have SAW administrative rights, click the Delete button to start the deletion process.

To view a log file or debug log file for a process, click the URL in the Name column of the table in the Log File Summary view. A web page containing the file appears. A header on the page identifies the file name and its size.

If you have SAW administrative rights, you can delete a log file or debug log file by clicking the trash can icon in the Delete column of the table in the Log File Summary view.

See Also

- *Setting up Security Access for SAW* in the *Server and Workstation Administration Guide* for more information about assigning SAW administrative rights

Active Log Files View

The Active Log Files view displays log file and debug log file information only on those processes currently running on the server.

Server Administration Workbench (SAW) for Enterprise Servers

Host: Port:

View: sawEnt

1 zombie process(es) on the Server; Please refer to the process table

Server Information:
 Server Status: Up for: 39 days, 2 hours, 2 minutes and 33 seconds
 Version: B9 - SP0.0 Disk Used: 65% Log Size: 120 k Log Files: 9 Running Processes Log Size: 120 K

Deleting files bigger than

#	file name	process name	last modification time	file size	Delete
1	idedebug_2726.log	/u15/shahin/b9_sp0/system/bin32/jdenet_n	Wed Mar 10 10:18:47 MST 2004	30935	
2	idedebug_2729.log	CALL OBJECT KERNEL	Mon Feb 09 15:42:59 MST 2004	30750	
3	jde_2728.log	REPLICATION KERNEL	Mon Feb 09 14:31:42 MST 2004	22536	
4	jde_2726.log	/u15/shahin/b9_sp0/system/bin32/jdenet_n	Wed Mar 10 10:18:47 MST 2004	15684	
5	idedebug_6098.log	jdenet_n	Tue Mar 16 16:13:54 MST 2004	9205	
6	jde_2748.log	SECURITY KERNEL	Tue Mar 09 09:59:40 MST 2004	4873	
7	jde_6098.log	jdenet_n	Tue Mar 16 16:13:54 MST 2004	4637	
8	jde_2747.log	SAW KERNEL	Mon Feb 09 14:31:45 MST 2004	1400	
9	jde_2729.log	CALL OBJECT KERNEL	Mon Feb 09 14:31:41 MST 2004	365	

PrintQueue Log Files View

The PrintQueue Log Files view has the same features as the Log Files view, but it only displays logs in the PrintQueue directory.

Disk Usage View

If you click Disk Usage in the dialog box of the Enterprise Server Monitor, the Disk Usage Summary view appears.

Server Administration Workbench (SAW) for Enterprise Servers

Host:

Port:

View:

sawEnt

 **1 zombie process(es) on the Server; Please refer to the process table**

Server Information:

Server Status: Up for: 39 days, 2 hours, 3 minutes and 36 seconds

Version: B9 - SP0.0

Disk Used: 65%

Log Size: 120 k

Log Files: 9

Running Processes Log Size: 120 K

#	disk name	kilo bytes	used	free	% used	mount
1	/dev/vg00/lvol3	311296	59264	250128	19	/
2	/dev/vg00/lvol1	298928	49728	219304	18	/stand
3	/dev/vg00/lvol12	8912896	160696	8683856	2	/var
4	/dev/vg00/lvol11	516096	27335	459296	6	/var/tmp
5	/dev/vg00/lvol10	2048000	602385	1355286	31	/var/adm/sw
6	/dev/vg00/lvol9	262144	54568	194666	22	/var/adm/sa
7	/dev/vg00/lvol8	2048000	1117383	878828	56	/var/adm/patch
8	/dev/vg00/lvol7	4096000	1563136	2513120	38	/usr
9	/dev/vgdata01/lvuxtools	5406720	2297607	2916738	44	/usr/local/uxtools
10	/dev/vgdata01/lvtools	819200	208974	572107	27	/usr/local/tools
11	/dev/vgdata01/lvusers	516096	83447	405639	17	/users
12	/dev/vgdata01/lvu17	10240000	4139480	5915682	41	/u17
13	/dev/vgdata01/lvu16	10240000	4951762	5123018	49	/u16
14	/dev/vgdata01/lvu15	10240000	6581744	3548594	65	/u15
15	/dev/vgdata01/lvu14	10240000	7022948	3116640	69	/u14
16	/dev/vgdata01/lvu13	10240000	4950482	5126200	49	/u13
17	/dev/vgdata01/lvu12	10240000	5027244	5049886	50	/u12

The Disk Usage Summary view displays the following information on the disks that make up the Enterprise server that you are monitoring:

- Disk name
- Volume ID
- Used space (in kilobytes)
- Free space (in kilobytes)
- Percentage of space used
- Type of disk mount

A green line in the table indicates the disk on which the enterprise server is running. If a line is red, 80 percent or more of the space on the disk has been used.

INI Settings View

The INI Settings view displays all the sections of the INI files along with most of their associated variables and values. This view does not display sensitive user information, such as Username and Password. The view allows you to display individual sections or a particular variable (INI File Key) by typing the section or variable in the appropriate field, and clicking the Submit button (>>).

Package Build Files View

The Package Build Files view displays the Package Build directory and allows you to navigate through all subdirectories, drill down to individual files, and click a file to view its contents. The Package Build directory is defined by the PackagePath setting in the [Install] section of the JDE.INI file.

Environment Variables View

The Environment Variables view displays all the environment variables associated with an Enterprise Server. You can search for a particular variable by typing the variable name in the Environment Variable name field.

Troubleshooting Call-Object Processes

When a call-object process dies, EnterpriseOne automatically saves the state of the process and its associated information. You can inspect this information by using the sawent.html page and typing in the server name and port number. The call-object process name will be displayed in the Enterprise Service Monitor table, even if it is no longer active. By clicking the call-object process name, you can view the saved state of the process and its associated information.

Note

This feature is only available with SP 18 and later.

Monitoring the JAS Servers from the Web

The Web Server Monitor provides a unified interface that administrators can use to monitor processes on the Web server and business functions running on enterprise servers connected to the Web server. Using the Web Server Monitor, you can monitor the following, in real time:

- Web server users and their use of the server
- JDENET connection pools to enterprise servers
- Business functions running on enterprise servers connected to the Web server
- Drivers used by the Web server to connect to data sources
- Log and debug files

The Web Server Monitor provides a continuous stream of Web server data that can help you troubleshoot and tune performance. The Web Server Monitor lets you to monitor the server when you are running EnterpriseOne on a HTML client. Thus, even with a client configuration that lacks the Server Administration Workbench executable, you can monitor the Web server simply—by typing in the server's Universal Resource Locator (URL).

Web Server Monitor Technical Considerations

You use the Web Server Monitor only to monitor Web server activities. You cannot use it to monitor Enterprise servers, although you can monitor the CallObject tasks that a Web server submits to an Enterprise server. For detailed monitoring of enterprise servers, use the Server Administration Workbench.

The Web Server Monitor will monitor only one server at a time, although you can change your monitoring from one Web server to another by typing in the name of a different server.

► To access the Web Server Monitor

Open your web browser and enter a URL with the following:

- Name of the EnterpriseOne web server
- Directory where the Web Server Monitor files are located

For example: `http://owweb1/jde/saw/sawWeb.html`

The Web Server Monitor form appears.

Configuring the Web Server Monitor

Using the Configure view of the Web Server Monitor, you can set the values of four monitoring parameters or accept the monitor's default values for the parameters:

- CallObject Timeout, in minutes. The default setting is 1.5 minutes.
- Maximum Host Pool, in number of enterprise server connections. The default setting is 50 connections.
- Maximum JAS Log Retrieve Size, in kbytes. The default setting is 50 kbytes.
- Maximum JAS Debug Log Retrieve Size, in kbytes. The default setting is 50 kbytes.

In addition, the Configure view contains a Flush OCM button. If the Object Configuration Manager (OCM) mappings on the server change, you can click this button to flush the OCM cache. Subsequent caching will reflect the new OCM mapping.

Note

The parameter values that you set in the Configure view are not persistent. You must reset them each time that you use the Web Server Monitor.

► To configure the Web Server Monitor

1. On Web Server Monitor, choose Config from the pull-down menu in the Views field.
2. Choose or enter values in the following fields:
 - CallObject Timeout

- Maximum Host Pool
 - Maximum JAS Log Retrieve Size
 - Maximum JAS Debug Log Retrieve Size
3. Click Refresh.
 4. To flush the OCM cache, click Flush OCM.

After you click Flush OCM, the following sentence appears above the button: "The OCM cache has 0 entries."

Monitoring the Web Server System

The Web Server Monitor form initially displays system summary information for the Web server. The System Summary view offers a quick synopsis of the system's operations at the point in time that you access it. The page will refresh automatically.

The System Summary view displays the following server operations data:

- Length of time the system has been running
- Number of users currently connected to the Web server
- Number of Web server users who have been active in the last 10 minutes
- Information about enterprise servers connected to the Web server
- Statistical data on tasks performed by the virtual machine software

In addition, the System Summary view displays a warning message if the log file exceeds 1 megabyte in size.

The most detailed information in the System Summary view relates to the enterprise servers connected to the Web server. The following table summarizes the information displayed in the table in the System Summary view:

System Summary Table Parameter

Description

Host	Name of the Enterprise server connected to the Web server and port number.
Max Waited	The number of users waiting. For example, if the maximum number of connections allowed is 10 (as defined in the server's jde.ini file), you see no value for the Max Waited parameter until the eleventh user tries to connect. The Max Waited parameter value is then 1. If another user tries to connect, the value is 2, and so on.
Waiting	The number of users currently waiting for a socket connection. The parameter has no value until the number of users trying to connect exceeds the number of socket connections allowed by the server's jde.ini file.
Current Size	The number of socket connections in the connection pool.
Busy	The number of socket connections currently in use.
CallObjs	The total number of business functions that have run on the enterprise servers while the system has been up.
SysErrors	The total number of system errors logged from business functions running on the Enterprise servers.
Timeouts	The total number of business functions that timed out while running on the enterprise servers.

To navigate to another view from the System Summary view or from any other view in the Web Server Monitor, click the scroll button in the upper left-hand corner of the form. A drop-down menu displays each available view in the Web Server Monitor, as listed and described below:

Web Server Monitor View	Information Displayed
User List	Data about individual Web server users.
CallObject Info	Data about business functions running on each Enterprise server connected to the Web server.
Log Files	JAS.log and JASdebug.log file information, and WebSphere application server information written to the stdout and stderr files.
Environment	Properties of the platform's virtual machine, including its full version and path.
Config	Configuration parameters for the Web Server Monitor (you can change these).
JDENET Pool	Data about pools of Enterprise server connections to the Web server.
JDBC Pool	Data about database drivers and pools of connections to the database.
Virtual Clients	Data about clients using the HTML interface. The parameter represents a subset of the User List parameter; the User List includes clients using both the HTML and JavaApplet interfaces.

Outstanding Requests	List of all users waiting for a response from the Web server.
Thread List	List of all threads, thread groups, and thread pools running on the Web server.

See Also

- ❑ *Monitoring JDENET Connection Pools* in the *Server and Workstation Administration Guide* for more information on connections to the enterprise server
- ❑ *Monitoring Business Functions Running on Connected Enterprise Servers* in the *Server and Workstation Administration Guide* for more information on monitoring business functions running on an enterprise server connected to the Web server

Monitoring Web Server Users

The System Summary view displays aggregate information about Web server users. The User List view displays data about individual users connected to the Web server. The list of users displayed in this view includes users who are currently logged on. The Web Server Monitor updates the user data each time you click the Refresh button.

You can use the User List view to see the following user data, which is not displayed in the System Summary view:

- The maximum number of users allowed on the server
- The maximum number of users who have connected to the server at any point during a session
- The number of user sessions serviced by the server
- The number of minutes a session can be idle before it times out
- The users who are currently active, that is, those who have not timed out of a session
- Active users who have accessed the Web server in the last 10 minutes
- The virtual clients associated with each user

Note

Starting with SP22, PeopleSoft introduced the Multiple Application Framework (MAF). MAF enables each user to open multiple applications in separate browsers. Each of these applications is associated with a unique virtual client. SAW allows an administrator to monitor all the virtual clients, and groups them by user. You can view virtual clients by clicking “View Virtual Clients” above the table on the User List view. The virtual clients appear as records under each user session.

PeopleSoft.

Server Administration Workbench (SAW) for JAS Servers

Host: denxpi32 Port: 85

Views: User List Refresh sawJas

User List

JAS Server: denxpi32:85 Up for: 10 days, 20 hours, 30 minutes and 30 seconds

Max Users Allowed: 100 Max Users Connected: 5 Session Serviced: 106 Session Timeout: 60

Current Users: 5 Active Users Last 10 Minutes: 3

[View With Virtual Clients](#)

#	User Name	Environment	Display Mode	Computer Name	Login Time	Idle Time	Host-PID	Session ID
1	JS8441257	STGAWU	HTML	den-js8441257.jdedwards.com	03/19/2004 12:37 PM	10 Minutes	DENXP112-3088	3690578305780770816
2	JV5443249	STGAWU	HTML	den-jv5443249.jdedwards.com	03/19/2004 12:40 PM	11 Minutes	DENXP112-6204	2693743573961615360
3	RP6126376	STGAWU	HTML	den-rp6126376.jdedwards.com	03/19/2004 12:52 PM	1 Minutes	DENXP112-3952	3108466456162497536
4	RP6126376	STGAWSC7	HTML	den-rp6126376.jdedwards.com	03/19/2004 12:51 PM	1 Minutes	OWNTS1-9900	3910850935703691264
5	RP6126376	STGAWU	HTML	den-rp6126376.jdedwards.com	03/19/2004 12:34 PM	18 Minutes	DENXP112-3096	1473803373290828800

A table in the User List view displays parameters and values for each user session. The following table summarizes the information displayed in the table in the User List view:

Parameter in Table of User List View	Description
User Name	EnterpriseOne login ID.
Environment	The path code and OCM mappings that define the particular environment in which the user is working, such as PDEVNIS2.
Display Type	The way that the EnterpriseOne Web page is displayed, such as HTML.
Computer Name	The name of the user workstation.
Login Time	The time and date the user connected to the EnterpriseOne Web server.
Idle Time	The amount of time, in minutes, that a user has been inactive during an EnterpriseOne session. The parameter value is Active for any user who has performed an action in EnterpriseOne within the last minute; a row is shaded green for users active in the last five minutes.
Host	The Enterprise server to which the user has connected during the session. This parameter can be blank.
Agent ID	The workstation browser type.
Session ID	A unique numerical identifier for each user session.

The table provides user information that can be useful in troubleshooting. For example, you might check the Agent ID parameter to see if the user has a browser that is out of date. You might also correlate problems with business functions running on an Enterprise server with a particular environment.

► To monitor web server users

1. In the System Summary view of the Web Server Monitor, click the scroll button and choose User List.

The Web Server Monitor displays the User List view.

2. Click the Refresh button to get the latest user information.

The Web Server Monitor displays information on users connected to the web server.

3. Click on any table column heading to sort the column data.

Monitoring JDENET Connection Pools

The Web Server Monitor allows you to monitor socket connections made from the Web server across the network to an Enterprise server, where business functions run.

The JDENET Connection Pool Information view displays connection and messaging data for each enterprise server connected to the Web server. The information is displayed in a table, as follows.

PeopleSoft.

Server Administration Workbench (SAW) for JAS Servers

Host: denxpi32 Port: 85

Views: JDENET Pool Refresh

Web Server denxpi32:85

Server Status Up for: 10 days, 19 hours, 5 minutes and 28 seconds

Server Time Fri Mar 19 11:29:05 MST 2004

Host	MaxPool	MaxWaited	Waiting	Current Size	Bytes Tx	Bytes Rx		
DENXPI12:6080	50	0	0	1	178.9K	318.5K		
	ID	port	inUse	bytesTx	bytesRx	msgTx	msgRx	inUseTime
	3925	2149	false	183,147	326,189	204	192	n/a
Host	MaxPool	MaxWaited	Waiting	Current Size	Bytes Tx	Bytes Rx		
OWNTS1:6080	50	0	0	1	2.8K	264		
	ID	port	inUse	bytesTx	bytesRx	msgTx	msgRx	inUseTime
	4235	1270	false	2,912	264	14	3	n/a
Host	MaxPool	MaxWaited	Waiting	Current Size	Bytes Tx	Bytes Rx		
HP9000B:6012	50	0	0	1	40.1K	420.3K		
	ID	port	inUse	bytesTx	bytesRx	msgTx	msgRx	inUseTime
	3991	49413	false	41,030	430,360	40	40	n/a

Note that connection data displayed in the JDENET Connection Pool Information view and the System Summary view correlate. For example, the Current Size parameter value in the System Summary view represents the current number of socket connections made between the Web server and an Enterprise server. The Connection Pool Information view also displays the Current Size parameter value.

The Connection Pool Information view also displays the maximum number of connections that can be made between the Web server and the Enterprise server, as well as the number of users waiting for a connection. Information about the volume of data and messages transmitted and received is contained in the following set of parameters:

- Bytes Tx, the value of which represents the number of bytes transmitted from the connection socket
- Bytes Rx, the value of which represents the number of bytes received from the connection socket
- msgtx, which represents the number of JDENET messages transmitted from the connection socket
- msgrx, which represents the number of JDENET messages received from the connection socket

► **To monitor connection pools**

1. In the System Summary view of the Web Server Monitor, click the scroll button and choose JDENET Pool.

The Web Server Monitor displays the JDENET Connection Pool Information view.

2. Click the Refresh button to get the latest connection information.

The Web Server Monitor displays connection and messaging information for each enterprise server.

Monitoring Business Functions Running on Connected Enterprise Servers

The Web Server Monitor maintains an ongoing, cumulative list of information about each business function that runs on an Enterprise server connected to the Web server.

The CallObject Information view displays data for each business function that has run on a connected enterprise server during a Web server session. The Web Server Monitor separates the business function data by Enterprise server and presents the data in lists that can be sorted.

Server Administration Workbench (SAW) for JAS Servers

Host: Port:

Views:

CallObject Information

Web Server denxpi32:85
 Server Status **Up for: 10 days, 19 hours, 8 minutes and 51 seconds**
 In-progress
 none

History

Host	CallObjs	SysErrs	Timeouts					
DENXPI12:6080	3,098	44	30					
CallObject	Called	firstTime	maxTime	minTime	avgTime	sysErrs	timeouts	
AcctRecFSBeginDoc	3	1,656	235	94	164	0	0	
AcctRecFSClearStack	2	16	16	16	16	0	0	
AcctRecFSEditLine	2	1,203	250	250	250	0	0	
AccumulateAmount	10	0	16	0	7	0	0	
AccumulateTotalAmount	326	0	32	0	9	0	0	
AdditionalEmployeeInformation	1	438	438	438	438	0	0	
BatchReviseOnExit	5	47	297	62	152	0	0	

Note that the Web Server Monitor does not perform a periodic flush of the list of business functions. Items in the list continue to accumulate until a server session ends. To make searching easier, you can sort the items in the list by clicking any heading. For example, you can display the list in descending order by amount of time required to run the business function.

The Web Server Monitor highlights (in red) the total number of business function errors for each Enterprise server and the total number of errors for each business function.

A table displays business function data for each Enterprise server. Each table displays the following general and cumulative information at the top:

- Server name
- Total number of business functions run on the server
- Total number of system errors
- Number of timeouts that occurred during the running of a business function

The table also displays cumulative data about each type of business function that has run during a server session. The following table summarizes the information displayed for each business function in the CallObject Information view:

Parameter in Table of CallObject Information View	Description

CallObject	The name of the business function.
Called	The total number of times the business function was called for processing on the enterprise server.
First time	The time, in milliseconds, that a business function took to run the first time it was called.
maxTime	The maximum time, in milliseconds, that a business function took to run when it was called.
minTime	The minimum time, in milliseconds, that a business function took to run when it was called.
avgTime	The average time, in milliseconds, required to run a business function.
sysErrors	The total number times a business function produced an error. If this parameter has a value, it appears in red.
Timeouts	The total number of times that a business function timed out before completion. If this parameter has a value, it appears in yellow.

The CallObject Information view also displays the number of business function processes in progress. This number can be important for troubleshooting purposes. A large number of in-progress processes may indicate that the enterprise server is running slowly and alert you that you need to investigate the possible causes.

► **To monitor business functions running on connected enterprise servers**

1. In the System Summary view of the Web Server Monitor, click the scroll button and choose CallObject Info.
The Web Server Monitor displays the CallObject Information view.
2. Click the Refresh button to get the latest connection information.
The Web Server Monitor displays connection and messaging information for each enterprise server.
3. To find business function errors, choose an Enterprise server and click the sysErrors column heading of the table.
If the calls with errors are not grouped at the top of the table after you click the column heading, click the heading again.
4. Click any of the other column headings to sort the parameter values.

Monitoring Database Connection Pools

The JDBC Connection Pool Information view allows you to monitor Web server connections to the database and the drivers that facilitate the connections.

From the JDBC Connection Pool Information view, you can view data on each available database connection, including:

- Last used, the value of which represents the number of seconds since a connection was last used
- First access, the value of which represents the number of seconds that have elapsed since a connection was first created

The JDBC Connection Pool Information view contains a Pool Cleaner button that; when you click it, starts a thread that goes through all the connections in pools and immediately runs a statement to determine the healthiness of each one. If the Pool Cleaner program finds a connection that does not work, it replaces the connection automatically.

You can also view additional information about an individual database connection by clicking the name of a pool key. The pool keys represent the database connections that meet the following criteria:

- Valid connection string, which is the URL for the database
- Valid database user ID
- Valid database password

The JDBC Connection Information view displays connection information for the pool key you chose.

In addition to the connection string, database user ID, and database password, the Connection Information view displays:

- Product information, which includes the database and version and the driver and version
- Connection properties, which list the database-specific properties used to establish the database connection
- Supports, which are the attributes of the connection

Note

A message in the Connection Information view indicating that a connection could not be established means either that the connection criteria have not been met or that the data source is unavailable. You might use this information for debugging if you are having trouble running applications.

Finally, you can view the list of database drivers that are available to the Web server by clicking the Loaded Drivers heading.

See Also

- ❑ *The jde.ini File* in the *System Administration Guide* for more information on database connections

► To monitor database connection information

1. In the System Summary view of the Web Server Monitor, click the scroll button and choose JDBC Pool.
2. On Jdbc Connection Pool Information, click Refresh to get the latest additions to the files.
3. To view information on an individual database connection, click the name of a connection string under the Pool Key column.

4. To view information on database drivers used to make connections, click Loaded Drivers.
5. If you are experiencing database problems, scroll to the bottom of the view and click Pool Cleaner.

Monitoring Virtual Clients

A new enhancement of the ERP Web product is Multiple Application Framework (MAF). MAF allows you to open multiple applications at the same time, in separate browser windows. As each application is opened, a separate virtual client is created. For example, a user logged into the Web client, with multiple applications open, has a unique virtual client entry for each application. Server Administration Workbench (SAW) displays each virtual client as associated with a distinct user session.

The SAW Web application allows the user to view information that is specific to a particular Web server and port. The information is similar to that in SAW, but it is specific to a server:port.

The virtual client records can be viewed in the User List View. Under the Location column, the name of the application and the form can be seen, along with the name and IP address of the web server where the virtual client exists. There is also an Action column, which displays 'View Detail.' Click on View Detail to view more specific information about the virtual client.

Property	Value
Virtual Client Properties	
Session ID	2693743573961615360
Virtual	OWVirtual@5957a8a6
Location	P0092W0092D1079725251551
Thread Name	No Thread Assigned
Max. Open Browsers	10
Parent URL	/servlet/com.jdedwards.runtime.virtual.LoginServlet
URL	null
Server Poll Interval	3000
Stack ID	1
Status	Ready
Page Timeout	300000
Virtual Client Environment Properties	
Environment	STGAWU
User Language	en
Language Preference	
User Id	JV5443249
User Name	com.jdedwards.services.login.UserName@1034a886
Virtual Client User Properties	
User Name	JV5443249
Environment	STGAWU
Display Mode	HTML
Computer Name	den-jv5443249.jdedwards.com
Login Time	03/19/2004 12:40 PM

From the Virtual Client Detail form, you can view the following.

- The Location row displays the name of the application and the form, as well as the name and IP address of the web server where the virtual client exists
- The Max. Open Browsers row displays the maximum amount of open browsers that each user is allowed. This value initially is defined in the JAS.INI
- Virtual Client Environment Properties and Virtual Client User Properties list information specific to the environment, user name, display mode, and user's computer name

This information is good for pinpointing where, when, and with whom a problem occurred.

You can select Virtual Clients from the drop down list. This view displays all of the virtual clients that are currently active.

Virtual Clients

JAS Server: denxpi32:85 Up for: 10 days, 19 hours, 29 minutes and 58 seconds

Max Users Allowed: 100 Max Users Connected: 5 Session Serviced: 101 Session Timeout: 60

Current Users: 0 Active Users Last 10 Minutes: 0

[Work With INI Settings](#)

#	User Name	Session ID	OW Virtual	Location	Thread Assigned	Action
1	Not Found	3273745803080583168	OWVirtual@15f828a3	No Forms in Stack	No Thread Assigned	View Detail
2	Not Found	2057540301528747008	OWVirtual@296928a7	No Forms in Stack	No Thread Assigned	View Detail
3	Not Found	2057540301528747008	OWVirtual@2566a8b3	No Forms in Stack	No Thread Assigned	View Detail

You can view and modify certain JAS.INI settings by clicking the hyperlink at the top of the grid 'Work With INI Settings'.

Note that the initial values for these settings are defined in the JAS.INI and are read when the web instance is started. You can modify these values through SAW without changing the JAS.INI file. The new values are immediately active, and maintain their status until the web instance is stopped. Once the instance is restarted, the JAS.INI values become active again.

Work with INI Settings for Virtual Clients

Max. number of open browsers: 10

Multiple browser enabled: true

Interactivity level: high

Cancel:

The Max. number of open browsers field defines how many browser sessions (JDE applications) each user will can have open at any one time. For example, a user logs into the web client, opening one browser. The user launches an application, opening a second browser. This number of browsers can be limited by setting this value.

The Multiple browser enabled field can be set to True or False. By this value being set to true, the users take advantage of the Multiple Application Framework (MAF), and can launch multiple applications in separate browsers at one time. By setting this value to false, the behavior of the web client is similar to that of the pre-SP22 client, where the user can launch only one application at a time, and the application is run in the same browser session.

Note

For more information about these settings, refer to the Web Server Installation guide for your platform.

► To run SAW Web

1. To run SAW Web, enter the following in your web browser:

`http://WebServerName:port/jde/saw/SawWeb.html`

2. In the SAW Web window, select User List from the drop down list.

The application displays a list of all the users that are currently logged into this particular web server and port. There is a hyperlink at the top of the grid that shows the virtual clients that are assigned to each of the user sessions. Under each user session is a record for each virtual client.

3. To view the virtual clients, click View With Virtual Clients.

4. To hide these records, click View Without Virtual Clients.

Under the Location column, the name of the application and the form can be seen, along with the name and IP address of the web server where the virtual client exists. There is also an Action column, which displays View Detail.

5. Click on View Detail to view more specific information on the virtual client.

The Location row displays the name of the application and the form, as well as the name and IP address of the web server where the virtual client exists.

The Max Open Browsers row displays the maximum amount of open browsers that each user is allowed. This value is initially defined in JAS.INI.

Virtual Client Environment Properties and Virtual Client User Properties list information specific to the environment, such as user name, display mode, and the user's computer name.

6. Choose Virtual Clients from the drop down list. This displays all of the virtual clients that are currently active, including the User Name and Location. There is also an option to View Detail.

7. At the top of the Virtual Clients form, click the Work With INI Settings hyperlink.

This form allows you to modify settings. Note, the initial values for these settings are defined in JAS.INI, and are read when the web instance is started. By modifying these values through SAW, JAS.INI is not actually changed. The new values are effective immediately, and maintain their status until the web instance is stopped. Once the instance is restarted, the JAS.INI values are again active.

Multiple Browser Enabled can be set to True or False. When True, users can take advantage of MAF and may launch multiple applications in separate browsers at one time. When False, the behavior of the web client is similar to that of the pre-SP22 client, where the user can launch only one application at a time and the application is run in the same browser session.

Max. Number of Open Browsers defines how many browser sessions (JDE applications) each user is allowed to open at any one time.

Monitoring Web Server Log Files

You can monitor Web server log files from the Web Server monitor. From the Log File view, you can view web server files, view logs, and change logging parameters.

The screenshot shows the PeopleSoft Server Administration Workbench (SAW) for JAS Servers. The interface includes a header with the PeopleSoft logo and a 'back' link. Below the header, there's a section for 'Server Administration Workbench (SAW) for JAS Servers' with fields for 'Host' (denxpi32) and 'Port' (85), and a 'Views' dropdown set to 'Work with Log files' with a 'Refresh' button. A 'sawJas' label is in the top right. The main area has tabs for 'Log Summary', 'JAS Log', 'DEBUG Log', 'stdout', 'stderr', 'RT Log', and 'RT Debug'. Below these tabs, there's a 'Java Logging System' section with a 'Work with Property Object' dropdown and a 'Web Server Files' section with a 'JAS.INI' dropdown. The 'Application Server Log Directory' is set to 'Z:\WebSphere\AppServer'. The 'Server Status' is 'Up for: 10 days, 19 hours, 36 minutes and 44 seconds'. The 'JAS Log' section shows 'Log size: -1 bytes'. The 'DEBUG Log' section shows 'Log size: -1 bytes' and 'Trace is turned on', 'NET Trace is turned off', and 'RT Trace is turned off'. The 'RT Log & Debug' section shows 'RT Log size: 53.4K bytes' and 'RT Debug size: 0 bytes'.

The Log File view provides the following menus:

- Log File buttons
- Java Logging System menu
- Web Server Files menu

Log File Buttons

The yellow buttons at the top of the screen allow you to access log files and a log file summary. A description of each of the buttons follows:

- Log Summary
Click this button to view the size of all logs and the last time they were written.
- JAS.Log
Click this button to view information on Java Application Server functions and records server errors.
- Debug.log
Click this button to view detailed information about the tasks running on the Java Application Server, including the SQL statements that were used.
- stdout

Click this button to view all of the outputs written by Java applications, including instantiation of servlets and business function processes running on the WebSphere Application Server.

- stderr

Click this button to view the error output from the main WebSphere servlet Java process.

- RT Log

Click this button to view routine runtime events.

- RT Debug Log

Click this button to view runtime events for system functions, business functions, table I/O, and other events.

You can also enable the Net Trace function, which offers a detailed view of the JDENET messages. You must enable JASdebug.log to enable Net Trace.

Caution

Enabling both JASdebug.log and Net Trace produces very large files, which could degrade server performance. Do not configure the Web Server Monitor to write these files unless you are performing a specific debugging task.

You configure the Web Server Monitor to determine the maximum size of the JAS.log and JASdebug.log files that can be retrieved.

You can also delete JAS.log and JASdebug.log files. However, if you do so, the files are deleted from the server, not from your workstation, and they cannot be retrieved. Do not delete the files unless you have administrative privileges.

► To monitor web server log files

1. In the System Summary view of the Web Server Monitor, click the scroll button and choose Log Files.

The Web Server Monitor displays the log file summary.

2. Click Refresh to get the latest additions to the files.
3. Click the JAS Log or Debug Log buttons to view the entire jas.log or jasdebug.log file.
4. To troubleshoot the log file, click inside the text file and perform a keyword search on "error."

► To enable Debug.log and Net Trace

1. On the web server, locate the sawLogButtons.jsp file, which is typically in the \E810\internet\dist\jdewww\saw\ directory.
2. Open the sawLogButtons.jsp file, locate the following parameters and change the values to "true," as follows:
 - boolean includeDeleteButtons = true

- `boolean includeTraceOnOffButtons = true`
3. Click Save.

These parameter changes will also enable the Delete JAS Log and Delete Debug Log buttons in the Log File view of the Web Server Monitor.
 4. In the Log File view, click the Debug Log button.

The Debug Log and Net Trace options should both be visible, as should the Delete buttons for both JAS Log and Debug Log.
 5. If you want to enable Net Trace, choose the Net Trace option.

You can enable Net Trace only if you have enabled the Debug Log option.
 6. Make sure the Delete JAS Log and Delete Debug Log files are visible in the Log File view.
 7. Click Refresh to update the file.

► To delete log files

1. In the System Summary view of the Web Server Monitor, click the scroll button and choose Work with Log Files.
2. Click Refresh to get the latest additions to the files.
3. To delete the JAS.log files, click Delete JAS Log.
4. To delete the JASdebug.log files, click Delete Debug Log.

Java Logging System Menu

The Java Logging System pull-down menu (located on the left side of the Log File view) includes the following options:

- Work with Property Object
- Work with Components
- Work with Log Files

Work with Property Object

Select Work with Property Object to display the properties of all the settings in the `jdelog.properties` file. In this view, you can display and change the properties of each of the keys in the `jdelog.properties` file.

When you click the Submit button next to the Add New Component field, the program displays a form for the new component. Type a name for the new component and select values for the other fields. Click the Submit button to accept your changes.

- To modify a component

When you click the component name in the Component column, the program displays fields for all properties of the component. Enter the new value for each field of the selected component. When you are done, click the Submit button to save the changes.

- To delete a component

Click Remove in the Action column.

Work with Log Files

Select Work with Log Files to display a list of all the log files and their associated properties. To view the file, click its name in the File Name column.

#	File name	Current size	Level	Append	Max Backup Index	Max File Size	Format
1	Z:\WebSphere\AppServer\installedApps\EA_JDEdwards_2.ear\webclient.war\logs\lderoof_382004.log	5858595	ERROR	FALSE	1	10MB	com.jdedwards.
2	Z:\WebSphere\AppServer\installedApps\EA_JDEdwards_2.ear\webclient.war\logs\las_382004.log	5858595	ERROR	FALSE	1	10MB	com.jdedwards.

Web Server Files Menu

The Web Server Files Menu allows you to view any of the following files associated with a particular Web Server:

- JAS.INI
- Oracle TNS
- MSSQL TNS
- xmlconfig.xml
- plugin-cfg.xml
- admin.config
- initial_setup.config
- xmlconfig.xml

- Application Server log files

To view one of these files, select the file from the Web Server Files pull-down menu, and click the Submit button.

Note

To view the Application Server log files, you must enter the path to these files in the Application Server Log Directory before you click the Submit button. The Application Server log files view displays the log files directory that you enter, and allows you to navigate through all subdirectories, drill down to individual files, and click a file to view its contents.

Troubleshooting with the Web Server Monitor

The available Web Server Monitor views offer data that can be useful to administrators in spotting potential problems. The following table lists Web Server Monitor parameters that you can monitor for troubleshooting purposes:

Parameter	Web Server Monitor View(s)	Parameter Meaning	Possible Problem Indication
MaxWaited	System Summary/JDENET Pool	Number of users demanding socket connections exceeds the number defined in jas.ini if this parameter has a value.	Graphic user interface presentation might be too slow.
Log file size	System Summary	Warning message appears only if file size exceeds 1 mb.	Numerous error messages might be in the log
Agent ID	User List	User's browser.	Browser might be out of date.
In Progress	CallObject Information	Many business functions currently running on enterprise server if the parameter value for the "In Progress" message is anything other than "None."	Enterprise server is running slowly, possibly because of problems with business functions.
Connection could not be established	Connection Information (JDBC Pool)	Message appears only if the server has a problem connecting to the database.	Connection parameters might be wrong or the data source is unavailable.
java.class.path	Environment	Virtual machine path.	Troubleshoot path if processing problems occur.
java.fullversion	Environment	Version of virtual machine.	User might need virtual machine upgrade if processing problems are occurring.

Locations of Key Parameter Values for Web Server Monitor Settings

The parameter values displayed by the Web Server Monitor are defined in various locations in the EnterpriseOne system.

The following table displays selected Web Server Monitor parameters, the Web Server Monitor view in which each is displayed, and the EnterpriseOne location where each Web Server monitor parameter is typically defined.

Web Server Monitor View	Parameter	EnterpriseOne Property	EnterpriseOne File Where Property Is Defined
System Summary	Heap Memory Allocated by VM (Java Virtual Machine)	java.mx	C:\WebSphere\AppServer\properties
Configuration	CallObject Timeout	[JDENET] enterpriseServerTimeout=	jas.ini file
Configuration	Maximum Host Pool	[JDENET] maxPoolSize=	jas.ini file
System Summary	Host		Object Configuration Manager table (F986110)
System Summary	Port (displayed together with Host parameter)	[SERVER] serviceNameConnect	jas.ini file
User List	Max users allowed	[OWWEB] MAXUser=	jas.ini file
User List	Session Timeout	[CACHE] UserSession=	jas.ini file
Jdbc Connection Information	Jdbc URL	[JDBC URL] constructed as follows: <environment>=<jdbcurl type owner user password>	jas.ini file

Monitoring System XPIe Servers from the Web

The System XPIe Server Monitor allows you to use the Web to monitor an XPIe Broker Server. The server-monitoring APIs for ActiveWorks are available in Java code, and EnterpriseOne uses them to provide the XPIe Server Monitor.

To monitor an XPIe server, you must enter a valid host name and port number in the XPIe Server Monitor workspace. From the dialog box in the workspace, you can choose from five views:

- Broker Summary
- Event Types
- Client Groups
- Client States
- Broker Logs

Broker Summary View

The Broker Summary view allows you to monitor information about one or more brokers running on the Broker Server. The broker provides essential system services, such as receiving, sending, and queuing events. Events are messages sent to and received by resources in the system, including client workstations and other servers.

To view broker summary information, enter a host name and port number in the XPIe Server Monitor workspace, and then choose Broker Summary from the dialog box and click the Monitor button.

The Broker Summary view appears.

The Broker Summary view contains a table with various parameters, the values of which provide information about one or more brokers running on the server. The following table identifies the parameters and offers a brief explanation of each one:

Parameter in Table of Broker Summary View	Parameter Meaning
Number	The number of the broker.
Broker Name	The name of the broker. The default name is Broker # 1.
Broker Host	The Broker Server on which the broker is running.
Territory	A set of brokers that share information about event types and client groups.
Description	A full description of the broker, provided when the broker is installed on the server.
Event Types	Types of messages received and sent by the broker.
Client States	Information about a client maintained by the broker. Clients connect to brokers.
Client Groups	A list of all client groups on the broker. A client group is a set of properties shared by broker clients.

The Event Types, Client States, and Client Groups parameter columns contain magnifying glass icons you can click to view more detailed information about each one. Each of these parameters has its own view, which you can access from the dialog box in the Broker Summary view or in the XPIe Server Monitor workspace.

The Broker Summary view also displays:

- Number of non-SSL (Secure Sockets Layer) connections
- Highest number of non-SSL connections

- Number of SSL connections
- Highest number of SSL connections
- Server disk space, in megabytes

You can view details about the server in which the broker is running by clicking the Server Statistics URL.

See Also

- ❑ *Event Types View, Client States View, and Client Groups View* in the *Server and Workstation Administration Guide* for more information about the Broker Summary Table parameters

Event Types View

An event type defines the properties of an event, including the data fields that the event carries, the event's unique name, and a storage type, which defines how the broker stores the event.

The word "events," when it is used in discussing the XPIe Server Monitor, corresponds to the word "messages," as that word is used in discussing the Enterprise Server Monitor. Events are messages exchanged by resources in the system. For example, an event might be processing a purchase order. To process a purchase order, the client and broker communicate.

When you choose Event Types from the XPIe Server Monitor workspace and click the Monitor button, the Event Types view appears.

The Event Types view contains a table with various parameters, the values of which provide information about the event types handled by the broker. The following table identifies the parameters and offers a brief explanation of each one:

Parameter in Table of Event Types View	Parameter Meaning
Event Name	The unique name of the event.
Description	A description of the function of each event.
Storage Type	An event attribute that determines how the event is stored in the broker. Storage types are: • Guaranteed, which means that events are stored on disk using a logged commit. • Persistent, which means that events are stored on disk using operating asynchronous input/output. • Volatile, which means that events are stored in memory.
Time to Live	The amount of time that an event type can exist in the broker.
Fields	The names and types of each data field within the event type.

You can view information about the data fields for each event type by clicking the magnifying glass icon in the Fields column. The Event Types Fields view appears; it identifies each field name for the event type and the field type, such as string.

Client Groups View

A client group is a set of broker clients with particular properties. For example, a client group defines the server on which clients access the broker.

When you choose Client Groups from the XPIe Server Monitor workspace and click the Monitor button, the Client Groups view appears.

The Client Groups view contains a table with various parameters, the values of which provide information about the client groups that provide control of client access to the broker. The following table identifies the parameters and offers a brief explanation of each one:

Parameter in the Client Groups View	Meaning
Client Group Name	The name of the client group. Each group has a specific set of properties defined using the ActiveWorks Manager.
ACL	The Access Control List, which is a list of SSL certificates that define the entities with permission to access the broker or create a client within a client group.
Can Publish	The event types that a client group can publish.
Can Subscribe	The event types that a client group can subscribe to.
Statistics	Additional statistics on the client group, such as how long the group has been on the server and how many events have been published by clients that belong to the group.

Click the magnifying glass in the ACL, Can Publish, Can Subscribe, and Statistics columns to view additional information about each of these parameters.

Client States View

A client state is information maintained by the broker about a client connected to the server. For example, the client group to which a client belongs is a client state.

When you choose Client States from the XPIe Server Monitor workspace and click the Monitor button, the Client States view appears.

The Client States view contains a table with various parameters, the values of which provide information about the clients connected to the server. The following table identifies the parameters and offers a brief explanation of each one:

Parameter in the Client States View	Meaning
Identifier	A unique identifier for the client connected to the server.
Client Group	The client group to which the client belongs.
App. Name	The name of the application that describes the client's connection to the server.
Access Label	Indicates, if appropriate, the value of the access label required for a client to connect to the client

	group to which it belongs.
Authenticator	The name of the certification authority that issued the certificate if SSL is enabled for the client.
Can Share	Indicates whether state sharing is enabled. If state sharing is enabled, the number of sessions allowed for a client can be set.
High Seq.	The highest published sequence number used by the client.
Max Shared Connections	The maximum number of server connections that the client can share.
Owner Name	The client's owner's user name.
Sessions	Information about a client session. Click the magnifying glass icon in the Sessions column to view detailed session information.

Broker Log View

You can use the XPIe Server Monitor to view a broker log containing information about events that have run on the server during a specified time. You can view all log messages, or you can limit your view to warnings, alerts, or information messages.

When you choose Broker Log from the XPIe Server Monitor workspace and click the Monitor button, the Broker Log view appears.

From the Broker Log, you can specify:

- The starting date from which you want to view log files
- The maximum number of entries you want to view
- The type of file you want to view

Click the View Log button to see the log file entries.

Monitoring EnterpriseOne with SAW on the iSeries

The Server Administration Workbench (SAW) for the iSeries uses the menu structure that is standard to the iSeries interface. Through this interface, SAW provides access to the following information:

Application	Description
Work with Servers	This option allows you to view EnterpriseOne services running on different ports on the same server or on other servers.
Work with Server Processes	This option provides the ability to start, stop, and view EnterpriseOne services for the local port and host only. Otherwise, this option allows viewing only of EnterpriseOne services. You can view information for other servers that run on the same port as the host server, but you can only start and stop services that run on the host server.
Work with Server Resources	This option allows you to view information about general EnterpriseOne resources on a server, such as locks and resource users, and disk space resources, such as CPU utilization and disk availability.

Work with Server Log Files This option provides a list of all log files for EnterpriseOne network, kernel, and batch processes. From this menu, you can view and delete the log files.

Work with Server INI File This option starts the Source Entry Utility (SEU) to allow you to view and change settings in the local jde.ini file on the local port and host only. Otherwise, this option is not displayed. You cannot access jde.ini files on remote servers.

Note

The name of the current server and the port for the host server appear at the top of every form in SAW.

Accessing SAW

You can access the Server Administration Workbench (SAW) from the command line.

Prerequisite

- ❑ SAW must reside in your EnterpriseOne system library in the library list before you can access SAW from the command line.

► To access SAW

On the command line, enter SAW.

The SAW main menu appears.

Monitoring Remote Servers

The Work with Servers application allows you to change the server and port that you monitor using Server Administration Workbench (SAW). EnterpriseOne services do not need to be running on the local port and host to run SAW.

► To monitor remote servers

From the SAW main menu, choose Work with Servers.

On Work with Servers, enter a new server name in the field.

When a valid server name and port are entered, the Server Administration Workbench screen displays the new server name and port.

Starting, Stopping, and Displaying ERP Server Processes

The Work with Server Processes form allows you to start or stop the local EnterpriseOne server and to display the EnterpriseOne processes currently running for any server on the network. You can only

start or stop services for the local server on the local port where Server Administration Workbench (SAW) was invoked. To start or stop services on remote servers, you must work with the server directly.

Note

The menu options for starting and stopping EnterpriseOne server processes do not appear when you switch the host or port using the Work with Servers option.

► To start EnterpriseOne processes on a server

From the SAW main menu, choose Work with Server Processes.

On Work with Server Processes, choose Start EnterpriseOne Processes.

If EnterpriseOne processes have already been started, this option does not appear on the form.

► To stop EnterpriseOne processes on a server

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Stop EnterpriseOne Processes.

Caution

Use caution when you stop EnterpriseOne processes on a server to avoid interrupting important server processes.

If EnterpriseOne processes are not active or if the host or port were not the default when starting SAW, this option does not appear on the form.

► To display EnterpriseOne server processes

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Display EnterpriseOne Processes.
2. On Display EnterpriseOne Processes, you can view information about kernel processes, network processes, and queue processes. You can also display and delete logs for the EnterpriseOne processes. When you display EnterpriseOne processes for the host server, you can end the processes and display detail information about the processes.
3. You can filter, sort, and view different information on this form by using key commands and function keys. Press F1 to see a list of these options.
4. If you are viewing processes for the host server, enter 5 in the Opt field in front of a process to view detail information about the process.

► To display and delete log files for server processes

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Display EnterpriseOne Processes.
2. On Display EnterpriseOne Processes, enter one of the following values in the Opt field in front of a process:
 - 7
This value allows you to view the log file.
 - 8
This value allows you to view the debug log file.
 - 9
This value allows you to delete all log files for the process.

Note

If you view log files for a remote server, you see only the data available at the time when you open the log. As a process continues, the system might add data to the log after you open the file. Reopen the log file at a later time to determine whether the system added new data.

► To end individual server processes

Note

You can only end EnterpriseOne processes for the host server. You cannot end processes for remote servers.

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Display EnterpriseOne Processes.
2. On Display EnterpriseOne Processes, enter 4 in the Opt field in front of the process that you want to stop.
SAW stops the process and then refreshes the form.

Displaying Server IPC and Disk Space Resources

The Work with Server Resources form provides information about current IPC resources and disk space information on the server. The IPC resources available for viewing are:

Resource	Description
Message queues (MSG)	A message queue allows one process to pass packets to another process.
Single Write Multiple Read	SWMR is a lock that allows only one writer to change data, but multiple readers can access the data. The system grants reader locks if no outstanding writer locks exist. The system

(SWMR)	grants writer locks if no outstanding reader locks exist.
Shared memory (SHM)	Shared memory is a segment of memory that appears in the address space of more than one process. For example, network processes might use a shared memory segment to monitor what processes exist and the status of those processes.
Mutual exclusion (MTX)	A mutual exclusion lock, or a mutex lock, allows only one process to hold the lock. This is the most common type of lock.

If EnterpriseOne is not running on the local port and local host and you try to display local IPC resources, you will quit SAW.

► To display IPC resources

From the SAW main menu, choose Work with Server Resources.

1. On Work with Server Resources, choose Display EnterpriseOne Resources.
2. On Display EnterpriseOne Resources, you can filter and sort information on this form by using key commands and function keys. Press F1 to see a list of these options.

► To display disk space resources

From the SAW main menu, choose Work with Server Resources.

On Work with Server Resources, choose Display Disk Space Resources.

If you display disk space resources for the host server, disk space resource information appears in the Display System Status (DSPSYSSTS) utility.

If you display disk space resources for a remote server on the network, the Display Disk Space Resources form appears to display the disk space resource information. You can sort information on this form by using key commands and function keys. Press F1 to see a list of these options.

Displaying, Printing, and Deleting Server Log Files

You can use the Work with Server Log Files option in Server Administration Workbench (SAW) to access all log files available on a specific server.

► To display, print, and delete server log files

From the SAW main menu, choose Work with Server Log Files.

The Work with Server Log Files displays all log files for network, kernel, and batch processes. SAW measures the size of the log files in Kilobytes.

1. You can sort information on this form by using key commands and function keys. Press F1 to see a list of these options.
2. In the Opt field in front of a process, enter one of the following values:

- 4

This value allows you to delete the log file.

- 5
This value allows you to view the log file.
- 6
This value allows you to print the log file.

Note:

If you view log files for a remote server, you see only the data available at the time when you open the log. As a process continues, the system might add data to the log after you open the file. Reopen the log file at a later time to determine whether the system added new data.

Modifying the Server jde.ini File

Server Administration Workbench (SAW) provides access to the jde.ini file for the server that you are using. You can view and modify settings in the jde.ini file *only* on the local host and local port. You cannot administer the jde.ini file on remote servers.

► **To modify the server jde.ini file**

From the SAW main menu, choose Work with Server INI File.

This option starts the Source Entry Utility (SEU) editor. Within this editor, you can view and modify the local jde.ini file.

The SEU editor appears with the jde.ini file displayed.

1. View and modify the jde.ini file as necessary and then press F3.
The Exit form for the SEU editor appears.
2. Save your modifications to the jde.ini file by entering Y in the Change/Create Member field.
You can disregard the modifications by entering N in this field.

Monitoring EnterpriseOne with SAW on UNIX

The Server Administration Workbench (SAW) on UNIX provides access to the following information:

- Work with Servers
This option allows you to view EnterpriseOne services running on other servers on any port.
- Work with Server Processes

This option provides the ability to start, stop, and view EnterpriseOne services. You can view information for other servers that run on any remote host server, but you can only start and stop services that run on the local host server.

- **Work with Server Resources**

This option allows you to view information about general EnterpriseOne resources on a server, such as locks and local resources, and disk space resources, such as disk availability.

- **Work with Server Log Files**

This option provides a list of all log files for EnterpriseOne network, kernel, and batch processes. From this menu, you can view and delete the log files.

- **Work with Server INI File**

This option starts the Source Entry Utility (SEU) to allow you to view and change settings in the local JDE.INI file. You cannot access JDE.INI files on remote servers.

Note

The name of the current server and the port for the host server appear at the top of every form in SAW.

The recommended terminal to use when you monitor your configuration with SAW is xterm with a minimum of 80 columns and 24 lines (default).

Accessing SAW

You can access Server Administration Workbench (SAW) from the command line.

► **To access SAW**

On the command line, enter the following command:

```
saw.sh
```

Monitoring Remote Servers

This application allows you to change the server that you monitor using Server Administration Workbench (SAW). The services on the remote server can run on any port number.

► **To monitor remote servers**

From the SAW main menu, choose Work with Servers.

You can either use the arrow keys to move focus to the appropriate option and press Enter or enter the number of the option on the command line.

On Work with Servers, enter a new server name and port number.

When you enter a valid server name, the name of the server replaces the previous server in the upper left corner of the form.

If you enter a server name that does not exist on the current port, or if the network fails to make a connection, SAW stops the process after fifteen seconds. You need to reset the server using Work with Servers.

Starting, Stopping, and Displaying Server Processes

The Work with Server Processes form allows you to start or stop the local EnterpriseOne server and to display the processes currently running for any server and port number on the network. You can only stop kernel and queue services on the local server. To stop services on other servers, you must work with the server directly.

► To start an EnterpriseOne server

From the SAW main menu, choose Work with Server Processes.

You can either use the arrow keys to move focus to the appropriate option and press Enter or enter the number of the option on the command line.

On Work with Server Processes, choose Start EnterpriseOne Processes.

If the server is already running, the Stop EnterpriseOne Processes option appears in bold, while the Start EnterpriseOne Processes option appears in a normal font.

► To stop an EnterpriseOne server

From the SAW main menu, choose Work with Server Processes.

On Work with Server Processes, choose the Stop EnterpriseOne Processes option.

If the server is not currently running, the Start EnterpriseOne Processes option appears in bold, while the Stop EnterpriseOne Processes option appears in a normal font.

Displaying, Printing, and Deleting Server Log Files

You can use the Work with Server Log Files option in SAW to access all log files available on a given server.

► To display EnterpriseOne server processes

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Display EnterpriseOne Processes.

The Display EnterpriseOne Processes form appears. On this form you can view detailed information about kernel processes, network processes, and queue processes. You can also display and delete logs for the EnterpriseOne processes. When you display EnterpriseOne processes for the local host server, you can end the processes.

You can modify the appearance of this form by using key commands and function keys. Press F1 to see a list of these options.

2. If you are viewing processes for the host server, press 5 to view detail information about the selected process. The following screen is an example of detailed information for a kernel process:

All other options on the Display EnterpriseOne Process help form are valid for both local and remote servers.

Note

You only need to press the key to perform the command. For example, press Shift + G or Shift + H and you automatically go to the end of the list or the beginning of the list respectively.

► To display and delete log files for server processes

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Display EnterpriseOne Processes.
2. On Display EnterpriseOne Processes, choose a process, then do one of the following:
 - Press 7 to view the log file.

Allows you to view the jde.log for the chosen process. If the size of the file you want to view is larger than the maximum allowed size, the system displays a confirmation form before opening the file.
 - Press 8 to view the debug log file.

Allows you to view the jdedebug.log for the chosen process. If the size of the file you want to view is larger than the maximum allowed size, the system displays a confirmation form before opening the file.
 - Press F4 to view the jdedebug file using vi.
 - Press 9 to delete all log files for the process.
 - Press A to delete all log files for all processes.

On the confirmation form that appears, enter Yes.
 - Press R to change the maximum time allowed for a UBE process to run.

The value is entered in minutes. The default value is 2880 minutes (48 hours). If a UBE process runs for more than this value, the process is marked with the letter R and will be underlined in the process list. This value is only used in SAW to monitor the execution time for UBEs and does not interfere with the execution of the UBE. UBEs are not stopped by SAW if their execution time exceeds this time limit.
 - Type c to change the maximum number of collisions allowed for a kernel process.

If the number of collisions exceeds this number, the process is marked with letter C and underlined in the process list. This value is used only in SAW to monitor the number of collisions for a kernel process and does not interfere with the execution of the process. Kernel processes with the number of collisions set larger than this limit are not stopped by the SAW interface.

A collision occurs when a kernel is processing a manual commit request and the server receives additional manual commit requests for the same process. If a kernel's number is much larger than other kernels, this information can be a warning for that kernel. The warning can mean the kernel process is taking too long to process the commit request and other requests are waiting for the same kernel. Kernel process collisions are harmless and normal. However, it is not normal to have a high number of collisions for a kernel while additional manual commit requests wait for the same kernel. For collisions, the number of outstanding kernel requests does not increment. The kernel process pulls out the request and inserts it into its own list.

Note

When you view a log file, press the Space bar for the next page of the file. Type q to stop viewing the log.

► To stop individual kernel processes

From the SAW main menu, choose Work with Server Processes.

1. On Work with Server Processes, choose Display EnterpriseOne Processes.
2. On Display EnterpriseOne Processes, choose the process to stop and press 4.
3. Click Yes to confirm.

Option 4 is valid only for the local server. This option allows you to stop a kernel process. This option is not active for a remote server or for jdenet processes.

SAW removes the process and then refreshes the form.

Displaying Server IPC and Disk Space Resources

The Work with Server Resources form provides information about current IPC resources and disk space information on the server. The IPC resources available for viewing are:

- Message queues (MSG)

A message queue allows one process to pass packets to another process on the same host.

- Single Write Multiple Read (SWMR)

SWMR is a lock that allows only one writer to change data, but multiple readers can access the data. The system grants reader locks if no outstanding writer locks exist and grants writer locks if no outstanding reader locks exist.

- Shared memory (SHM)

Shared memory means that a segment of memory appears in the address space of more than one process. For example, network processes might use a shared memory segment to monitor what processes exist and the status of those processes.

- Mutual exclusion (MTX)

A mutual exclusion lock, or a mutex lock, allows only one process to hold the lock. This is the most common type of lock.

► **To display IPC resources**

From the SAW main menu, choose Work with Server Resources.

On Work with Server Resources, choose Display EnterpriseOne Resources.

The Display EnterpriseOne Resources form appears.

You can modify the appearance of this form by using key commands and function keys. Press F1 to see a list of these options.

► **To display disk space resources**

From the SAW main menu, choose Work with Server Resources.

On Work with Server Resources, choose Display Disk Space Resources.

The Display Disk Space Resources form appears.

SAW measures disk space in KB. You can modify the appearance of this form by using key commands and function keys. Press F1 to see a list of these options.

Note

The underlined disk resource represents the EnterpriseOne disk resource for the current server.

► **To display, print, and delete server log files**

From the SAW main menu, choose Work with Server Log Files.

The Work with Server Log Files displays all log files for currently running network, kernel, and batch processes including log files for previously run processes. SAW measures the size of the log files in KB. Press F1 to see a list of options for this screen.

On Work with Server Log Files, choose a log file, and then do one of the following:

- Press 4 to delete one or more log files on the server.

To get a list of options explaining how you can remove files, press the F1 key. When log files are removed, they are marked with <Deleted> in the log list at the end of each line, but they are not actually deleted until you refresh the form (F5). Until you press F5, you can undo a remove if necessary. Once you press F5 to refresh the form, the files are deleted from the server.

- Press 5 to view the log file.
- Press 6 to print the log file.
- Type b to remove all log files bigger than the current size limit.

The current size limit appears at the top of the form. You can undo the remove by typing b again before refreshing the form. Once you refresh the form (F5), the files you marked are deleted from the server.

- Enter l to set the size limit for retrieving log files.

The number is entered in kilobytes. Log files can be big on the server and choosing to view them, especially if you are monitoring a remote server, can be time consuming. Before viewing or transferring a log file, SAW checks the log file's size. If the log file exceeds the size limit, SAW verifies that you do want to view the file.

Note

When you view a log file, press the Enter key for the next page of the file. Type q to stop viewing the log. You can press F4 to view the file using the vi editor.

Working with the Server jde.ini File

SAW provides access to the local jde.ini file. You can view and modify settings in the jde.ini file *only* on the local server. You cannot perform administration tasks for the jde.ini file on remote servers.

► To modify the server JDE.INI file

From the SAW main menu, choose Work with Server INI File.

1. On Work with Server INI file, choose a section and then press Enter.

A form appears that lists the settings for the JDE.INI section.

If you don't know the section in which a variable resides but you know the name of the variable, press F7 to perform a search on the variable. SAW returns a list of sections in which the variable resides.

2. Choose a setting and then press F6.
3. On Changing Value for Current Variable, enter a new value for the setting and then press any key to return to the settings list.
4. Press the left arrow key to return to the sections list.

► To edit the server JDE.INI file

From the SAW main menu, choose Work with Server INI File.

1. On Work with Server INI file, Press F8.

The vi editor opens to display the full JDE.INI.

2. Change the JDE.INI file as necessary and then save and quit vi.

► **To turn on or turn off debugging**

From the SAW main menu, choose Work with Server INI file.

1. Work with Server INI file, press F5.

The Turning debug On/Off form appears, informing you that debugging was either turned on or off.

2. Press F5 again to reset debugging to the previous setting.

► **To change the port number of the server**

From the SAW main menu, choose Work with Server INI File.

1. On Work with Server INI file, press F6.
2. On Changing port number value, enter a new port number.

The JDE.INI file immediately reflects the changed port number, but you must stop and restart the server to access information based on the new port number.

Monitoring EnterpriseOne with the Knowledge Module

BMC Software has a product called PATROL, which allows companies to monitor and detect errors across large networks and a variety of third-party software packages, including PeopleSoft EnterpriseOne. The PATROL console offers a single interface for data monitoring and collection across an entire distributed enterprise.

If you are using PATROL to monitor your system, you can use the EnterpriseOne Knowledge Module (OWKM) to monitor the processes on your EnterpriseOne server. OWKM uses a command line interface (CLI) to call JDESAW APIs, which retrieve server information and return the data to the CLI. The CLI, in turn, sends the information to OWKM, which displays the information on the PATROL console.

The following table lists the components of the OWKM solution and summarizes the role of each component:

Component	Role
JDESAW	Dynamic link library containing SAW APIs, which retrieve data about OW servers running locally or remotely.
OWKM	Program that sends a request from an EnterpriseOne administrator to a CLI, receives the returned server data, and displays the data in the PATROL console.
CLI	Standalone, running program that receives an EnterpriseOne administrator's request from OWKM, calls JDESAW APIs, receives server data from the JDESAW APIs, formats the data, and sends it back to OWKM.

Administrators can access the following server information on the PATROL console:

- Process status (active or inactive)
- Server and print queue log files, their size and the percentage of space used

- Total number of server processes
- CLI status of data, including partial or corrupted data
- Kernel process status, outstanding requests, and processing time
- Network process status, processing time, and server connections
- Number of resources for each type of interprocess communication resource
- Free space, used space, and the percentage of disk space used
- Disk space usage over time displayed in a graph

The PATROL console, as it is used to monitor EnterpriseOne servers, displays icons for each server component. You can double-click these icons to monitor each component. In addition, you can choose OWKM commands for each component to gather additional data about each one.

Operating System-Specific Considerations

You can monitor EnterpriseOne servers using OWKM and the PATROL console regardless of the operating system you use. However, the monitoring setups for each operating system differ slightly.

The following table summarizes the operating system considerations for monitoring EnterpriseOne servers with OWKM and PATROL.

Operating System	Consideration
iSeries	Because PATROL has not been developed for iSeries systems, you monitor the iSeries server by using PATROL and an EnterpriseOne client install from a Windows workstation.
Windows	To monitor EnterpriseOne servers running on Windows, you can either install PATROL on the monitored server and monitor the local server, or use a remote workstation with an EnterpriseOne install of PATROL.
UNIX	To monitor EnterpriseOne servers running on UNIX, you can either install PATROL on the monitored server and monitor the local server, or use a remote workstation with an EnterpriseOne install of PATROL.

Adding OWKM Files to the PATROL Environment

You add OWKM files to the PATROL environment automatically, on either a Windows or UNIX system, by running an installation script.

► To add OWKM files to the PATROL environment automatically

Run the installation script.

- For Windows, run the script called OW_install_km.bat with the name PATROL_DIRECTORY as the argument of the script:

```
cd: c:\810\system\owpatrol\owkm\psl
```

```
OW_install_km.bat <PATROL_DIRECTORY>
```

- For UNIX, run the script called `OW_install_km.sh` with the name `PATROL_DIRECTORY` as the argument of the script:

```
cd c:/u10/peoplesoft/E810/system/owpatrol/owkm/psl

OW_install_km.sh <PATROL_DIRECTORY>

cd /u10/peoplesoft/E810/system/bin32

chmod 6555 owptrl_cli
```

Loading OWKM Files

You load the OWKM files so that the `SYSTEM_SETUP` icon appears in the PATROL console. Once EnterpriseOne and PATROL are set up, you will use the `SYSTEM_SETUP` icon to access information about your servers. This procedure needs to be done only once to configure your system.

► To load OWKM files

1. On the PATROL console, choose File, then Load KM.
2. Choose `OW_ALL.kml`.

The `SYSTEM_SETUP` icon should appear in the PATROL console. If you accidentally repeat this procedure at a later time, choose Skip All to bring up the `SETUP` icon.

Adding a Host

Before you can monitor EnterpriseOne servers using OWKM and PATROL, you must use the PATROL console to enter information about the host machine. After you have entered the required information, an icon for the host appears in the PATROL console. You can monitor information about the host by double-clicking its icon.

► To add a host

1. On the PATROL console, choose Add from the Hosts menu.
2. On Add Hosts, complete the following fields:
 - Host Name
Enter the name of the local workstation.
 - Computer Class
From the list, choose the name of the operating system for the local workstation.
 - Connection Mode
Enter the network protocol for the agent. If you started PATROL in operator mode, this field is disabled. If you started PATROL in developer mode, you can check this option to change the alarms and other features.
 - Protocol
Enter the mode of communication with the PATROL console. Use the default value.

- Port
Enter the port number used by PATROL to run on the system. It is not the same as the port number used by the EnterpriseOne server. If there is a conflict with the default value, choose an unused port number.
 - Username
Enter your network user name.
 - Password
Enter your network password.
 - Verify Password
Reenter your network password.
3. Click OK.
 4. On the PATROL console, double-click the PatrolMainMap icon or expand the node.
An icon labeled with the name of the host that you added appears in the PATROL console.

Adding an EnterpriseOne Server

Before you can monitor EnterpriseOne servers using OWKM and PATROL, you must enter information about each server. After you enter the information, an icon for each server appears in the PATROL console.

► To add an EnterpriseOne server

1. On the PATROL console, right-click the MainMap icon, and then choose the following:
 - KM Commands
 - Configure
 - EnterpriseOne Monitored Servers
2. On Configure EnterpriseOne Monitored Servers, choose the Add option and then click Execute.
3. On Add EnterpriseOne Monitored Server, complete the following fields:
 - Host Name
Enter the name of the local or remote server.
 - Port Number
Enter the server's port number.
 - Web Server
Choose this option if you are monitoring a web server.
 - EnterpriseOne Userid
Enter your EnterpriseOne ID.

- EnterpriseOne Password

Enter your EnterpriseOne password.

- JDE Base Path

Enter the path to the EnterpriseOne installation directory on your local workstation. The CLI, which communicates with any local or remote server that you monitor, is installed under this directory. For example, for Windows, you might enter d:\E810\system. For UNIX, you might enter /u15/E810/system.

4. Click OK.

After you add a server to be monitored, OW_SYSTEM is added to the MainMap icon, indicating that you have configured at least one server.

Configuring OWKM

After you have added EnterpriseOne servers for monitoring, you can configure OWKM from the PATROL console. You configure OWKM to set, for example, how frequently the CLI retrieves data from the server or the maximum log size you want to retrieve.

Some OWKM configuration settings are persistent, meaning that they remain until you change them. Others are non persistent, meaning that they revert to default status each time you restart PATROL.

Persistent settings are:

- Max Log Size to Retrieve
- Max Log Lines to Display
- Send Alarm Event When Server Bounced
- UBE Alarm Timeout

Non persistent settings are:

- Data Request Timeout
- Log Request Timeout
- Preserve Non-Error Log Files
- Use Defaults

The following table lists and briefly describes each OWKM configuration parameter:

OWKM Configuration Parameter	Explanation
Max Log Size to Retrieve	Determines how many bytes from the bottom of the log file the CLI retrieves from the server. Range is 100 to 2,000 kbytes. Default value is 100 kbytes.
Max Log Lines to Display	Determines how many lines of a log file will be displayed in the PATROL console. Range is 100 to 10,000 lines. Default value is 5,000 lines.
Send Alarm Event When Server Bounced	Sends an alarm whenever a server you are monitoring is bounced. Default value is off.

UBE Alarm Timeout	Sends an alarm if a UBE runs longer than the value you set. Range is 1 to 2,880 minutes. Default value is 60 minutes.
Data Request Timeout	Sets the number of seconds to wait to retrieve data from the server. Value range is 1 to 300 seconds. Default value is 30 seconds.
Log Request Timeout	Sets the number of seconds to wait to retrieve log files from the server. Value range is 1 to 300 seconds. Default value is 180 seconds.
Preserve Non-Error Log Files	Creates a JDEDEBUG log file each time the CLI retrieves data from the server. Use for debugging purposes only, as the CLI process will create a large number of files in your system. Default value is off.
Use Defaults	Retrieves and sets default values for each configuration parameter after you choose the option and click OK. Default value is off.

► To configure OWKM

1. On the PATROL console, right-click the OW_SYSTEM icon, and then choose the following:
 - KM Commands
 - Configure
 - EnterpriseOne KM (CLI)
2. On EnterpriseOne KM Configuration, make any necessary changes to the following configuration settings:
 - Data Request Timeout
 - Log Request Timeout
 - Preserve Non-Error Log Files
 - Max Log Size to Retrieve
 - Max Log Lines to Retrieve
 - Send Alarm Event When Server Bounced
 - UBE Alarm Timeout
 - Use Defaults
3. If you want to retrieve the default values for each setting, choose the Use Defaults option and click OK.
4. Review all configuration settings and click OK.

Enterprise Server Components Monitored by OWKM

After you add a server to be monitored by OWKM and PATROL, the PATROL console displays the name of the server. You can expand the server name node or double-click the icon to reveal the monitored components. A yellow, blinking component icon indicates warning status. A red, blinking component icon indicates alarm status. For instance, if a disk runs out of space, the disk icon blinks red. If you expand the Disks icon node, the disk that has run out of space also blinks red.

The following table lists and briefly describes each monitored enterprise server component and briefly describes what PATROL and OWKM monitor for each one:

Component	Explanation
CLI_Status	Indicates if the CLI is communicating with the server. If the status is OK, communication is occurring properly. If there is a CLI error or the server is down, the CLI_Status icon blinks to indicate the problem.
Detailed Status	Provides summary information of all the other components.
Disks	Displays disk usage on the server.
INI	Allows users to view the EnterpriseOne server jde.ini file. Server Administration Workbench (SAW) administrators can edit the file, if necessary.
Logs	Displays all the log and debug log files generated by EnterpriseOne server processes.
Processes	Displays information about EnterpriseOne processes monitored on each server, including business functions, jdenet, queue kernel, replication, SAW, Scheduler, and security.
Resources	Displays information about EnterpriseOne interprocess communication resources, such as shared memory, message queues, and semaphores.

The PATROL console also displays subcomponents of the Disks, Logs, Processes, and Resources components. You can expand the node for each of these components and double-click a subcomponent to view information about it. For example, if you expand the Disks node, you reveal nodes for each disk on the server. If you double-click an icon for a particular disk, you can view detailed status information about the disk, as well as data about the percent used, free space, and used space on the disk.

Monitoring Enterprise Servers Using KM Commands

The KM Commands menu item is an EnterpriseOne-specific addition to PATROL features. The menu appears when you right-click some icons in the PATROL console. KM Commands give you additional monitoring options for selected server components.

Note

You can choose any available enterprise server component for monitoring by right-clicking the server icon and choosing KM Commands and View.

The following table lists and summarizes the KM Commands options for each enterprise server component in the PATROL console:

Component	KM Commands Parameters	Comments
Server	- View - Refresh Parameters	- View allows user to choose an individual server component and view information about it. - Refresh Parameters gets server information immediately rather than

	• Stop Monitoring • Detailed Status	• waiting the normal time of one minute. • Stop Monitoring deletes the server from the list of servers to be monitored by PATROL and OWKM. • Detailed Status displays information on Disks, Logs, Processes, and Resources components in one form.
Disks	• View • Detailed Status	• View allows user to view a list of disks and the percentage used for each. Choose a disk and click View Details for host, port, path, mount point, and usage data.
INI	• View • Edit	• View allows users to view only the server's jde.ini file. • Edit allows users with Server Administration Workbench (SAW) administrator privileges to edit the server's jde.ini file.
Logs	• View • Detailed Status	• View displays lists of log files, debug files, server logs, and print queue logs. Users can sort, select, and delete logs. Choose an individual log and click Execute to view the log. • Detailed Status displays log type, file count, and disk usage information.
Processes	• View • Detailed Status • Clear Missing Process Alarms	• View displays a list of processes with options to view details, view error log, view debug log, or delete logs. • Detailed Status displays process ID, type, and name of each process. • Clear Missing Process Alarms allows an administrator to delete blinking (alarm) icons for failed processes.
Resources	• View • Detailed Status	• View displays a list of resource names and types. Choose a resource and click View Details to view additional information about the resource, including the processes used, the time, and the state of the process, such as unlocked. • Detailed Status displays each resource name and type.

Monitoring Enterprise Server Components

After you have configured OWKM, you are ready to use the PATROL console to monitor EnterpriseOne enterprise server components.

Monitoring Enterprise Server Components from the PATROL Console

You can monitor the components of each enterprise server that you have added for monitoring.

► To monitor enterprise server components from the PATROL console

1. On the PATROL console, double-click the OW_SYSTEM icon.
2. Double-click an enterprise server icon to view the components monitored by PATROL and OWKM.
3. Double-click one of the following component icons to view its status:
 - CLI_Status
 - Detailed Status

- Disks
- Logs
- Processes
- Resources

Editing the Enterprise Server jde.ini File

If you have Server Administration Workbench (SAW) administrator privileges, you can edit the enterprise server's jde.ini file from the PATROL console. Without administrator privileges, you are permitted only to view the file.

► To edit the enterprise server jde.ini file

1. On the PATROL console, double-click an enterprise server icon.
2. Right-click the INI icon and choose the following:
 - KM Commands
 - Edit

The jde.ini file appears.

3. Edit and save the file.

Monitoring Enterprise Server Log Files

You can access and view the log files for each enterprise server that you monitor from the PATROL console. You can also sort and select the files using options in the Server Log Files form.

► To monitor enterprise server log files

1. On the PATROL console, double-click an enterprise server icon.
2. Right-click the Logs icon and choose the following:
 - KM Commands
 - View
 - All Server Logs
3. On EnterpriseOne Server Log Files, choose one of the following options to change the view of the list of logs:
 - View Log
 - Sort by Size
 - Sort by Date
 - Sort by Name
 - Select by Number
 - Reverse Selections
 - Clear Selections

4. To view a log file, select the log, choose the View Log option, and click Execute.
The EnterpriseOne web server log file or debug file appears in the PATROL console.
5. To search the file, right-click inside the file, choose Find and enter a word to search for, such as "Error," and a direction to search, either up or down.

Note

Remember that the maximum size of the file you can retrieve when you use this form is set when you configure OWKM. If the size of the file you want to retrieve exceeds this size, you must reconfigure OWKM and increase the value of this parameter.

Viewing Log Files for an Enterprise Server Process

You can monitor an enterprise server process and view the related log files.

► To view log files for an enterprise server process

1. On the PATROL console, choose a server whose processes you want to monitor and double-click the icon.
2. Double-click the Processes icon.
3. Right-click a process and choose the following:
 - KM Commands
 - View
 - Log File or Debug Log File

PATROL displays the contents of the log file or debug log file.

Deleting Log Files for an Enterprise Server Process

You can delete the log or debug log files for an EnterpriseOne enterprise server process. You should do this periodically to free up space on your disk.

► To delete log files for an enterprise server process

1. On the PATROL console, select a process whose log file you want to delete.
2. Right-click the process icon and choose the following:
 - KM Commands
 - Delete
 - Log File or Debug Log File
3. On Confirmation, click Yes to delete the file.

Viewing Enterprise Server Disk Information

You can view disk information for each server you monitor from the PATROL console.

► To view enterprise server disk information

1. On the PATROL console, double-click the icon for the server whose disk information you want to view.
2. Right-click a disk icon and choose the following:
 - KM Commands
 - View
3. On EnterpriseOne Disks, select the name of a disk and then click View Details.

The Disk Detailed Status form appears, containing the following disk information:

- Host
 - Port
 - Name
 - Volume
 - Used space
 - Free space
 - Percent used
 - Mount Point
4. Click OK to exit the form.

Note

Clicking Detailed Status after KM Commands launches the Detailed Status form, which allows you to view disk name, mount point, used space, free space, and percent used.

Changing Alarm Ranges

Some server component settings have alarm ranges that are adjustable by administrators. The alarm settings determine when a component will blink yellow, indicating a warning status, or blink red, indicating an alarm status. To change the alarm status, you must run the PATROL console as a developer, not as an operator.

► To change alarm ranges

1. Log on to PATROL as a developer.
2. On the PATROL console, click the KM tab.
3. Expand the Knowledge Module and Application Classes nodes.

4. Choose an Application Class, such as OW_FILESYS, and expand the node.
5. Expand the Parameters node.
6. Double-click a parameter.
7. On Parameter Properties, click the Alarm Ranges tab.
8. Change the alarm ranges by making any necessary adjustments to the following controls:
 - Border, which is the range of valid values for the alarm
 - Alarm 1, which sets the range for a warning
 - Alarm 2, which sets the range for an alarm
9. Click Apply.

You can use PATROL and OWKM to monitor web servers as well as enterprise servers. When you add a web server to be monitored, the PATROL console again displays the name of the server, and you can view the monitored components by clicking the server icon node or by double clicking the icon. The PATROL console again indicates a warning status for a component with a yellow, blinking icon, and an alarm status for a component with a red, blinking icon.

Monitoring Web Server Components

You can use BMC PATROL and OWKM to monitor Web server components, including logging files, users, and connection pools.

Web Server Components Monitored by OWKM

The following table lists each monitored Web server component and briefly describes what PATROL and OWKM monitor for each one:

Component	Explanation
CLI_Status	Indicates if the CLI is communicating with the server. If the status is OK, communication is occurring properly. If there is a CLI error or the server is down, the CLI_Status icon blinks to indicate the problem.
Detailed Status	Provides summary information for all the other components.
Server Status	Indicates whether the Web server you are monitoring is running or not.
Connection Pools	Provides information on all the EnterpriseOne enterprise servers that users connect to through the Web server.
Users	Displays information on all the users connected to the Web server.

The PATROL console also displays subcomponents of the Connection Pools and Users components. You can expand the node for each of these components and double-click a subcomponent icon to view more information about it. For example, if you expand the Connection Pools node, you reveal nodes for each enterprise server that users connect to through the Web server. If you double-click an icon for a particular enterprise server, you can view detailed status information about the enterprise server, as well as data about the number of business functions running on that server.

Monitoring Web Servers Using KM Commands

The KM Commands menu item is an EnterpriseOne-specific addition to PATROL features. The menu appears when you right-click some icons in the PATROL console. KM Commands give you additional monitoring options for selected Web server components.

Note

You can choose any available Web server component for monitoring by right-clicking the server icon and choosing KM Commands and View.

The following table lists and summarizes the KM Commands options for each Web server component in the PATROL console:

Component	KM Commands Parameters	Comments
Server	- View - Delete - Set Logging Flags - Refresh Parameters - Stop Monitoring - Detailed Status	- View allows user to choose an individual Web server component and view information about it. - Delete allows user to delete the log file and debug log file for the monitored Web server. - Setting Log Flags allows user to enable or disable logging for JAS and network files. - Refresh Parameters gets server information immediately rather than waiting the normal time of one minute. - Stop Monitoring deletes the server from the list of servers to be monitored by PATROL and OWKM. - Detailed Status displays information on Users and Connection Pools components in one form.
Connection Pools	- View - Detailed Status	- View displays list of enterprise servers connected to the Web server. Choose an enterprise server and click View Details for a list of all business functions running on the enterprise server. - Detailed Status displays in one form business function data for all enterprise servers connected to the Web server.
Users	- View - Detailed Status	- View: Displays list of users logged on to the Web server, by user ID and machine name. Select a user and click View Details for additional user data. - Detailed Status displays in one form data about each user logged on to the Web server, including user ID, Web server host, enterprise server host, environment, log-in time, browser name, and machine name.

Monitoring Web Server Components from the PATROL Console

You can monitor the components of each Web server that you have added for monitoring.

► To monitor web server components from the PATROL console

1. On the PATROL console, double-click the OW_SYSTEM icon.
2. Double-click a web server icon to view the components monitored by PATROL and OWKM.

3. Double-click one of the following component icons to view its status:
 - CLI_Status
 - Connection Pools
 - Detailed Status
 - Server Status
 - Users

Monitoring Web Server Log Files

A separate component icon for Web server logs in the PATROL console does not exist. You access the files from the Web server icon.

► To monitor web server log files

1. On the PATROL console, right-click the server icon and choose the following:
 - KM Commands
 - View
 - Logs
 - Log File or Debug File

The EnterpriseOne web server log file or debug file appears in the PATROL console.

2. To search the file, right click inside the file, choose Find and enter a word to search for, such as "Error," and a direction to search, either up or down.

Deleting Web Server Log Files

You can delete EnterpriseOne Web server log and debug logs. You access the command to delete the files from the Web server icon.

► To delete web server log files

1. On the PATROL console, right-click the server icon and choose the following:
 - KM Commands
 - Delete
 - Log File or Debug File
2. On Confirmation, click Yes if you are sure you want to delete the file.

Setting Logging Flags

The Set Logging Flags parameter allows you to enable or disable the log files for the JAS server and for the network. Note that you cannot view the Web network files from the PATROL console. When you view the log files for the server, the PATROL console simply notifies you that the network trace is turned on.

► To set logging flags

1. On the PATROL console, right-click the server icon and choose the following:
 - KM Commands
 - Set Logging Flags
2. On Set Web Log Flags, choose the following options if you want to enable JAS and network server logging:
 - Enable Java Server Log
 - Enable Net Server Log
3. Click OK.

An Information form notifies you whether or not the flag was set successfully.

Monitoring Web Server Users

You can monitor the users who are currently logged on to EnterpriseOne Web servers. You can choose an individual user and view data about that user, or you can view, in a single form, data about all the logged-on users.

► To monitor web server users

1. On the PATROL console, double-click the web server icon.
2. Right-click the Users icon and choose the following:
 - KM Commands
 - View
3. On EnterpriseOne Web Users, choose a user and click View Details.
4. On User Detailed Status, click OK.

To view in a single form data on all of the users, choose KM Commands and Detailed Status.

Monitoring Web Server Connection Pools

You can monitor the connection pools of enterprise servers that are connected to the EnterpriseOne Web server. You can choose an individual enterprise server and view business data about business functions running on that server, or you can view, in a single form, business function data about all the enterprise servers connected to the Web server.

You monitor the following types of information about enterprise servers connected to Web servers:

- Names of business functions that have run on an enterprise server during a server session
- Minimum time required to run a business function
- Maximum time required to run a business function
- Average time required to run a business function
- Number of timeouts that occurred during the running of a business function
- Number of system errors that occurred during the running of a business function

- Number of application errors that occurred during the running of a business function
- Number of times a business function was called during a server session

► **To monitor web server connection pools**

1. On the PATROL console, double-click the web server icon.
2. Right-click the Connection Pools icon and choose the following:
 - KM Commands
 - View
3. On EnterpriseOne Connection Pools, choose an enterprise server and click View Details.
4. On Connection Pools Detailed Status, click OK.

To view in a single form data on all the enterprise servers connected to the EnterpriseOne web server, choose KM commands and Detailed Status.

EnterpriseOne on Windows Terminal Server Edition

Windows Terminal Server Edition (TSE) provides an excellent solution for EnterpriseOne in a WAN environment. TSE allows you to set up multiple terminal server client machines that need only contain the TSE client software. You can use less powerful machines to function as terminal server clients. These clients connect to a machine set up with TSE software. Multiple users can simultaneously connect to the same terminal server to run EnterpriseOne.

Understanding Windows Terminal Server Edition

Windows Terminal Server Edition (TSE) is a multi-user extension to the Microsoft Windows family of operating systems. TSE allows users to share an application that resides on the terminal server. The terminal server performs all the processing for an application, and then sends a picture of the screen to the client terminal. Only keystrokes and mouse movement occur at the terminal. These movement commands travel through the network to the server, which returns the modified screen to the terminal.

The following provides an example of the TSE process flow:

Step 1: EnterpriseOne client applications execute on terminal server.

Step 2: The terminal server sends the video for the user interface across any connection.

Step 3: The terminal server client displays the user interface.


Step 4: The terminal server client sends actions, such as keystrokes and mouse movement, back to the terminal server where processing occurs.

By sending only the information necessary to re-create the screen and convey mouse and keyboard events, TSE provides LAN-like performance over WAN and dialed connections.

TSE allows you to set up multiple users to work with a single client installation of EnterpriseOne. By sharing a single copy of EnterpriseOne on the terminal server, you reduce the costs of deployment and administration.

The following illustration provides an example of an EnterpriseOne configuration with a terminal server:

PeopleSoft EnterpriseOne Configuration: Terminal Server



Note

Sun-Solaris can be added to the list of enterprise servers that can run in an EnterpriseOne configuration with a terminal server.

Incorporating Citrix MetaFrame with TSE

TSE provides multi-user technology that uses a presentation protocol called Remote Desktop Protocol (RDP). RDP, based on the International Telecommunications Union T.120 protocol, is a viable option if you plan to use only Win16/Win32 clients in an uncomplicated configuration.

If your network comprises multiple platforms and requires optimum performance, Citrix offers a product called Citrix MetaFrame with multi-user technology that provides additional functionality to TSE, such as load-balancing and the support of more client platforms. MetaFrame uses the Independent Computing Architecture (ICA) presentation protocol on which Citrix based WinFrame, a successful multi-user product for Windows 3.51.

Note

If you start EnterpriseOne as a specified application through ICA, you cannot view the jde.ini, jde.log, and jdedebug.log files.

The following table lists the capabilities of TSE RDP and MetaFrame:

Capability	TSE RDP	MetaFrame
Client Platforms	TSE client software runs on the following platforms: - Windows 16-bit - Windows 32-bit - Some RDP-equipped Windows terminals - Windows CE	MetaFrame runs on the following platforms: - DOS - Windows 16-bit - Windows 32-bit - X-Term - MacIntosh - Solaris - Windows CE - Some ICA-equipped network computers - Some Internet browsers as a plug-in
Network Topologies	TSE RDP supports the TCP/IP standard.	MetaFrame supports the following standards: - IPX - SPX - PPP - NetBIOS
Load-balancing	Windows 2000 supports load balancing.	You can purchase an option for MetaFrame that provides load-balancing capabilities.
Encryption	N/A	You can purchase an option for MetaFrame that provides the encryption of ICA traffic.
Capability	TSE RDP	MetaFrame
CCPDD:cut/copy/paste/drag/drop	Windows 2000 supports cut, copy, paste, drag, and drop.	MetaFrame allows CCPDD between the session window and the underlying Windows desktop.

Device Mapping	TSE RDP allows you to map local devices for printing through a work-around.	MetaFrame allows you to map devices local to the TSE client from the terminal server. For example, you can locally map hard drives, fax modems, and printers.
Session Shadowing	Windows 2000 supports remote control.	With MetaFrame loaded, TSE supports an administration tool called session shadowing. Session shadowing helps administrators audit remote sessions. You might also use session shadowing for video conferencing and in a support desk role.

TSE Restrictions in Multi-user Mode

EnterpriseOne is "TSE-aware." TSE-aware means that when you deploy EnterpriseOne as a client on a terminal server, the software automatically recognizes the terminal server and configures itself to run in multi-user mode. Because of the configuration required by multi-user mode, the following restrictions apply to terminal server users:

- Disabled development with Form Design Aid, Report Design Aid, and Table Design Aid, including Object Management Workbench check-in and check-out capabilities

Note

The restriction on development does not prevent new versions of existing applications, but only the modification of current applications and the creation of new applications. The reason for this restriction is that Microsoft recommends against using Visual C++ remotely through TSE under any circumstances. Also, a complete set of development specifications (500+ MB) defeats the purpose of a thin client.

- Disabled local processing for batch applications

All batch applications process on a separate batch server to avoid an impact to performance on the terminal server.

- Disabled Just-In-Time-Installation (JITI)

Because you deploy a full client package to the terminal server, JITI is not recommended.

- Disabled intensive specification file access

Applications that intensively access specification files, such as the Universal Table Browser application, are not active due to the strain put on the terminal server when the application retrieves data. Also, file level locking could prevent access to data in the specification files for other users.

A standard EnterpriseOne configuration and a terminal server configuration are not mutually exclusive within an enterprise. You can mix a standard configuration with a terminal server configuration to maximize your overall performance. For example, you can use a standard EnterpriseOne configuration over a LAN, and use a terminal server configuration to support remote sites across a WAN.

Network Considerations

The terminal server must reside on the same local area network (LAN) as the enterprise server or database server, or both. Include one normal EnterpriseOne client on the LAN to verify performance and function. Normal EnterpriseOne LAN requirements apply.

For a wide area network (WAN), you must use a 56KB, or faster, line.

Performance Considerations

When you add any ICA session, change the Window Colors display properties to 16 color mode or to the lowest setting that your software and hardware allow. You can access display properties from the Control Panel.

Also, for your Citrix setup, you need to choose the Compress data stream and Cache bitmaps to disk options. See the appropriate Citrix documentation for more information about how to modify these settings.

Setting Up EnterpriseOne on the Terminal Server

Because EnterpriseOne is "TSE-aware," running on a terminal server is almost identical to running a standard client.

When performing certain processes, such as creating log files and running UBEs, EnterpriseOne checks whether it resides on a terminal server or standard client. If it detects a terminal server, the software automatically switches to multi-user mode.

In multi-user mode, EnterpriseOne processes data while simultaneously protecting data integrity and maintaining performance on the terminal server. Multi-user mode also masks any of the multi-user activity from a user so that a terminal server session of EnterpriseOne looks no different than a standard EnterpriseOne client session.

► To set up EnterpriseOne on the Terminal Server

1. Install Windows Terminal Server Edition on the machine that you will use as your terminal server.

Refer to the Microsoft documentation for information about how to set up Windows Terminal Server software.

2. Install Microsoft Terminal Server Client software onto the machines that you will use as terminal server clients.

You can connect to the terminal server from the Terminal Server Client option on the Programs menu. This file also resides in the Terminal Server Client subdirectory in the Program Files directory. Refer to Microsoft documentation for information about how to set up Terminal Server Client software.

3. Depending on the type of database you use, you might need to install software so that your client can properly connect with the server where the database resides.

Note

You must use Add/Remove Programs on the Control Panel to install applications on the terminal server. During installation, make sure that you choose the All users begin with common application settings option on the Change User Option dialog box.

Install the following software on the terminal server as necessary.

- Oracle for Windows
 - SQL Server client
 - Client Access
 - DB2 Connect
4. Install a full package of EnterpriseOne. You need to install a full package because EnterpriseOne on the terminal server is multi-user. If you install a partial package, multiple users will simultaneously experience Just-In-Time-Installation, which will negatively affect performance on the terminal server.
-

Note

Use Add/Remove Programs on the Control Panel to install EnterpriseOne on the terminal server. During installation, make sure that you choose the All users begin with common application settings option on the Change User Option dialog box.

After you perform these steps, you should be able to successfully run EnterpriseOne from terminal server client machines.

See Also

- *Hardware and Software Requirements* in the *EnterpriseOne Installation Guide* for information about additional requirements or service packs needed to run software packages

Troubleshooting EnterpriseOne on Windows Terminal Server Edition

This chapter details situations that can exist when you run EnterpriseOne on a Microsoft Windows Terminal Server Edition (TSE).

Troubleshooting: UBE Output Security on TSE

An EnterpriseOne TSE user sends PDF files by default from the Enterprise server to the local \\E810\PrintQueue directories using the Work With Servers (P986116) application. Users choose View PDF from the Row menu of the Submitted Job Search form. Because the files are saved to a user's local PrintQueue directory, another user can view the PDF file in Windows Explorer or in Adobe Acrobat.

You can relocate the PrintQueue directory by adding the following section to the jde.ini file:

[NETWORK SETTINGS]

OutputDirectory=C:\WTSRV\Profiles\USERNAME\Windows

Server administrators need to make the jde.ini file modification of each user on each TSE so that the PDF output points to each individual's user profiles directory. With the PrintQueue directory located in the user profiles directory, the PDF files are protected by Windows security. Only server and system administrators have access to the files.

Troubleshooting: Submit UBE Locally to TSE

You cannot locally submit UBEs to run on the TSE because of resource constraints such as CPU power.

You can run UBEs on a dedicated TSE or during hours when no other EnterpriseOne users are using the machine. Complete the task to submit a UBE locally and run it on the TSE.

► To submit a UBE locally and run it on the TSE

From the System Administration Tools menu (GH9011), choose Logical Data Sources (P986115).

1. On Logical Data Sources, choose the TSE machine name with System as the data source and click Select.
2. On Work with Data Sources, click Add.
3. On Data Source Revisions, complete the following fields:
 - Data Source User
Enter DB to specify a Local Data Source.
 - Data Source Name
Enter TSE Local.
 - Data Source Type
Enter M to specify MSDE/OLEDB.
 - DLL Name
Enter JDBODBC.DLL.
 - Database Name
Enter System Local.
 - Server Name
Enter LOCAL.
 - Platform
Enter LOCAL.
4. When you run the UBE on the TSE, on the Work with Batch Versions from, choose the report you want to run and click Select.

5. On Version Prompting, choose Advanced from the Form menu.
6. On Advanced Version Prompting, choose the Override Location option and click OK.
7. On Version Prompting, choose the Data Selection option and click Submit.
8. On JDE Data Sources, choose TSE Local as the data source and click Select.

Troubleshooting: Import/Export with Microsoft Excel

Importing a Microsoft Excel spreadsheet into an EnterpriseOne grid intermittently fails when users are running TSE.

PeopleSoft recommends using Microsoft Office 2000 on TSE. To install Microsoft Office 2000 in a TSE environment, access the following URL:
<http://www.microsoft.com/Office/ORK/2000/Two/3053.htm>. You install a single of Microsoft Office 2000 on the TSE. Multiple users then connect to the server and run Microsoft office from the server.

Troubleshooting: Specification Files are Locked

In the following circumstances, EnterpriseOne users get a message box to indicate that a specification file is currently unavailable:

- When another EnterpriseOne session on the same TSE machine performs a data dictionary Just-In-Time-Installation (JITI).
- When another EnterpriseOne session receives a TSE run-time error dialog box. Generally, this type of error occurs when a memory violation occurs.

In both cases, the specification file or files are locked. The specification files will be unlocked when either the TSE completes the JITI process or a user closes the TSE run-time error dialog box.

To prevent the specification files from being locked when EnterpriseOne performs a JITI, reduce the frequency that EnterpriseOne performs JITIs. Complete the task appropriate to your release number to change the frequency of JITIs:

- Reduce JITI frequency (B73.2 and B73.3)
- Reduce JITI frequency (B73.3.1, B73.3.2, and B73.3.3)

► To reduce JITI frequency (B73.2 and B73.3)

1. On a non-TSE workstation, create the business function B98CRTGL.
2. Go to Design and type GenGbltblSpec in the Function Name field.
3. Choose the row and then choose Parameters from the Row menu.
4. Type D9800330, click Find, and then click Select.
5. Click OK.
6. On Business Function Design, click OK.
7. Copy B98CRTGL.c to the \b7\PRODB732\source directory and B98CRTGL.h to the \b7\PRODB732\include directory.

Important

The contents of the specification files are different for B732 and B733. Make sure that you copy the correct specification files for your release.

8. On Business Function Source Librarian, click Build.
Verify that the project configuration in BusBuild is Optimize.
9. Create a report (UBE) without a business view. On this report, create a group section that calls B98CRTGL (business function created in step 1).

Note

Parameter passing is not necessary.

10. Create a version of the UBE and then run the version locally on the workstation. Copy the full glbltbl.ddb and glbltbl.xdb to the TSE machines.

The UBE should take less than an hour, during which time, you will see that the JITI process is running on the machine. The UBE generates the full GLBLTBL specification in the ..\b7\PRODB732\spec directory.

► To reduce JITI frequency (B73.3.1, B73.3.2, and B73.3.3)

1. Run the Generate global table spec (R98CRTGL) batch application on a non-TSE EnterpriseOne client to generate full GLBLTBL specification files.
2. Copy the full glbltbl.ddb and glbltbl.xdb to the TSE machines.

Troubleshooting: User Cannot Restart EnterpriseOne

Occasionally, when a memory violation occurs in an EnterpriseOne TSE session, the terminal server prevents the user from restarting EnterpriseOne. The administrator must then sign onto the terminal server and end the OEXPLORER.exe process from the Task Manager. After the administrator ends the process, the user can sign on to EnterpriseOne again.

When a run-time exception occurs, EnterpriseOne should immediately exit. To instruct EnterpriseOne to immediately exit in this situation, change the EXCEPTION_Enabled setting in the jde.ini to False, as shown:

<pre>[INTERACTIVE RUNTIME] EXCEPTION_Enabled=False</pre>

Troubleshooting: Logging Off Versus Disconnecting

Users should always log off their TSE session rather than disconnecting. Logging off shuts down all processes completely for the user.

Troubleshooting: Shortcuts Do Not Work in E-Mail Messages

Workflow provides the ability to send shortcuts to EnterpriseOne applications via e-mail messages. This function does not work when the e-mail application, such as Microsoft Outlook, is not currently active on the terminal server. When the e-mail application invokes the shortcut, the operating system attempts to launch the shortcut on the local machine and not on the terminal server.

Run the e-mail application and EnterpriseOne on the same terminal server machine.

Troubleshooting: Data Selection and Sequencing Criteria Lost

The following situation occurs when two or more users are signed on to the same terminal server using the same pathcode.

The first user submits a batch application from Batch Versions, changes the data selection criteria, and then stops at the printer screen. The second user then goes into Batch Versions to submit the same version of the batch application, changes the data selection criteria, and then stops at the printer screen. When the users click OK to send the batch application to the enterprise server for processing, the data selection criteria for the second user overrides the selection criteria for the first user.

A modification to batch processing in EnterpriseOne now saves data selection and sequencing criteria in memory rather than in specification files.

Troubleshooting: Run-Time Error Occurs During Server Connection Test

The Server Administration Workbench (SAW) application receives a run-time error when it performs a server connection test. This situation occurs when the user who performs the connection test does not possess the authority to access the pinging mechanism on the target machine.

Troubleshooting: EnterpriseOne Development Tools Are Disabled

EnterpriseOne development tools are disabled on the terminal server. Currently, PeopleSoft instructs customers to perform all development on non-TSE machines.

Troubleshooting: Users Experience Problems Accessing EnterpriseOne

Only administrators can run EnterpriseOne. This situation is a result of the way EnterpriseOne was installed on the terminal server.

As the Administrator, you should use the Add/Remove Programs application on the Control Panel to install EnterpriseOne on the terminal server. During installation, make sure that you choose the "All users begin with common application settings" option on the Change User Option dialog box. This option ensures that the terminal server maintains EnterpriseOne specific files, such as the jde.ini file, across user profiles.

Troubleshooting: Log Path is Incorrect

The log path in the jde.ini for individual users is incorrect.

The EnterpriseOne installation program sets the WTSLogs setting to False. Users should change this setting to True after the installation and before any users run EnterpriseOne.

When the WTSLogs setting is True, the output log directories for each user point to the home directory of the user rather than to the root directory of the drive. The output log directories settings are also defined in the jde.ini file.

Troubleshooting: Shortcut Path is Incorrect

The shortcut path on the terminal server is incorrect. This situation occurs when the terminal server uses the server-based profiles.

The resolution involves "Shortcuts Created Under TSE 4.0 Resolve to UNC Paths." For more information, see the following web address:

<http://support.microsoft.com/support/kb/articles/Q195/8/87.ASP>

Troubleshooting: Only One User Can Log in to EnterpriseOne

The main EnterpriseOne window fails to appear after entering the password for all other users.

For B73.3.2 with Service Pack 10 or greater, place the EnterpriseOne command line switch "/NoLogo", located in the EnterpriseOne shortcuts, on the desktop and on the Start menu if not already present. Separated by a space, append the text to the end of the line in the "Target" edit box of the shortcut properties window. Do not include the quotation marks. This will prevent the display of the splash screen. There is no resolution for EnterpriseOne version B73.3.2 Service Pack 9.

EnterpriseOne on a Cluster

High availability clusters provide redundancy of software and hardware so that a single point of failure will not interrupt service. If a failure occurs, the clustering software automatically detects the problem and shifts to an alternate machine without ending processes and interrupting your enterprise.

Clustering allows EnterpriseOne processes running on a machine that fails to continue running without interruption on a second machine. The second machine has a setup that supports the given processes. Essentially, EnterpriseOne "moves" to the alternate machine without requiring you to restart a process that was active on the machine that failed.

Note

Each node in the cluster must have the appropriate software and hardware to ensure that processing moves successfully from server to server.

Make sure you understand the clustering software and the tasks necessary to implement the software on a given platform.

Hp-UX Clustering

Hewlett-Packard provides two mutually exclusive software products to manage high availability clusters:

- Hewlett-Packard Multi-Computer/ServiceGuard (MC/ServiceGuard)
- Hewlett-Packard Multi-Computer/LockManager (MC/LockManager)

You must use MC/LockManager when you use Oracle Parallel Server (OPS).

You will set up only one of these products for your HP-UX cluster.

EnterpriseOne requires a named IP address for workstations to connect with a server. With the Hewlett-Packard clustering software, you can assign a floating IP address that can move from node to node within the cluster. You should enter this IP address into the WINS or DNS database so that workstations can access the address. If your enterprise servers are not using DNS to resolve host names, you must also add the floating IP address to the `/etc/hosts` file on each node in the cluster where EnterpriseOne might run.

Note

If you do not assign a floating IP address, then whenever EnterpriseOne moves to another node in the cluster, the workstations will be unable to connect with the servers.

Configuring Oracle Parallel Server (OPS)

Oracle Parallel Server (OPS) allows concurrent database access from multiple nodes in a cluster. If you use OPS, you must install MC/LockManager, not MC/ServiceGuard.

► To set up Oracle Parallel Server and MC/LockManager

This task describes how to set up OPS and MC/LockManager. For more information, refer to the Hewlett Packard documentation on setting up OPS and MC/LockManager.

1. Install Oracle client software on each node in the cluster where it might run. The EnterpriseOne database should be created on shared disks, or on a machine outside the cluster, so that it can be accessed from multiple nodes.
2. Create a package in MC/ServiceGuard with no services.

This package should be set up with an associated IP address so that it can be reached from any node where it might run. This package should also specify the shared volume group on which the EnterpriseOne database will reside.
3. Edit the package control script to add the Oracle startup and shutdown commands. The following code sample provides an example of the *customer_defined_run_cmds* function from a package control script:

```
function customer_defined_run_cmds
{
# ADD customer defined run commands.
export ORACLE_HOME=/u01/app/oracle/product/8.0.5
export ORACLE_SID=jde1
export ORAENV_ASK=NO
. $ORACLE_HOME/bin/oraenv
su oracle -c '$ORACLE_HOME/bin/lsnrctl start'
su oracle -c '$ORACLE_HOME/bin/svrnmgrl' <<EOF1
connect internal
startup
exit
EOF1
test_return 52
}
```

4. You can use these same steps to enter the Oracle shutdown commands into the *customer_defined_halt_commands* section of the package control script.

Setting Up an Oracle Package for MC/ServiceGuard

If you use MC/ServiceGuard, you should set up a package for Oracle. An Oracle package allows the Oracle processes to move from one node to another when a node fails or during scheduled maintenance.

You do not need to perform this step if you use MC/LockManager.

► To set up an Oracle package for MC/ServiceGuard

1. Install Oracle on each node in the cluster.

Create the EnterpriseOne database on shareable disks so that multiple nodes can access the database.

2. Create a package in MC/ServiceGuard with no services.

You should set up this package with an IP address so that any node on the cluster can access and run the package. This package should also specify the shared volume group where the database will reside.

3. Edit the package control script to add the Oracle startup and shutdown commands.

The following code sample provides an example of the `customer_defined_run_cmds` function from a package control script:

```
Function customer_defined_run_cmds
{
# ADD cusomter defined run commands.

Export ORACLE_HOME=/u01/app/oracle/product/8.0.5

Export ORACLE_SID=jde1

Export ORAENV_ASK=NO

.$ORACLE_HOME/bin/oraenv

su oracle '$ORACLE_HOME/bin/svrnmg1'<<EOF1

connect internal

startup

exit

EA0F1

Test return 52

}
```

You can use the same steps to enter the Oracle shutdown commands into the `customer_defined_halt_commands` section of the package control script.

Setting Up an EnterpriseOne Package

The standard EnterpriseOne Enterprise server software requires minimal modifications to function in a cluster. These modifications include the following items:

- Enterprise server `jde.ini`

- owenv script file in the \$SYSTEM/bin32 directory
- Package control script

► **To set up an EnterpriseOne package for a cluster**

1. In the server jde.ini file, locate the [CLUSTER] section, then change the following setting:
[CLUSTER]

PrimaryNode=Package IP name

Variable Value	Description
EnterpriseOne Package IP Name	This value represents name given to the IP address that is associated with the EnterpriseOne package that you are creating.

2. Change the owenv script file in the \$SYSTEM/bin32 directory. The owenv script file contains the settings for various UNIX environment variables required by EnterpriseOne.
3. Create a package using SAM.

Note

You can also use the command line to create a package. See HP documentation for details.

For EnterpriseOne, set up a package with an associated floating IP address, but with no services. This setup is necessary because the cluster manager needs to start services without environment variables under the root user. The software will not run properly unless you set environment variables. This setup also allows you to utilize the installation defaults and the start and end scripts provided by PeopleSoft.

Depending on the needs of your enterprise, you might want to install EnterpriseOne on a shared volume group. This setup allows multiple nodes in a cluster to access a single version of EnterpriseOne, but only one node at a time. This setup also allows you to easily update EnterpriseOne through server package installations.

4. Using SAM, modify the package control script on each node to start and stop EnterpriseOne.

Note

You might need to vary the control scripts for different nodes in the cluster to configure different volume group names or path names. In this case, you must edit the scripts on each node individually instead of using SAM.

The following code sample provides an example of the customer_defined_run_cmds function from a package control script:

```
function customer_defined_run_cmds
{
# ADD customer defined run commands
```

```
# wait 60 seconds for Oracle to come up
sleep 60
. /home/jde/owenv
su jde << EOF1
mv $OWHOME/log/jde*.log $OWHOME/log/oldlogs
cd $SYSTEM/bin32
RunOneWorld.sh
EOF1
test_return 51
}
```

The following code sample provides an example of the `customer_defined_halt_cmds` function from a package control script:

```
function customer_defined_halt_cmds
{
# ADD customer defined halt commands.
. /home/jde/owenv
su jde << EOF2
cd $SYSTEM/bin32
EndOneWorld.sh
sleep 15
rmics.sh
EOF2
test_return 52
}
```

The following list provides explanations for these functions:

- `sleep 60`

The “run” function first waits 60 seconds for Oracle processes to start. HP states that you should set the *PKG_SWITCHING_ENABLED* parameter to *NO* for applications that access OPS. This setting prevents these applications from starting before OPS is active. If you use the sleep command in your script, you do not need to modify this setting. You can remove the sleep command from the script when you use the control script for a backup node with OPS running.

`./home/jde/owenv`

This line runs the *owenv* script provided by PeopleSoft to set up UNIX environment variables. The *owenv* script resides in the *\$SYSTEM/bin32* directory. You must edit this script to ensure that the correct setup exists for all necessary environment variables for EnterpriseOne and Oracle. In these examples, the script has been moved to the home

directory of the jde user. The script might need to move to the home directory if you use a different SID to access Oracle from different nodes.

su jde

This line switches to the user ID that owns the EnterpriseOne processes. If you omit this line, the root user owns the EnterpriseOne processes.

mv \$OWHOME/log/jde*.log \$OWHOME/log/oldlogs

This line moves any logs in the log directory to a backup log directory, which you create. This command is particularly important if your EnterpriseOne instance resides on a shared disk where a “failed over” instance of EnterpriseOne will use the same physical disk space as the failed instance. You might consider adding the command *rm \$OWHOME/log/oldlogs/** before this line to clean out any older versions of logs.

RunOneWorld.sh; EndOneWorld.sh

These are the standard start and stop scripts that PeopleSoft provides for the UNIX enterprise server.

The directory that contains the package control script also contains the control.sh.log file, which contains the results of starting and stopping a package. This file is the first place to check if problems arise when you start or stop a package. In particular, it will contain any output or error messages from the customer-defined commands you might enter.

Maintaining Multiple Instances of EnterpriseOne

When you run multiple instances of EnterpriseOne in a clustered environment, you must consider several factors. Even though each instance might begin on a separate node, a situation might arise for which multiple instances need to run on the same node. When this happens, communication to each EnterpriseOne instance must occur on a different port number, or service name, and each instance must use a different range of IPC keys. The following parameters in the jde.ini file control these settings:

```
serviceNameListen=Service Name or Port Number
serviceNameConnect=Service Name or Port Number
[JDEIPC]
startIPCKeyValue=Numeric Value
```

Variable Value	Description
Service Name or Port Number	The service name parameters can use an actual port number or the name of a service that you enter into the /etc/services file.
Numeric Value	The IPC key values should differ by at least 1000 between any two EnterpriseOne instances.

Troubleshooting: HP-UX Clustering

The sections below address specific problems with HP-UX clustering.

Problems with Oracle Parallel Server (OPS)

Complete the following steps if you experience problems with OPS:

- Verify that the cluster software is operational. OPS requires the cluster software to start before OPS can start.
- Verify that DLM is enabled in the DLM configuration. Oracle Group Management Services (OGMS) will not start if DLM is disabled.

EnterpriseOne Does Not Start

Complete the following steps if EnterpriseOne fails to start:

- When you start EnterpriseOne using the package control script, first check the control script log for errors. Look for errors in the script that occur before the RunOneWorld.sh command.
- Check the log directory for log files. If none reside, verify that the EnterpriseOne processes exist in the proper directory and that you correctly set the *\$SYSTEM* environment variable.
- If the log file names are in all capital letters, the *\$JDE_BASE* environment variable might be set incorrectly. If you incorrectly set this environment variable, the process will be unable to locate the jde.ini file.
- Verify whether an entry exists in the /etc/hosts table for the floating IP address. If no entry exists, jdenet_n will start, but all other processes will return the following message in the log: 239–gethostbyname returned Connection refused.
- If no entry for the floating IP address exists that the workstation can reference, it will fail to connect and return the following message in the log: 11001–gethostbyname returned 11001 (WSAHOST_NOT_FOUND): The host was not found.

Problem with Workstation Connection to an EnterpriseOne Server; Endnet Works Improperly on the Server

You must associate an IP address with the EnterpriseOne package. The package must be operational for the IP address to be active. Otherwise, workstations will not connect to the server and endnet will not work properly on the server.

EnterpriseOne Does Not Work From the Package Control Script

Oracle must be operational and the owenv must reference the proper SID for EnterpriseOne to work from within the package control script.

Package Does Not Switch to the Backup Node upon Failure or Removal from the Cluster

You must enable automatic switching in the package failover options. If you do not enable this setting, the package will not switch to the backup node when the node fails or you remove the node from the cluster. If you do not want the package to switch, for example, you might want to stop EnterpriseOne, you can disable this flag and then halt the package.

Package Halt Fails

If EnterpriseOne does not end cleanly during a package halt, the package halt might fail. This could occur if Oracle is not operational or if EnterpriseOne cannot access the database. You might need to change the test condition in the package control script, or add commands to search for remaining J.D> Edwards processes and end them.

Placement of the owenv File

Generally, the owenv file should not reside on the shared disk. Different environment settings, particularly ORACLE settings, might exist depending on which node you run a package. If you placed the EnterpriseOne bin32 directory on a shared disk, move the owenv file to another directory.

HACMP for AIX Clustering

This chapter explains how to set up a cluster environment using High Availability Cluster Multi-Processor (HACMP) software. If a failure occurs, HACMP provides a transparent recovery for critical applications. You can configure a cluster using any RS/6000 processor and a variety of network adapters and disk subsystems to satisfy your LAN, disk capacity, and performance requirements.

Shared File Considerations

Be careful when you delete or write to shared files. You might want to move old log files rather than delete them. If you move a package running on a shared file system from one node to another, the new instance of EnterpriseOne references the logs and files from the old instance.

How HACMP Works

HACMP for AIX (Version 4.2) allows customers to automatically detect system failures and recover users, applications, and data on backup systems, minimizing downtime to minutes or seconds. In addition, using HACMP for AIX virtually eliminates planned outages, since users, applications and data can be moved to backup systems during scheduled system maintenance. HACMP Version 4.2 adds new features such as the Cluster Single Point of Control (CSPOC) and Dynamic Reconfig, which allow the system administrator to add users, files, and security functions without stopping mission-critical jobs.

HACMP provides several configuration options, including the following:

- Idle standby for up to seven processors being backed up by a single processor

- Rotating standby for up to seven processors backed up by a standby processor in a predefined or contention takeover sequence
- Mutual takeover for up to eight processors backing each other up by sharing the application workloads
- Concurrent access for up to eight processors working on the same jobs and sharing the same data.

The configuration flexibility of HACMP allows customers to choose the cluster topology and database manager that best suits the requirements of their computing environment. IBM states that HACMP can support both concurrent and parallel data access within a common cluster. HACMP also operates with the new Parallel Database Products, such as IBM's DB2 Parallel Edition and Oracle 8 Parallel Server.

Several components make up the HACMP environment, including the following:

- Nodes
Nodes are the core of an HACMP cluster. A node is a processor that runs the AIX operating system, HACMP, and the mission-critical software. Software execution can be spread over several nodes for system load balancing. In the event of a failover, HACMP executes customer-defined scripts that will establish environments and start specific software packages on a standby node.
- Shared external disk
Shared external disks are disks that are physically connected to multiple nodes. The shared disks store mission-critical data, which is shared among processes running on separate nodes.
- Networks
Networks are the independent components of HACMP. TCP/IP is the protocol with which HACMP was designed to function. It has been tested with ethernet, token ring, and Fiber Distributed Data Interface (FDDI) topology.
- Network adapters
- Clients

Installation Considerations

Installing the HACMP cluster requires that you create logon accounts and use the Oracle Standard Enterprise Database Management System (DBMS).

User Logon Accounts

A major consideration when setting up HACMP and EnterpriseOne is user accounts. When these accounts are created, they are given unique user IDs and unique role IDs. When a node fails over to another node, these unique IDs are matched to names in the `/etc/passwd` and `/etc/role` files. If no matches occur, the unique user and role values are then used as IDs, which can create problems with access and security.

To avoid a problem, before starting the EnterpriseOne installation or configuration, create all user accounts and roles on all nodes that will be used in the cluster environment. Use the same unique

number for all users and the same unique number for all roles. The easiest way to do this is to use the add user/role function found in the HACMP extension of SMIT.

If EnterpriseOne is already installed, use the existing user and role ID numbers to create accounts on the nodes that are defined in the resource group.

Oracle Database

Oracle Standard Enterprise DBMS is used in the HACMP configuration explained in this chapter. The resource control scripts are coded to start and stop the database using standard Oracle program calls. These scripts can be easily modified to allow for changes in database start and stop procedures as well as the introduction of commands for Oracle Parallel Server. To minimize installation and configuration problems, have the database administrator review the commands in the control scripts to ensure that they are correct for your installation.

Creating Group and User Accounts

By performing this operation using the add group/user option in the HACMP component of SMIT, all user accounts and group assignments are synchronized across all nodes. This ensures that when the resource disk volume groups remount on the failover system, the user and group IDs match.

► To create group and user accounts

1. Verify that HACMP is running on all nodes within the resource group.
2. Select a unique ID number that can be assigned to the new group and user that you want to create.
3. Verify your selection by searching all password and group files on the node where the new user and group will be created.
4. Enter the following command on the command line:

```
# smit hacmp
```
5. From the menu, choose Cluster System Management, then Cluster Users & Groups, then Groups, and then Add a Group to the Cluster.
6. Choose the resource group to which you want to add the new group. The resource group identifies the nodes that will need to be updated. Next, add a group called EnterpriseOne and assign it the unique ID number that you chose.
7. Press Enter.
8. From the Cluster Users & Groups panel, choose Users and then Add a User to the Cluster.
9. Choose the same resource group that you chose above for adding a group.
10. Add a user "psft," assign it the pre-selected unique ID number, and choose the EnterpriseOne group.
11. Repeat these steps for the Oracle sign on, creating the group "dba".

Setting Up EnterpriseOne for HACMP

The standard EnterpriseOne Enterprise server software requires minimal modifications to function in a cluster, including editing the Enterprise server JDE.INI file, the owenv script, the start resource control script, and the stop resource control script.

► To edit the owenv script

1. In Windows Explorer, go to the /\$SYSTEM/bin32 directory and open the owenv file.
2. Edit the bold lines shown below:

```
#!/bin/ksh
## set OWHOME to point to the base install path for EnterpriseOne
export OWHOME=/ow2/PeopleSoft/E810
## set ENVIRON to the path code from which you want to run business functions
export ENVIRON=MSTR
## set up the path to your EnterpriseOne system and path code
export SYSTEM=$OWHOME/system
export APPDEV=$OWHOME/$ENVIRON
## set JDE_BASE to the location of your JDE.INI file
export JDE_BASE=$SYSTEM/ini/aix
## set up the Oracle environment
export ORACLE_HOME=/u01/app/oracle/product/8.0.5
export ORACLELIB=$ORACLE_HOME/lib
## the remaining variables point to libraries and executables
export SHLIB_PATH=$SYSTEM/lib:$APPDEV/bin32:$ORACLELIB:$SYSTEM/libv32
export LD_LIBRARY_PATH=$SHLIB_PATH
export PATH=$PATH:$SYSTEM/bin32
```

3. Save and close the file.

► To edit the start resource control script

1. In Windows Explorer, go to the /\$SYSTEM/bin32 directory and open the StartResource.sh file.
2. Edit the bold lines as shown below:

```
# _____
#Global Variables
# _____
#
#export PATH=$PATH
```



```
#
Set environment variables
#
./usr/sbin/cluster/scripts/owenv
#
loop=0
StartUpError=false
ORACLEPROCESS="oracle"
OWStartupDir=$SYSTEM"/bin32"
StartUpLog=$OWStartupDir"/OWStartup.log"
ORACLE_UID=oracle
APP_UID=psft
LOGFILES=$OWStartupDir"/jde*.log"
```

The first bold line executes the script that sets various required EnterpriseOne environment variables.

Note

This script, as well as others, will be relocated into the scripts directory as described in the Control Scripts section under *Creating an Application Server* in the *Server and Workstation Administration Guide*.

ORACLE_UID and APP_UID are the login ID names for Oracle and EnterpriseOne. These are used in the script so that the respective applications are started with the proper application ownership.

Note

This script is delivered with these IDs undefined. If the script is executed, an error message will be generated.

3. Save and close the file.

► To edit the stop resource control script

1. In Windows Explorer, go to the /\$SYSTEM/bin32 directory and open the StartResource.sh file.
2. Edit the bold line as shown below:

```
#
#Global Variables
#
export PATH=$PATH:.
#
# Set environment variables
#
./usr/sbin/cluster/scripts/owenv
LogicalVolumn=/ow2
ShutdownDir=$SYSTEM"/bin32"
StartupLog=$OWStartupDir"/OWStartup.log"
APP_UID=psft
```

Similar to the script modification described in the previous procedure, the APP_ID needs to have the login ID of the application owner. In this example, psft owns the application.

3. Save and close the file.

Creating an Application Server

The application server is a method that invokes predetermined actions of applications. The server is called in the cluster startup or shutdown sequence and executes predefined scripts depending on what activity is occurring. As part of the cluster installation and configuration process, an application server must be created and the access path to the control scripts defined.

Prerequisite

- ❑ The control scripts that are included with EnterpriseOne are located in the /\$SYSTEM/bin32 directory. These scripts are generic and can be modified as needed for your requirements. Before you can create an application server, you must move the scripts to a non-shared disk directory.

► To move the control scripts

Log in as root and enter the following commands:

```
#export SYSTEM=< EnterpriseOne system directory path>
```

Where EnterpriseOne system directory path is the path to your system directory. An example of the path is /ow2/PeopleSoft/810/system.

```
# cd /usr/sbin/cluster
# mkdir scripts
# cd scripts
```

Be sure to include the periods, preceded by a single space, in the following commands:

```
# cp $SYSTEM/bin32/StartResource.sh .
# cp $SYSTEM/bin32/StartResource.sh .
# cp $SYSTEM/bin32/owenv .
# chmod 755 *
```

Repeat this step on all failover nodes and FTP over the modified script files.

► To define an application server

1. From the command line, enter the following command:

```
# smit hacmp
```

2. From the menu, choose Cluster Configuration, then Cluster Resources, then Define Application Servers, and then Add an Application Server.
3. Complete the following fields:
 - Server Name
Enter OneWorldSrv. This adds a label to the resource server that controls the starting and stopping of EnterpriseOne.
 - Start Script
Enter the path of the StartResource.sh script - for example, /usr/sbin/cluster/scripts/StartResource.sh.
 - Stop Script
Enter the path of the StopResource.sh script - for example, /usr/sbin/cluster/scripts/StopResource.sh.
4. Press Enter.

► To define cluster resources

This procedure explains how to let HACMP know that you have defined an application server. This is so that HACMP will know to use the server during a cluster-related event. Within the cluster parameters display is a field in which this resource is defined.

1. From the command line, enter the following:

```
# smit hacmp
```
2. From the menu, choose Cluster Configuration, then Cluster Resources, and then Change/Show Resource for a Resource Group.
3. Choose the resource group - for example, EnterpriseOne.
4. On Configure Resources for a Resource Group, in the Application Server field, enter the name of the application server that you defined in the previous procedure.

Maintaining Multiple Instances of EnterpriseOne

Several considerations exist when you run multiple instances of EnterpriseOne in a clustered environment. Even though each instance might begin on a separate node, a situation might arise for which multiple instances need to run on the same node. When this happens, communication to each EnterpriseOne instance must occur on a different port number or service name, and each instance must use a different range of IPC keys. The following parameters in the JDE.INI file control these settings:

[JDENET]	
ServiceNameListen=	Service Name or Port Number
ServiceNameConnect=	Service Name or Port Number

[JDEIPC]	
StartIPCKeyValue=	Numeric Value

Variable Name	Description
Service Name or Port Number	The service name parameters can use an actual port number or the name of a service that you enter into the /etc/services file.
Numeric Value	The IPC key values should differ by at least 1000 between any two EnterpriseOne instances.

Troubleshooting: AIX Clustering

This section explains how to work through several basic problems that can occur with AIX clustering.

EnterpriseOne Does Not Start

Follow these steps if EnterpriseOne fails to start:

- When you start EnterpriseOne using the control script, first check the control script log for errors.
- Check the log directory for log files. If no log files reside in the directory, verify that the EnterpriseOne processes exist in the proper directory and that you correctly set the \$SYSTEM environment variable.
- If the log file names are in all capital letters, the \$JDE_BASE environment variable might be set incorrectly. If so, the process will not be able to locate the JDE.INI file.

To resolve this, verify whether an entry exists in the /etc/hosts table for the floating IP address. If no entry exists, jdenet_n will start, but all other processes will return the following message in the log:

Error 239 -- gethostbyname returned Connection refused

- If no entry for the floating IP address exists that the workstation can reference, the workstation fails to connect and returns the following message in the log:

```
11001 -- gethostbyname returned 11001 (WSAHOST_NOT_FOUND): The host was not found
```

Note

See the task *To edit the enterprise server JDE.INI file* in the *Server and Workstation Administration Guide*.

Resource Does Not Switch to the Backup Node upon Failure

Verify that the control scripts are in the correct directory and set to be executable. Verify that the application server name is correct in the resource group.

Resource Shutdown Request Fails

Check the OWShutdown script log to determine why the shutdown request failed. This log is located in the `/SYSTEM/bin32` directory.

- If EnterpriseOne does not end cleanly during a failover, the node will not fail over. This failure might occur if Oracle is not operational or if EnterpriseOne cannot access the database. You might need to change the test condition in the resource control script, or add commands to search for leftover EnterpriseOne processes and end them.
- If the script is failing during the unmount of the file system, you may need to add a delay between the fuser command and unmount command. The fuser command should remove any processes accessing the file system. Adding a one-second delay will allow this command to complete before the unmount is attempted.

► To edit the enterprise server JDE.INI file

1. Open the enterprise server's JDE.INI file.
2. Edit the following line in the [CLUSTER] section so it looks like this:

```
[CLUSTER]
```

```
Primary Node=Resource IP Name
```

Where Resource IP Name represents the IP address associated with the EnterpriseOne server. This IP is not a UNIX server IP, but is a separate address given to this resource. This address must be defined on all servers to which this resource might failover.

3. Save and close the JDE.INI file.

Sun Solaris Clustering

This chapter explains how to set up EnterpriseOne to work with Sun's clustering software. This software provides higher availability for your applications because it allows you to recover almost

instantaneously from a power failure or hardware problem. It also allows applications to be available during scheduled downtime.

This documentation assumes that you have successfully installed Oracle and the SUNClustering software. If you are having trouble with either of these products, you should contact Oracle or Sun as needed.

Requirements for Clustering on Sun

You must have a disk that is accessible to all machines in the cluster, and this disk must be large enough to accommodate your EnterpriseOne installation. If you wish to place the database on the same cluster as well, the database file must also be placed on a shared disk accessible to all machines in the cluster (although not necessarily on the same shared disk as the one that the EnterpriseOne server is on).

The SUNClustering 2.2 or greater API is recommended.

Clustering Scripts and How to Modify Them

Four cluster-specific scripts are delivered with EnterpriseOne:

- SunStartResource.sh
- SunStopResource.sh
- owenv
- SunOracleMgr.sh

These scripts can be found in the system/bin32 directory under the base EnterpriseOne installation directory.

SunStartResource.sh

This script runs whenever a node in the cluster starts the EnterpriseOne service. It must be registered with the SUNClustering software and should handle everything that needs to happen when the EnterpriseOne service starts or is switched from one node to another.

See Also

- ❑ *Registering EnterpriseOne with SUNClustering in the Server and Workstation Administration Guide* for more information about SunStartResource.sh

► To modify the SunStartResource.sh script

1. Under Global Variables, there is a call to `/suncldata/peoplesoft/E810_sp1/system/bin32/owenv`. Change `/suncldata/peoplesoft/E810_sp1/` to the same path that OWHOME was set to in the owenv script.
2. Set the APP_UID to the user name that EnterpriseOne is to be run under.
3. Set ORACLE_UID to the user that is to run Oracle if the database is on the same cluster as EnterpriseOne.

4. If you are using Oracle, uncomment the section under "Check for ORACLE running" and "Check to see if ORACLE started/running."

SunStopResource.sh

This script runs whenever a node in the cluster stops the EnterpriseOne service. It is registered with the SUNClustering software, and handles shutting down various processes and any cleanup that needs to happen when the EnterpriseOne service is stopped.

► To modify the SunStopResource.sh script

1. Under Global Variables, change the call to `/suncldata/peoplesoft/E810_sp1/system/bin32/owenv` to the same thing that is in the `SunStartResource.sh` script.
2. Set `APP_UID` and `ORACLE_UID` to the users running EnterpriseOne and Oracle, respectively. These values will be the same as in `SunStartResource.sh`.
3. Under the "Shutdown EnterpriseOne " section, set `LOGDIR` to be the location where all the log files are located.

owenv

This script sets various UNIX environment variables that are needed by EnterpriseOne. The script is called from within the `SunStartResource.sh` and `SunStopResource.sh` scripts.

► To modify the owenv script

1. Set `OWHOME` to be the base directory of EnterpriseOne, for example, `/suncldata/peoplesoft/E810_sp1/`.
2. Set `ENVIRON` to the pathcode that you are using, for example, `PROD` or `CRP`.
3. Once `OWHOME` and `ENVIRON` are set, `SYSTEM`, `APPDEV`, AND `JDE_BASE` should be correct.
4. Set `ORACLE_HOME` to be the location of the ORACLE installation on the machine, for example, `/suncldata/app/oracle/product/8.0.5`.
5. Set `ORACLE_SID` as needed.

SunOracleMgr.sh

This script is necessary if you are running the database on the same cluster as the EnterpriseOne server. If you are not running the database on the same cluster, you can ignore this section.

► To modify the SunOracleMgr.sh script

1. In the Setup Global Variables section of the script, set `ORACLE_UID` to the user ID that is used to start Oracle.
2. Set `ORACLE_HOME` to the appropriate value for your Oracle installation.

Registering EnterpriseOne with SUNClustering

You must register EnterpriseOne with the SUNClustering software.

► To register EnterpriseOne with SUNClustering

1. To register EnterpriseOne with SUNClustering, log in as the root user and enter the following command:

```
/opt/SUNWcluster/bin/hareg -r[service name] -m start_net=[absolute path of SunStartResource.sh] -mstop_net=[absolute path of SunStopResource.sh]
```

Where service name can be anything you want, but you may want to make it something easy to remember like "peoplesoft" since you will need to use that name when modifying EnterpriseOne registry with SUNClustering.

2. Enter the following command with no options:

```
/opt/SUNWcluster/bin/hareg
```

If the line containing the service name you assigned to EnterpriseOne contains "off," then enter the following command:

```
/opt/SUNWcluster/bin/hareg -y [service name]
```

This sets the data service to "on" which allows the data service to switch physical hosts when appropriate.

3. To test whether the cluster switches, enter the following command:

```
/opt/SUNWcluster/bin/haswitch
```

Also, try pulling the plug on the current active host.

Note

You should *not* use "reboot" or "shutdown" on the primary node as a test as doing so will result in an error and not in a switchover.

Windows Clustering

This section provides information about the installation of EnterpriseOne on Microsoft Cluster Services (MSCS).

Note

For more information about how to install and set up a Microsoft cluster, see the appropriate Microsoft documentation.

Prerequisites

- ❑ Partition your disk array to logically divide the software components. Typically, you will set up the following partitions:
 - A partition that contains the cluster software
 - A partition that contains the EnterpriseOne software
 - A partition that contains the database management system (DBMS) software and database if these reside on the cluster system
- ❑ If you will be using Microsoft clustering in conjunction with a DBMS, such as Oracle Fail Safe, SQL Server, or UDB DB2, consult the installation documentation for the appropriate DBMS before installing Windows enterprise server.
- ❑ Configure your network, which includes setting up connections among servers, workstations, and printers.

Note

Microsoft cluster server software only supports TCP/IP. Microsoft cluster server software does not support the use of Dynamic Host Configuration Protocol (DHCP) for the assignment of IP addresses.

Although you need only one network card in each node, you should use two cards to ensure recoverability. One network card will communicate with the public network, and the second card will connect between nodes. This setup allows your cluster to remain active when the primary node loses the network connection. If you use only one network card, when a node loses the network connection that node also loses the connection to other nodes in the cluster.

- ❑ If the database and EnterpriseOne are both running on the cluster, they can be configured to run on separate nodes. To accomplish this, separate EnterpriseOne resources and database resources into different groups. Also, be sure that EnterpriseOne and the database do not share disk resources. EnterpriseOne resources will need to be in the same group as the cluster network name and cluster IP address. This can be the cluster group.

When the EnterpriseOne and database groups are in separate groups, the database group must be online before bringing the EnterpriseOne resources online.

If you do not require EnterpriseOne and database resources to run on separate nodes, place all database and EnterpriseOne resources in the cluster group.

Troubleshooting: Changing an IP Address

After moving the cluster equipment to another area, PeopleSoft had to change the IP address for cluster nodes and any virtual machines connected to them. In order to perform this task, PeopleSoft had to uninstall Microsoft Cluster Services (MSCS), which required the removal of all resources and groups already defined. It also required uninstalling database management system (DBMS) cluster software, which included SQL Server Cluster Manager and Oracle Failsafe. The final steps were uninstalling MSCS, changing the physical IP address in the nodes, reinstalling MSCS using the new IP address for the virtual machine, and reinstalling SQL Server Cluster Manager and Oracle Failsafe.

Troubleshooting: Reinstalling MSCS

If you have already installed a version of Microsoft Cluster Server (MSCS) software and you need to uninstall it, you must use the Add/Remove Programs tool to uninstall MSCS.

Setting Up EnterpriseOne on a Windows Cluster

The following task explains how to set up EnterpriseOne on a Windows cluster. The Windows cluster uses a two-server configuration.

► To set up EnterpriseOne on a Windows cluster

1. Install the supported Microsoft Windows operating system for clustering as listed in the PeopleSoft Minimum Technical Requirements (MTR) web site on the Knowledge Garden.

The nodes can be domain controllers or member servers, but we do not recommend that you run EnterpriseOne or the database on a domain controller.

2. Install Microsoft Cluster Server (MSCS) software on each node using an account that has administrator authority.

Refer to the appropriate Microsoft documentation for specific instructions on the installation of MSCS. Access the cluster to verify that MSCS is installed.

Note

If you already installed a version of Microsoft Cluster software, before you reinstall you must uninstall the software using Add/Remove programs.

3. If you want to manage the cluster remotely, for example, at your desk rather than physically at the cluster, you can install Cluster Administrator on workstations.

Refer to the appropriate Microsoft documentation for specific instructions about installing the Cluster Administrator.

4. Install EnterpriseOne on the cluster disk using the cluster name as the enterprise server name. All nodes will share a single copy of EnterpriseOne.

5. Install the EnterpriseOne network and queue services on each node. Make sure that the EnterpriseOne partition (that is, cluster disk resource) is owned by the node on which you are installing the network and queue services.
6. Go to Settings, Control Panel, Services, and then Startup to assign a generic EnterpriseOne user ID and password for the network and queue services to allow workstations to connect to the EnterpriseOne server.
7. In the server jde.ini file, all server parameters except the database system settings should point to the name of the cluster. Set the database system settings to the name of the database server.

Adding EnterpriseOne Network Resources to the Group

The following task explains how to add EnterpriseOne on a Windows cluster. The Windows cluster uses a two-server configuration.

► To add the EnterpriseOne network resources to the cluster group

1. From the Cluster Administrator main menu, choose New and then Resource from the File menu.
2. On New Resource, complete the following fields for the JDE network resources, and then click Next.
 - Name
 - Description
 - Resource Type
Choose Generic Service from the list.
 - Group Name
Choose the cluster group.
3. Do not turn on the option to Run this resource in a separate Resource Monitor.
4. On Possible Owners, verify that the possible owners match the preferred owners, and then click Next.
5. On Dependencies, choose the EnterpriseOne physical disk resource and the cluster network name, click Add, and then click Next.

You establish this dependency to ensure that the resources are available when the network service starts. Network and queue resources will not function if the drive is unavailable.
6. On Generic Service Parameters, do the following, and then click Next:
 - Type the JDE network service name as it appears in the NT Services applet.
 - Leave the Startup Parameters field blank.
7. On Registry Replication, click Finish.
8. On Cluster Administrator, click OK.

Starting and Stopping the EnterpriseOne Resources on the Cluster

Important

In a cluster environment, use Cluster Administrator to start and stop EnterpriseOne network and queue services. EnterpriseOne resources should be started and stopped in the same order as in the Windows Services applet: start the network service first and then the queue service. When stopping, stop the queue service first and then the network service.

► To start and stop the EnterpriseOne resources on the cluster

1. To start network or queue services or both services, in Cluster Administrator, right-click the EnterpriseOne resource and choose Online.
2. To stop network or queue services or both services, in Cluster Administrator, right-click the EnterpriseOne resource and choose Offline.

Testing EnterpriseOne Connections on the Cluster

Important

EnterpriseOne network and queue services must be set up before you can test connections on the cluster.

See Also

- *Starting the Windows Enterprise Server* in the *EnterpriseOne Installation Guide* for information about testing connections on the cluster

► To test EnterpriseOne connections on the cluster

1. Sign on to Windows.
2. Using Cluster Administrator, verify that all EnterpriseOne resources are online for the EnterpriseOne group.
3. Run porttest to verify that EnterpriseOne is installed correctly on the cluster server. Enter the following commands to run porttest:

```
cd \PeopleSoft\ddp\xxxx\System\Bin32
```

```
porttest userID password environment
```

Where *xxxx* is the specific release of EnterpriseOne (for example, 810), *userID* is a valid EnterpriseOne user, *password* is the password for that user, and *environment* is the environment you are verifying. These parameters are case-sensitive. Enter the porttest command for each environment.

The porttest program initializes a user and an environment, assuming that EnterpriseOne is correctly installed and configured. The program should display messages indicating the selection of records out of an EnterpriseOne table. If it does not, review the jde_####.log file specified in the jde.ini file on the enterprise server.

4. Sign on to an EnterpriseOne workstation, and then run a batch application.
For example, submit the Business Unit Master List - All Companies report (R0006P). Verify that the report processed successfully on the UBE server.
5. Using the Cluster Administration tools, stop EnterpriseOne on the first node as follows:
 - a. Right-click the node name and choose Cluster Service.
All groups will failover to the second node.
6. Verify that all resources are online on the second node, and repeat these steps.

See Also

- ❑ *Working with Server Jobs* in the *System Administration Guide* for more information about verifying batch processing

Setting Advanced Options for EnterpriseOne Resources

The following task explains how to set up advanced options for EnterpriseOne on a Windows cluster. The Windows cluster uses a two-server configuration.

► To set advanced options for EnterpriseOne resources

1. From the Cluster Administrator main menu, open the group that contains the resource you want to modify.
2. Choose a resource, and then choose Properties from the File menu.
3. Click the Advanced tab, and choose one of the following options:

- Do Not Restart
- Restart

If you want to disable failover, choose Do Not Restart.

4. Review the values on the form. For more information about these values, see the appropriate Microsoft clustering documentation.

Setting Failover and Failback Parameters

The following task explains how to set up fail parameters for EnterpriseOne on a Windows cluster. The Windows cluster uses a two-server configuration.

► To set failover and failback parameters

1. From the Cluster Administrator main menu, choose the appropriate group, and then choose Properties from the File menu.

2. Click the Failover tab, and review the following fields:

- Threshold

This field determines the number of failovers attempted by MSCS.

- Period

This field determines the period of time during which MSCS attempts failovers.

3. Click the Failback tab and review the following options:

- Prevent Failback
- Allow Failback

The setting for this option can be immediate or during a set period of time.

For more information about these parameters, see the appropriate Microsoft clustering documentation.

Backing Up EnterpriseOne Tables

A well-planned backup strategy is essential to protect your enterprise's information assets. Rigorously following the backup strategy will provide insurance against data lost by acts of nature, hardware or software failure, or human error. Your backup strategy must balance the level of protection you need against the physical constraints of your system, such as information storage capacity.

PeopleSoft recommends that your backup strategy include the following:

- Perform a full system backup whenever data is at risk, such as when you are installing or upgrading software. In this circumstance, at least back up the database completely.
- Each night, back up changed objects, such as tables and EnterpriseOne objects.
- Each week, back up the deployment server, enterprise servers, and the full database.

Note

You should outline and implement your backup strategy before you begin the Prototype phase of implementation.

Understanding Backup Requirements for Servers

When you perform a backup on a server, you can back up either the entire server or only the changed objects and data. You do not need to perform a complete backup of the server nightly. Only directories that change daily require daily backups.

Deployment Server

EnterpriseOne on the deployment server includes the following items:

- EnterpriseOne directory (all subdirectories and contents)
- jde.ini file on c:\winnt
- Services file on c:\winnt\system32\drivers\etc
- Registry export file
- EnterpriseOne files in the root directory (c:\):
 - jdeapp.ddp
 - jdeapp.xdp
 - jdeauth.dda
 - jdeauth.xda
 - jdemod.ddm

- jdemod.xdm
- jdesec.dds
- jdesec.xds
- jdecode.ddm
- jdecode.xdm

If you modify objects, build new packages, or update the Access database delivered during a workstation installation, create backups of the PD810, DV810, and PY810 directories. If you modify help files, create a backup of the HELPS directory. If your media objects reside on the deployment server, create a backup of the MEDIA OBJ directory.

If important data, such as system data, resides on your deployment server, create nightly backups of the EnterpriseOne data sources (Oracle or SQL Server). For example, if your central objects or Object Management Workbench resides on the deployment server, create a nightly backup.

Enterprise Server

EnterpriseOne on the enterprise server runs on the iSeries, UNIX, or Windows operating systems. You back up key libraries on the iSeries and key files on the UNIX and Windows operating systems.

iSeries

The EnterpriseOne iSeries® libraries comprise the following:

Note

Shut down the database before you create any backups.

- All EnterpriseOne system libraries
 - JDEOW
 - SYS810
 - E810SYS
 - SVM810
- EnterpriseOne data dictionary library
 - DD810
- EnterpriseOne Object Management Workbench library
 - OL810
- All EnterpriseOne production libraries (This example is for pristine and production)
 - PD810
 - PY810
 - PRODDTA
 - PRSTDTA

- All EnterpriseOne business data libraries
 - PRODDTA
 - CRPDTA
 - PRSTDTA
 - TESTDTA
- All EnterpriseOne control libraries
 - PRODCTL
 - CRPCTL
 - TESTCTL
 - PRSTCTL
- All EnterpriseOne versions libraries
 - PD810DNT
 - PY810DNT
 - PS810DNT
 - DV810DNT
- IFS (Integrated File System) libraries
 - PD810
 - PY810
 - PS810
 - TS810
 - DV810
- IBM libraries that require backups
 - QCPA
 - QGPL
- Central objects on the deployment server in Oracle or Microsoft SQL Server database

UNIX

The EnterpriseOne UNIX® system and database files comprise the following:

Note

Shut down the database before you create any backups using Backup Manager. If you export or import using Data Manager, you do not need to shut down the database.

- System files

Create backups of all host files under the PeopleSoft/E810 directory. For example, /u03/PeopleSoft/E810/*.

- Database files

Create backups of all data files that reside in the EnterpriseOne tablespaces.

Use the Oracle Data Manager Tool on the deployment server to make a .dmp file of the desired database, and then back up the .dmp file on tape or hard disk.

Windows

The EnterpriseOne Windows® system and database files comprise the following:

Note

Shut down the database before you create any backups.

- System files

PeopleSoft\ddp\E810 directory.

- Oracle database files

Create backup files for all data files that reside in the EnterpriseOne tablespaces

Use the Oracle Data Manager Tool on the deployment server to make a .dmp file of the desired database, and then back up the .dmp file on tape or hard disk.

- Microsoft SQL Server database files

Create backup files for all tables that reside in the EnterpriseOne databases.

Use the SQL Server Database/Object Transfer tool on the enterprise server to copy the desired tables or database (for example, PSFT810) to a backup database.

Note

PeopleSoft recommends that you use the backup tool provided by the RDBMS vendor.

EnterpriseOne Tables and Object Owner IDs

The following tables list EnterpriseOne tables by type and with the associated object owner IDs.

Note

If any of the control table merges fail or if the specification merge fails, you might need to restore the tables to a pre-merge condition and run the merge again. Follow the restore instructions for your database.

System Tables

F00053

F000531

F000532

F0092

F00921

F00924

F0093

F0094

F00941

F00942

F00945

F00946

F00948

F00950

F00960

F98101

F986101

F98611

F986115

F986116

F98613

F986150

F986151

F986152

F98616

F986161

F986162
F986163
F986164
F986165
F98701
F98800D
F98810D
F9882
F98825
F9883
F9885
F9886
F9887
F9888
F98881
F98882
F98885
F98887
F9889
F98891
F98892
F98980
F98CONST
F98DRENV
F98DRLOG
F98DRPCN
F98DRPUB

F98DRSUB
F98EVDTL
F98EVHDR
F98MOQUE
F98OWSEC
F98TMPL
F98VAR

Object Owner

sys810

Object Management Workbench Tables

F00165
F9860
F9861
F9862
F9863
F9865

Object Owner

obj810

Data Dictionary Tables

F00165
F9200
F9202
F9203
F9207
F9210
F9211

Object Owner

dd810

Server Map Tables

F986101

F98611

F986110

F986111

F986113

F98DRPCN

F98DRLOG

Object Owner

svm810

Control Tables

F0002

F00021

F0004

F0004D

F0005

F0005D

F0082

F00821

F00825

F00826

F0083

F0084

Object Owners

Control Tables - PROD: prodtl

Control Tables - CRP: crpctl

Control Tables - TEST: testctl

Control Tables - PS810: prstctl

Versions Tables

F983051

F98306

Object Owners

Versions - PD810: PD810

Versions - PY810: PY810

Versions - DV810: DV810

Versions - PS810: PS810

Central Objects

F980011

F980021

F983051

F98306

F98710

F98711

F98712

F98713

F98720

F98740

F98741

F98743

F98745

F98750

F98751

F98752

F98753

F98760

F98761

F98762

F98950

Object Owners

Central Objects - PD810: pd810

Central Objects - PY810: py810

Central Objects - DV810: dv810

Central Objects - PS810: PS810

Business Data

Business Data - PROD: proddta

Business Data - CRP: crpdta

Business Data - TEST: testdta

Business Data - PS810: prstdta

Backing Up EnterpriseOne Tables on Servers

Depending on your platform and database type, the procedures you need to perform for backups vary.

Prerequisites

- ❑ If you are using SQL Server or Oracle, verify that you have enough disk space for your backup copy before you begin the backup.
- ❑ If you are using SQL Server, verify that the Select Into/Bulk Copy option on the Options form is turned on for the database into which you will transfer objects. Double-click the database in the tree structure to access the Options form.

► To create a backup for iSeries

1. On a tape drive, back up the following libraries, depending on which path codes you have installed:

Library name	Description
SYS810	System library

SVM810	Server Map
OL810	Object Librarian
DD810	Data Dictionary
COPY810	Central Objects - Prototype
COPS810	Central Objects - PS810
COPD810	Central Objects - PROD
CODV810	Central Objects - DEV
PRODDTA	Production Business Data
PRODCTL	Production Control Tables
CRPDTA	Prototype Business Data
CRPCTL	Prototype Control Tables
TESTDTA	Test Business Data
TESTCTL	Test Control Tables
PRSTDTA	Pristine Business Data
PRSTCTL	Pristine Control Tables
PY810DNT	Versions for CRP
PD810DNT	Versions for PROD
PS810DNT	Versions for PRST
DV810DNT	Versions for DEV
E810SYS	Server system library
JDEOW	PeopleSoft Installation
PY810	Server modules - Prototype
PY810FA	Package Library - Prototype
PS810	Server modules - PS810
PS810FA	Package Library - PS810
PD810	Server modules - PROD

PD810FA	Package Library - PROD
DV810	Server modules - DEV
DV810FA	Package Library - DEV

2. Back up the following IFS structure with the subdirectories:

Library name	Description
PSFT810	Logging directory
E810SYS	Kernel spec and XML
PY810	Spec files for Prototype
PS810	Spec files for PSFT
PD810	Spec files for PROD
DV810	Spec files for DEV
PeopleSoft	Contains the spec files for each path code. \PeopleSoft\PACKAGES\PY810FA\SPEC*.* \PeopleSoft\PACKAGES\PS810FA\SPEC*.* \PeopleSoft\PACKAGES\PD810FA\SPEC*.* \PeopleSoft\PACKAGES\DV810FA\SPEC*.*

► To create a backup for Oracle on UNIX or Windows

1. From the Oracle Enterprise Manager Tool, open Data Manager and choose Export from the Data menu.
2. Type the name for your export utility .dmp file.
Click the Browse button to choose the directory where your .dmp file will reside.
3. Click Next.
4. On the Object Selection form, choose the objects you want to back up, and then click Next.

Note

Objects chosen in the tree on the Data Manager form appear in the Selected Objects form. You can move objects between forms using the arrow buttons or by dragging and dropping.

To export objects, expand the Available Objects tree and choose the item to export. Use the arrows to move objects to and from the Selected Objects form.

5. On the Tuning form, you can choose to generate a log file, if needed.
6. Click Next.

Note

Choose the Generate Log File option and enter a log file name or use Browse to choose a log file.

7. On the Advanced Options form, take the default values or choose the desired options, and click Next.
8. On the Summary form, verify that all of the chosen objects and options are correct.
9. Click Finish to begin exporting objects.

A message window opens that displays information about the progress of the export process.

10. When the export process is completed, you will receive the following message: "Export terminated successfully without warnings."

If errors or warnings exist, check your log file to review the export process.

► To create a backup for SQL Server

1. From SQL Enterprise Manager, choose Database/Object Transfer from the Tools menu.
2. On the Database/Object Transfer form, choose a destination server and database on which to create backup copies of your tables.

Note

The source server and the destination server can be the same, but the database must be different.

3. Keep all default settings and then click the Start Transfer button.
The Database/Object Transfer tool moves your objects.
4. Perform either of the following to verify whether the backup was successful:
 - When the process completes the transfer, click the View Logs button to review the transfer process.
 - Run a SELECT statement to verify that your backup tables transferred to the new database with data.

► To restore a backup file for Oracle on UNIX or Windows

1. From the Oracle Enterprise Manager Tool, open Data Manager and from the Data menu, and choose Import.
2. Type the name of your import utility .dmp file.
3. Click Next.
4. On the Object Selection form, choose the objects you want to restore and click Next.

The Importable Objects tree contains the objects that are importable in the file you specified. To move the object to the Selected Objects tree, choose an object in the tree and click the down arrow.

Note

When the .dmp file is on a remote machine, Data Manager uses the Console job and event system to retrieve the file before displaying the data through the Import Wizard. The Remote Import page of the Import Wizard has a status line at the top of the page that shows the progress of data retrieval. The Oracle Enterprise Manager Console must be running.

Three conditions can be displayed: Job Submitted, Job Started, and Job Completed.

Important

Data retrieval must complete successfully before beginning the import operation.

The Selected Objects/Available Objects tree contains the objects to be imported. To remove an object from the list, choose the object and use the up arrow or drag and drop.

5. Click Next.
 6. On the Associated Objects form, accept the defaults and click Next.
 7. On the Tuning form, you can choose to generate a log file, if needed.
 8. Click Next.
-

Note

Choose the Generate Log File options and enter a log file name or use Browse to choose a log file.

9. On the Advanced Options form, choose the Increment Type. If you followed the instructions to create a backup, choose None for Increment Type and click Next.
 10. On the Summary form, verify that all chosen objects and options are correct.
-

Important

You must drop the existing objects in the database that you want to restore or the import process will fail.

11. Click Finish to begin importing objects.
12. When the import process is completed, you will receive the following message: "Process terminated successfully with no warnings."
If errors or warnings exist, check your log file to review the export process.
13. Perform a SELECT statement to verify that your backup tables are populated with data.

► To restore a backup file for iSeries

Restore the libraries and IFS directories that you backed up from tape.

SQL Server Parameters

The following material on SQL Server parameters is reprinted from *SQL Server Books Online - Transact SQL Reference 6.0* with permission from Microsoft Corporation.

Variable Name	Description
database_name	Specifies the database. If the table being copied is in your default database, this parameter is optional.
owner	The owner's name. This name is optional if you own the table being copied. If no owner is specified and you do not own a table of that name, the program will not execute.
table_name	Specifies which database table to copy.
in out	Specifies the direction of the copy. The in option copies from a file into the database table; the out option copies to a file from the database table.
datafile	The full path of an operating system file when copying a table to or from multiple diskettes; datafile is a drive specifier only (such as C:).
/n	Performs the copy operation using the data's native (database) data types as the default. This option does not prompt for each field because it uses the default values.
/u login_id	The user ID.
/p password	A user-specified password. If the /p option is not used, bcp prompts for a password. If the /p option is used at the end of the command line without any password, bcp uses the default password (NULL).
/s servername	Specifies which SQL Server to connect to. The servername is the name of the server machine on the network. This option is required when you perform bcp from a remote machine on the network.

► To restore a backup file for SQL Server

Verify that the Select Into/Bulk Copy option on the Options form is turned on for the database into which you will transfer objects. Double-click the database in the tree structure to access the Options form.

1. From SQL Enterprise Manager, choose Database/Object Transfer from the Tools menu.
2. On the Database/Object Transfer form, choose a destination server and database from which to transfer backup copies of your tables.

Note

The source server and the destination server can be the same, but the database must be different.

3. Turn off the Transfer All Objects option, but keep all of the other default settings.

4. Click the Choose Objects button, choose the objects that you want to transfer, and then click OK to return to the Database/Object Transfer form.
5. Click the Start Transfer button.

The Database/Object Transfer tool moves your objects.

6. Perform either of the following to verify whether the backup was successful:
 - When the process completes the transfer, click the View Logs button to review the transfer process.
 - Run a SELECT statement to verify that your backup tables transferred to the new database with data.

► **To restore a backup file for SQL Server on Windows**

Verify that the Select Into/Bulk Copy option on the Options form is turned on for the database into which you will transfer objects. Double-click the database in the SQL Enterprise Manager tree structure to access the Options form.

1. Generate scripts for the tables you want to restore and then drop the tables.
2. Use SQL to recreate the scripts for the tables.
3. From the command line, type the following command:
`bcp [[database_name.]owner .] table_name(in|out) datafile /n /u /p /s`
4. Perform a SELECT statement to verify that data populates your backup tables.

SnapShot


The SnapShot application enables you to install multiple versions of EnterpriseOne on a single workstation. Use SnapShot to rename the installation directory, named \E810 by default, and to move files and registry settings to the renamed directory. Once you rename the EnterpriseOne installation directory, you can install a new version of EnterpriseOne without overwriting the previous version. SnapShot also allows you to switch between the versions without requiring you to perform any administrative tasks.

Using SnapShot

You must restore the SnapShot file to a true EnterpriseOne installation before you can run the software from that installation. You cannot begin an instance of EnterpriseOne from the SnapShot file.

The SnapShot process follows:

PeopleSoft EnterpriseOne Snapshot Process



The following tasks refer to saving and restoring an instance from the SnapShot file:

- To save EnterpriseOne in a SnapShot file
- To restore a SnapShot
- To delete a SnapShot
- Running SnapShot from the command line

Prerequisite


- ❑ Before you run SnapShot, copy SnapShot.exe to a directory outside of the \E810 directory. If you try to run SnapShot from inside the \E810 directory with EnterpriseOne open or with an EnterpriseOne directory open, you will receive an error message.

► **To save EnterpriseOne in a SnapShot file**

1. Run SnapShot.exe.

The default directory for SnapShot is \E810\SYSTEM\Bin32.


The EnterpriseOne SnapShot form appears.



On this form, you can review a list of previously created SnapShot files, create a new SnapShot of EnterpriseOne, restore a SnapShot file to a working installation of EnterpriseOne, or delete a SnapShot file.

2. Click Save.

The Make a New SnapShot form appears.



On this form, you can define the name of your SnapShot file and determine the folder name where your SnapShot file will reside.

3. Click OK to complete the SnapShot.

Your new SnapShot file appears in the list box on the EnterpriseOne SnapShot form.

You can now install a new version of EnterpriseOne either by installing a new package on the workstation or restoring another SnapShot file.

► To restore a SnapShot

1. Run SnapShot.

The EnterpriseOne SnapShot form appears. On this form you can review a list of previously created SnapShot files, create a new SnapShot of EnterpriseOne, restore a SnapShot file to a working installation of EnterpriseOne, or delete a SnapShot file.

2. From the list box, choose a SnapShot to restore and then click Restore.

SnapShot restores your EnterpriseOne installation and removes the SnapShot file name from the list box.

Note

If a viable installation of EnterpriseOne exists on the workstation, SnapShot prompts you to save the current installation to a SnapShot directory before you restore the alternative installation.

► To delete a SnapShot

1. Run SnapShot.

The EnterpriseOne SnapShot form appears. On this form you can review a list of previously created SnapShot files, create a new SnapShot of EnterpriseOne, restore a SnapShot file to a working installation of EnterpriseOne, or delete a SnapShot file.

2. From the list box, choose a SnapShot to delete and then click Delete.

SnapShot deletes the directory that stores your SnapShot file and removes the SnapShot file name from the list box.

Running Snapshot from the Command Line

Run SnapShot from the command line for batch processing and to implement shortcuts that automatically save and restore SnapShot files.

► To run SnapShot from the command line

Enter one of the following commands in the command line:

- To save the current installation of EnterpriseOne in a SnapShot file:

SnapShot [/SSnapShot name]

You do not need to specify a SnapShot name when you save an installation. If you do not specify a name when you save EnterpriseOne in a SnapShot file, the current version of EnterpriseOne will be used to define the name for the SnapShot, such as 810.

- To restore a SnapShot file:

SnapShot [/RSnapShot name]

You do not need to specify a SnapShot name when you restore an installation. If you do not specify a name when you restore a SnapShot file, you must specify a target directory in which to locate the SnapShot file.

Note

For SnapShot names that contain embedded spaces, enclose the entire name in quotation marks.

Depending on your needs, you can also use the following switches in the command line:

- **/U**

This switch disables the user interface, including warnings and errors. Use the following syntax:

```
SnapShot [/SSnapShot name] [/U]
```

```
SnapShot [/RSnapShot name] [/U]
```

- **/D**

This switch informs SnapShot not to rename the installation directory. Use this switch when EnterpriseOne resides in uniquely named directories. Use the following syntax:

```
SnapShot [/SSnapShot name] [/D]
```

```
SnapShot [/RSnapShot name] [/D]
```

- **/A**

The Disable Autosave switch turns off the autosave feature in SnapShot. By default, when you restore a previous EnterpriseOne installation, SnapShot automatically saves the current active EnterpriseOne installation. When you turn on the Disable Autosave switch, SnapShot overwrites the current active installation with the restored EnterpriseOne SnapShot file.

```
SnapShot [/RSnapShot name] [/A]
```

- **/T**

This switch designates the target directory where you save your SnapShot file. If you do not specify a SnapShot name in the /R switch, you must specify a target directory. For a target directory with an embedded space, enclose the entire name in quotation marks. Use the following syntax:

```
SnapShot [/R] [/Ttarget dir]
```

Generating Serialized Objects for the EnterpriseOne Web Server

To run the web server, the server must have access to a set of serialized EnterpriseOne JAS objects. These objects can be generated directly from EnterpriseOne objects using the appropriate set of EnterpriseOne specifications.

Generating EnterpriseOne serialized objects requires a specific machine configuration. While it is possible to configure a web server to generate EnterpriseOne serialized objects, the recommended method is to dedicate a separate generation machine for this process. Navigate the following path on the Knowledge Garden for a list of all requirements for the generation machine:

Home - Support - Software Engineering - Platform Technologies -
Minimum Technical Requirements

If you upgraded to EnterpriseOne from a previous release and customized your EnterpriseOne objects, you should first test your custom modifications and then generate serialized JAS objects from the upgraded path code.

Note

PeopleSoft strongly recommends that you generate all objects so that JAS is working with the same specs as your EnterpriseOne Windows clients. Although we ship pre-generated objects for all platforms, the ability to successfully generate objects is key to implementing the web server. Since generating JAS objects is equivalent to building a package for EnterpriseOne clients, you must be able to generate objects to ensure a complete and successful implementation.

Complete the following tasks to install eGenerator and generate EnterpriseOne JAS objects from a set of EnterpriseOne objects.

To generate serialized objects for the Java server, your developers must install and use the eGenerator, a PeopleSoft tool that turns EnterpriseOne specifications into Java code, which allows you to access EnterpriseOne applications in HTML. The EnterpriseOne forms and applications that you generate using the tool are HTML objects. EnterpriseOne stores the objects in a database and retrieves them at runtime.

Note

The eGenerator software can change with each Tools Release. To ensure you have the latest documentation, see the chapter *Generating Serialized Objects for the EnterpriseOne Web Server* in the *Web Server Installation Guide* for your platform from the Update Center.

Prerequisites

- ❑ Ensure that you have installed and are running an EnterpriseOne client that is running at the same Tools Release level as the Web Server. In this task, you will configure the EnterpriseOne client as the generation machine. For instructions about installing an

EnterpriseOne client, see the *EnterpriseOne Installation Guide for Windows NT-Based Systems*.

Installing eGenerator

This topic discusses the tasks you follow to install the eGenerator:

Prerequisites

- ❑ Verify that the paths indicated in the `gen.bat` file on your generation machine are valid for your configuration. If needed, change them to point to the correct local jar files on your generation machine.
- ❑ For Unix and Windows: Complete this task if you are using DB2 UDB on the Enterprise Server. On the Generation machine, open the `db2cli.ini` file and comment out or remove the `LobCachSize` parameters under database aliases associated with EnterpriseOne.

This file is typically located under `DB_HOME\SQLLIB\`.

Setting the Default Storage Parameter (Windows and Unix)

The default storage parameter on your EnterpriseOne database might not allow enough space to transfer all the standard EnterpriseOne Java objects. To avoid this problem, from the following two tasks, complete the one that is relevant to your database.

► To set the default storage parameter (Oracle only)

Log in to the Oracle database and enter the following command:

```
Alter tablespace xx9T default storage (maxextents unlimited);
```

Where "xx" is the environment and "9" is the release of the software to which you are installing the Java objects. For example, use PY9T for the 8.9 Prototype environment.

► To set the default storage parameter (SQL only)

Open the SQL database and set your database Maximum File Size parameter to "Unrestricted filegrowth."

Complete this step for each environment (for example, PY9) to which you are installing the Java objects.

Setting Up the Win32 Client to Generate Serialized Objects

Once you install an EnterpriseOne client, you can use this machine to generate serialized objects.

► To set up the Win32 client to generate serialized objects

1. In Windows Explorer, find the directory that stores the `jde.ini` file (usually `c:\winnt`).
2. In the [Interactive Runtime] section of the `jde.ini` file, add the following parameter and value:

WebAdmin=1

Setting the value for this parameter to 1 allows you to view and use the eight tabs in the eGenerator, thus enabling the generation of all serialized objects for the default user.

3. Click Save.

Note

Any changes to the jde.ini file will be overwritten when you deploy a full package to the Generation machine. To preserve these settings, you can copy the jde.ini file to a backup directory. After you deploy a new package, you can copy the customized settings for the Java generation (listed below) from the [Interactive Runtime] section of the backup file to the updated jde.ini file. Do not overwrite the entire updated jde.ini file with the backup file after a package build. Doing so can delete updated settings that are required by the new package.

Configuring the Generation Machine

The JAS build is automatically installed in the x:\E810\system\generator directory when you install the EnterpriseOne client. The \generator directory is preconfigured and contains most of the files you need to run EnterpriseOne properly. However, you need to complete the following tasks to set up a few of the configuration files.

► To copy the genapp.ini file to the default Windows directory

1. On the generation machine, navigate to the x:\E810\system\generator directory, and locate the genapp.ini file.
2. Copy the genapp.ini file to the default Windows directory on your generation machine.

The default Windows directory for Windows 2000 is x:\WINNT

► To copy the jas.ini file and tnsnames file to the generation machine

1. On the JAS server, navigate to the following directory:

For WebSphere 4.0:

```
/PeopleSoft/JAS/EA_JS_<port number>_<machine name>.ear/webclient.war/WEB-INF
```

For WebSphere 5.0:

```
<WAS_HOME>/installedApps/<machine name>/EA_JS_<port number>_<machine name>.ear/webclient.war/WEB-INF
```

2. Copy the following files to the x:\E810\system\Generator\WEB-INF directory on the generation machine:
 - jas.ini
 - tnsnames.ora (Oracle only)

Note

The above directories are the default locations of the files, but the directory paths might be different on your machines.

(Windows and Unix) The tnsnames file might reside in a different place on the JAS server. To find the location of this file, open the jas.iniJDBj.ini file and note the value (directory path) of the following setting:

tns setting in the [JDBj-Oracle] section (Oracle only)

Preparing EnterpriseOne for Serialized Objects

Before you generate serialized Java objects in EnterpriseOne, you must populate the Data Dictionary and global specs on the Java generation machine.

► To populate the data dictionary

1. On the Java generation machine, log on to a valid environment in EnterpriseOne.
2. In the Fast Path, type BV and press Enter.
3. On Batch Applications - [Work with Batch Versions - Available Versions], type R92TAM in the Batch Application field and click Find.
4. Choose the XJDE001 version and click Select.
5. On Batch Applications - [Version Prompting], click Advanced on the Form menu.
6. On Batch Applications - [Advanced Version Prompting], choose Override Location and click OK.
7. On Batch Applications - [Version Prompting], click Submit.
8. On Batch Applications - [PSFT Data Sources], choose Local and click Submit.

The application starts a UBE that populates databases stored in the x:\E810 directory. This process runs for an hour or more.

9. When the process is complete, copy the following files from the x:\E810 directory to the \spec directory under the path code of the environment you will use to generate the serialized Java objects (for example, x:\E810\PY810\spec):
 - dddict.xdb
 - ddtext.xdb
 - dddict.ddb
 - ddtext.ddb

► To populate the global specs

1. In the Fast Path, type BV and press Enter.

2. On Batch Applications - [Work with Batch Versions - Available Versions], type R98CRTGL in the Batch Applications field and click Find.
3. Choose the XJDE001 version and click Select.
4. On Batch Applications - [Version Prompting], click Advanced on the Form menu.
5. On Batch Applications - [Advanced Version Prompting], choose Override Location and click OK.
6. On Batch Applications - [Version Prompting], click Submit.
7. On Batch Applications - [JDE Data Sources], choose Local and click Select.
8. Click OK to accept the default output.

The application starts a UBE that runs for an hour or more.

Modifying gen.bat

eGenerator includes a bat file (gen.bat) that contains a list of standard configuration settings. You must modify these settings for your environment.

Note

You cannot use old versions of gen.bat (SP20 or lower) with new versions of eGenerator. If you upgrade to a new Tools Release, make sure you use the gen.bat file that is included with the new eGenerator software.

► To modify gen.bat

1. From Windows Explorer, navigate to the x:\E810\system\Generator directory on your generation machine.
2. Right click gen.bat and choose Edit.
3. In the gen.bat configuration files, change the following settings:

Note

For iSeries, ensure that the OWResource.jar file is in the location specified in GEN_CLASSPATH.

Parameter	Setting
set JDK=	Path to the location of the JDK application used by WebSphere. This path is typically x:\WebSphere\AppServer\java.
set JAS INI=	Path to the location of the jas.ini file. (Required for direct generation only.) Typically, this path is x:\E810\system\Generator\WEB-INF.
set DBDRIVERS=	Paths and file names of the database JDBC drivers. (Required for Direct generation only.) Enter the path to the JDBC driver for your particular database: For Oracle: classes12.zip For SQL 2000: msbase.jar, mssqlserver.jar, msutil.jar For DB2/UDB: db2java.zip For DB2 (iSeries): jt400.jar
Safe Mode	Adding the safe mode switch in gen.bat makes parallel generation more stable and less prone to hanging when using IBM JDK1.3.1. When the safe mode is activated, the Generator defaults to Sequential Generation during "Generate All," and inserts one object at a time to the serialized database. Safe Mode reduces the number of concurrent threads from 11 to 1, and enhances the stability, although the "safe mode" is 30 percent slower than regular parallel generation. To activate Safe Mode, add the following bolded text to the command line in the gen.bat file: <pre>%JDK%\bin\java %GENERATOR_OPT% -Dmode=safe -Ddefault_path=%JAS_INI% -classpath %GEN_CLASSPATH% com.jdedwards.runtime.generator.Generator %1</pre>
JASLOGGER=	WEB-INF\lib\log4j.jar Delete the following comment and replace it with the path to the folder containing the jdelog.properties file: <folder containing jdelog.properties>. For example, if the default location for the jdelog.properties file is the WEB-INF directory, then the JASLOGGER setting is:

Parameter	Setting
	WEB-INF\lib\log4j.jar;\WEB-INF

Note

If your path name to the jas.ini file or the database driver contains a space character, you need to wrap the entire path with double quotes. For example, for Windows and WebSphere 5.0, set JAS INI= "<WAS_HOME>\installedApps\<machine name>\EA_JS_<port number>_<machine name>.ear\webclient.war\WEB-INF"

Generating Serialized Objects

You need to generate all serialized objects after you initially promote the JAS build. The eGenerator, once installed on your Win32 client with a JAS build, can generate objects of compiled Java bytecode, and you have access to all EnterpriseOne objects when you run your Java server.

► To generate serialized objects

From your generation machine, log off EnterpriseOne.

1. Under the /JAS directory on your generation machine, run gen.bat.

This batch file calls an initialization file called genapp.ini. You use this file to configure specification applications for serialized object generation.

Note

Ensure that the genapp.ini file is located in the WINNT directory of your generation machine before you run gen.bat.

2. Log on to a Java environment in EnterpriseOne (for example, JPY810).

Note

Verify that the system displays the message "Login to the server successful" in the DOS window before you proceed. This message can take up to 10 minutes to display.

3. On Java & HTML Generator Version 2.0, complete the following field by entering the name of your JAS Server:
 - Web Server Name

Note

If you are running multiple Web server ports on your generation machine, type machine name:port, where "machine name" is the name of your generation machine and "port" is the port number.

4. Click Advanced Settings.
5. On Log Files, type the name and path to the Status Log and Error Log files and click OK.
6. Choose the "Generate all objects" option.
7. Choose the following options and click OK:

- Business Views
- Tables

The system generates all EnterpriseOne business views and tables.

8. Verify that the DOS window indicates the generation process has completed successfully.
9. Choose the "Generate all objects" option.
10. Choose all options except Business Views and Tables, and click OK.

The generation process takes five to ten hours, depending on the speed of your generation machine.

11. Verify that the DOS window indicates the generation process has completed successfully.

Note

If you are upgrading the Java server from a previous release, stop and restart the virtual Web host after the generation process is complete. For details, see the *EnterpriseOne Web Server Installation Guide (Windows-Based Systems)*.

Logging In

Using eGenerator, you can generate EnterpriseOne objects in one of the following ways:

- Generate to the Web Server
- Bypass the Web server and generate directly to a database

Each method of generation has a specific way to log into eGenerator. Select the task that corresponds to the way you want to generate objects.

Prerequisites

- Ensure that the following files and folders reside on your workstation:
 - gen.bat

- OWResource.jar
 - xalan.jar
 - xerces.jar
 - images.jar
 - classes folder
- For direct generation to a database, you also need the following files:
- jas.ini
 - jdbj.ini
 - Database drivers
 - tnsnames.ora (Oracle only)

jas.ini Settings

Verify that the following sections in the jas.ini file are configured correctly. If necessary, change them to match your configuration.

- [LOGS]
Verify that the paths to the log files point to a valid directory.
- [OWWEB]
If you are using an Oracle database, verify that the “OracleTNS=” setting points to the location of the tnsnames.ora file.

JDBj.ini Settings

Verify that the following sections in the JDBj.ini file are configured correctly.

- [JDBj-SPEC DATA SOURCE]
Change the settings to match your configuration.

Note

For more information on these settings, refer to Appendix D: Parameters and Values for the jdbj.ini file in the Web Server Installation Guide.

- [JDBj-ORACLE]
Even if you are not using Oracle as the database, a path to this database must be present in the following setting:

tns=c:\E810\system\Generator\tnsnames.ora

The path and tnsnames.ora file are required for the eGenerator Diagnostic tool to work. If you are not using Oracle, use a text editor to create a blank file called tnsnames.ora and place it in the above path.

Running the eGenerator Diagnostic Tool

eGenerator includes an application that is automatically launched every time eGenerator is started. This application is a diagnostic tool that checks the configuration of eGenerator and reports incorrect settings. The diagnostic tool categorizes errors into two types: fatal and non-fatal. If a fatal error is detected, the application displays an error message and does not allow eGenerator to launch. If a non-fatal error is detected, the application displays an error message but allows you to continue with the Generation process.

Note

The auto diagnostic tool can be suppressed by launching gen.bat with the -nodiag parameter.

When you launch gen.bat, the diagnostic tool automatically performs the following checks on the general parameters, and it displays a message asking if you want to validate the configuration for direct generation. If you press y, the program also diagnoses the direct generation parameters.

The auto diagnostic tool performs the following operations:

For General Use

Operation	Fatal Error
Validates the location of the JDK driver.	X
Reads the classpath and checks that all classpath files are in their stated directories.	X
Verifies that the genapp.ini and jde.ini files are in the default Windows directory (c:\winnt).	
Verifies that jnigen.dll is in the system/bin32 folder on the Generation machine.	X
Checks if the R92TAM and R98CRTGL reports must be run to populate the Data Dictionary and global table specs.	X
Verifies that the jnigen.dll file matches the java part of the Generator.	X
Verifies that the serialized object tables have the correct indices.	

For Direct Generation Only

Operation	Fatal Error
Validates the location of the database drivers.	X
Generates a Web Code Level Object, which is used by JAS at startup to determine if its code matches the generated objects.	X
Checks the location of the jas.ini file and validates the following key parameters in the file:	

Operation	Fatal Error
- The location of the tnsnames.ora and tnsnames.sql files. - The location of the jas logs.	

For JAS Generation only

Operation	Fatal Error
Verifies that the Tools Release of the JAS code on the generation machine matches the Tools Release of the JAS code on the Web Server.	X

Generate to the Web Server

Complete this task if you want to generate objects to the web server. To bypass the web server and generate objects directly to a database, skip this task and complete the next task.

► To generate to the web server

Access your generation machine.

- Under the x:\E810\system\Generator directory on your generation machine, run gen.bat.
This batch file calls an initialization file called genapp.ini. You can use this file to configure specific applications for serialized object generation.

Note

If you do not want to run the diagnostic program, include the parameter -nodiag in the command statement when you run gen.bat. See *Running the eGenerator Diagnostic Tool* for more information about the autodiagnostic tool.

Ensure that the genapp.ini file is located in the WINNT directory on your generation machine.

A sample genapp.ini file is included in the section called *Configuring the genapp.ini File*.

- On EnterpriseOne Sign On, log on to a Java environment in EnterpriseOne (for example, JPY810).

Tip

After you click OK to log on to an EnterpriseOne environment, it can take up to two minutes for EnterpriseOne to complete the logon process.

- On eGenerator, complete the following field:

- JAS Server Name

Type the name of the JAS server that you want to log in to.

The system gets the database path from your jas.ini file. eGenerator generates all objects to this web server.

Note

If you are running multiple web server ports on your generation machine, type *machine name:port*, where *machine name* is the name of your generation machine and *port* is the port number.

4. If you want to generate specs from a remote location, click the Remote Spec Folder option and then type the machine name and path into the field to the right of this option.

For example, `\\machinename\c$\8.9\DV810\SPEC`.

Tip

You can use this field to point to the latest specs on the deployment server and generate objects from the most recent, undeployed specs.

5. Click Connect.

The system connects to the web server that you indicated. A message appears on the status bar when login is complete.

Bypass the Web Server and Generate Directly to a Database

Select direct generation by clicking the Direct Generation check box at the bottom of the Connect to Web Server screen of the eGenerator.

In this mode, the generator initializes a "mini" JAS Server instance on the generation machine. During generation, the generator will bypass the JAS Server and store the objects directly to the database.

Pros:

- Can generate objects before setting up the Java server.
- Faster for all forms of generation.
- Easier to maintain.

Cons:

- Harder to set up initially.

► **To generate objects directly to a database**

1. Launch Gen.bat.

Note

If you do not want to run the diagnostic program, include the parameter `-nodiag` in the command statement when you run `gen.bat`. See the section called *Running the eGenerator Diagnostic Tool* for more information about the autodiagnostic tool.

2. On EnterpriseOne Sign On, log in to a Java environment (for example, JPY9).
3. Click the Direct Generation option.

In Direct Generation Mode, the eGenerator locates the database by looking for the server in the JDBJ-SPEC DATA SOURCE section of the `jdbj.ini` file, which is located on the local disk.

4. Click Connect.

The system bypasses the web server and connects directly to the database specified in the `jdbj.ini` file.

Configuring eGenerator

You can set up a number of eGenerator options for your particular environment. These options are described in the following topics.

Setting eGenerator Options

Before generating objects, you define a variety of parameters that affect how the eGenerator functions.

► **To set eGenerator options**

1. From the pull-down menu, select Options-> Advanced Settings.
2. On Advanced Settings, select the options appropriate for your generation requirements:

- **Logging**

Specify whether to log the generation process by clicking the checkboxes next to the Status and Error logs. You can also enter the location of the log files. Logging is crucial for troubleshooting problems with the generation process.

- **Generation Modes**

Specify if the generation of All Objects will be done in parallel mode or sequential mode by checking the appropriate option.

Note

For more information about Generation Modes, see *Generation Modes* in the *Server and Workstation Administration Guide*.

- **Report Licensing**
Click the checkbox to generate Report Licensing after every report.
- **User Option**
You can generate a public version of the applications or a personal version by choosing the appropriate option. Personal versions are available only to the person who created the versions.
- **Application Entry Point**
Choose this option to generate Application Entry Point information after every application.

Generation Modes

The eGenerator has the following modes for generating web objects:

- Sequential Generation
- Parallel Generation

Sequential Generation

In this mode, web objects are generated one after the other when you select the Generate-All Objects option.

Pros:

- Most stable mode of generation
- Requires the least system resources (for example, memory)
- Safe to use when JITI (Just in Time Install) is occurring

Cons:

- Slowest mode of generation

Parallel Generation

This is the default mode of generation. You can select this mode of generation by clicking Advanced Settings, and then clicking Parallel Generation.

In this mode of generation the web objects are generated concurrently when you select the Generate-All Objects option

Pros:

- Much faster than Sequential Generation, especially on multi-processor workstations

Cons:

- Cannot be used when the user doesn't have the full set of spec files because JITI might occur
- If a JITI occurs during parallel generation, the spec files can be corrupted

Choosing the HTMLOverrides File

Choose the HTML Overrides file to generate HTML objects with advanced features, such as multi-line edit.

► To choose the HTMLOverrides File

1. Launch Gen.bat.
2. On EnterpriseOne Sign On, log in to EnterpriseOne.
3. On eGenerator, complete the following field:

- HTMLOverrides File

Type the path and name of the htmloverrides.ini file, or click Select File and browse for the file in the pop-up window.

When you run eGenerator and generate applications, the generator will process the htmloverrides.ini file.

Choosing Languages

You can generate applications and reports in any language supported by the eGenerator application.

Note

You must install the Language Tools Release for each language you want to use.

► To choose languages

1. On eGenerator, select Options->Languages from the pull-down menu.
2. Click the checkbox next to each language you want to use.
3. Click OK.

Configuring the genapp.ini File

The genapp.ini file contains the list of applications that you want generated in a mode other than the default mode. Mode 1, the default mode, is for Windows. You list applications that you want generated in another mode in the genapp.ini file. Mode 2 is for Java and Mode 3 is for HTML.

The following file is a sample genapp.ini file. eGenerator reads this file. You can use this file to configure specific applications for generation.

```
# genapp.ini   Last Revised 3/15/02
# Application is the name of the application to generate
# Mode is the FDA mode used to create the new "version" of the application.
# Attach a mode to a menu from menu design.  This determines which mode of
# the application is displayed.  The mode has nothing to do with whether or
# not it gets generated in HTML or Java - they are always generated for both.
# An application can only be specified for a mode once.
# Footer is the name of the menu you want displayed at the bottom of the
# application.  Footer menus are optional.

P4015:2:G42314
P4015:3:G42314
P4210:3:G4231
P4210:2:G42314
P4101H:3:G42314
P41202:2:G42314
P41202:3:G42314
P03B2002:2:G42314
P03B2002:3:G42314
P42050:2:G42314
P42050:3:G42314
P4015W:2:G42314
P4015W:3:G42314
P40215:2:G42314
P40215:3:G42314
P41829W:2:G42314
P41829W:3:G42314
P42232:2:G42314
P42232:3:G42314
P4006:2:G42314
```

P4006:3:G43S11
P42045:2:G42314
P42045:3:G42314
P01012:2:G42314
P01012:3:G43S11
P4310:2:
P4310:3:G43S11
P43214:2:
P43214:3:
P4334:2:
P4334:3:G43S11
P0411:2:
P0411:3:G43S11
P34301:2:
P34301:3:G43S11
P40ITM3:2:G42314
P40ITM3:3:G42314
P4960:2:
P01013:2:G42314
P01013:3:G42314
P3460:2:G42314
P3460:3:G42314
P43214:3:G43S11
P41201:3:G43S11
P43230:3:G43S11
P44200:3:G43S11
P430114:3:G43S11
P43100:2:G43S11
P43100:3:G43S11
P0150R:3:G42314
P0150S:3:G42314
P055011:2:G05BESS1
P055011:3:G05BESS1
P4915:3:G43S11
P4960:3:G43S11
P17500:2:G1715

P17500:3:G1715

P17501:2:G1715

P17501:3:G1715

P48201:2:G1715

P48201:3:G1715

P17714:2:G1715

P17714:3:G1715

P1723:2:G1715

P1723:3:G1715

P4947:2:G1705

P4947:3:G1705

P32942:3:

P053020:3:

P986116 and P986110B are being kept in genapp.ini to prevent any
potential issues with previous releases.

P986116:2:

P986116:3:

P986110B:2:

P986110B:3:

#These apps were added for the project DONUT (6/22/01)

P0092SS:3:

P01012SS:3:

P4334SS:3:

P4311S:3:

P4312S:3:

P4310SS:3:

P41204:3:

P04111:3:

P3462:3:

P4210SS:3:

P4947S:3:

P03B2003:3:

P43230SA:3:

```
P42230:3:
P42235:3:
P42240:3:
P34301:3:

#SAR 5336769
P0005S:3:
P0006S:3:
P0101S:3:
```

Generating All Standard Serialized Objects

The eGenerator, in combination with a JAS build, can generate a complete set of Java Serialized Objects. These serialized objects allow you access to all EnterpriseOne objects when you run your web server. If you are installing the web server or upgrading to a new tools release, you will likely want to generate a complete set of serialized objects.

Running the TAM Analyzer

The TAM Analyzer enables eGenerator to generate all objects without saving them to the database, which allows you to quickly test the generation process and detect corrupt specs by checking the log files. To use the TAM Analyzer, you must activate logging for both the Status log and the Error log.

► To run the TAM Analyzer

1. From the pull-down menu, select Options->Advanced Settings.
2. Click the check boxes next to Error Log and Status Log, and click OK.
3. From the pull-down menu, select File -> Tam Analyzer.
The Generator will start generating all the objects.
4. When the generation is complete, check the Status Log and Error Log for corrupt specs.

Generating All EnterpriseOne Objects

Complete the following task to generate a complete set of EnterpriseOne objects.

► To generate all EnterpriseOne objects

1. From the pull-down menu, select File and then Core Objects.

Note

For more information about generating core objects, see the *Web Server Installation Guide*.

2. If the core objects generate successfully, select Generate-> All Objects from the pull-down menu.
3. On Generate All Objects, check all of the objects listed.
4. Click Start.

eGenerator checks the database connection before it attempts to generate objects. If it cannot connect to the database, eGenerator displays an "invalid connection" message.

- a. If an "invalid connection" message displays, click Stop and verify the connection to the database.

If the last Generate-All session did not conclude successfully, the "Resume last failed session" message appears.

- b. Select one of the following options:

- To resume the previous generation process, click RESUME OLD SESSION.

We recommend you choose this option, which restarts the previous process from the point of failure. Once the previous process completes, you can repeat this task to generate the new set of objects.

- To cancel the previous generation process and generate the new objects, click CREATE NEW SESSION.

Select this option if you are currently generating a complete new set of all object types.

In either case, an object set will start generating. The progress of object generation is displayed on the screen. The blue progress bar indicates the progress of the generation process, and the green progress bar indicates the progress of that objects are inserted into the database.

If the database insertion of objects is not complete when the generation is done, the Generator will display the number of objects remaining in the queue. Once all objects are inserted into the database, the program displays a message that the generation completed successfully.

Executable Files on the Workstation

The following tables present an alphabetical list of executable files in the workstation system/bin32 directory. Each table includes a description of the executable file and instructions for running it.

EnterpriseOne Linked Executable Files

The executables in the following table are:

- Called by other EnterpriseOne programs
- Called by the EnterpriseOne kernel
- Have no value if run independently of EnterpriseOne
- Will not run unless called by EnterpriseOne

Executable	Description	Call Details
Ap22.exe	EnterpriseOne uses this program to display spreadsheets in a dialog box. This executable is obsolete and has no function in SP10 and beyond.	Obsolete.
BLC2Text.exe	EnterpriseOne uses this program to read workstation JDEBLC spec files and generate a text file with details about each Business Function source file listed in the spec file.	Called by an internal Business Function build program that is not shipped to customers.
Dir2txt.exe	This program takes a path and a text file name as arguments and places the directory name of the highest branch in the path into the text file.	Called from makefiles generated by BusBuild.
Drilldwn.exe	EnterpriseOne uses this utility when generating Balance Auditor functions in Tabular reports.	Called by the EnterpriseOne UBE kernel.
DSArguments.exe	EnterpriseOne uses this utility program to create a CID argument when attempting to connect to an EnterpriseOne Data Source.	Called by EnterpriseOne kernel.
GBLib.exe	The EnterpriseOne process BusBuild uses this program to determine if object files exist.	Called from makefiles generated by BusBuild.
Guimole.exe	EnterpriseOne uses this program to create a bridge between the workstation and the iSeries server to allow green screens to be displayed through EnterpriseOne.	Called by EnterpriseOne kernel.

Executable	Description	Call Details
InstMon.exe	EnterpriseOne uses this program during log-on to install update packages if an update package is selected.	Called by EnterpriseOne kernel.
JDEGenEx.exe	This program generates a list of exports for each dll.	Called from makefiles generated by BusBuild.
Rtt.exe	PeopleSoft ships this program for use by business partners only. EnterpriseOne uses this program to build resource files for language translation. The risk is that the existing resource files could be confused with the newly generated files. The user would have to intentionally continue through multiple screens for this to happen.	Do not use this program.
Servermon.exe	This obsolete PeopleSoft internal tool was created before the SAW interface for monitoring the server. It sent e-mail alerts to specified users when events occurred on the servers.	Obsolete. Called by an internal script.
Ubemon.exe	This program monitored long running UBEs and reported their completion. This program is obsolete and was disabled in SP10.	Obsolete.
Vdt.exe	This Business View Design Tool creates business views when called from Object Librarian or Object Management Workbench (OMW).	Called from Object Librarian or Object Management Workbench (OMW).
owptrl_cli.exe	owptrl_cli is the communication bridge between BMC Patrol Monitoring Tool and EnterpriseOne Enterprise and Web Servers. owptrl_cli converts SAW data to BMC data so the Agent will understand. The Agent sends a request to the owptrl_cli (for example, give me the list of all processes running on HP9000B port 6012); the owptrl_cli returns detailed information about all the processes running on the Server in a format known by the Agent. In this way, BMC can monitor EnterpriseOne servers.	Called by BMC Patrol Agent with a predefined argument list.
DbidCapture.exe	This utility program is called by Autopilot or EventCapture to capture database IDs necessary for Autopilot to access EnterpriseOne tables. It is not an end-user program and has no purpose apart from	Called by Autopilot or EventCapture.

Executable	Description	Call Details
	Autopilot or EventCapture.	
ubeprint.exe	ubeprint.exe is not for direct customer use, although the EnterpriseOne product suite does use it.	Called by EnterpriseOne kernel.
genver.exe	genver creates the win32 version information for the build process.	Called by EnterpriseOne kernel.
poda.exe	Processing Option Design Aid. All design tools are client side only. OMW passes the executable a set of parameters similar to RDA, TDA, and BDA.	Called from OMW when you design a processing option.
RDA.exe	Report Design Aid. All design tools are client side only. OMW passes the executable a set of parameters similar to PODA, TDA, and BDA.	Called from OMW when you design a batch application. RDA can also be opened without command line parameters.
guimole.exe	A secondary executable called to pass parameters into the WorldVision session.	Called by WorldVision.
FDA.exe	Form Design Aid (FDA) is used to create interactive applications. FDA is currently configured to run on a fat client.	Called from the design window in OMW for an application.
JdeCabExtract.exe	JdeCabExtract creates self-extracting.exe files.	PeopleSoft internal tool.
JdeCompress.exe	JdeCompress creates PeopleSoft-compatible cabinet files.	PeopleSoft internal tool.
krnlspec.exe	Generates jdekrnl.xdb and jdekrnl.ddb specs from the pristine database.	PeopleSoft internal tool.
netmon.exe	netmon is the first (primitive) monitoring tool for EnterpriseOne Enterprise Servers. It sends requests to different Server Kernel processes to verify that they are running. It is obsolete and was replaced by SAW.	Obsolete.
pssg.exe	An obsolete file not called by any EnterpriseOne applications.	Obsolete.
GLBUILD.exe	GLBUILD was replaced by Busbuild.exe. It was used to build the business functions.	Obsolete.
krnlspec.exe	Used to generate jdekrnl.xdb and jdekrnl.ddb specs from the pristine database.	PeopleSoft internal tool.

Executable	Description	Call Details
XFDA.exe	XFDA.exe only exists in EnterpriseOne. It is the previous FDA executable that was restructured and saved for testing purposes. It might be included in some beta releases of EnterpriseOne, but will not be released with the GA version.	Test only. Obsolete.

Standalone Executable Files

You can run the following executable files directly from either the command line or through Windows Explorer.

Executable	Description	Run Instructions
JDECOMConnector2.exe	This program sets up COM connections to the server using the COM Connector product and only works in that context. Contact Customer Support for full documentation.	Run from the command line with a -regserver option.
LogViewer.exe	This program employs a user friendly interface to view and modify plain ASCII EnterpriseOne files such as: • jdedebug.log • jde.log • olt.log • jde.ini	Double-click the executable.
MOConv.exe	This utility converts all records to use a period as the decimal separator. Use this utility when records are entered into a single table using both commas and periods as decimal separators. This utility is driven by the MOConv.ini file.	Exit EnterpriseOne and then double-click the executable.
Nettest.exe	This utility tests basic JDENET connectivity using the “netecho” function against an Enterprise server. Enter the name of the Enterprise server in the Host Name box and press Send. The returned data indicates success or failure.	Double-click. Required argument: • Enterprise Server Name

Executable	Description	Run Instructions
Regdlls.exe	This program adds the following dynamic link libraries (DLLs) to the registry: • jdetapitest.dll • jdetapicomtek.dll Register these DLLs before using RunTAPI.exe.	Run from a DOS window in the EnterpriseOne system/bin32 directory.
RunTAPI.exe	This program controls interoperability between EnterpriseOne and telephone switching systems. It is a snap-in (harness) to ComTech's CTI Server objects. It requires jde.ini file changes and Regdlls.exe before it can be run. Contact Customer Support for full documentation.	Double-click.
SABridge.exe	This Object Export Facility displays the names of the EnterpriseOne objects along with their descriptions and corresponding product codes.	Double-click the executable.
SnapShot.exe	This program manages multiple workstation installations on the same PC. You can install a new instance of EnterpriseOne by clicking Save to store the current workstation installation in a newly named location. Click Restore to toggle between the current and saved versions.	Exit EnterpriseOne and then double-click the executable.
VerifyOCM.exe	This program reads the OCM tables from the database and verifies that the mappings in OCM are correct.	Run from a DOS window. Required arguments: • EnterpriseOne user • EnterpriseOne password • EnterpriseOne Environment
Vercheck.exe	This program displays, on one screen, the properties of all the files in a directory. The properties are the same as those displayed when you right-click a file and choose Properties.	Open a DOS window, change the directory to the desired target, and double-click the executable.

Executable	Description	Run Instructions
GenCOM.exe	This program generates COM wrappers for the business functions specified in the script.	Run GenCom.exe from the command line with the name of the script file.
GenCORBA.exe	Creates CORBA wrappers around EnterpriseOne business functions. This is a command line utility that requires a script file as an input. GenCORBA generates COBRA interfaces for EnterpriseOne business functions.	Syntax: GenCORBA[options] [libraries] For example: GenCORBA /Cat /UserID Devuser1 /Password Denuser1 /Environment ADEVHPO2 CAEC
GenJava.exe	GenJava provides access to EnterpriseOne business functions by generating pure Java interfaces to them.	Run GenJava. Syntax: GenJava [options] [libraries] For example: GenJava /Cat /UserID Devuser1 /Password Denuser1 /Environment ADEVHPO2 CAEC
LaunchUBE.exe	LaunchUBE.exe is used to launch the UBE job stand-alone (not going through EnterpriseOne). It replaces the User Interface of UBEPrint.exe.	Double-click the executable or start via the command line.
oexplore.exe	"It is obsolete" used to start EnterpriseOne FAT client.	Double-click the executable.
Autopilot.exe	Autopilot is the centerpiece of all automated testing tools. Using Autopilot, a person can script EnterpriseOne applications to run automatically and save the scripts to run many times. Autopilot is used throughout the company and by many customers for a wide variety of purposes.	Normally started from a desktop icon or from the Start menu without command line options.

Executable	Description	Run Instructions
EventCapture.exe	EventCapture is a small program that can be activated alongside EnterpriseOne (in lieu of Autopilot) to capture performance and debugging information. EventCapture is often used instead of Autopilot because it is simpler and quicker than creating an Autopilot script for a single use. With EventCapture, the user drives EnterpriseOne; with Autopilot, the script drives EnterpriseOne.	Normally started from a desktop icon or the Start Menu without command line options.

APTestMgr.exe	Autopilot Test Manager is used to run multiple Autopilot scripts in a batch, and to manage batches for repeated execution. It has some ability to summarize results, and it is frequently used for regression testing.	Normally started from a desktop icon or from the Start Menu without command line options.
VSMerge.exe	Visual ER Compare Tool is used to compare and merge Event Rules (ER) for EnterpriseOne Applications, Reports, Table Conversions, NERs (Named Event Rules), and TERs (Table Event Rules). It also can be used to compare and merge C Business functions.	You can launch Visual ER Compare tool from OMW or from the command line.
VSMEditor.exe	VSM Editor is a rarely used GUI tool that creates .VSM files. VSM files are "super scripts" that name one or more virtual Autopilot scripts to be executed in succession by VAPPlayer.	Normally run by double-clicking vsmeditor.exe in bin32.
VirtualRunner.exe	VirtualRunner is a GUI tool for controlling multiple VAPPlayer processes on a single workstation.	Run the tool from a shortcut on the desktop or in the Start Menu. This tool does not use command line arguments.
vapplayer.exe	Virtual Autopilot Player allows you to simulate multiple concurrent EnterpriseOne users on a single workstation. It is used primarily for concurrency testing during development and for performance and scalability testing of EnterpriseOne applications. VAPPlayer requires a proper vap.ini (initialization) file. VAPPlayer has many command line arguments, which are optional if vap.ini is fully utilized. See documentation for details. VAPPlayer has no user interface. It produces output in log files.	VAPPlayer can be run from a command line, from the VirtualRunner graphical user interface, or from the Mercury LoadRunner (third-party) software console.
VAPPerf.exe	VAPPerf is better known as EnterpriseOne Analyzer; it has been renamed Analyzer.exe in EnterpriseOne. It is a powerful instrument used to analyze performance data and other	Create a desktop icon or Start Menu icon. No command line arguments are used.

	debugging information generated by an EnterpriseOne application run under Autopilot or EventCapture.	
UTBrowse.exe	UTB is a tool used for viewing the records in tables. We also use it to view local EnterpriseOne object specifications. UTBrowse.exe uses the following two libraries in the bin32 directory: • datautils.dll • envtool.dll.	Type UTB in the EnterpriseOne fast path or click the executable.
tda.exe	Use TDA is to modify EnterpriseOne tables.	On the Command Line, type tda.exe - <i>idtablename</i>, where <i>tablename</i> is the name of the table you want to modify. For example, F0101. You do not need to run EnterpriseOne before running tda.
tc.exe	tc.exe opens the EnterpriseOne Table Conversion Design Tool. This tool is used to design EnterpriseOne Table Conversion batch applications.	Double-click the executable, or execute it from the command line using the optional parameter idXXXX, where XXXX is the name of an existing Table Conversion object.
Tamvrfy.exe	tamvrfy checks the integrity of all the tam files listed in the tamvrfy.lst.	Double-click the executable.
tamtool.exe	tamtool can perform the following functions: • Recreate a tam file • Copy a tam file • Print index information • Print the index key • Verify the tam file	Run from the command line.
tampack.exe	tampack.exe is a backup utility in case tamftp.exe does not work for the customer. tampack.exe has about half the functionality of tamftp.exe. tampack.exe is included with the workstation and the deployment server. tampack.exe creates a translated copy of	You must run tampack.exe from a DOS shell and pass in parameters.

	TAM files (RDASPEC.DDB, GBRSPEC.DDB, etc.) on the PC. The translated copies are known as pack files. Once the program is finished, the user can manually run ftp.exe to transfer them to a remote enterprise server. Once the pack files are on a remote enterprise server, the user can unpack them on the enterprise server.	
ServerAdministrationWorkbench.exe	Starts the SAW interface.	Double-click the executable.
tamftp.exe	tamftp.exe comes with the workstation and the deployment server. tamftp.exe transfers TAM files (RDASPEC.DDB, GBRSPEC.DDB, and so on) from the PC to a remote enterprise server that is operational.	You must run the program from a DOS shell and pass it parameters.
pdf2pdl.exe	pdf2pdl is an MFC application that converts PDF files into files containing the printer-specific protocol language for a selected printer. This application is only intended to allow development to troubleshoot problems with a customer's EnterpriseOne output. The tool can help solve configuration problems.	Double-click the executable.
pdfcompare.exe	Displays the objects in the PDF document as a list and compares them.	Double-click the executable.

Troubleshooting the Workstation

Use this section as a general guide for basic troubleshooting techniques on the EnterpriseOne workstation. In order to troubleshoot problems, you will need a thorough understanding of the interactive error messages, EnterpriseOne Message Center, logging process, and associated log files.

This section provides solutions to the following problems that you might encounter on the workstation:

Error Messages

- Error message appears at the bottom of a form.
Press F8 or click the Display Errors button to view an error description.

Report Batch Process

- Report shows no data. It shows only the report headers and the text "No Data Selected."
- Batch process shows errors on the report.
- Batch process gives unexpected data on the report.

Environment Issues

- Works when the batch process or business function runs locally but not when it runs on the enterprise server.
- For store-and-forward operation, data entered to the local database is not moved to the server as expected.
- Tables are missing.

Data Source Setup Problems

- Unable to connect to the enterprise server environment.
- Data is displayed incorrectly on the interactive form or batch report.

See Also

- ❑ *Running the Verify Object Configuration Manager* in the *Configurable Network Computing Implementation Guide* for instructions about running the Verify OCM report

Understanding Error Messages

EnterpriseOne provides a variety of error messages. This chapter explains the nature of some of these error messages to help you correct the errors or determine the steps that caused the error.

Troubleshooting: Looking At Error Message Details

When you encounter an error, right-click the error in the error message window and choose Detail to provide additional information about the error. This information provides the source file and the

source line that turned the error on. If you try to set up an Item/Branch record in P41026 with an invalid item number, you will receive error number 0267 (Item Number Invalid).

When indicating the source file that generated the error, the system provides the entire path of the source location. In the above example, the source file is c:\E810\MSTR810\X4101.C, and business function X4101 created the error. The other pieces of the path are directory names. The important information in this example is the file with the .C extension (X4101.C).

If the detail for your error includes the name of your source file, you can identify the process that the file performs to determine what might occur to cause an error. For example, the name of the source file might include system code that indicates the process performed by the file. The process might attempt to run in a module that is not fully functional. The cause of the error might be a constant set to perform a function that is currently unavailable. Once you turn off the constant, you avoid the error.

Note

If you see a source file description that begins with c:\E810\SYSTEM, the error did not occur through a business function. Possibly, the error occurred through an event rule or the tool, while automatically triggering a data dictionary edit.

Look above the source line for conditional statements that determine whether to turn on the error message. Look for table names to see if the program attempts to retrieve data. Look for other programs that the program might call. Also, read the programmer comments included in the source, which might provide a literal explanation for why the code issues an error.

Also look at the data item that caused the error. The data item represents a control on the form. If you get a Blanks Invalid error without an indication of what field you left blank, look at the data item in this error detail box to see which control triggered the error. The field that contains the error might be a hidden field. For example, if you process a transaction that requires a supplier number determined by an Item/Branch combination decided by EnterpriseOne (not by a value you define on the form) but no supplier number exists for the Item/Branch combination, the software returns the Blanks Invalid error. The field for supplier number does not appear on the form, so the cause of the error is not readily apparent. The data item might alert you to the hidden field and help you resolve the error.

Troubleshooting: Error Messages Generated By Applications

These error messages are maintained in the data dictionary and are intentionally set to inform the user of a problem. The error might indicate that the setup is incorrect or that the user is attempting an invalid action. Examples of these kinds of errors include Record Invalid and Blanks Invalid. Some generic errors lack applicable descriptions; techniques for troubleshooting these errors will be discussed.

Troubleshooting: Frequent Generic Error Messages

Some error messages are too generic to immediately explain an error. Examples are Null Pointer and File Can Not Be Accessed. The full descriptions of these error messages do not provide much information as to how to resolve the problem.

To troubleshoot generic errors, retrace the exact steps that led to the original error. The goal is to reproduce the error. If you cannot duplicate an error condition, then the application is accessing

different lines of code than it did when the error occurred. Also look at the information in the error detail box, such as the source file, the source line, and the data item.

Troubleshooting: Memory Violations

Memory violations occur when you encounter memory leaks in an application. A memory leak is a bug that prevents a program from freeing memory that the program no longer needs. The program continues to consume more memory until no memory remains, and the program crashes. EnterpriseOne applications set aside memory while they run. When the application no longer needs that memory, the application should free the memory for other applications to use. When an application does not properly free memory or when an application attempts to use invalid memory, you receive a memory violation.

Use the following techniques to troubleshoot these errors:

- Look at the `jdedebug.log` to find information about the processing that occurred at the time of the error, such as programs called and tables accessed.
- Follow the exact steps that led to the error to reproduce the memory violation. If you cannot duplicate the violation, then the application is accessing different lines of code than it did when the violation occurred. Also look at the information in the error detail box, such as the source file, the source line, and the data item. For UBEs, if the UBE uses a business function that causes memory violations, the UBE will simply stop. In this case, the `ube.log` is the only way to find out what failed.

Troubleshooting: Form and Grid Add Failures

The two error messages that follow indicate that an attempt to add a new record to the database failed. The first message indicates that an add within a fix/inspect form failed. The second message indicates that an add within a grid failed. If you receive these errors, you could be attempting to add a duplicate record.

- Attempt to add form record failed.
- Attempt to add grid record failed.

The `jde.log` is a helpful reference when these errors occur. In general, it includes detailed information about the table into which the user attempted to add a duplicate record.

Troubleshooting: Communication Failure

When submitting batch processes to a server, you might receive an error telling you that a communication failure has occurred.

When you submit a batch job to a server, you are first asked if you would like to install the specifications. If the job is submitted successfully, EnterpriseOne reverts to the initial form.

Troubleshooting the Production Workstation

The troubleshooting procedures that you use for a workstation depend on whether the workstation is a production or development machine. Production machines contain only EnterpriseOne applications, so the scope of problems that can occur is limited. In addition to containing pre-built applications, development machines are equipped with EnterpriseOne and third-party tools. These tools enable developers to create, modify, compile, generate, and troubleshoot EnterpriseOne applications.

Performing Preliminary Troubleshooting

As a system administrator, you can perform preliminary troubleshooting on your production workstation to verify the nature of the problem. You will also want to isolate problems to a user's particular workstation and environment.

► To perform preliminary troubleshooting

1. Determine whether you can consistently duplicate the problem.
2. If you can duplicate the problem, restart the current application.
3. If the error recurs, restart EnterpriseOne.
4. If the error recurs, reboot the workstation.

These steps clear any memory or caching problems with the workstation.

Troubleshooting Interactive Application Problems

In general, when you are running an interactive application, the system displays errors at the bottom of a form. The system highlights the fields with errors in red. You can choose Details on an error message to see information about where the error was set. For example, if the error resulted from within a business function, the system displays the business function and line number where the error was set.

If the errors cannot be resolved through the error messages received in the application, check the error messages in the log files for additional information.

If an application has stopped running, you might need to create or retrieve a new set of specifications for that application. You can overwrite a single application by building a partial package and deploying that package.

See Also

- ❑ *Working With Packages* in the *Package Management Guide* for details about building partial packages

► To troubleshoot interactive application problems

1. Choose one of the following to see the text of the message:
 - Display Errors from the Help menu selection
 - Display icon on the toolbar
 - F8

2. To see the full description of an error message, right-click and choose Full Description.
The system displays a full glossary of the error and includes information for resolving the issue.

Troubleshooting Batch Process Problems

A user might encounter several problems when attempting to run a batch process. For example, the output might show only the report headings or it might print a message such as “No Data Selected.” If the result of a batch process is no data, several factors could be causing the problem.

This section contains the following tasks, which provide solutions to the above problems.

► To troubleshoot batch processes resulting in no data

1. Verify that the data selection on the batch process is appropriate and that data should result.
Data selection on an item that has no data, such as an inactive company, or an incorrect value will result in a batch process with no data.
2. Check the Work Center to see if the batch process resulted in an error.
Most error messages are not printed on the report but are sent as an email message to the user who submitted the report.

These messages will give the user an example of why the batch process ended without producing the desired results. For example, when executing a G/L post that ends in error, the report will print only the report headings. All error messages are sent to the Work Center.

Upon exiting to the Work Center, the user receives error messages and a glossary description that indicate why the batch process resulted in no data. Some error messages include hot links that will link the user directly to the appropriate interactive application to correct the error.

3. If checking the data selection and correcting any errors does not resolve the issue, turn on the applicable logs and continue with the following steps.
 - a. Run your batch process and locate the batch process log for the report you ran.
EnterpriseOne names this log with the following convention:

report name_version_other identifiers.log

This log is located in the local directory under \E810\PrintQueue. If you ran report R04431, it would appear in your local directory, 810\PrintQueue, as follows:

R04431_XJDE0001_D960823_T104512.log

- b. View the log file using the EnterpriseOne Log Viewer or an ASCII editor such as Notepad or Wordpad.
Inspect the log for errors or failures of any kind. Also examine the SQL statements created by the batch process and verify that they should result in data on the report. The batch process log is your main source for debugging batch processes. However, you can look in your jde.log and jdedebug.log for errors or failures of any kind.

- c. Verify that data exists in the tables for the database that you are accessing.
Use the Universal Table Browser tool to view the database table.

See Also

- ❑ *Enterprise Workflow Management Guide* for more information about the Work Center
- ❑ *Universal Table Browser* in the *System Administration Guide* for more information about how to directly view the database table
- ❑ *Troubleshooting: Local Data Availability Problems* and *Troubleshooting: Enterprise Server Data Availability Problems* in the *Server and Workstation Administration Guide* for more information on troubleshooting batch processes

► To troubleshoot batch processes showing errors on the report

Some batch processes will give error messages directly on the report. These messages should include both the short description and error message number. You can view the full description of the error by opening the message in Data Dictionary Design.

From the Data Dictionary Design menu (GH951), choose Error Messages.

Description	Data Item	Language	System Code	Product Code Reporting

1. On Error Messages, complete the following field:
 - Glossary Group
2. Complete the following optional fields:
 - Language
 - Alpha Description

3. On the grid, enter values in the Data Item field and click Find to narrow the search to your particular error code.

For example, enter 0002 to search for the data item associated with the Record Invalid error condition.

4. To see an extended description of the error, choose Glossary from the Row menu.

► To troubleshoot batch processes showing unexpected data on the report

1. Verify that the data selection on the batch process is correct and should result in the data output expected.
2. Turn on the batch process log and run the report.
3. Examine the report process flow and SQL statements to see why the data output on the report is selected.

► To troubleshoot batch processes ending in an error when submitted on the server

The default processing location for batch jobs is the server. If a job gives incorrect results or ends in error when run on the server, the problem could lie with the batch process or with the server.

1. Rerun the report, but override the location to execute on the workstation rather than the server.

You should be aware that if this is a very large report, the processing may take a significant amount of time. You may want to select less material to speed up the processing time.

2. Verify if the outcome is the same as when the report was executed on the server. If so, use the other troubleshooting procedures for batch processing to resolve the issue.

See Also

- ❑ *Troubleshooting the Enterprise Server* in the *Server and Workstation Administration Guide* for more information about reports that do not execute correctly on the server

Troubleshooting: Local Data Availability Problems

Inspect the local database at `\E810\pathcode\data\JDELocal_PD810.mdf` to verify that data exists in the tables the batch application is accessing.

► To troubleshoot local data availability problems

1. To find the calling SQL statement, open the batch process log.
EnterpriseOne names this log using the following convention: report name_version_other identifiers.log. This log is located in your local directory, `\E810\PrintQueue`.
2. Highlight the SQL statement, right-click, and copy the contents to the clipboard.
3. To view data in the local database, open the Universal Table Browser (UTB) and retrieve the table that the batch application is accessing from the local data source.
4. Use the information you copied from the SQL statement to query the table in UTB.

If this action causes the expected records to be found, the data you specified in the data selection matches your SQL statement, which means that data selection is not the cause of the problem.

Troubleshooting: Standalone Installation of EnterpriseOne

If you find you cannot perform a force checkout on a PC running a standalone installation of EnterpriseOne, it is because the software cannot determine your system name for a standalone installation.

The solution is to disable the DNS name in Windows.

Troubleshooting: Enterprise Server Data Availability Problems

If your workstation is running a report against any enterprise server database, such as Oracle, SQL Server, or DB2/400, you need to check the database to see if your SQL statement can find data in the tables. With the help of a database administrator, you can execute the same SQL statement against the server database to verify that the expected data exists in the tables.

As an alternative or in addition to these procedures, you can also use the Universal Table Browser to verify table structure and data availability.

If you do not find any data in the tables for the environment against which you are running, then your SQL statement might be incorrect or the table is empty. Check the data selection and processing options, and verify that they are selecting data that is in the tables. If you do not have data in the tables to match what you are searching, then you will get unexpected results or no data on your report.

For example, if you leave your processing options blank (even though that may be a valid entry for an EnterpriseOne batch process), the process might be searching for blank values or for all values. If your data selection is selecting on a company that does not have any records, then the report batch process does not find any records.

Troubleshooting: Printing Problems

Most printing errors are written to the batch process log. However, some errors might show on reports or be visible in another form. For example, the report prints an error message, prints in the wrong font, or prints landscape instead of portrait.

The following printing problems can occur:

- Your batch application produces error messages on the report, for example, Invalid Company Number.
- Your report batch process shows the wrong font on the report.

Check the report properties of the version you just ran. Also, for the section that is not printing the correct font, check the section properties for the font. If the font is correct, then try printing to a different printer. Otherwise, try using another workstation to see if the font that is being sent to the printer is not interpreted correctly.

- Your report batch process prints portrait instead of landscape or landscape instead of portrait.
Check the report properties of the version you just ran and verify that the properties are correct.

Troubleshooting: .DLL Problems on a Production Workstation

Problems with workstation .DLL files are indicated if you receive an error message such as the following:

CALLBSFN.DLL Load Lib failed

Such a message might indicate that the object does not exist on the workstation. Use a tool such as Explorer to verify if the file exists. You can find consolidated .DLLs in the \E810\path code\bin32 directory.

If the .DLL does not exist on the workstation or if it exists but you continue to get the error even after restarting EnterpriseOne, you can get the correct parent .DLL by reinstalling EnterpriseOne on the workstation from the deployment server. Another option is to copy the parent .DLL from the deployment server package location or another functioning workstation. This option will be successful if the business functions built into the parent .DLL are the same on the workstation that you are copying to as they are on the one you are copying from. Use caution when copying .DLLs. A workstation installation is the preferred method.

Troubleshooting: Data Source Setup Problems

If errors are received when attempting to sign on to an EnterpriseOne environment, a possible cause is an incorrect data source setup on the workstation. Some indications of incorrect setup are as follows:

- A form continues to request a user ID, a password, and a data source even after valid ones are entered.
- Data is displayed incorrectly on an interactive application.
- Messages in the logs refer to problems connecting to data sources or concerning incorrect passwords.

See Also

- ❑ *Running the Verify Object Configuration Manager in the Configurable Network Computing Implementation Guide*

► To troubleshoot data source setup problems

1. From the Control Panel, verify that the ODBC settings are correctly defined and that the data source exists. The proper settings vary by data source.
2. If other users will log on to the same workstation, verify that the data sources are set up as system data sources rather than user data sources.

Data sources set up as user data sources must be set up for each user accessing EnterpriseOne on the workstation.

Troubleshooting the Development Workstation

The troubleshooting procedures that you use on a workstation depend on whether the workstation is a production or development machine. Production machines contain only EnterpriseOne applications, so the scope of the problems that can occur is limited. In addition to containing pre-built EnterpriseOne applications, development machines are equipped with EnterpriseOne and third-party tools. These tools enable developers to create, modify, compile, generate, and troubleshoot EnterpriseOne applications.

You can perform troubleshooting procedures to isolate and resolve a problem with an EnterpriseOne development workstation.

Troubleshooting: .DLL Problems on a Development Workstation

Problems with workstation .DLL files are indicated if you receive an error message such as the following:

CALLBSFN.DLL Load Lib failed

Such a message might indicate that the object does not exist on the workstation. Use a tool such as Explorer to verify if the file exists. You can find consolidated .DLLs in the directory *810\path code\bin32* and *810\system\bin32*.

If the .DLL does not exist on the workstation or if it does exist but you continue to get the error even after restarting EnterpriseOne, the workstation has a problem with the build of one or more consolidated .DLLs. You can rebuild libraries or .DLLs using the BusBuild application from Windows Explorer. The path to busbuild.exe is *810\system\bin32\busbuild.exe*.

► To troubleshoot .DLL problems on a development workstation

Use this procedure if you are receiving the error on a specific business function that cannot be found in the parent .DLL.

1. Verify that the correct parent .DLL for the business function being executed is being referenced when you receive the error.
2. If the wrong parent .DLL is referenced, choose Synchronize JDEBLC from the Tools menu within BusBuild to correctly synchronize the parent .DLLs.
3. Attempt to rebuild the business function from the BusBuild.exe.
The rebuild should include the business function in the parent .DLL.
4. To verify which business functions are part of a parent .DLL, choose Dumpbin from the Tools menu within Busbuild.

This option lists all of the business functions included in the parent .DLL.

Troubleshooting: Event Rule Problems

When encountering problems with event rules on an interactive or batch application, several tools are available to help resolve the problem.

- Review the event rules attached to the application or batch process for obvious problems such as disconnected assignments or incorrect parameters passed to business functions.
- When generating the application, a compile error log is generated, which documents errors in the event rules. Review this log for errors within the Event Rules.
- The Debug Application within EnterpriseOne allows you to debug the event rules for an application or batch process.

See Also

- *Debugging* in the *EnterpriseOne Development Tools Guide* for more information about the debug application

Troubleshooting: Business Function Problems

You might be having business function problems if you are getting unexpected results or getting a .DLL error when you execute a business function.

Microsoft Visual C++ allows you to debug a business function. You can use this tool to step through the logic and inspect variables, which often helps you detect the error.

See Also

- *Debugging* in the *EnterpriseOne Development Tools Guide* for more information about debugging business functions
- *Business Function Problems* in the *Configurable Network Computing Implementation Guide* for more information about debugging business functions

Working with the Workstation Log Files

You should be familiar with the various logs used to troubleshoot problems on the workstation. By using these logs and the procedures outlined in this chapter, you can troubleshoot problems with interactive applications, batch applications, or business functions running locally on the EnterpriseOne workstation. Determine whether you can duplicate the problem consistently or whether it is intermittent.

Do not leave the debugging logs active when the logs are not in use. The logs consume disk and processor resources, and, therefore, affect performance.

Note

If you do not use data replication in your configuration, ignore error messages that refer to the following tables in the JDE.LOG and JDEDEBUG.LOG:

- F98DRPUB
 - F98DRENV
 - F98DRSUB
 - F98DRLOG
 - F98DRPCN
-

Global Tables

Each EnterpriseOne workstation uses global tables (glbltbl.xdb and glbltbl.ddb) to write disk cache files containing internal session-specific and workstation-specific information. For example, information about data dictionary tables and business views is cached. By maintaining a history of this cached session information, individual workstations will improve runtime performance based on their usage.

If you are doing application development, you might need to delete the global tables to see the results of your changes. This is because the system looks first to the disk cache to read certain table information. The information contained in the disk cache might not be synchronized with your current development. You cannot edit the contents of the disk cache.

PeopleSoft recommends that normal startup of EnterpriseOne in a production environment NOT include the deletion of these global tables. These files should only be deleted as a troubleshooting technique or a development aid.

If the global table files do not exist when EnterpriseOne is started, they will be created. If they do exist, they will be appended, not overwritten. The files are located in the 810/*pathcode*/spec directory.

In general, on EnterpriseOne workstations, logs are classified in the following categories:

- Logic processing
- Batch processing
- Application development (compiling and generating)
- Object Management Workbench transactions

Logic Processing Logs

You use two major log file sources for troubleshooting processing faults on the workstation:

- `jde.log`
This log shows fatal errors. Jde.log can track any fault that might occur within EnterpriseOne.
- `jddebug.log` (JDEDEBUG on iSeries)
This log tracks API calls and SQL statements as well as other messages. You can use this file to determine the point in time when normal execution stopped. The system does not use `jddebug.log` to track errors. Instead, this log is used to track the timing of EnterpriseOne processes.

Application Development Logs

For EnterpriseOne workstations in application development environments, you can use the following logs to identify faults in processing related to compiling and generating applications and business functions:

- `compile_error.log`
The `compile_error.log` contains compile errors for event rules. You can use this log to view event rules that might not properly compile and execute. These include Named Event Rules, Table Event Rules, and event rules embedded in applications.
- `jdecpy.log`
This log is produced each time you run the copy table program (`cpytbl.exe`). Copy table error messages and IDs are logged. This log also indicates if any inserts failed that could indicate a possible error.
- `sql.log`
You use this log to view exactly what is being sent through the ODBC driver. This is not an EnterpriseOne log; another software vendor provides this log process.

For workstations in production environments, you can use the following log to identify faults in EnterpriseOne silent installation:

- `jdeinst.log`
If you use the silent installation process for EnterpriseOne installations on a workstation, you can use this log to view the status of the silent installation.

Using Log Files to Troubleshoot Strategies

You can create a normal (successful) jde.log by logging on to EnterpriseOne and then immediately logging off. Use this log of successful startup statements to compare against logs that have a problem.

If you know your problem is not related to startup, you can clear and save the log without exiting EnterpriseOne. When you re-create the problem, the contents of the log should only contain errors that occurred since you cleared the log.

You can also rename the log to indicate the kind of problem. For example, you might delete the jde.log and then run a report that causes an error condition. You could rename the jde.log to report.log.

Another alternative is to add comment lines to the jde.log indicating the sequence of events you are performing. For example, you might be running an application that you know causes an error. Just before you run the application, you could edit the jde.log to add a comment line stating that you are about to start the suspect application.

Most error messages in the jde.log have a unique number assigned to them. You can view an extended description of the error, including possible causes and resolutions, by searching on the error number in the Error Messages application (P92002).

See Also

- ❑ *Understanding Error Messages* in the *Server and Workstation Administration Guide* for more information about the Error Messages application

Viewing Log Files

You can view EnterpriseOne log files from within any application. If you want to view log files outside of EnterpriseOne, you can use a text editor like Notepad or Wordpad.


► To view log files

1. From within any EnterpriseOne application, right-click to open the pop-up menu.
2. On the pop-up menu, choose the View System Log option.
3. On Log Viewer, choose Open from the File menu to locate and open an EnterpriseOne log file.

You can also use the View menu selection to choose log files. If you have viewed log files previously, the File menu keeps a history of those files.

Example: Log Viewer

This example, jde.log, shows errors due to a missing table:



Setting Up the Workstation jde.log

You can use the workstation jde.log as a general purpose log to track fatal errors generated by EnterpriseOne processing. The jde.log tracks any fault that might occur within EnterpriseOne. When you are looking for startup errors, you should read the jde.log from the top down. For other errors, you should read from the bottom up.

The workstation jde.log is created (if it does not exist) or overwritten (if it already exists) at the start of every EnterpriseOne session.

► To set up the workstation jde.log

1. Locate the workstation jde.ini file.

The EnterpriseOne setup program places this file in the working Windows directory; for example, c:\Windows\jde.ini. If you are unsure of your workstation's working Windows directory, use the Find command to locate the jde.ini file.

2. Use an ASCII editor (like Notepad or Wordpad) to open the file.
3. In the [DEBUG] section, verify or change your settings for the job file variable as shown below:

[DEBUG]

JobFile=job file

Where JobFile=job file specifies the location and name of the jde.log file. The default value is c:\jde.log.

Note


You can disable the jde.log on the workstation by specifying a blank or invalid value for JobFile. If you delete or disable (comment out) the JobFile parameter, EnterpriseOne automatically creates and writes to a log file called jde.log in the c:\ directory of the workstation.

4. Save the changes and close the jde.ini file.

Example: Workstation jde.log

This example log shows errors due to a missing table. The highlighted lines indicate the unique error message number that can be used to search for an extended description in the Data Dictionary Error Message application.

Most entries in the jde.log are significant, and you should examine them closely. Developers also use this information to indicate problems with an application that need to be addressed.



Setting Up the Workstation jdedebug.log

The workstation jdedebug.log file contains messages relating to API calls and SQL statements, as well as other messages. You can use this file to determine the point in time when normal execution stopped. The system does not use jdedebug.log to track errors. Instead, it uses this log to track the timing of EnterpriseOne processes.

You can use jdedebug.log to find out where a process has ended. For example, log data can include what the ODBC was trying to connect to, the SQL statement that was being executed for a specific table, and if memory has been freed.

If your process failed and you have logging turned on, look in your jdedebug.log for the following messages:

- Not Found
- Failure

Also, look at the end of the log to see what process executed last. In general, important lines in the log are:

- SELECT

The SELECT lines indicate which table you are selecting. The log tells you in which library (for the iSeries) or environment (for the non- iSeries) the table resides. You should verify that the selected libraries and environments are correct.

- ODBC Version

The ODBC lines indicate whether you are having problems connecting to the driver.

► To set up the workstation jdedebug.log

1. Locate the workstation jde.ini file.

The EnterpriseOne setup program places this file in the working Windows directory; for example, c:\Windows\jde.ini. If you are unsure of your workstation's working Windows directory, use the Find command to locate the jde.ini file.

2. Use an ASCII editor (like Notepad or Wordpad) to open the file.
3. Verify or change the name of the jdedebug.log file. The location and file name of the jdedebug.log file is defined by the following setting in the jde.ini file:

[DEBUG]

DebugFile=location and name

Where DebugFile=location and name specifies the location and name of the jdedebug.log file. The default value is c:\jdedebug.log.

4. Enable or disable the logging of events to the jdedebug.log file by setting the following parameter in the [DEBUG] section:

[DEBUG]

Output=output parameter

Where the valid values for Output=output parameter are:

NONE	No trace information is written to jdedebug.log.
FILE	Database and runtime trace information is written to the file specified by the DebugFile= parameter in the [DEBUG] section.
EXCFILE	Runtime trace information is written to the file specified by the DebugFile= parameter in the [DEBUG] section.
BOTH	Trace information is written to both jde.log and jdedebug.log.

Note

The primary method of disabling the jdedebug.log is by using the Output parameter. However, if you set Output=FILE and you leave the DebugFile value blank (or specify an invalid location), EnterpriseOne still performs debug tracing but does not write the results to any jdedebug.log file. If you delete or disable (comment out) the DebugFile parameter, EnterpriseOne automatically creates and writes to a log file called jdedebug.log in the c:\directory of the workstation.

5. Set the level of debugging information that you want written to the jdedebug.log file. The debug level is determined by the following parameter in the [DEBUG] section:

[DEBUG]

Output=debug level

Where the setting for LEVEL=debug level is described as follows:

This parameter controls the debug level. You can specify any combination of allowable values using commas as delimiters. The default setting is LEVEL=BSFN,EVENTS.

Valid values are:

- EVENTS

Traces the start and stop of events.

- BSFN

Traces when business functions are entered and when they return.

- **SF_x**

Traces when system functions execute. The x variable is any allowable system function value. Valid values are listed below:

GRID

PARENT_CHILD

GENERAL

MESSAGING

WORKFLOW

WORKFLOW_ADMIN

MEDIA_OBJ

CONTROL

For example, LEVEL=SF_CONTROL. In addition, you can specify multiple system functions by separating them with commas. For example, LEVEL=SF_GRID,SF_CONTROL. You can also specify numeric values:

- 1 Traces critical errors. This is the default level. That is, whether you specify this value or not, the system will always trace critical errors.
- 2 Traces critical errors. This is the default level. That is, whether you specify this value or not, the system will always trace critical errors.
- 3 Traces statements as the software enters and exits each event. Specifying this value is the equivalent of specifying the EVENTS value.
- 4 Traces main messages that the software sends to a controlling parent of a child. These messages concern the processing functions such as the grid.
- 5 Provides a detailed trace report of every function that the software calls in the interactive runtime module. This setting is only applicable to developers of the runtime module.

6. Save the changes and close the jde.ini file.

Example: jdedebug.log (Workstation - Output Not Enabled)

This example shows a sample of jdedebug.log when the output is not enabled on the workstation.

4294849873/4294749677 **** jdeDebugInit -- output disabled in INI file

Setting Up the Batch Process Log

You can use the batch process log to identify faults in EnterpriseOne processing related to batch processes. The batch process log resides in the \E810\PrintQueue directory. The log file name is *batch_process.log*, where *batch_process* represents the report name, version name, date, and time.

Based on the setting of the UBESaveLogFile parameter in the [UBE] section of the jde.ini file, this log file is deleted or saved on successful completion of batch processes. This log file shows different types of messages that can help in tracking errors in the batch process. The messages are:

- Section Level Process
- Object Level Process
- ER Level Process
- DB Level Process

The batch process log can contain ER references, batch process flow, and SQL statements, among other messages. You can use the batch process log file to determine when normal execution stopped.

Prerequisite

- ❑ See *Working with the Enterprise Server jdedebug.log File* in the *Server and Workstation Administration Guide* for more information about turning on the debug log file

Batch Processes: General Event Flow

The batch process log file shows you the process flow in batch processes. This flow is completed in the following steps:

- When batch processes complete a section, starting with the INIT section, a business view is opened. After the INIT Section log, you should see a SQL statement.
- After INIT Section, the batch engine calls Adv Section to retrieve a record.
- After the retrieve, batch engine processes the Do Section Processing.
- From Do Section, each object is processed in Init Object - Do Object - End Object order.
- After Do Object message, you should see Printed value in the log. ER events are logged in a different event level.

► To set up the batch process log

1. Locate the workstation jde.ini file.

The EnterpriseOne setup program places this file in the working Windows directory; for example, c:\Windows\jde.ini. If you are unsure of your workstation's working Windows directory, use the Find command to locate the jde.ini file.

2. Use an ASCII editor (such as Notepad or Wordpad) to open the file.
3. Set the level of batch report debugging information that you want written to the batch process log file, and set whether you want the file to be saved. These settings are controlled by the following parameters in the [UBE] section:

[UBE]

`UBEDebugLevel=debug level`

`UBESaveLogFile=0/1`

Where `UBEDebugLevel=debug level` specifies the level of UBE debug logging. The default value is 0, and valid values are:

- 0 No message in a log file
- 1 Error messages, and log entry and section level messages
- 2 Object level messages (plus Level 1 messages)
- 3 ER messages and database mapping messages (plus Level 1 and 2 messages)
- 4 SQL statements (plus Level 1, 2, and 3 messages)
- 5 Batch process function calls and printed output values (plus Level 1, 2, 3, and 4 messages)
- 6 Batch process function calls and printed output values (plus Level 1, 2, 3, 4, and 5 messages)

and `UBESaveLogFile=0/1` specifies whether the `<batch_report>.log` file will be saved. Valid values are:

- 0 The `<batch_report>.log` file is not saved.
- 0 The `<batch_report>.log` file is saved in the workstation's EnterpriseOne print queue directory (810\PrintQueue).

4. Save the changes and close the `jde.ini` file.

Troubleshooting: Working with the Compile Error Log

For EnterpriseOne workstations in development environments, use this log to identify faults in EnterpriseOne processing that are related to compiling and generating applications and business functions. This log for compiled event rules provides an account of event rules (Named Event Rules, Table Event Rules, and applications) that do not properly compile and execute. EnterpriseOne generates this log file every time the Code Generator program (`cg.exe`) is run and errors occur with compiled event rules.

The `<compile_error>` portion of the log file name refers to a variable value for the name of the event rules being compiled. For example, a name of a log file for compiling NER N3200780 is `N3200780.log`. The error log from an application containing compiled event rules replaces the first letter of the application name with an E; for example, P0101 generates an error log named `E0101.log`.

Use this log when errors have occurred within the Code Generator while you were compiling an application, Named Event Rules, or Table Event Rules. When this happens, a message box appears beneath the EnterpriseOne Code Generation form with the source member and the problem description. You can use the log file to keep a record of such problems. The compile error log resides in the log folder under the path code portion of the 810 directory tree; for example, `c:\E810\PD810\LOG`.

Troubleshooting: Working with jdecpy.log

The system produces output for jdecpy.log each time the copy table program (cpytbl.exe) is run on the workstation. In general, the file contains records of those tables that were successfully copied from the local database to the chosen server. This log also indicates if any inserts failed. Such failures indicate a possible error. This log is automatically stopped after cpytbl.exe completes.

The jdecpy.log resides in the root directory of the workstation, usually in c:\. EnterpriseOne automatically generates this log every time you run cpytbl.exe. The log is created or overwritten each time it runs.

Once you use jdecpy.log to determine that a copy table error has occurred, you should refer to the jde.log. If a table does not copy properly, the detail of the error text is written to jde.log. The jde.log contains the actual error message and message ID. The message ID relates to the line prefix numbers in the jdecpy.log. This ID will help you locate the applicable error text written to the jde.log.

The following unsuccessful jdecpy.log example shows a message ID of 116/75 for a copy failed message. The corresponding detail error text, which is written to the jde.log, also contains the same message ID, 116/75, as shown in the excerpt of the jde.log in the example.

Example: Successful Copy of jdecpy.log

A successful copy generates a log similar to the one shown below:

```
91/104   STARTING NEW TABLE COPY
91/104   File F0002
91/104   Creating - F0002
91/104   Created Table
91/104   Copied 20 records for table F0002           from DS Access32 to DS
SQLSERVER JDFDTA71 - from Env PRODLOC to Env PRODLOC.
91/104   0 Inserts failed - check jde.log
91/104   DONE
91/104
91/104   STARTING NEW TABLE COPY
91/104   File F00021
91/104   Creating - F00021
91/104   Created Table
91/104   Copied 2 records for table F00021          from DS Access32 to DS
SQLSERVER JDFDTA71 - from Env PRODLOC to Env PRODLOC.
91/104   0 Inserts failed - check jde.log
91/104   DONE
```

Example: Unsuccessful Copy of jdecpy.log

The following example shows how the copy to the Oracle data source failed:

```
116/75
116/75   STARTING NEW TABLE COPY
116/75   File F0002
116/75   Creating - F0002
116/75   Copy failed
```

By looking in jde.log for this specific error condition (excerpt of the jde.log file shown in the example below), you can further examine the Oracle error. In this case, the associated message ID from the jdecpy.log (116/75) shows that the Oracle error is ORA-01017. The Oracle documentation for error ORA-01017 explains that the table copy failed because the user ID and password were invalid.

```
**** Module: ORACLE    Fri Jan 31 14:05:49 1997        FILE:
      \B7\SYSTEM\JDBDRV\ORACLE\DBINITCN.C LINE: 133
                        116/75      Error - ORA-01017: invalid username/password;
logon denied
```

Troubleshooting: Working with the sql.log

You can use sql.log to view exactly what is being sent through the ODBC driver. This log is not an EnterpriseOne log; another software vendor provides this log process. For workstations, sql.log resides in the default root directory of the workstation, usually in c:\. You can direct the output to any file in any location. In general, instead of using the sql.log, you can use the jdedebug.log because it also tracks SQL statements.

In sql.log, the important lines to search for are:

- **SELECT * FROM**
- SQLBindCol
- Table not found

If you experience a problem with the ODBC settings or can't connect to an EnterpriseOne ODBC database, turn on logging for jde.log, jdedebug.log, and sql.log. Duplicate the problem, check jde.log or jdedebug.log to view the ODBC error messages, and check the end of sql.log to determine the last process. The majority of ODBC problems occur when the following processes are called:

- Process SQL Statements
- Receive Results

Troubleshooting: Reading sql.log

In sql.log, the important lines for you to search are:

- **SELECT * FROM** (bolded in the following example)
- SQLBindCol
- Table not found

Verify that you are reading the correct table. For example, in the sql.log example, a line exists for every column in the selected table, which indicates that the correct table is selected.

If you are having difficulty reading the table, verify that the table has the correct number of columns. If you have added columns to the table and you cannot locate the correct number of columns, you need to configure the table. This information is also provided in jde.log.

► To turn on sql.log

1. From the Windows Control Panel, choose 32bitODBC.
2. On Data Sources, click Options.

Caution

Leave the Stop Tracing Automatically option turned on. Because this log grows rapidly, PeopleSoft recommends you stop the trace in this way to preserve disk space resources and CPU cycles.

Ensure that Trace ODBC Calls is turned off when you are not debugging. The log files can consume large amounts of disk space as well as CPU cycles.

► To troubleshoot ODBC problems using sql.log

1. Ensure the data source names are set up correctly (as system data sources) and that a driver has been set up in the 32bitODBC in Control Panel.
2. Make sure that Client Access has the correct parameters.
3. Ensure that the library to which you are pointing is set up correctly.
4. Look for the following ODBC error messages in jde.log and jdedebug.log:
 - Table not in library
If the table specified couldn't be found in the specified location, go to the appropriate DBMS and attempt to locate the table.
 - If the table does not exist, you must generate the table.
 - If the table exists but has been moved, you must change your data source to point at the new library.
 - Not Binding Column Data Types
This error message means that the row is in use and that another program has a lock on that data. As a result, you cannot use this row until the program that is currently using it releases it.

See Also

- ❑ The applicable chapters in the *EnterpriseOne Installation Guide* for information about Client Access parameters

Working with the jdeinst.log

Use jdeinst.log to view the status of the EnterpriseOne silent installation. The silent installation mode allows you to submit a workstation installation request through command line arguments. EnterpriseOne creates a log file that records error conditions encountered during the silent installation, and it indicates if the silent installation was successful. The jdeinst.log file is located in the root directory of the workstation.

Example: jdeinst.log

This example shows a jdeinst.log indicating a successful EnterpriseOne silent installation completed on a workstation. Comments are included.

<pre>File Warning: Directory d\E810\PROD exists, contents are not erased. Congratulations: Install completed!</pre>

Troubleshooting the Enterprise Server

Using the techniques presented below, you can troubleshoot batch applications and business functions that process on the EnterpriseOne enterprise server. Platform-specific procedures are presented in other sections of this guide.

General Information for Troubleshooting Enterprise Servers

You might encounter the following types of general problems on an EnterpriseOne enterprise server. The information presented applies to all operating systems.

- Communication failure when submitting a UBE or when trying to run business function logic on the server
- Error message appearing at the bottom of a form (press F8 or click Bitmap to view an error description)

You should be familiar with the various logs used to troubleshoot problems on the server. Using these logs, you can troubleshoot batch applications and business functions that are executing on the enterprise server.

Types of Enterprise Server Log Files

In general, logs on EnterpriseOne enterprise servers are classified as either logic processing logs or batch processing logs.

Logic Processing Logs

You can use the following two major log file sources for troubleshooting processing faults on the Enterprise server:

- `jde.log`
This log shows fatal errors. It can track any fault that might occur within EnterpriseOne.
- `jddebug.log`
This log tracks API calls and SQL statements as well as other messages. You can use this file to determine the point in time when normal execution stopped. The system does not use `jddebug.log` to track errors; instead, this log is used to track the timing of EnterpriseOne processes.

Batch Processing Logs

You can use the following log to identify faults in EnterpriseOne processing related to batch processes:

- Batch process log

Use batch process logs to view the process flow in batch processes. This log can contain event rule (ER) references, batch application process flow, and SQL statements, as well as other messages.

You can also view logs using the Work With Servers application.

See Also

- ❑ *Work with Servers Program* in the *EnterpriseOne System Administration Guide* for more information about viewing logs using the Work with Servers application

Viewing Enterprise Server Logs from the Workstation

You must log on to the server to view logs for the server. You can also view portions of log files from the workstation that initiated the calls to the server.

► To view server logs from the workstation

1. In the [DEBUG] section of the Enterprise server jde.ini file, set the ClientLog parameter to 1. This setting enables the server to send logs to workstations. For example:

```
[DEBUG]
```

```
ClientLog=1
```

2. In the [DEBUG] section of the Workstation jde.ini file, set the ServerLog parameter to 1. This setting enables the workstation to receive log information from the Enterprise server. For example:

```
[DEBUG]
```

```
ServerLog=1
```

See Also

- ❑ *Troubleshooting: Enabling and Disabling jde.log* in the *Server and Workstation Administration Guide*

Working with the Enterprise Server jde.log File

You can use the Enterprise server jde.log to track fatal error messages generated by batch applications and business functions that are executing on the Enterprise server. The jde.log tracks any fault that might occur within EnterpriseOne. When you are looking for startup errors, you should read the jde.log from the top down. For other errors, you should read from the bottom up.

If jde.log is enabled, a uniquely identified log file is created each time you start an EnterpriseOne job (including EnterpriseOne startup) on the Enterprise server. These logs are associated with an Enterprise server process ID (Job Number for iSeries).

The process ID (Job Number for iSeries) is appended to the file name, before the .log extension, with an underscore character (for example, jde_442.log).

jde.log File Creation

The Enterprise server jde.log file is created (if it does not exist) or overwritten (if it exists) at the start of every EnterpriseOne session. For a Windows Enterprise server jde.log file, EnterpriseOne appends new information to the end of the jde.log.

Troubleshooting: Enabling and Disabling jde.log

Normally, the Enterprise server should be set to enable the jde.log and disable the jddebug.log. Valid combinations are shown below for the jde.ini parameter setting for enabling or disabling server logs.

Enable jde.log

```
[DEBUG]
Output=NONE
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Enable jde.log and jddebug.log

```
[DEBUG]
Output=FILE
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Disable jde.log

```
[DEBUG]
Output=NONE
JobFile=blank/invalid location/name (1)
DebugFile= blank/invalid location/name (2)
```

Note

Files and members generated by the jde.log will be located in JobFile. EnterpriseOne uses the following naming conventions:

jde_*process_ID*.log

Where jde is the file or member name prefix, *process_ID* is a uniquely named process ID, and log is the file or member suffix or extension.

For non-iSeries Enterprise servers, files generated by the jddebug.log will be located in DebugFile. EnterpriseOne uses the following naming conventions:

jddebug_*process_ID*.log

Where `jddebug` is the file name prefix, `process_ID` is a uniquely named process ID, and `log` is the file suffix or extension.

Caution

Verify whether the paths for the `JobFile` and the `DebugFile` settings are valid. If the paths for these settings are invalid, EnterpriseOne creates no logs at all.

For iSeries Enterprise servers, the members generated by `jddebug` will be located in `DebugFile`. EnterpriseOne uses the following naming conventions:

`jddebug_process_ID`

Where `jddebug` is the file name prefix and `process_ID` is a uniquely named process ID.

Troubleshooting: Recommendations for the Enterprise Server `jde.log`

You can create a normal (successful) `jde.log` by signing on to EnterpriseOne and then immediately signing off. Use this log of successful startup statements to compare against logs that have a problem.

You can also rename the log to indicate the nature of the problem. For example, you might delete the `jde.log` and then run a report that causes an error condition. Then you could rename the `jde.log` to `report.log`.

If you are the only user running an instance of EnterpriseOne, you can add comment lines to the `jde.log` indicating the sequence of events you are performing. For example, you might be running an application that you know causes an error. Before you run the application, you could edit the `jde.log` to add a comment line stating you are about to start the suspect application.

Troubleshooting: Recommendations for Setting Up Server Locations

EnterpriseOne recommends that you create a separate directory on your enterprise server for logs. You should set up the `jde.ini` file to explicitly direct log files to that directory. For `jde.log`, the location and name of the log file are controlled by the following default setting:

<pre>[DEBUG] JobFile=jde.log</pre>

Note

Files generated by the `jde.log` are located in JobFile. EnterpriseOne uses the following syntax for naming files:

`jde_process_ID.log` (`jde_jobnumber.log` for iSeries)

Where `jde` is the file/member name prefix, `process_ID` is a uniquely named process ID, and `log` is the file suffix or extension.

If you do not specify a location, EnterpriseOne places the log files in the directory where you ran the EnterpriseOne startup executable (the default). On a UNIX machine, if you start EnterpriseOne with the following commands and if logging is enabled, the system places the log files in the `/u13/PeopleSoft/E810/system/bin32` directory:

```
cd /u13/PeopleSoft/E810/system/bin32
RunOneWorld.sh
```

If you start EnterpriseOne with the following commands and if logging is enabled, the system places the log files in the `/usr/PeopleSoft` directory because that is the working directory:

```
cd /usr/PeopleSoft
/u13/PeopleSoft/E810/system/bin32/RunOneWorld.sh
```

If you set up your UNIX machine to automatically start EnterpriseOne when the machine is booted, it is especially important that you specify the full path of the log file in the `jde.ini` file.

Naming Conventions for `jde.log`

EnterpriseOne processes create logs as `jde_processID.log` (`jde_JobNumber.log` for iSeries), where *processID* is the process ID of the process that creates the log.

Non-iSeries EnterpriseOne processes move logs for batch jobs to the PrintQueue directory and rename them as *report_version_date_time.log*, where *report* is the report name and *version* is the version name; for example, `R014021_XJDE0001_D990312_T161854215.log`.

► To set up the Enterprise server `jde.log`

1. Locate the Enterprise server `jde.ini` file (JDE member for iSeries).
 - For iSeries Enterprise servers, the JDE member is located under the file `.INI`, which is located within the *releaseSYS* library. The variable *release* is the EnterpriseOne release level (for example, `E810SYS`).
 - For UNIX Enterprise servers, you can locate the `jde.ini` file by examining the UNIX environment variable `JDE_BASE` (the variable name is case-sensitive). For example, you might enter the following command:

```
echo $JDE_BASE
```

In general, this file is located in a directory called "ini" in the directory tree under the base install directory; for example, `/u04/PeopleSoft/E810`. In this case, the directory with

the JDE.INI file (the file name is case-sensitive) is
JDE_BASE=/u04/PeopleSoft/E810/ini.

- For Windows enterprise servers, you can locate the jde.ini file in the EnterpriseOne installation directory under system\bin32.
2. Use a text editor to open the jde.ini file.
 3. In the [DEBUG] section, verify or change your settings for the *job file* variable:

[DEBUG]

JobFile=job file

Where JobFile=job file specifies the location and name of the jde.log file/member. The default value is jde.log.

4. Enable or disable the logging of errors to the jde.log file by setting the following parameter in the [DEBUG] section:

[DEBUG]

LogErrors=0/1

Where LogErrors=0/1 is a parameter that controls whether the logging function (for both jde.log and jddebug.log) is enabled. Valid values are:

- 0 Disabled. Do not write errors to the jde.log file/member or events to the jddebug.log file (JDEDEBUG member for iSeries).
- 1 Enabled. Write errors to the jde.log file and, if enabled, events to the jddebug.log file (JDEDEBUG member for iSeries).

5. Save your changes and close the file.

See Also

- ❑ *Server Administration - iSeries* in the *Server and Workstation Administration Guide* for a complete list of libraries and their contents

Example: Enterprise Server jde.log

The following example of the jde.log from the Enterprise server shows errors caused by signon tables that were not properly closed after fetching data. Normally, the only way this can happen is if a business function program did not close the table. Therefore, generated code applications cannot have this problem.

Most entries in the jde.log file are significant, and you should examine them closely. This information is also used by developers to indicate problems with the application that need to be addressed.

Working with the Enterprise Server jddebug.log File

You can use the Enterprise server jddebug.log to determine the point in time when normal execution stopped. The system does not use jddebug.log to track errors. Instead, it uses this log to track the

timing of EnterpriseOne processes. The log contains API calls and SQL statements as well as other messages.

You can use `jddebug.log` to find out where a process ended. For example, log data can include what the ODBC was trying to connect to, the SQL statement that was being executed for a specific table, and if memory has been freed.

If `jddebug` is enabled, each `jdenet_n` job and batch process started on a server creates a uniquely identified `jddebug.log`. These logs are associated with an Enterprise server process ID. Each time EnterpriseOne is started on the Enterprise server and each time a batch process job is executed on the Enterprise server, a new `jddebug.log` is created.

For Enterprise servers, the process ID (Job Number for iSeries) is appended to the file name with an underscore character before the `.log` extension. For example, the file name might be `jddebug_442.log`. The Enterprise server `jddebug.log` is created (if it doesn't exist) or overwritten (if it exists) at the start of every EnterpriseOne session. For a Windows Enterprise server `jde.log` file, EnterpriseOne appends new information to the end of `jde.log`.

Note

Server administrators are responsible for clearing and deleting `jde.log` and `jddebug_*.log` files from the Enterprise server.

Troubleshooting: Reading the `jddebug.log`

If your process failed and you have logging turned on, look in your `jddebug.log` for the following messages:

- Not Found
- Failure

Also, look at the end of the log to see what task was executed last. In general, important lines in the log are:

- SELECT

The SELECT lines indicate which table you are selecting. The log tells you where the table resides. For the iSeries, this location will be a library. For non- iSeries servers, this location will be an environment. You should verify that the selected libraries and environments are correct.

- ODBC Version

The ODBC lines will indicate whether you are having problems connecting to the driver.

Troubleshooting: Enabling and Disabling `jddebug.log`

Normally, the Enterprise server should be set to enable the `jde.log` and disable the `jddebug.log`. Valid setting combinations for enabling or disabling server `jddebug.log` are shown below.

Enable jdedebug.log

```
[DEBUG]
Output=FILE
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Enable jde.log and jdedebug.log

```
[DEBUG]
Output=BOTH
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Disable jdedebug.log

```
[DEBUG]
Output=NONE
LogErrors=0
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

The [DEBUG] section of the jde.ini file contains the files and members generated by the jde.log. EnterpriseOne uses the following naming conventions:

jde_<pid>.log

Where jde is the file or member name prefix, <pid> is a uniquely named process ID, and log is the file or member suffix or extension.

For Enterprise servers, the files generated by the jdedebug.log will be located in the jde.ini file. EnterpriseOne uses the following naming conventions:

jdedebug_<pid>.log (jdedebug_<JobNumber>.log)

Where jdedebug is the file name prefix, <pid> is a uniquely named process ID, and log is the file suffix or extension.

Troubleshooting: Recommendations for the Enterprise Server jdedebug.log

You can create a normal (successful) jdedebug.log (JDEDEBUG for iSeries) by logging on to EnterpriseOne and then immediately logging off. Use this log of successful start up statements to compare against logs that have a problem.

You can also rename the log to indicate the nature of the problem. For example, you might delete the jdedebug.log and then run a report that causes an error condition. Then you could rename the jdedebug.log to report.log.

Another alternative is to add comment lines to the `jddebug.log` that indicate the sequence of events you are performing. For example, you might be running an application that you know causes an error. Before you run the application, you could edit the `jde.log` to add a comment line stating that you are about to start the suspected application.

Troubleshooting: Recommendations for Setting Up Server Locations

PeopleSoft recommends that you create a separate directory on your Enterprise server for logs. You should set up the `jde.ini` file to explicitly direct log files to that directory. For `jddebug.log`, the following setting controls the location:

```
[DEBUG]
```

```
DebugFile=jddebug.log
```

For Enterprise servers, the files generated by the `jddebug.log` will be located in `DebugFile`. EnterpriseOne uses the following naming conventions:

```
jddebug_ process_ID.log (jddebug_ JobNumber.log for iSeries)
```

Where `jddebug` is the file name prefix, *process_ID* is a uniquely named process ID, and `log` is the file suffix or extension.

By default, EnterpriseOne places the log files in the directory where you ran the startup executable. For example, on a UNIX machine, if you start EnterpriseOne with the following commands:

```
cd /u13/PeopleSoft/E810/system/bin32 RunOneWorld.sh
```

and assuming that logging is enabled, the system places the log files in the `/u13/PeopleSoft/E810/system/bin32` directory. Similarly, on a UNIX machine, if you start EnterpriseOne with the following commands:

```
cd /usr/PeopleSoft /u13/PeopleSoft/E810/system/bin32 RunOneWorld.sh
```

and assuming that logging is enabled, the system places the log files in the `/usr/PeopleSoft` directory. This is the working directory. If you set up your UNIX machine to automatically start EnterpriseOne when the machine is booted, it is especially important that you specify the full path of the log file.

Naming Conventions for `jddebug.log` on the Enterprise Server

EnterpriseOne processes create logs as `jddebug_ process_ID.log`, where *process_ID* (Job Number for iSeries) is the process ID of the process creating the log. For example, a batch report running on a UNIX server as process 123456 would produce a file named `jddebug_123456.log`.

► To set up the Enterprise server `jddebug.log`

1. Locate the Enterprise server `jde.ini` file (JDE member for iSeries).
2. Use a text editor to open the `jde.ini` file.
3. In the `[DEBUG]` section, verify or change your settings for the *debug file* variable:

```
[DEBUG]
```

DebugFile=*debug file*

Where *debug file* specifies the name of the jdedebug.log file (JDEDEBUG member for iSeries). For non-iSeries Enterprise servers, the default value is jdedebug.log. For iSeries Enterprise servers, the default value is JDEDEBUG.

4. Enable or disable the logging of events to the jdedebug.log file (JDEDEBUG member for iSeries) by setting the following parameter in the [DEBUG] section:

[DEBUG]

LogErrors=0/1

Output=output parameter

Where LogErrors=0/1 is the parameter that controls whether the logging function (for both jde.log and jdedebug.log) is enabled. Valid values are:

- 0 Disabled. Do not write errors to the jde.log file/member or events to the jdedebug.log file (JDEDEBUG member for iSeries).
- 1 Enabled. Write errors to the jde.log file and, if enabled, events to the jdedebug.log file (JDEDEBUG member for iSeries).

And where Output=output parameter has the following valid values:

NONE	No trace information is written to jdedebug.log (JDEDEBUG member for iSeries).
FILE	Database and runtime trace information is written to the file/member specified by the DebugFile= parameter in the [DEBUG] section.
EXCFILE	Runtime trace information is written to the file/member specified by the DebugFile= parameter in the [DEBUG] section.
BOTH	Trace information is written to both jde.log and jdedebug.log (JDEDEBUG member for iSeries).

5. Save the changes and close the jde.ini file.

Troubleshooting: Working with the Batch Process Log File

Whenever you run a batch process requested from a workstation, an individual log file is created in the EnterpriseOne print queue directory (810\PrintQueue) on that workstation. For any batch process request issued from a workstation, this file is created even if you have specified that the batch process report is to run on the Enterprise server. For batch processes requested from a server, the jdedebug.log file is created on the server in the print queue directory.

Based on the setting of the UBESaveLogFile parameter in the [UBE] section of the jde.ini file, this log file is deleted or saved on successful completion of batch processes. This log file shows different types of messages that can help track errors in the batch process. The messages are:

- Section Level Process

- Object Level Process
- ER Level Process
- DB Level Process

The batch process log can contain ER references, batch process flow, and SQL statements, among other messages. You can use the batch process log file to determine when normal execution stopped.

The batch process log file shows you the process flow in batch processes. The following example describes the event flow within the batch engine and provides sample messages that would be written to the log at each point in the event flow, assuming UBEDebugLevel is set to 6. Note that each message written to the log file displays the error level of that message in brackets. For example, -UBE--[2]—indicates a section-level message.

When a UBE processes a section, it begins by opening the business view for that section within the INIT section event. As a result, a SELECT statement will follow in the logging directly after the INIT section for each section.

```
--UBE--[2]-- 355/392 Process Init Section
```

```
--UBE--[2]-- 355/392 InitSection for Business Unit Report Driver
```

```
--UBE--[2]-- 355/392 InitSection for Business Unit Report LBH
```

```
--UBE--[4]-- 355/392 SELECT T0.MCMCU, T0.MCSTYL, T0.MCLDM, T0.MCCO,
T0.MCAN8, T0.MCCNTY, T0.MCADD5, T0.MCFMOD, T0.MCDL01, T0.MCDL02,
T0.MCDL03, T0.MCDL04, T0.MCRP01, T0.MCRP02, T0.MCRP03, T0.MCRP04,
T0.MCRP05, T0.MCRP06, T0.MCRP07, T0.MCRP08, T0.MCRP09, T0.MCRP10,
T0.MCRP11, T0.MCRP12, T0.MCRP13, T0.MCRP14, T0.MCRP15, T0.MCRP16,
T0.MCRP17, T0.MCRP18, T0.MCRP19, T0.MCRP20, T0.MCRP21, T0.MCRP22,
T0.MCRP23, T0.MCRP24, T0.MCRP25, T0.MCRP26, T0.MCRP27, T0.MCRP28,
T0.MCRP29, T0.MCRP30, T0.MCPECC, T0.MCALS, T0.MCALCL, T0.MCSBLI,
T1.CCCO, T1.CCNAME, T1.CCRCD FROM F0006 T0,F0010 T1 WHERE (
T1.CCCO=T0.MCCO ) ORDER BY T0.MCCO ASC,T0.MCMCU ASC
```

After INIT Section, the engine calls Advance Section to retrieve a record from the SELECT statement.

```
--UBE--[2]-- 355/392 Process Adv Section
```

```
--UBE--[2]-- 355/392 Processing Adv Section for Page Header
```

After the retrieve, the engine performs the DO Section processing. This includes any event rules attached to the DO Section event.

```
--UBE--[2]-- 355/392 Process DO Section
```

```
--UBE--[2]-- 355/392 Processing DO Section for Page Header
```

```
--UBE--[4]-- 355/392 --ER: Line(1): Loading Data Structure for BSFN
```

```
--UBE--[4]-- 355/392 --ER: Line(1): Processing BSFN : GetCompanyAndReportDesc
```

```
--UBE--[4]-- 355/392 --ER: Line(1): Done Processing BSFN : GetCompanyAndReportDesc
```

--UBE--[4]-- 355/392 --ER: Line(1): Unloading Data Structure for BSFN

--UBE--[4]-- 355/392 --ER: Line(1): Done Processing ER BSFN

Within DO Section, each object is processed and eventually printed in INIT, DO, and END object order.

--UBE--[3]-- 355/392 Process Init Object

--UBE--[3]-- 355/392 Processing Init Item SystemTime in Section Page Header

--UBE--[3]-- 355/392 Process DO Object

--UBE--[3]-- 355/392 Processing Do Object SystemTime in Section Page Header

--UBE--[6]-- 355/392 Printing Object Value = 14:35:46

--UBE--[3]-- 355/392 Process End Object

--UBE--[3]-- 355/392 Process Init Object

--UBE--[3]-- 355/392 Processing Init Item SystemDate in Section Page Header

--UBE--[3]-- 355/392 Process Do Object

--UBE--[3]-- 355/392 Processing Do Object SystemDate in Section Page Header

--UBE--[6]-- 355/392 Printing Object Value = 3/6/00

--UBE--[3]-- 355/392 Process End Object

After all the objects for a section have been processed, the engine calls Process Last Object and then begins processing for the next section in the report.

--UBE--[3]-- 355/392 Processing Do Object ModelAccountsandConsolid in Section Page Header

--UBE--[6]-- 355/392 Printing Object Value = MD

--UBE--[3]-- 355/392 Process End Object

--UBE--[3]-- 355/392 Process Last Object

--UBE--[2]-- 355/392 Process End Page Header Section

--UBE--[2]-- 355/392 Process Do Section

--UBE--[2]-- 355/392 Process Do Section for Business Unit Report Driver

When all sections have been processed, if the report finishes without errors, the following messages are displayed at the end of the log:

--UBE--[6]-- Successfully Finishing Engine

...

UBE Job Finished Successfully.

The level of detail provided by the batch process log is controlled by the UBEDebugLevel parameter of the jde.ini file. Valid values for UBEDebugLevel are shown below:

Setting	Description
0	No error messages
3	Object-level messages
4	Event rule messages and SQL statements (plus levels 1-3)

► **To set up the <batch_process>.log file**

1. Locate the workstation jde.ini file.

The EnterpriseOne setup program places this file in the working Windows directory (for example, c:\WINNT40\jde.ini). If you are unsure of your workstation's working Windows directory, use the Find command to locate the jde.ini file.

2. Use an ASCII editor (such as Notepad or Wordpad) to open the file.
3. Set the level of batch report debugging information that you want written to the batch process log file and whether you want the file to be saved.

These settings are controlled by the following parameters in the [UBE] section:

[UBE]

UBEDebugLevel=debug_level

UBESaveLogFile=0/1

Where UBEDebugLevel=debug_level specifies the level of UBE debug logging. The default value is 0. Valid values are:

- 0 No error messages
- 1 Warnings and high-level information
- 2 Section-level messages (plus Level 1 messages)
- 3 ER messages and database mapping messages (plus Level 1-2 messages)
- 4 SQL statements (plus Level 1-3 messages)
- 5 Database output (plus Level 1-4 messages)
- 6 Batch process function calls and printed output values (plus Level 1-5 messages)

And where `UBESaveLogFile=0/1` specifies whether the `batch_report.log` file will be saved. Valid values are:

- 0 The `batch_report.log` file is not saved.
 - 1 The `batch_report.log` file is saved in the workstation's EnterpriseOne print queue directory (810\PrintQueue).
4. Save the changes and close the `jde.ini` file.

Troubleshooting General Problems

This section provides general troubleshooting information for the Enterprise server.

General Problems

You can troubleshoot general Enterprise server problems using the Server Administration Workbench (SAW), a PeopleSoft executable that enables you to monitor server components, processes, and resources.

► To troubleshoot general problems

1. Use SAW to verify that you are looking at the correct port and the server is operational on that port.

SAW runs only on workstations and Windows Enterprise servers. On other platforms, you can use the tool to view network connections to Enterprise servers.

2. Verify the `netTrace` setting in the Enterprise server's `jde.ini` file:

[JDENET]

`netTrace=0/1` (disabled/enabled)

When the variable `netTrace=0`, EnterpriseOne does not generate Net log information. When `netTrace=1`, EnterpriseOne generates Net log information.

Note

Using SAW for Windows, you can turn logging on or off for a particular kernel process.

3. Return to EnterpriseOne and duplicate the problem.

The trace facilities write debugging information to the `jde.log` and `jdedebug.log` files.

4. After running the business function again, look at the `jde.log` files on the server. Search for the following message (you must search for lower case):

`jdenet_n process`

If you cannot find this message, bring the server down and back up. If you do find this message, look at the jde.log file with the same process ID as the net process.

5. Verify that the user is running in the correct environment or path code; for example PD810 or DV810.

If this environment is not set up on the server, you receive errors on the workstation jde.log as well as the Enterprise server jde.log.

6. In the jde.logs on the Enterprise server, look for a JDENET_SendMSg Failed Error=12 message.

This message means that the JDENET server is down and you must restart it.

7. In the jde.log file on non-iSeries Enterprise servers, look for any “Unable to connect to Oracle” messages. Search on ORA-.

If you find messages, they indicate problems connecting to Oracle. You get an indication of an Oracle connection problem if, in a business function, you select find/browse, data is not found, and no errors are received from the application. You need help from an Oracle database administrator at this point. To debug this problem, see the section in this document about sql.log.

8. Look in the jdexxx.log file (where xxx is the ID of the process that created the log) on the server for the following message:

Could not find symbol in the <BSFN dll name>

If present, this message might mean that the business function did not build on the enterprise server.

9. If you have not found a problem indicating why you are unable to run an application on the Enterprise server, you will need to debug it on the server.

Note

For Windows Enterprise servers, if you cannot identify a problem by reading the log, you need to put the business function through debug on the server. This action requires knowledge of C++ and how to debug. See Microsoft's documentation for Debugging C++.

See Also

- ❑ *Monitoring EnterpriseOne with SAW on Windows in the Server and Workstation Administration Guide* for information about logging on/off for a particular kernel process

Troubleshooting: Communication Problems

When you submit an application to an Enterprise server through an override of the master business function set in Object Configuration Manager, you might experience communication problems with the Enterprise server. The business function then runs locally on the client workstation. EnterpriseOne displays a window to inform you that the business function is running in a new location.

► To troubleshoot communication problems

Note

Use this procedure if EnterpriseOne displays a window to inform you that a business function is running in a new location.

1. Check your `jde.ini` on the workstation to make sure the JDENET service name (port number) is correct and valid.

This port number must match the settings in the server `jde.ini` file, and the EnterpriseOne server must be running to successfully submit reports or to run business logic on a server. Security services and transaction management services also require the EnterpriseOne server to be running.

[JDENET]

`serviceNameListen=service name`

`serviceNameConnect=service name`

Where `serviceNameListen=service name` specifies the communications service port on the TCP/IP network. EnterpriseOne uses this port address to listen for requests on the network. Using a file called "services," you can associate the port number with a unique name. The default value is `jde_server` (port number 6003).

Where `serviceNameConnect=service name` specifies the communications service port on the TCP/IP network. EnterpriseOne uses this port address to connect to the network. Using a file called "services," you can associate the port number with a unique name. The default value is `jde_server` (port number 6003).

2. On your workstation, exit EnterpriseOne and turn logging on in the `jde.ini`.
 - a. Run the application on the server again, and then check the `jde.log` file to see if any of the following errors are logged:

JDENET_SendMsg Failed Error=8

This error can mean you are not using the correct TCP/IP service port or that the Enterprise server does not have that JDENET listing.

JDENET_SendMsgFailed Error=5, 11, or 12

These errors can mean that the message is being sent to the correct port, but the Enterprise server's JDENET is down.

3. From within SAW, change the port address to determine if both the workstation and server are using the same port.

Typically, network administrators run SAW using a workstation's EnterpriseOne software, although you can run a standalone version on Windows server platforms.

4. Check the services file on your workstation (located in your *operating system directory*\System32\drivers\etc for Windows).

Ensure that a blank line exists at the end of the file and that you have the service name mentioned in Step 1 (for example, `jde_server`) going to the correct port address on the server. Verify the port address with the server administrator.

5. If you receive a Communication Failure message, try resubmitting the application. A time-out may have occurred.

[JDENET]

`netTrace=0/1` (*disabled/enabled*)

6. Look in the log file for the following message:

Could not find symbol in the *<BSFN dll name>*

Troubleshooting: Deadlock Collision Problems

When a business function performs a manual commit inside a CallObject process, no other business function can access the tables that the original manual commit locks. Deadlock collisions occur when other business functions attempt to access a table or other resource locked by the original manual commit.

To handle the deadlock collision, the CallObject process puts the processing request in a queue. The number of deadlock collisions grows during the life of a CallObject kernel process and provides evidence of how many times a request has been placed in a queue because the needed resources were locked.

The longer the time a CallObject process requires to run a request, the greater the number of deadlock collisions that are likely to occur. If the number of deadlock collisions grows quickly, the likelihood is that too many users are attempting to connect to the same CallObject process. This, in turn, might be an indication that you do not have enough CallObject processes running on the Enterprise server.

You can monitor the number of deadlock collisions for each process running on the Enterprise server using the Server Administration Workbench (SAW). If you observe the value for the Lock Conditions parameter in the SAW form increasing rapidly, you can increase the number of CallObject kernel processes on the server and monitor the effect of that action.

See Also

- ❑ *Monitoring Process Statistics for Enterprise Servers* in the *Server and Workstation Administration Guide* for details on using SAW to monitor processes running on the enterprise server

Troubleshooting: Server Map Problems

If you change the Object Configuration Manager or the Data Source Master files in the Server Map data source, you can test the changes using the PeopleSoft program PORTTEST. This test is designed to validate your environments.

See the section specific to your platform type for more information about PORTTEST.

Troubleshooting the iSeries Enterprise Server

This chapter explains how to troubleshoot problems that can occur on an iSeries Enterprise server. When troubleshooting, follow these guidelines:

- Try to narrow the definition of any problem that you have, particularly when communicating the issue to someone, such as PeopleSoft Worldwide Customer Support Services. For example, rather than reporting that the batch application failed, explain how the batch application failed. The more specific your information, the faster the problem can be solved. Rather than reporting that “The report had the wrong data,” say that “The batch status is E.”
- When communicating an error message to someone, be sure to include all parts of the error message exactly as they appear in the log file or on the screen. Parts of the message that might not seem important might actually hold the key as to why an error occurs. Also, distinguish between characters that might be misinterpreted, for example, the capital letter “O” and the numeral zero.
- As soon as you notice an error, examine the log files. Messages near the end of the log files will probably reveal the most important information about the cause of the error.
- Before you restart EnterpriseOne on the server, either delete or move all the files from the log directories. Refer to the JDE.INI file for the locations of the log files.
- When you first try to get EnterpriseOne running, verify that you have logging turned on. Examine the JDE.LOG and JDEDEBUG files carefully.
- Carefully examine the JOBLOGs and JDE.LOG files of the EnterpriseOne jobs to ensure that authorities and OCM are set correctly. Look for messages such as the following in the JDE.LOG files

```
JDB3100011 - Failed to get location of table F983051 for environment PD810
```

Look for messages similar to the following in the JOBLOGs:

```
File F98306 not found in library PRODDTA.
```

You might want to temporarily modify the job description of the EnterpriseOne user profile to always write the joblog until you are satisfied that all setup is correct.

Important

To complete the resolutions provided for these issues, you must sign on to the Enterprise server using an account that has administrative privileges.

See Also

- ❑ *Understanding the EnterpriseOne iSeries Architecture and Process Flow for iSeries* in the *Server and Workstation Administration Guide* for a better understanding of the flow of the iSeries EnterpriseOne server processes

Troubleshooting iSeries Enterprise Server Installation

This section explains topics that might create issues during the installation of an iSeries enterprise server.

Troubleshooting: Library Installation Verification

Issue	Resolution
You want to verify that the correct libraries and data dictionary items are installed on the iSeries.	See the list of libraries and data dictionary items and descriptions of their contents.

Note:

For a description of data dictionary items and iSeries libraries, see *Understanding the iSeries Library Structure for EnterpriseOne* in the *Server and Workstation Administration Guide*.

Troubleshooting: Database Table Configuration

Issue	Resolution
Strange database results or errors imply that Object Configuration Manger (OCM) is not set up correctly. For example, you see the following message in the JDE.LOG file: JDB3300011 - Failed to get location of table F983051 for environment PD810	• Verify that environments set up in the OCM are correct. • Review the description of how OCM is used by EnterpriseOne in EnterpriseOne Initialization. • Run the VerifyOCM program to ensure that the OCM tables are set up correctly. You must have one valid environment available to run VerifyOCM.

Troubleshooting: Setting up the iSeries .INI File

Issue	Resolution
You cannot find the .INI file	Find it in IFS. The file should be located in the /<release>/ ini directory. For example, /E810sys/ini/JDE.INI.
You need more information on using the iSeries .INI file	Review the notes and descriptions of .INI settings.

Note

For a description of the .INI settings, see *Troubleshooting the JDE.INI File* in the *Server and Workstation Administration Guide*.

Troubleshooting: Finding the Log Files

Issue	Resolution
You cannot find the log files	In B73.3 and subsequent releases, the logging is performed to the iSeries Integrated File System (IFS). The naming convention is similar to that of the UNIX enterprise servers. That is, the default names of the files are JDE_AS400JobNumber.log, JDEDEBUG_AS400JobNumber.log, and JDETS.LOG, where AS400JobNumber is the iSeries Job Number of the job that generated the file. The system creates these files automatically, but the path to the files must exist before logging begins. The created log file directory should match the JOBLOG and JDELOG settings in the JDE.INI file. The path to the log files stored in the IFS can be created by performing successive calls to the iSeries command MKDIR. For example, to create the path /PSFT810/LogFiles, enter the following command: • MKDIR DIR('/PSFT810') DTAAUT(*RWX) OBJAUT(*ALL) Followed by: • MKDIR DIR('/PSFT810/LogFiles') DTAAUT(*RWX) OBJAUT(*ALL) Logging might be turned off in the .INI. Turn on logging in the .INI using the following settings in the [DEBUG] section: <pre>[DEBUG] LogErrors=1 Output=FILE</pre> Where variable names and descriptions are as follows: LogErrors valid values are: 0 = Do not generate logs. 1 = Create logs. Output valid values (always in upper case) are: - NONE = Do not write debug messages to any output device. - FILE = Write messages to log files.
Not enough relevant information is written to the log files	Additional logging information may need to be turned on in the .INI. Set the following keys in the .INI for additional information to be output to the log files: <pre>[JDENET] netTrace=1 [JDEIPC] ipcTrace=1 [DEBUG] TAMTraceLevel=1 [UBE] UBEDebugLevel=6 [TCEngine]</pre>

TraceLevel=10

Where variable names and descriptions are as follows:

netTrace valid values are:

0 = Do not generate JDENet error messages (that is, communication between platforms).

1 = Generate JDENet error messages.

ipcTrace valid values are:

0 = Do not generate Interprocess Communication (IPC) error messages (that is, communication between processes on a single platform).

1 = Generate IPC error messages.

TAMTraceLevel valid values are:

0 = Do not generate Table Access Management (TAM) error messages (that is, regarding specification files).

1 = Generate TAM error messages.

UBEDebugLevel valid values are:

0 = Do not generate batch application error messages.

1-6 = Generate increasingly detailed error messages (1 indicates the least specific message; 6 indicates the most detailed message).

TraceLevel valid values are:

0 = Do not generate Table Conversion (TC) error messages.

1-10 = Generate increasingly detailed error messages (1 indicates the least specific message; 10 indicates the most detailed message).

Note

Because NetTrace and ipcTrace messages are written to the debug log associated with that job, the [DEBUG] section of the jde.ini file requires the Output=FILE setting.

Troubleshooting: Testing with PORTTEST

In general, turn on logging when running PORTTEST. Review the JDE.LOG and JDEDEBUG members generated by running PORTTEST. Also review the iSeries job log generated by running PORTTEST. These logs provide valuable information about your EnterpriseOne iSeries configuration and setup.

Issue

Resolution

An error with OCM

Verify that the OCM is correct for the environment. Disable the security server

occurred.	in the JDE.ini file and make sure that porttest runs successfully. For this work, you must log on with a User ID that has administrative privileges.
An error with the security server occurred.	The EnterpriseOne network might not be running. Clear the Interprocess Communication (IPC) structures using the EnterpriseOne iSeries CLRIPC command and restart EnterpriseOne. If you have different versions of EnterpriseOne running, make sure that they are on different ports and have different values for startIPCKeyValue. In the [JDEIPC] section of the JDE.INI file. Also, note that the different versions of EnterpriseOne should have different EnterpriseOne libraries, database files, and IFS directories Successful running of CLRIPC should result in the appearance of no messages on the screen. If messages appear as a result of CLRIPC, one or more jobs (including an interactive job that ran PORTTEST) might have locked some of the IPC shared memory. Determine which job locked shared memory and end it. Try logging off of a session in which you ran PORTTEST and running CLRIPC. If all attempts fail, you may change the .INI setting [JDEIPC] startIPCKeyValue to at least 1000 more or less than the current setting. Log off and back on again to ensure the new value is read. Attempt CLRIPC again, and restart EnterpriseOne if CLRIPC is successful.
An error with the security server occurred	The EnterpriseOne network might be running as a service under one library list and you are trying to run PORTTEST under another library list. Display all the libraries in the current library list and correct the list if the displayed library list is wrong. Then run PORTTEST. If the library list is correct, the problem could be because the activation group under which your job is running on the iSeries may retain some of the information from previous attempts. Log off, log on, and run PORTTEST again.
An error with the security server occurred	The supplied user name or password might not match any names or passwords in the EnterpriseOne security table. Try one of the following: • Run PORTTEST with a valid user name and password. • Add the given user name and password to the EnterpriseOne security table.
You get the following message on the screen:	You might not have included the correct number of arguments to PORTTEST. Use the following arguments:
Invalid parms PORTTEST <USER> <PWD> <ENV>	User - A valid EnterpriseOne user ID. Password - Password for the EnterpriseOne user ID. Environment - A valid EnterpriseOne environment.
Fewer than 99 F0902 records are written to the screen by PORTTEST	A possible PORTTEST failure. Examine the log files. • Fewer than 99 records might exist in the F0902 table. This is not an error, but you should review the log files for any errors. • The F0902 database table might not be accessible. Verify that you can query the F0902 table using SQL. Use the STRSQL command on the iSeries.
An error initializing	The environment might not be set up correctly. Any errors in the affected .INI keys or

the environment occurs in the log file database tables could cause the EnterpriseOne initialization to fail. The environment that PORTTEST uses is passed as a command line argument.

Note

See *Understanding EnterpriseOne Initialization for iSeries in the Server and Workstation Administration Guide* for information about how EnterpriseOne programs use OCM.

Troubleshooting: Running JDENET

Issue	Resolution
NETWORK dies immediately.	IPCs might not have been cleared before starting EnterpriseOne (that is, starting JDENET using the EnterpriseOne iSeries command STRNET). End EnterpriseOne (ENDNET). Clear IPCs (using the CLRIPC command) and restart EnterpriseOne. The startIPCKeyValue in the .INI file could be used by another version of EnterpriseOne. Try one of the following: • Change the startIPCKeyValue and restart the software. This problem is not easily evident by examining the log files or reviewing error messages. Symptoms of the problem include: • You attempt to run more than one version of EnterpriseOne on the iSeries. • One environment can be successfully started by itself. A second environment cannot be successfully started (that is, the JDENET_N job ends almost immediately after starting) for the second version. • Look in the JDE_xxx and JDEDEBUG_xxx files for specific error messages. • Determine if PORTTEST runs correctly. If not, correct those problems, and then try restarting EnterpriseOne using STRNET. • The configuration for the local host name, local domain name, and IP address might be incorrect. In the command line, enter CFGTCP to access the Configure TCP/IP form. Choose option 12 (Change local domain and host names) and verify the settings for the local domain name and the local host name (for example, YOURCOMPANY.COM and SRVR1 respectively). Then choose option 10 (Work with TCP/IP host table entries) and verify that two names exist in connection with the IP address for the iSeries. One name is a combination of the local host name and the local domain name (for example, SRVR1.YOURCOMPANY.COM). The other name is just the local host name (for example, SRVR1). • Verify that the Relational Database Directory name is set up correctly. Use the WRKRDBDIRE command to verify that the name of the *LOCAL database is the same as the server. If they are different, refer to the iSeries Configuration guide to determine how to set this up correctly.
An error initializing the environment occurs in the log file.	• Examine the issues in this section about PORTTEST. • Determine if PORTTEST runs correctly. If not, correct those problems, and then try restarting EnterpriseOne using STRNET.

Troubleshooting: Testing EnterpriseOne by Submitting a Report

Issue	Resolution
You get the following message:	You might see a message referencing an error of 11, indicating a time out occurred because the server was started after the client was run. Try resubmitting the report.
Communication Failure with <server name>	A time out might have occurred because of heavy network traffic or server load. Increase the time out value for the JDENETTimeout setting in the [NETWORK QUEUE SETTINGS] section of the jde.ini file on the workstation. The wrong communications port might have been used. Verify that the serviceNameListen value in the [JDENET] section of the jde.ini file on the workstation matches the serviceNameConnect value in the [JDENET] section of the jde.ini on the server. In addition, the serviceNameConnect value in the client's jde.ini must match the serviceNameListen in the jde.ini on the server. Other communications problems might exist. Run SERVERADMINISTRATIONWORKBENCH.exe (found in the system\bin32 directory on the workstation). This program shows only the machines on the specified port (also known as "service") that are running EnterpriseOne (either client or server). Use the following information to track down the problem: • If the remote machine is visible, a time-out probably occurred. Rerun the report. • If the remote machine is not visible, try to ping the remote machine using the name of the machine. • If the ping fails, try to ping the remote machine using its IP address. • With this information, determine if the client and server agree on the IP address for the server. • If none of these steps identify the problem, a general network error probably occurred (for example, the network is down or a machine is disconnected). Track it down.
The report does not show any data.	No data might exist in the database for the report that you are running, or you do not have access to the data. Try the following: • Select a different report to verify that some reports do produce data. • Verify the database contains data that should be included in the report. Add data if necessary. • Change the processing options for the report. • Change the OCM and/or data sources to reference the correct library. • If the report is launched on the server, make sure that the vertical tables in the server OCM match those of the OCM for the workstation. If no data is found, it could be because: • No data exists. • The processing options are incorrect. • The OCM for either the client or server is pointing to the wrong data source. • The data sources for either the client or server are pointing to the wrong database. • The SQL statement is incorrect (possibly due to a program bug). • The database drivers are out of date.
The report does not show any data.	An error might have occurred with the report. Review the jddebug.log and jde.log files for errors.

An error initializing the environment occurs in the log file.	The environment might not be set up correctly. Try the following: • Check for errors in .INI keys or database tables that can cause an initialization failure. • Stop EnterpriseOne and determine if PORTTEST runs correctly. If not, correct the problems, and then rerun EnterpriseOne manually.
You get the following message:	The server might have two network cards, which can confuse JDENET when the net communications are initialized between the client and server. One machine tries to connect using one network card, and the other machine connects using the other network card.
Communication Failure with <server name>	The hosts file on the server should list two different IP addresses for the server: one for each network card. The solution for the error involves setting the NetHostName field in the [JDENET] section of the JDE.INI to one of the names for the server given in the hosts file.
This error occurs sometimes on the workstation	JDENET then uses the IP address associated with the given network card.
Restarting JDENET_N sometimes gets rid of the error	
SAW (running on the workstation) cannot always see the server	
You can ping the server from the workstation	

Note

See *EnterpriseOne Initialization for iSeries* in the *Server and Workstation Administration Guide* for information about how the software uses OCM.

Troubleshooting: Shutting Down JDENET

Running the iSeries command CLRIPC immediately after shutdown (that is, after running the iSeries command ENDNET) each time you shut down will help you avoid most restart problems.

Troubleshooting: E-Mail and PPAT

Issue	Resolution
The batch application, server package installation, or table conversion log file (in the PrintQueue directory) shows the message:	The particular user may not be found in the Address Book table (F0101). Add the user to the Address Book table (F0101).

DoSendMessage Error: User 5600427 does not exist in the

address book file (F0101).

Troubleshooting: Multiple Release Setup

This section explains how to troubleshoot problems that can occur with multiple releases on the iSeries.

Issue	Resolution
When you try to run multiple releases of EnterpriseOne at the same time, conflicts seem to occur between each release.	Each installed release of EnterpriseOne may not have its own unique set of keys in the .INI. Change the following keys in one or both .INI files: [JDEIPC] startIPCKeyValue [JDENET] serviceNameListen serviceNameConnect Variable names and descriptions: startIPCKeyValue An integer value that indicates an arbitrary starting memory offset for interprocess communications. For multiple instances of EnterpriseOne server, be sure that the differences between these values are 1000 or more. The default value is 5000. Note: IBM Opti-Connect and Opti-Mover products use the IPC shared memory address 9999. Avoid setting the jde.ini file setting IPCStartKey to a starting value using the range of 9000 to 9999. serviceNameListen Port through which JDENet listens for communications attempts. The default is jde_server (translated using the "services" file). Each instance of the EnterpriseOne server needs to communicate with EnterpriseOne clients through different ports. serviceNameConnect Port through which JDENet tries to initialize connections with other platforms. The default is jde_server (which is translated using the "services" file). Each instance of EnterpriseOne server needs to communicate with EnterpriseOne clients through different ports. Also, verify that each version of EnterpriseOne has a unique set of libraries and database files. This is done using the ApplicationPathAddendum setting in the JDE.INI file.

Troubleshooting: JDBNET

This section explains how to troubleshoot problems that can occur with JDBNET.

Issue	Resolution
You do not know how JDBNET is used	JDBNET processes database requests using a client and server. It can also be configured to process server-to-server requests. This is, one server functions as a JDBNET client and the other as a JDBNET server. JDBNET eliminates the need for database-specific network software. All database requests are transported to the JDBNET server, processed in a local database, and the results are transported back to the JDBNET client.
You get an error that the data source on the JDBNET server is not found	The correct data source on the JDBNET server may not exist. Create a data source on the server that will be used by JDBNET. This is a normal configuration for a server data source that can be accessed by JDENet running on that server. Note the data source name (OMDATP) that will be used for the JDBNET client configuration.
You get an error that the data source on the JDBNET client is not found	The correct data source on the JDBNET client may not exist. Use the P98611 application to create a JDBNET data source in the F98611 table using the following information: • Data source name (OMDATP field) - Used to access tables as specified in the F986101 table. • Server name (OMSRVR field) - Identifies the JDBNET server. • Database name (OMDATB field) - Matches exactly the data source name (that is, the OMDATP field) to be used by the JDBNET server. • All other columns must match the values in the corresponding columns of the server data source. Set this data source as an active override in the F986101 table for all tables that will be accessed through JDBNET.
JDBNET does not transfer any data	The network may not be running. End EnterpriseOne, clear IPC (via the iSeries CLRIPC command), and restart EnterpriseOne.
JDBNET does not transfer any data	The JDBNET server and client may not be using the same server port number. Modify the serviceNameListen and serviceNameConnect fields in the [JDENET] section of both the JDBNET jde.ini files on the server and on the workstation. These values must match on both the JDBNET server and JDBNET client.

Troubleshooting: Interprocess Communications

This section explains how to troubleshoot problems that can occur with Interprocess Communication (IPC).

Issue	Resolution
EnterpriseOne jobs cannot communicate with one another with the following symptoms: PORTTEST fails	This could be because the iSeries release is pre-V4R2. In these releases, damaged IPC message queues might result when you end EnterpriseOne jobs using the command ENDJOB* IMMED. • Use the *CNTRLD option to end an iSeries job.

Caution

The security server on the iSeries fails
UBE submission fails

If you activated ipcTrace in the [JDEIPC] section of the server jde.ini file, an error similar to the following should appear in the JDEDEBUG.log:

IPC2100017 createIPC Msgq (name Port6005) failed, errno=3484: A damaged object was encountered

You might still have damaged IPC message queues if the iSeries-controlled ending times out.

- Run the following program to verify whether a damaged message queue exists. You must have V4R1 PTF# SF45946.
CALL QPOZIPCS PARM('-aqE')

This program generates a spool file called IPCS that contains information about message queues on the system. Look for the following output:

```
KEY    MODE
0x00000000 -----
0x00000000 --RW-----
0x00000000 --RW-----
0x00000000 --RW-----
0x00001234 D-RW----RW-
```

In this example, the message queue 0x00001234 is damaged. To fix, stop and restart JDENET using the following commands:

ENDNET

CLRIPC

STRNET

Also, if the ipcTrace setting in the [JDEIPC] section of the jde.ini file on the server is not set, activate the setting and run PORTTEST to determine whether any message queues are damaged. Look for the word "damage" in the JDEDEBUG.log file.

Caution: Some of your message queues might be damaged even if the JDEDEBUG.log file does not indicate that any damage exists.

Troubleshooting: JDE.INI File

This chapter explains how to troubleshoot problems that can occur with the JDE.INI file. The following notes apply to the .INI file in the E810SYS library:

- It is composed of several sections. The section names are enclosed in square brackets - for example, [JDENET].
- Within each section are one or more keys or settings. The key name is on the left side of the equals sign, and the value of the key is on the right side.
- Do not include spaces in the names or values of the keys unless you know that a space is required. Do not include spaces immediately before or after the equals sign.
- Keys may be commented out by adding a semicolon (;) at the start of the key name.
- PeopleSoft recommends that you place any incidental comments on a separate line above the key to which the comment applies. Be sure to include a preceding semicolon. Comments can be included at the end of the keys' values, but these comments can be wrongly interpreted if they are not separated from the keys' values by enough white space. Because the amount of

white space needed between the keys' values and the comments is not strictly defined, PeopleSoft recommends that you do not place comments after the values of the keys.

- The section and key names are not case sensitive.
- Many key values are case sensitive.
- Although all of the following may mean to turn a feature on, they may not be interchangeable as values in the .INI. Use a value that is comparable to the default value provided in the original .INI. Also, many values are case sensitive. If you have any questions about valid values, contact PeopleSoft Worldwide Customer Support Services.
 - YES
 - ON
 - TRUE
 - 1
- Likewise, the following values mean to turn a feature off. They are not necessarily interchangeable as values in the .INI.
 - NO
 - OFF
 - FALSE
 - NONE
 - 0

If you are told by PeopleSoft Worldwide Customer Support Services to modify a key that does not exist, you can add the key. Just be sure that it is in the correct section.

See Also

- ❑ *The Jde.ini File* in the *System Administration Guide* for more information about troubleshooting problems that can occur with the JDE.INI file

Troubleshooting the UNIX/Linux Enterprise Server

This section describes some typical problems that you might encounter and their solutions. When troubleshooting, follow these guidelines:

- Check your logs. Many times, the logs point to the problem. As soon as you notice an error, examine the log files. Messages near the end of the log files will probably reveal the most important information about the cause of the error.
- Try to narrow down the definition of any problem that you may have, particularly when communicating the issue to someone, such as PeopleSoft Worldwide Customer Support Services. For example, rather than reporting that the batch application failed, explain how the batch application failed. The more specific your information, the faster the problem can be solved. For example, rather than reporting that “The report had the wrong data,” say that “The batch status is E.”
- When communicating an error message to someone, be sure to include all parts of the error message exactly as they appear in the log file or on the screen. Parts of the message that may

not seem important may actually hold the key as to why an error occurred. Also, distinguish between characters that might be misinterpreted; for example, the capital letter O and the numeral zero (0).

- Before you restart EnterpriseOne on the server, either delete or move the jde_xxx.log and jdedebug_xxx.log files (where xxx is a number). Do not rename the log files because it is easier to work with logs that use the standard naming convention (jde_xxx.log and jdedebug_xxx.log). If you need to save the log files until the problem is solved, then create a temporary directory and move the files.
- Clear the log directory regularly to avoid filling the file system. If the file system fills up, then the specification files can become corrupted.
- Always keep a backup of the specification files handy in case they become corrupted. Specification files should be backed up regularly for easy recovery of specification installs. If spec files have to be replaced, all specification installations will be lost if backups are not kept.
- To find problems that occur due to server failure, go to the system/bin32 directory:
 - `grep -n "failed" *log* > problems.txt`
The file problems.txt will contain a list of errors with the file and line number.
- Remember that UNIX is case-sensitive: jde.ini is not the same file as JDE.INI.

Caution

To complete the resolutions provided for these issues, you must sign on to the UNIX enterprise server using an account that has administrative privileges.

Troubleshooting the UNIX Enterprise Server Installation

This section provides details of the following topics that might create issues during the installation of a UNIX enterprise server.

Troubleshooting the JDE.INI File

To locate the JDE.INI file, search in the system/bin32 subdirectory. For example, /u01/PeopleSoft/E810/ini/JDE.INI. The following notes apply to the JDE.INI:

- It is composed of several sections. The section names are enclosed in square brackets; for example, [JDENET].
- The environment variable \$JDE_BASE should contain the location of the JDE.INI file.
- If you copy the JDE.INI file to other directories (for example, the \$SYSTEM/bin32 directory), the EnterpriseOne programs might read the wrong JDE.INI file. This error occurs because some programs might look for the JDE.INI file in the current directory before looking at the JDE_BASE environment variable.
- Each section contains one or more keys. The key name is on the left side of the equal sign, and the value of the key is on the right side.

- Do not include spaces in the key names or key values unless you know that a space is required. Do not include spaces immediately before or after the equal sign.
- Keys can be commented out by adding a semicolon (;) at the start of the key name.
- PeopleSoft recommends that you place incidental comments on a separate line above the key to which the comment applies. Be sure to include a preceding semicolon. Comments can be included at the end of the key value, but these comments can be incorrectly interpreted if they are not separated from the values of the keys by sufficient white space. Because the amount of white space between the values of the keys and the comments is not strictly defined, we recommend that you do not place comments after the values of the keys.
- Section and key names are not case sensitive.
- Many key values are case sensitive.
- Although all of the following may mean to turn on a feature, they may not be interchangeable as values in the JDE.INI. Use a value that is comparable to the default value provided in the original JDE.INI. Also, many of these values are case sensitive. If you have any questions about valid values, contact PeopleSoft Worldwide Customer Support Services.

YES

ON

TRUE

1

The following values mean to turn a feature off. They are not necessarily interchangeable as values in the JDE.INI.

NO

OFF

FALSE

NONE

0

If you are told by PeopleSoft Worldwide Customer Support Services to modify a key that does not exist, you can add the key. Ensure that the key is in the correct section and entered with the correct spelling and case.

Troubleshooting: Copying EnterpriseOne to a Server

If you cannot copy files from the deployment server to the temporary directory on the Enterprise server, this could be because ftp cannot connect. See your system administrator.

Troubleshooting: Configuring Database Tables

If results or errors occur that imply that OCM is not set up correctly, review the description in this guide of how OCM is used by EnterpriseOne.

Troubleshooting: Setting Up a Printer

If reports do not print from a server, verify the name of the default printer. Send a simple text file to the default printer using the lp command. If you get an error similar to the following, then the printer is not configured on the server or is not online:

“lp: destination aPrinter non-existent”

Contact your system administrator for assistance.

For Linux, do not set up a print queue that translates files to postscript. Your Linux print queues that are used by EnterpriseOne should generally be “raw” print queues that simply redirect the output of the file to the printer.

Troubleshooting: E-Mail

If the report, server package installation, or table conversion log file (in the PrintQueue directory) shows the message "DoSendMessage Error: User 5600427 does not exist in the address book file (F0101)," the particular user might not be found in the Address Book table (F0101). Add the user to the Address Book table (F0101).

Troubleshooting: Multiple Release Setup

Each installed release of EnterpriseOne has its own JDE.INI in its ini directory. Point the user entries in the JDE.INI files to the directories of the log and other files. If the log files do not go to separate directories, change the appropriate keys in one or both JDE.INI files to point to unique directories for each installed instance of EnterpriseOne.

Troubleshooting: Finding the Report Files

If you cannot find the report output files, consider the following:

- The location is specified as the OutputDirectory key of the [NETWORK QUEUE SETTINGS] section in the JDE.INI on the server. If this key is not found, the location is the PrintQueue subdirectory of the EnterpriseOne base directory (for example, /u01/PeopleSoft/E810SYS/PrintQueue).
- The JDE.INI file on the workstation may have the SaveOutput key of the [NETWORK QUEUE SETTINGS] section set to FALSE. This is because a problem after the report has been printed. After the report is printed, then the record will be deleted, as will the .PDF file. Change the value of the SaveOutput key of the [NETWORK QUEUE SETTINGS] section in the JDE.INI on the workstation to TRUE.

Troubleshooting: JDBNET Server Not Found

If you get an error that the data source on the JDBNET server is not found, the correct data source on the JDBNET server might not exist. Create a data source on the server that will be used by JDBNET. This is a normal configuration for a server data source that can be accessed by JDENet running on that server. Note the data source name (OMDATP) that will be used for the JDBNET client configuration.

If you get an error that the data source on the JDBNET client is not found, the correct data source on the JDBNET client might not exist. Create a JDBNET data source in the F98611 table using the following information:

- Data source name (OMDATP field)
Used to access tables as specified in the F986101 table.
- Server name (OMSRVR field)
Identifies the JDBNET server.
- Database name (OMDATB field)
Matches exactly the data source name (that is, the OMDATP field) to be used by the JDBNET server.
- Shared library name (OMDLLNAME field)
Identifies the JDBNET client .DLL. (libjdbnet.sl on HP-UX, libjdbnet.so on AIX).
- All other columns must match the values in the corresponding columns of the server data source.

Set this data source as an active override in the F986101 table for all tables that will be accessed through JDBNET.

Troubleshooting: EnterpriseOne Testing

If porttest does not run successfully after startup:

- If you have Oracle or UDB running on the Enterprise server and the database and EnterpriseOne services are set to start automatically at system startup, EnterpriseOne services may start before the database is running completely. You must ensure that the database software is running before starting any EnterpriseOne processes.
- If EnterpriseOne loses the connection to the database because either the network or database went down, you should see some sort of network or database error in the log files.
- Stop the EnterpriseOne services, clear the logs, and then restart the EnterpriseOne services to see if the problem is resolved.

Troubleshooting the Windows Enterprise Server

This section describes some typical problems that you might encounter and their solutions. When troubleshooting, follow these guidelines:

- Narrow the definition of any problem that you might have, particularly when communicating the issue to someone, such as PeopleSoft Worldwide Customer Support Services. For example, rather than reporting that the batch application failed, explain how the batch application failed. The more specific your information, the faster the problem can be solved. For example, rather than reporting that “The report had the wrong data,” say that “The batch status is E.”
- When communicating an error message to someone, be sure to include all parts of the error message exactly as they appear in the log file or on the screen. Parts of the message that may

not seem important may actually hold the key to why an error occurs. Also, distinguish between characters that might be misinterpreted (for example, the capital letter "O" and the number zero "0").

- As soon as you notice an error, examine the log files. Messages near the end of the log files will probably reveal the most important information about the cause of the error.
- Before you restart EnterpriseOne on the server, either delete or move the jde_xxx.log and jdedebug_xxx.log files (where xxx is a number). Do not rename the log files; it is easier to work with logs that use the standard naming convention ("jde_xxx.log" and "jdedebug_xxx.log"). If you need to save the log files until the problem is solved, create a temporary directory and move the files there.
- Clear the log directory regularly to avoid filling the file system. If the file system fills up, then the specification files will become corrupt.
- Always keep a backup of the specification files in case they become corrupt. Specification files should be backed up regularly for easy recovery of spec installs. If specification files have to be replaced, all specification installations will be lost if backups are not kept.

Important

To complete the resolutions provided for these issues, you must sign on to the Windows Enterprise server using an account that has administrative privileges.

Troubleshooting the Windows Enterprise Server Installation

This section explains how to troubleshoot problems that can occur when installing the Windows Enterprise server.

Troubleshooting: Setting up EnterpriseOne Accounts

If you cannot set up any accounts in the User Manager program, the account you are logged into in Windows may not have the privileges to modify or add accounts. Log out of Windows and log back on under the Administrator account or an account in the Administrators group.

Troubleshooting: Copying EnterpriseOne to a Server

If you cannot copy files from the CD to the EnterpriseOne directory on the Enterprise server, verify that the CD is in the CD-ROM drive. Another cause is that one or more of the files to be copied is currently open on the CD:

- Close any files on the CD that are open.
- Close any applications that may have files open on the CD.

If one or more of the files that will be overwritten in the target directory is open:

- Close any files in the target directory that are open.
- Close any applications that may have files open in the target directory.

If the target disk is full:

- Delete or move files from the target disk.

- Copy EnterpriseOne to a different disk.

Troubleshooting: Configuring Database Tables

If the OCM is not set up correctly and errors occur, run the VerifyOCM program to ensure that the OCM tables are set up correctly.

Troubleshooting: Setting Up a Printer

If you cannot set up a printer:

- The printer may not be attached (local printer) or the print server may not be available (network printer). Attach to the local printer or determine why the print server is not available.
- The printer drivers may not be installed. Install the correct printer drivers.

Troubleshooting: Setting Up the JDE.INI File

If you cannot find the JDE.INI file:

- Search in the system\bin32 subdirectory in the EnterpriseOne tree. For example, z:\PeopleSoft\E810\ddp\system\bin32\jde.ini.
- Make sure you have access rights to the system\bin32 directory by logging on to Windows as a user who has administrative rights.

Troubleshooting: Finding the Log Files

If you cannot find the log files:

- Log files are listed in the DebugFile and JobFile keys in the [DEBUG] section of the jde.ini. If no paths are shown, the logs are in the system\bin32 directory. The log files are named according to the following scheme:

An underscore (_) and the process ID of the process that creates the log file are inserted before the period for example, jde_123.log or jdedebug_123.log for a process with an ID of 123.

The log file associated with the DebugFile key contains the sequence of EnterpriseOne events.

The default value for this key is jdedebug.log.

The log file associated with the JobFile key contains error messages that occur in EnterpriseOne.

The default value for this key is jde.log.

- When a batch application is run and the jde.ini on the workstation has [NETWORK QUEUE SETTINGS] SaveOutput=TRUE, the jde_xxx.log and jdedebug_xxx.log files for the runbatch that processed the batch application is copied to a file in the PrintQueue directory. The root name of the files are the same as the name of the PDF file. The extension is .jde.log

and .jdedebug.log. The duplication of these log files does not occur if the batch application runbatch.exe dies before duplication.

- Verify logging in the jde.ini is turned on using the following settings in the [DEBUG] section:

[DEBUG]

LogErrors=1

Output=FILE

Variables and their descriptions:

LogErrors

0 = Do not generate logs.

1 = Create logs.

Output

NONE = Do not write messages to any output device.

AUX = Write messages to a console window.

FILE = Write messages to log files.

BOTH = Write messages to log files and console window.

If not enough relevant information is written to the log files, this could be because additional logging information needs to be turned on in the jde.ini. Set the following keys in the jde.ini for additional output to the log files:

[JDENET]

netTrace=1

[JDEIPC]

ipcTrace=1

[DEBUG]

TAMTraceLevel=1

[UBE]

UBEDebugLevel=6

[TCEngine]

TraceLevel=10

Variables and their descriptions:

netTrace

0 = Do not generate JDENet error messages (that is, communication between platforms).

1 = Generate JDENet error message.

ipcTrace

0 = Do not generate Interprocess Communication (IPC) error messages (that is, communication between processes on a single platform).

1 = Generate IPC error messages.

TAMTraceLevel

0 = Do not generate Table Access Management (TAM) error messages (that is, regarding specification files).

1 = Generate TAM error messages.

UBEDebugLevel

0 = Do not generate batch application error messages.

1 = Generate increasingly detailed error messages (1 gives the least specific messages, whereas 6 gives the most detailed messages).

TraceLevel

0 = Do not generate Table Conversion (TC) error messages.

1-10 = Generate increasingly detailed error messages (1 gives the least detail, whereas 10 gives the most detail).

Troubleshooting: Testing with PORTTEST

If an error with the security server occurred:

- Verify the EnterpriseOne network is running either as a service or started from a command prompt.
- If the security server is inactive, or if it is active on a server and port that is different from the ones PORTTEST uses, do one of the following:

Start EnterpriseOne net on the server and port where PORTTEST is being run. The security server key in the [SECURITY] section of the jde.ini specifies the security server, and the serviceNameListen and serviceNameConnect settings in the [JDENET] section specify the ports.

Change the name of the security server or the names of the ports, or both, in the jde.ini file to point to the correct security server.

- Make sure that the EnterpriseOne network and PORTTEST are running under the same account:

To determine under which account PORTTEST is running, press the Control, Alt, and Delete keys at the same time.

If the EnterpriseOne network is running as a service, determine under which account it is running. To do this, choose the service in Windows's Control Panel, then go to Services and click on Startup.

For initial testing, you can stop the EnterpriseOne network service, open a Windows command prompt, cd to the system\bin32 directory, run jdenet_n without any parameters, and rerun PORTTEST. When finished, stop jdenet_n from the Windows Task Manager.

To run PORTTEST under the same account as the EnterpriseOne network service, log out of Windows, log into the same account under which the service is running, open a Windows command prompt, cd to the system\bin32 directory, and rerun PORTTEST.

- To make sure the supplied user name and password, or both, match names and passwords, or both, in the EnterpriseOne security table:

Run PORTTEST with a valid user name and password.

Add the given user name and password to the EnterpriseOne security table.

If you get the message "Invalid parms PORTTEST: <USER> <PWD> <ENV>", the correct number of arguments to PORTTEST may not have been included. Use the following arguments:

User - A valid EnterpriseOne account name.

Password - Password for the EnterpriseOne account.

Environment - A valid EnterpriseOne environment.

Fewer than 99 records are written to the screen by PORTTEST.

If PORTTEST failed, examine the log files.

If fewer than 99 records exist in the F0902 table, this is not an error. You should review the log files for errors.

If the F0902 table is not accessible, verify that you can query the F0902 table using SQL.

If an error initializing the environment occurs in the log file, the environment may not have been set up correctly. See the chapter Understanding the EnterpriseOne Initialization for Windows in this guide for more information about how EnterpriseOne programs use OCM. Any errors in the affected jde.ini keys or database tables could cause the EnterpriseOne initialization to fail. The environment that PORTTEST uses is passed as a command line argument.

Troubleshooting: Running EnterpriseOne Manually

If

- If the EnterpriseOne network is not running, start the EnterpriseOne network service.
- Verify the EnterpriseOne network is running by doing the following:

- The EnterpriseOne network should either be running as a service or from a Windows command prompt.
- If it is running as a service, determine under which account it is running. To do this, choose the EnterpriseOne network service in Windows's Control Panel, choose Services, and then choose Startup. Note the account name. If you are using Windows 2000, choose the EnterpriseOne network service in the Windows Control Panel, choose Services, and then choose Properties.
- If it is run from a command prompt, the network will be running under the Windows account into which you signed on. When you log off Windows, network processes started from a command prompt and all child processes will terminate.
- If the setup of some part of EnterpriseOne, such as the jde.ini file or OCM, is incorrect, determine if PORTTEST runs correctly. If not, correct those problems and then try running EnterpriseOne manually.

If an error initializing the environment occurs in the log file, the setup for some part of EnterpriseOne, such as the jde.ini file or OCM, may be incorrect. Examine the applicable problems under Testing with PORTTEST in this chapter. Determine if PORTTEST runs correctly. If not, correct those problems, and then try running EnterpriseOne manually.

Troubleshooting: Finding the Report Files

If you cannot find the report output files:

- Check the OutputDirectory key of the [NETWORK QUEUE SETTINGS] section in the jde.ini file on the server. If no location is shown, the files are in the PrintQueue directory of the EnterpriseOne base directory. For example, z:\PeopleSoft\E810\ddp\PrintQueue.
- Verify that SaveOutput in the [NETWORK QUEUE SETTINGS] section in the jde.ini file on the workstation is TRUE.

Troubleshooting: Testing EnterpriseOne by Submitting a Report

- If a time-out occurred because the EnterpriseOne server was started after the client, resubmit the report.
- If a time-out occurred due to heavy network traffic or server load, increase the time-out value in the jde.ini file on the workstation and resubmit the report. Use the JDENETTime-out setting in the [NETWORK QUEUE SETTINGS] section.
- If the wrong communications port is being used, do one of the following:
 - Verify that the serviceNameListen value in the [JDENET] section of the jde.ini file on the workstation matches the serviceNameConnect value in the [JDENET] section of the jde.ini file on the server. In addition, the serviceNameConnect value in the jde.ini file on the workstation must match serviceNameListen in the jde.ini file on the server. If the values of these keys are strings, the numeric value is retrieved from the services file in the c:\winnt\system32\drivers\etc directory (Windows; client or server).
 - The services file contains a list of strings and their corresponding port numbers. If the port that you are interested in is on the last line of the services file, be sure to include a return at the end of the line or else the string will not be translated to the corresponding port number.

- If the client is using Dynamic Host Configuration Protocol (DHCP) and the server does not have an entry for itself in its hosts file in the c:\winnt\system32\drivers\etc directory, add an entry for the server in the hosts file on the server.
- Run the Server Administration Workbench (SAW) program. This program shows only the machines on the specified port (also known as service) that are running EnterpriseOne (either client or server). Troubleshoot using the following:
 - If the remote machine is visible, a time-out probably occurred. Rerun the report.
 - If the remote machine is not visible, ping the remote machine using the name of the machine.
 - If the ping fails, ping the remote machine using its IP address.
 - If these pings fail, run SAW from the other machine (either client or server) and repeat the above steps.
 - With this information, determine if the client and server agree on the IP address for the server.
 - If none of these steps identify the problem, a general network error probably occurred (for example, the network is down or a machine is disconnected).
- The following situations can occur:
 - "Communications failure" error message on the workstation.
 - Restarting Network Service or jdenet_n sometimes gets rid of the error.
 - SAW (running on the workstation) cannot always see the server.
 - You can ping the server from the workstation.

These issues can occur because the server has two network cards, which confuses JDENET when the net communications are initialized between the client and server. One machine tries to connect using one network card, and the other machine connects using the other network card.

The hosts file on the server should list two different IP addresses for the server--one for each network card. Resolve the error by setting the NetHostName field in the [JDENET] section of the jde.ini to one of the names for the server given in the hosts file. JDENET then uses the IP address associated with the given network card.

- For the error "Cannot connect to printer" in the jde_xxx.log or the log file in the PrintQueue subdirectory:
 - If a general printing error occurred, try to print a text document from Notepad. Resolve any issues.
 - If no default printer is set up on the enterprise server, set up a printer using the task "Add a new printer" or "Modify an existing printer" in the System Administration Guide.
 - If you do not have privileges to the printer, define the owner as a local or network account. The type of account depends on the type of printer:

Local printer. The owner could be either a local or network account but either type must have privileges to access the printer.

Network printer. The owner must be a network account with access privileges.

To take ownership of a printer:

6. From the Windows Start menu, choose Settings, and then Printers.
7. Right-click the desired printer.
8. Choose Properties and then the Privileges tab.
9. Click Ownership and then Take Ownership.

If the printer drivers are not installed, see the section Database Driver Files in this guide for information about which drivers you need.

If the report printouts are in portrait mode but should be in landscape mode (or vice versa):

Verify that the orientation specified in RDA for the report is correct.

If the default printer is set to the wrong orientation, set the orientation using the following task:

10. From the Windows Start menu, choose Settings, and then Printers.
11. Right-click on the desired printer.
12. Choose Document Defaults.
13. Choose the desired default orientation.
14. Click OK.

All jobs sent to this printer using the current server will default to the selected orientation. Note that the report template or other programs may override this default orientation.

If you cannot change the printer orientation, you may not have the right to change the orientation. Log on to Windows in an account that has administrative rights for the printer. For a local printer, use an account that has administrative privileges. For a network printer, use an account given administrative privileges by a network administrator.

If the report does not show any data, the data may not exist in the database for the report that you are running or you do not have access to the data. Try one of the following:

Select a different report.

Add data to the database.

Change the processing options for the report.

Change the OCM and/or data sources to point to the correct database.

If the report is launched on the server, verify the vertical tables in the server OCM match those in the workstation OCM.

If you believe data should have been found:

Edit the report `jdedebug.log` found in the `PrintQueue` subdirectory.

Search for the SQL select statement used to retrieve data from the database. You must have some idea what data is being read to do this.

Copy the SQL statement.

Open the database's respective SQL command interface - for example, SQL Plus or ISQL_w.

Paste the SQL statement into the SQL command interface.

Submit the SQL statement.

If no data is found, one of the following may be true:

No data exists.

The processing options are incorrect.

The OCM for either the client or server is pointing to the wrong data source.

The data sources for either the client or server are pointing to the wrong database.

The SQL statement is incorrect (possibly due to a program bug).

The database drivers are out of date.

If an error occurred with the report, look in the `jde_XXX.log` for error messages.

If an error initializing the environment occurs in the log file, the environment may not be set up correctly.

Note

See *Understanding the EnterpriseOne Initialization for Windows* in the *Server and Workstation Administration Guide* for information about how EnterpriseOne programs use OCM. Any errors in the affected JDE.INI keys or database tables could cause the EnterpriseOne initialization to fail.

Stop EnterpriseOne and determine if PORTTEST runs correctly. If not, correct those problems and then run EnterpriseOne manually.

Troubleshooting: Stopping EnterpriseOne as Run Manually

If you need to stop the EnterpriseOne processes that you started from the command prompt, for example, `jdenet_n`, stop any of the following processes that are running:

- `Jdenet_n.exe`
- `Jdenet_k.exe`
- `Runbatch.exe`

- ipcsrv.exe

These additional processes, such as jdenet_k and runbatch, are started by jdenet_n and queue kernel. To stop all the EnterpriseOne processes:

15. Run the Windows Task Manager.
16. Choose the Processes tab.
17. Choose one of the running processes.
18. Click End Process.
19. Repeat for each process to be stopped.

If you do not have the rights to stop the processes:

20. Log on to Windows in an account that has rights to stop processes.
21. Stop processes via Visual C++:
 1. Run the Windows Task Manager.
 2. Choose the Processes tab.
 3. Choose one of the running processes.
 4. Click Debug Process. Visual C++ will start.
 5. Click on the "X" in the upper right-hand corner to close Visual C++. Do not save the project workspace. This should kill the runaway process.
22. Repeat these steps for each runaway process. If they still do not end, reboot the machine.

Troubleshooting: E-Mail

If the report, server package installation, or table conversion log file in the PrintQueue directory shows the message DoSendMessage Error:

- User 5600427 does not exist in the address book file (F0101). This could be because the particular user is not found in the Address Book table (F0101). Add the user to the Address Book table (F0101).

See Also

- *The Jde.ini File* in the *System Administration Guide* for more information about troubleshooting problems that can occur with the JDE.INI file

Troubleshooting Web Servers

This section describes some typical issues you might encounter when using WebSphere and Java Application Server (JAS). It also explains other issues you might encounter with web servers and how to track down problems by using the log files in SAW.

Troubleshooting: IIS and IBM HTTP Web Servers

If you need to configure with IIS and an IBM HTTP Server, refer to the installation documentation.

If you receive the message "Recursive error - page not found," you need to make sure IIS is running for a particular instance of JAS. IIS instances can be stopped easily, and the user may forget to restart them. To make sure IIS is running for the particular instance of JAS, verify IIS instance properties by selecting the appropriate instance, and then right-clicking and choosing Properties. Confirm that the correct paths are listed for the desired JAS code.

Troubleshooting: JAS

If no logs appear, verify that the [LOGS] setting in the jas.ini has logging turned on and points the log files to reside in the desired location (for example, ;log=d:\E810\internet\jas.log or ;debuglog=d:\E810\internet\jasdebuglog). If the log file paths are not correctly stipulated, the logs may be writing to a file located elsewhere.

If JAS seems slow, check to see whether jdbcTrace is set to TRUE or FALSE. If tracing is turned on or set to TRUE, the additional logging will dramatically slow JAS performance.

Troubleshooting: Serialized Database and Generation Issues

If you receive the message "Form is out of date...most likely needs to be regenerated," this error usually occurs because the specifications used to construct the serialized database do not match the JAS code. Ensure that the date the JAS code was written matches the date of the jdecom.dll that resides in the 810\system\bin32 directory of the generating machine.

Also be sure to register the jdecom.dll. After you run the regsvr32 jdecom.dll command, the eGenerator recognizes the jdecom.dll and uses it to fetch EnterpriseOne specs and convert them into Java serialized objects.

If the menu does not appear when the user signs on to EnterpriseOne, check for the following:

- [JDBC URL] section in jde.ini is set correctly or [JDBC DRIVERS] is set correctly. The [JDBC URL] points to the serialized database (the one you just set up).
- Bounce the WebSphere application server. Menus are cached, and by bouncing the server you clear the cached information.
- Ensure that the host database for serialized objects is running.

Troubleshooting: SQL Server Issues

If SQL Server process or Oracle process consumes excess CPU in a web server environment, the serialized objects for the web server are stored in either SQL server or the Oracle database. The web server must access these tables frequently when running an application. Indexes may be missing, which can cause severe performance problems.

Ensure that all existing EnterpriseOne indexes are created for tables F989998 and F989999. You should have one index for F989998 for columns WBJOBID and WBOID. You should also have one index for F989999 for columns WBUID, WBOID, WBLNGPREF. If these indexes do not exist in your database, generate them using Object Librarian.

Add a new index to the F989999. This index should include columns WBOID, WBUID, and WBJVER. Generate this index over the F989999 table.

Update statistics on both tables as follows:

- For Oracle, issue the following commands in SQL *Plus:
ANALYZE TABLE owner.F989999 COMPUTE STATISTICS

ANALYZE TABLE owner.F989999 COMPUTE STATISTICS
- For SQL Server, issue the following commands:
UPDATE STATISTICS owner.F989999

UPDATE STATISTICS owner.F989998

Improvements will vary depending on the number of users accessing the serialized database.

Troubleshooting: Problems Using Log Files

If you need to view logging information for the Java client, open the Java Console by choosing Java Console from the View menu in Internet Explorer. The Java Console displays all problems that the Java Virtual Machine on your client is having. Errors appear as uncaught exceptions in the console.

Note

You must have the appropriate Internet options turned on to view the Java Console.

To enable the Java Console in Internet Explorer, choose Tools, and then choose Internet Options. In Internet Options, click the Advanced tab, scroll down to the section titled Java VM, and choose the following options:

- Java Console enabled
- Java logging enabled
- JIT compiler for virtual machine enabled

If you need to troubleshoot errors in web applications:

- Verify that the problem is only a problem on the Web. Test the fat client version of the same application against the same Enterprise server that the web is using. Make sure that you use the same EnterpriseOne accounts and environments.
- Determine whether the problem happens in HTML, Java, or both. Since both Java and HTML use the Java runtime engine, they should behave the same. Some variation will exist based on the inherent differences between the Portal, HTML page processing and Java interactive processing, but underlying functionality and processing should be the same.
- Re-create the problem on the web server. (The logs will work in the Portal, HTML, and Java.)
- Open a separate Internet Explorer browser and use it to access the Web Server Monitor for the web server being used.

Note

See *Monitoring the JAS Servers from the Web* in the *Server and Workstation Administration Guide* for information about accessing the Web Server Monitor.

- Check the Standard Error Log (stderr.log) for errors. A common error you might see here is BSFN Failed. If you see this error, verify that the enterprise server is up and that the BSFN is not a T1 BSFN.
T1 refers to Type 1 business functions, which are client-only business functions. They cannot run on a server.
- Check the Standard Output Log (stdout.log) for more information. For example, you can view the time and date stamps from the errors found in both the Jas.log and the standard error log to find more detailed information about what was occurring at about the same time that the errors occurred.

If you need more information, enable Debug.log and set Net Trace, which you can do in the [LOGS] section of jas.ini file. Re-create the problem, view the Debug.log, and look for more information.

You can also use the Server Administration Workbench (SAW) to monitor web servers.

Note

See *Monitoring EnterpriseOne with SAW on Windows* in the *Server and Workstation Administration Guide* for more information.

Try to find SQL statement information. SQL statements can give you an idea of what values are being passed. Some common failures include:

- Form Interconnects are passing incorrect information.
Verify that the fat client is working correctly. Watch especially for null, blank, and zero problems, as well as special characters.
- String is too big.
Note carefully what you did to get this error.
- Null values are being passed.
Your SQL statement information search results in nothing being found. Check the SQL statements and make sure that correct values were passed. Determine where the failure occurred and make a note of it.
- The application stops responding.
Check logs for BSFN failures.

EnterpriseOne PeopleBooks Glossary

“as of” processing	A process that is run at a specific point in time to summarize item transactions.
52 period accounting	A method of accounting that uses each week as a separate accounting period.
account site	In the invoice process, the address to which invoices are mailed. Invoices can go to a different location or account site from the statement.
active window	The window that contains the document or display that will be affected by current cursor movements, commands, and data entry in environments that are capable of displaying multiple on-screen windows.
ActiveX	A technology and set of programming tools developed by Microsoft Corporation that enable software components written in different languages to interact with each another in a network environment or on a web page. The technology, based on object linking and embedding, enables Java applet-style functionality for Web browsers as well as other applications (Java is limited to Web browsers at this time). The ActiveX equivalent of a Java applet is an ActiveX control. These controls bring computational, communications, and data manipulation power to programs that can “contain” them—for example, certain Web browsers, Microsoft Office programs, and anything developed with Visual Basic or Visual C++.
activity	In Advanced Cost Accounting, an aggregation of actions performed within an organization that is used in activity-based costing.
activity driver	A measure of the frequency and intensity of the demands that are placed on activities by cost objects. An activity driver is used to assign costs to cost objects. It represents a line item on the bill of activities for a product or customer. An example is the number of part numbers, which is used to measure the consumption of material-related activities by each product, material type, or component. The number of customer orders measures the consumption of order-entry activities by each customer. Sometimes an activity driver is used as an indicator of the output of an activity, such as the number of purchase orders that are prepared by the purchasing activity. See also cost object.
activity rule	The criteria by which an object progresses from a given point to the next in a flow.
actual cost	Actual costing uses predetermined cost components, but the costs are accumulated at the time that they occur throughout the production process.
adapter	A component that connects two devices or systems, physically or electronically, and enables them to work together.
add mode	The condition of a form where a user can enter data into it.
advanced interactive executive	An open IBM operating system that is based on UNIX.
agent	A program that searches through archives or other repositories of information on a topic that is specified by the user.
aging	A classification of accounts by the time elapsed since the billing date or due date. Aging is divided into schedules or accounting periods, such as 0-30 days, 31-60

	days, and so on.
aging schedule	A schedule that is used to determine whether a payment is delinquent and the number of days which the payment is delinquent.
allegato IVA clienti	In Italy, the term for the A/R Annual VAT report.
allegato IVA fornitori	In Italy, the term for the A/P Annual VAT report.
application layer	The seventh layer of the Open Systems Interconnection Reference Model, which defines standards for interaction at the user or application program level.
application programming interface (API)	A set of routines that is used by an application program to direct the performance of procedures by the computer's operating system.
AS/400 Common	A data source that resides on an AS/400 and holds data that is common to the co-existent library, allowing PeopleSoft EnterpriseOne to share information with PeopleSoft World.
assembly inclusion rule	A logic statement that specifies the conditions for using a part, adjusting the price or cost, performing a calculation, or using a routing operation for configured items.
audit trail	The detailed, verifiable history of a processed transaction. The history consists of the original documents, transaction entries, and posting of records and usually concludes with a report.
automatic return	A feature that allows a user to move to the next entry line in a detail area or to the first cell in the next row in several applications.
availability	The expression of the inventory amount that can be used for sales orders or manufacturing orders.
available inventory	The quantity of product that can be promised for sale or transfer at a particular time, considering current on-hand quantities, replenishments in process, and anticipated demand.
back office	The set of enterprise software applications that supports the internal business functions of a company.
backhaul	The return trip of a vehicle after delivering a load to a specified destination. The vehicle can be empty or the backhaul can produce less revenue than the original trip. For example, the state of Florida is considered a backhaul for many other states—that is, many trucking companies ship products into the state of Florida, but most of them cannot fill a load coming out of Florida or they charge less. Hence, trucks coming out of Florida are either empty or produce less revenue than the original trip.
balance forward	The cumulative total of inventory transactions that is used in the Running Balance program. The system does not store this total. You must run this program each time that you want to review the cumulative inventory transactions total.
balance forward receipt application method	A receipt application method in which the receipt is applied to the oldest or newest invoices in chronological order according to the net due date.
bank tape (lock box) processing	The receipt of payments directly from a customer's bank via customer tapes for automatic receipt application.

base location	[In package management] The topmost location that is displayed when a user launches the Machine Identification application.
basket discount	A reduction in price that applies to a group or “basket” of products within a sales order.
basket repricing	A rule that specifies how to calculate and display discounts for a group of products on a sales order. The system can calculate and display the discount as a separate sales order detail line, or it can discount the price of each item on a line-by-line basis within the sales order.
batch job	A job submitted to a system and processed as a single unit with no user interaction.
batch override	An instruction that causes a batch process to produce output other than what it normally would produce for the current execution only.
batch process	A type of process that runs to completion without user intervention after it has been started.
batch program	A program that executes without interacting with the user.
batch version	A version of a report or application that includes a set of user-defined specifications, which control how a batch process runs.
batch/lot tracking	The act of identifying where a component from a specific lot is used in the production of goods.
batch/mix	A manufacturing process that primarily schedules short production runs of products.
batch-of-one processing	A transaction method that allows a client application to perform work on a client workstation, and then submit the work all at once to a server application for further processing. As a batch process is running on the server, the client application can continue performing other tasks. See also direct connect, store-and-forward.
binary large object (BLOB)	A collection of binary data stored as a single entity in a [file].
binder clip	See paper clip.
black products	Products that are derived from the low or heavy end of the distillation process—for example, diesel oils and fuel oils. See also white products.
blend note	Document that authorizes a blending activity, and describes both the ingredients for the blend and the blending steps that occur.
blend off	Reworking off-specification material by introducing a small percentage back into another run of the same product.
blind execution	The mode of execution of a program that does not require the user to review or change the processing options set for the program, and does not require user intervention after the program has been launched.
boleto	In Brazil, the document requesting payment by a supplier or a bank on behalf of a supplier.
bolla doganale	VAT-Only Vouchers for Customs. In Italy, a document issued by the customs

	authority to charge VAT and duties on extra-EU purchasing.
bookmark	A shortcut to a location in a document or a specific place in an application or application suite.
bordero & cheque	In Brazil, bank payment reports.
broker	A program that acts as an intermediary between clients and servers to coordinate and manage requests.
BTL91	In the Netherlands, the ABN/AMRO electronic banking file format that enables batches with foreign automatic payment instructions to be delivered.
budgeted volume	A statement of planned volumes (capacity utilization) upon which budgets for the period have been set.
bunkering	A rate per ton or a sum of money that is charged for placing fuel on board; can also mean the operation itself.
business function	An encapsulated set of business rules and logic that can normally be re-used by multiple applications. Business functions can execute a transaction or a subset of a transaction (check inventory, issue work orders, and so on). Business functions also contain the APIs that allow them to be called from a form, a database trigger, or a non-EnterpriseOne application. Business functions can be combined with other business functions, forms, event rules, and other components to make up an application. Business functions can be created through event rules or third-generation languages, such as C. Examples of business functions include Credit Check and Item Availability.
business function event rule	Encapsulated, reusable business logic that is created by using through event rules rather than C programming. Contrast with embedded event rule. See also event rule.
business object library	[In interoperability] The repository that stores EnterpriseOne business objects, which consist of Java or CORBA objects.
business unit	A financial entity that is used to track the costs, revenue, or both, of an organization. A business unit can also be defined as a branch/plant in which distribution and manufacturing activities occur. Additionally, in manufacturing setup, work centers and production lines must be defined as business units; but these business unit types do not have profit/loss capability.
business view	Used by EnterpriseOne applications to access data from database tables. A business view is a means for selecting specific columns from one or more tables with data that will be used in an application or report. It does not select specific rows and does not contain any physical data. It is strictly a view through which data can be handled.
business view design aid (BDA)	An EnterpriseOne GUI tool for creating, modifying, copying, and printing business views. The tool uses a graphical user interface.
buy-back crude	In foreign producing oil countries, that portion of the host government's share of "participation crude" which it permits the company holding a concession to "buy back."
CAB	In Italy, the bank branch code or branch ID. A five-digit number that identifies any agency of a specific bank company in Italy.

cadastro de pessoas físicas	Cadastro de pessoas físicas. In Brazil, the federal tax ID for a person.
category code	A code that identifies a collection of objects sharing at least one common attribute.
central object	A software component that resides on a central server.
central objects merge	A process that blends a customer's modifications with the objects in a current release with objects in a new release.
central server	A computer that has been designated to contain the originally installed version of the software (central objects) for deployment to client computers.
certificate input	See direct input.
certificate of analysis (COA)	A document that is a record of all of the testing which has been performed against an item, lot, or both, plus the test results for that item and lot.
change management	[In software development] A process that aids in controlling and tracking the evolution of software components.
change order	In PeopleSoft, an addendum to the original purchase order that reflects changes in quantities, dates, or specifications in subcontract-based purchasing. A change order is typically accompanied by a formal notification.
chargeback	A receipt application method that generates an invoice for a disputed amount or for the difference of an unpaid receipt.
chart	EnterpriseOne term for tables of information that appear on forms in the software. See forms.
check-in location	The directory structure location for the package and its set of replicated objects. This location is usually \\deploymentserver\release\path_code\package\packagename. The subdirectories under this path are where the central C components (source, include, object, library, and DLL file) for business functions are stored.
checksum value	A computed value that depends on the contents of a block of data, and that is transmitted or stored with the data to detect whether errors have occurred in the transmission or storage.
class	[In object-oriented programming] A category of objects that share the same characteristics.
clean cargo	Term that refers to cargoes of gasoline and other refined products. See also dirty cargo.
client access	The ability to access data on a server from a client machine.
client machine	Any machine that is connected to a network and that exchanges data with a server.
client workstation	A network computer that runs user application software and is able to request data from a server.
ClieOp03	In the Netherlands, the euro-compliant uniform electronic banking file format that enables batches with domestic automatic direct debit instructions and batches with domestic payment instructions to be delivered.

ClieOp2	In the Netherlands, the uniform electronic banking file format that enables batches with domestic automatic direct debit instructions and batches with domestic payment instructions to be delivered.
cluster	Two or more computers that are grouped together in such a way that they behave like a single computer.
co-existence	A condition where two or more applications or application suites access one or more of the same database tables within the same enterprise.
cold test	The temperature at which oil becomes solid. Generally considered to be 5 degrees F lower than the pour point.
commitment	The number of items that are reserved to fill demand.
common object request broker architecture	An object request broker standard that is endorsed by the Object Management Group.
compa-ratio	An employee's salary divided by the midpoint amount for the employee's pay grade.
component changeout	See component swap.
component object model (COM)	A specification developed by Microsoft for building software components that can be assembled into programs or add functionality to existing programs running on Microsoft Windows platforms. COM components can be written in a variety of languages, although most are written in C++, and can be unplugged from a program at runtime without having to recompile the program.
component swap	In Equipment/Plant Management, the substitution of an operable component for one that requires maintenance. Typically, you swap components to minimize equipment downtime while servicing one of the components. A component swap can also mean the substitution of one parent or component item for another in its associated bill of material.
conference room pilot environment	An EnterpriseOne environment that is used as a staging environment for production data, which includes constants and masters tables such as company constants, fiscal date patterns, and item master. Use this environment along with the test environment to verify that your configuration works before you release changes to end-users.
configurable network computing (CNC)	An application architecture that allows interactive and batch applications that are composed of a single code base to run across a TCP/IP network of multiple server platforms and SQL databases. The applications consist of re-usable business functions and associated data that can be configured across the network dynamically. The overall objective for businesses is to provide a future-proof environment that enables them to change organizational structures, business processes, and technologies independently of each other.
configurable processing engine	Handles all "batch" processes, including reporting, Electronic Data Exchange (EDI) transactions, and data duplication and transformation (for data warehousing). This ability does not mean that it exists only on the server; it can be configured to run on desktop machines (Windows 95 and NT Workstation) as well.
configuration management	A rules-based method of ordering assemble-to-order or make-to-order products in which characteristics of the product are defined as part of the Sales Order Entry

	process. Characteristics are edited by using Boolean logic, and then translated into the components and routing steps that are required to produce the product. The resulting configuration is also priced and costed, based on the defined characteristics.
configured item segment	A characteristic of a configured item that is defined during sales order entry. For example, a customer might specify a type of computer hard drive by stating the number of megabytes of the hard drive, rather than a part number.
consuming location	The point in the manufacturing routing where a component or subassembly is used in the production process. In kanban processing, the location where the kanban container materials are used in the manufacturing process and the kanban is checked out for replenishment.
contra/clearing account	A G/L account used by the system to offset (balance) journal entries. For example, you can use a contra/clearing account to balance the entries created by allocations.
contribution to profit	Selling price of an item minus its variable costs.
control table	A table that controls the program flow or plays a major part in program control.
control table workbench	During the Installation Workbench process, Control Table Workbench runs the batch applications for the planned merges that update the data dictionary, user defined codes, menus, and user overrides tables.
control tables merge	A process that blends a customer's modifications to the control tables with the data that accompanies a new release.
corrective work order	A work order that is used to formally request unscheduled maintenance and communicate all of the details pertaining to the requested maintenance task.
corrective work order	A work order that is used to formally request unscheduled maintenance and communicate all of the details pertaining to the requested maintenance task.
cost assignment	Allocating resources to activities or cost objects.
cost component	An element of an item's cost—for example, material, labor, or overhead.
cost object	Any customer, product, service, contract, project, or other work unit for which you need a separate cost measurement.
cost rollup	A simulated scenario in which work center rates, material costs, and labor costs are used to determine the total cost of an item.
costing elements	The individual classes of added value or conversion costs. These elements are typically materials, such as raw and packaging; labor and machine costs; and overhead, such as fixed and variable. Each corporation defines the necessary detail of product costs by defining and tracking cost categories and subcategories.
credit memo	A negative amount that is used to correct a customer's statement when he or she is overcharged.
credit notice	The physical document that is used to communicate the circumstances and value of a credit order.
credit order	A credit order is used to reflect products or equipment that is received or returned so that it can be viewed as a sales order with negative amounts. Credit orders

	usually add the product back into inventory. This process is linked with delivery confirmation.
cross segment edit	A logic statement that establishes the relationship between configured item segments. Cross segment edits are used to prevent ordering of configurations that cannot be produced.
crude oil assay	A procedure for determining the distillation curve and quality characteristics of a crude oil.
cumulative update	A version of software that includes fixes and enhancements that have been made since the last release or update.
currency relationships	When converting amounts from one currency to another, the currency relationship defines the from currency and the to currency in PeopleSoft software. For example, to convert amounts from German marks to the euro, you first define a currency relationship between those two currencies.
currency restatement	The process of converting amounts from one currency into another currency, generally for reporting purposes. It can be used, for example, when many currencies must be restated into a single currency for consolidated reporting.
current cost	The cost that is associated with an item at the time a parts list and routing are attached to a work order or rate schedule. Current cost is based on the latest bill of material and routing for the item.
customer pricing rules	In Procurement, the inventory pricing rules that are assigned to a supplier. In Sales, inventory pricing rules that are assigned to a customer.
D.A.S. 2 Reporting (DAS 2 or DADS 1)	In France, the name of the official form on which a business must declare fees and other forms of remuneration that were paid during the fiscal year.
data dictionary	A dynamic repository that is used for storing and managing a specific set of data item definitions and specifications.
data source workbench	During the Installation Workbench process, Data Source Workbench copies all of the data sources that are defined in the installation plan from the Data Source Master and Table and Data Source Sizing tables in the Planner data source to the System - release number data source. It also updates the Data Source Plan detail record to reflect completion.
data structure	A description of the format of records in a database such as the number of fields, valid data types, and so on.
data types	Supplemental information that is attached to a company or business unit. Narrative type contains free-form text. Code type contains dates, amounts, and so on.
datagram	A self-contained packet of information that is forwarded by routers, based on their address and the routing table information.
date pattern	A period of time that is set for each period in standard and 52-period accounting and forecasting.
DCE	See distributed computing environment.
DEB	See déclaration d'échange de biens.

debit memo	In Accounts Payable, a voucher that is entered with a negative amount. Enter this type of voucher when a supplier sends you a credit so that you can apply the amount to open vouchers when you issue payment to the supplier.
debit memo	A form that is issued by a customer, requesting an adjustment of the amount, which is owed to the supplier.
debit statement	A list of debit balances.
de-blend	When blend off does not result in a product that is acceptable to customers. The further processing of product to adjust specific physical and chemical properties to within specification ranges. See also blend off.
déclaration d'échange de biens (DEB)	The French term that is used for the Intrastat report.
delayed billing	The invoicing process is delayed until the end of a designated period.
delta load	A batch process that is used to compare and update records between specified environments.
denominated-in currency	The company currency in which financial reports are based.
deployment server	A server that is used to install, maintain, and distribute software to one or more enterprise servers and client workstations.
detail	The specific information that makes up a record or transaction. Contrast with summary.
detail information	Information that primarily relates to individual lines in a sales or purchase order.
direct connect	A transaction method in which a client application communicates interactively and directly with a server application. See also batch-of-one immediate, store-and-forward.
direct input	The system calculates the net units when you enter gross volume, temperature, and gravity or density. This data is generally entered during product receiving from the certificate that is prepared by an independent inspector.
direct ship orders	A purchase order that is issued to a third-party supplier who designates the destination as the customer. A direct ship sales order is also created for the customer. Direct ship orders occur when a product is not available from a company-owned or company-operated source, so the system creates an order to ship the product from a third-party source directly to the customer. Sometimes referred to as a drop ship or third-party supply.
direct usage	Consumption of resources that are attributable to specific production runs because the resources were directly issued to the schedule/order.
director	An EnterpriseOne user interface that guides a user interactively through an EnterpriseOne process.
dirty cargo	Term that refers to crude oil cargoes or other non-refined petroleum cargoes. See also clean cargo.
dispatch planning	Efficient planning and scheduling of product deliveries. Considerations include: Dispatch groups

	Scheduled delivery date Scheduled delivery time Preferred delivery date Preferred delivery time Average delivery time for that geographical location Available resources Special equipment requirements at the product's source or destination.
displacement days	The number of days that are calculated from today's date by which you group vouchers for payment. For example, if today's date is March 10 and you specify three displacement days, the system includes vouchers with a due date through March 13 in the payment group. Contrast with pay-through date.
display sequence	A number that the system uses to re-order a group of records on the form.
distributed computing environment (DCE)	A set of integrated software services that allows software which is running on multiple computers to perform seamless and transparently to the end-users. DCE provides security, directory, time, remote procedure calls, and files across computers running on a network.
distributed data processing	Processing in which some of the functions are performed across two or more linked facilities or systems.
distributed database management system (DDBMS)	A system for distributing a database and its control system across many geographically dispersed machines.
do not translate (DNT)	A type of data source that must exist on the AS/400 because of BLOB restrictions.
double-byte character set (DBCS)	A method of representing some characters by using one byte and other characters by using two bytes. Double-byte character sets are necessary to represent some characters in the Japanese, Korean, and Chinese languages.
downgrade profile	A statement of the hierarchy of allowable downgrades. Includes substitutions of items, and meeting tighter specifications for those products with wider or overlapping specification ranges.
DTA	Datenträgeraustausch. A Swiss payment format that is required by Telekurs (Payserv).
dual pricing	To provide prices for goods and services in two currencies. During the euro transition period, dual pricing between the euro and Economic and Monetary Union (EMU) member currencies is encouraged.
dynamic link library (DLL)	A set of program modules that are designed to be invoked from executable files when the executable files are run, without having to be linked to the executable files. They typically contain commonly used functions.
dynamic partitioning	The ability to dynamically distribute logic or data to multiple tiers in a client/server architecture.
economy of scale	A phenomenon whereby larger volumes of production reduce unit cost by distributing fixed costs over a larger quantity. Variable costs are constant; but fixed costs per unit are reduced, thereby reducing total unit cost.

edit mode	A processing mode or condition where the user can alter the information in a form.
edit rule	A method that is used for formatting user entries, validating user entries, or both, against a predefined rule or set of rules.
embedded event rule	An event rule that is specific to a particular table or application. Examples include form-to-form calls, hiding a field that is based on a processing option value, or calling a business function. Contrast with business function event rule. See also event rule.
employee work center	A central location for sending and receiving all EnterpriseOne messages (system and user-generated), regardless of the originating application or user. Each user has a mailbox that contains workflow and other messages, including Active Messages. With respect to workflow, the Message Center is MAPI compliant and supports drag-and-drop work reassignment, escalation, forward and reply, and workflow monitoring. All messages from the message center can be viewed through EnterpriseOne messages or Microsoft Exchange.
Emulator	An item of software or firmware that allows one device to imitate the functioning of another.
encapsulation	The ability to confine access to and manipulation of data within an object to the procedures that contribute to the definition of that object.
engineering change order (ECO)	A work order document that is used to implement and track changes to items and resulting assemblies. The document can include changes in design, quantity of items required, and the assembly or production process.
enhanced analysis database	A database containing a subset of operational data. The data on the enhanced analysis database performs calculations and provides summary data to speed generation of reports and query response times. This solution is appropriate when external data must be added to source data, or when historical data is necessary for trend analysis or regulatory reporting. See also duplicated database, enterprise data warehouse.
enterprise server	A computer containing programs that collectively serve the needs of an enterprise rather than a single user, department, or specialized application.
EnterpriseOne object	A re-usable piece of code that is used to build applications. Object types include tables, forms, business functions, data dictionary items, batch processes, business views, event rules, versions, data structures, and media objects. See also object.
EnterpriseOne process	Allows EnterpriseOne clients and servers to handle processing requests and execute transactions. A client runs one process, and servers can have multiple instances of a process. EnterpriseOne processes can also be dedicated to specific tasks (for example, workflow messages and data replication) to ensure that critical processes do not have to wait if the server is particularly busy.
EnterpriseOne web development computer	A standard EnterpriseOne Windows developer computer with the additional components installed: Sun's JDK 1.1. JFC (0.5.1). Generator Package with Generator.Java and JDECOM.dll.

	R2 with interpretive and application controls/form.
environment workbench	During the Installation Workbench process, Environment Workbench copies the environment information and Object Configuration Manager tables for each environment from the Planner data source to the System release number data source. It also updates the Environment Plan detail record to reflect completion.
equivalent fuel	A barrel of equivalent fuel supplies six million BTUs of heat. Fuel gas quantities are usually calculated as equivalent fuel barrels in economic calculations for refinery operations.
escalation monitor	A batch process that monitors pending requests or activities, and restarts or forwards them to the next step or user after they have been inactive for a specified amount of time.
ESR	Einzahlungsschein mit Referenznummer. A pay slip with a reference number.
event rule	[In EnterpriseOne] A logic statement that instructs the system to perform one or more operations that are based on an activity that can occur in a specific application, such as entering a form or exiting a field.
exit bar	[In EnterpriseOne] The tall pane with icons in the left portion of many EnterpriseOne program windows.
facility	An entity within a business for which you want to track costs. For example, a facility might be a warehouse location, job, project, work center, or branch/plant. Sometimes referred to as a business unit.
fast path	[In EnterpriseOne] A command prompt that allows the user to move quickly among menus and applications by using specific commands.
file handle	A temporary reference (typically a number) that is assigned to a file which has been opened by the operating system and is used throughout the session to access the file.
file server	A computer that stores files to be accessed by other computers on the network.
find/browse	A type of form used to: Search, view, and select multiple records in a detail area. Delete records. Exit to another form. Serve as an entry point for most applications.
firm planned order (FPO)	A work order that has reached a user defined status. When this status is entered in the processing options for the various manufacturing programs, messages for those orders are not exploded to the components.
fiscal date pattern	A representation of the beginning date for the fiscal year and the ending date for each period in that year.
fix/inspect	A type of form used to view, add, or modify existing records. A fix/inspect form has no detail area.
fixed quantity	A term that indicates the bill of material relationship between a parent item and its components or ingredients. When a bill of material component has a fixed quantity relationship to its parent, the amount of the component does not change

	when the software calculates parts list requirements for different work order quantities. Contrast with variable quantity.
flexible account numbers	The format of account numbers for journal entries. The format that you set up must be the three segments: Business unit. Object. Subsidiary.
form design aid (FDA)	The EnterpriseOne GUI development tool for building interactive applications and forms.
form exit	[In EnterpriseOne] An option that is available as a button on the Form Exit bar or as a selection in the Form menu. It allows users to open an interconnected form.
form interconnection	Allows one form to access and pass data to another form. Form interconnections can be attached to any event; however, they are normally used when a button is clicked.
form type	The following form types are available in EnterpriseOne: Find/browse. Fix/inspect. Header detail. Headerless detail. Message. Parent/child. Search/select.
form-to-form call	A request by a form for data or functionality from one of the connected forms.
framework	[In object-oriented systems] A set of object classes that provide a collection of related functions for a user or piece of software.
frozen cost	The cost of an item, operation, or process after the frozen update program is run; used by the Manufacturing Accounting system.
frozen update program	A program that freezes the current simulated costs, thereby finalizing them for use by the Manufacturing Accounting system.
globally unique identifier (GUI)	A 16-byte code in the Component Object Model that identifies an interface to an object across all computers and networks.
handle	[In programming] A pointer that contains the address of another pointer, which, in turn, contains the address of the desired object.
hard commitment	The number of items that are reserved for a sales order, work order, or both, from a specific location, lot, or both.
hard error	An error that cannot be corrected by a given error detection and correction system.
header	Information at the beginning of a table or form. Header information is used to

	identify or provide control information for the group of records that follows.
header information	Information that pertains to the entire order.
hover help	A help function that provides contextual information or instructions when a cursor moves over a particular part of the interface element for a predefined amount of time.
ICMS	Imposto sobre circulação de mercadoria e serviços. In Brazil, a state tax that is applied to the movement of merchandise and some services.
ICMS Substituto	Imposto sobre circulação de mercadoria e serviços substituto. In Brazil, the ICMS tax that is charged on interstate transactions, or on special products and clients.
ICMS Substituto-Markup	See imposto sobre circulação de mercadoria e serviços substituto-markup.
imposto de renda (IR)	Brazilian income tax.
imposto sobre produtos industrializados	In Brazil, a federal tax that applies to manufactured goods (domestic and imported).
imposto sobre services (ISS)	In Brazil, tax on services.
inbound document	A document that is received from a trading partner using Electronic Data Interface (EDI). This document is also referred to as an inbound transaction.
indented tracing	Tracking all lot numbers of intermediates and ingredients that are consumed in the manufacture of a given lot of product, down through all levels of the bill of material, recipe, or formula.
indexed allocations	A procedure that allocates or distributes expenses, budgets, adjustments, and so on, among business units, based on a fixed percentage.
indirect measurement	Determining the quantity on-hand by: Measuring the storage vessels and calculating the content's balance quantity. or Theoretically calculating consumption of ingredients and deducting them from the on-hand balance.
indirect usage	Determining what should have been used by multiplying receipt quantity of the parent times the quantity per statement in the formula, recipe, or bill of material. This transaction typically affects both consumption on schedule as well as issue from on-hand balances.
in-process rework	Recycling a semi processed product that does not meet acceptable standards. Further processing takes the product out of a given operation and sends it back to the beginning of that operation or a previous operation (for example, unreacted materials). Rework that is detected prior to receipt of finished goods and corrected during the same schedule run.
INPS withholding tax	Istituto Nazionale di Previdenza Sociale withholding tax. In Italy, a 12% social security withholding tax that is imposed on payments to certain types of contractors. This tax is paid directly to the Italian social security office.

inscrição estadual	ICMS tax ID. In Brazil, the state tax ID.
inscrição municipal	ISS tax ID. In Brazil, the municipal tax ID.
integrated toolset	Unique to EnterpriseOne is an industrial-strength toolset that is embedded in the already comprehensive business applications. This toolset is the same toolset that is used by PeopleSoft to build EnterpriseOne interactive and batch applications. Much more than a development environment, however, the EnterpriseOne integrated toolset handles reporting and other batch processes, change management, and basic data warehousing facilities.
integrity test	A process that is used to supplement a company's internal balancing procedures by locating and reporting balancing problems and data inconsistencies.
interbranch sales order	A sales order that is used for transactions between branch/plants other than the selling branch/plant.
Interoperability	The ability of different computer systems, networks, operating systems, and applications to work together and share information.
inventory pricing rule	A discount method that is used for purchases from suppliers and sales to customers. The method is based on effectivity dates, up-to quantities, and a factor by which you can mark up or discount the price or cost.
inventory turn	The number of times that the inventory cycles, or turns over, during the year. A frequently used method to compute inventory turnover is to divide the annual costs of sales by the average inventory level.
invoice	An itemized list of goods that are shipped or services that are rendered, stating quantities, prices, fees, shipping charges, and so on. Companies often have their invoices mailed to a different address than where they ship products. In such cases, the bill-to address differs from the ship-to address.
IP	See imposto sobre produtos industrializados.
IR	See imposto de renda.
IServer Service	Developed by PeopleSoft, this Internet server service resides on the Web server and is used to speed up delivery of the Java class files from the database to the client.
ISS	See imposto sobre serviços.
jargon	An alternate data dictionary item description that EnterpriseOne or PeopleSoft World displays, based on the product code of the current object.
java application server	A component-based server that resides in the middle-tier of a server-centric architecture and provides middleware services for security and state maintenance, along with data access and persistence.
JDBNET	A database driver that allows heterogeneous servers to access each other's data.
jde.ini	A PeopleSoft file (or member for AS/400) that provides the runtime settings that are required for EnterpriseOne initialization. Specific versions of the file or member must reside on every machine that is running EnterpriseOne, including workstations and servers.
JDE.LOG	The main diagnostic log file of EnterpriseOne. Always located in the root

	directory on the primary drive. Contains status and error messages from the startup and operation of EnterpriseOne.
JDEBASE Database Middleware	PeopleSoft proprietary database middleware package that provides two primary benefits: 1. Platform-independent APIs for multidatabase access. These APIs are used in two ways: a. By the interactive and batch engines to dynamically generate platform-specific SQL, depending on the data source request. b. As open APIs for advanced C business function writing. These APIs are then used by the engines to dynamically generate platform-specific SQL. 2. Client-to-server and server-to-server database access. To accomplish this access, EnterpriseOne is integrated with a variety of third-party database drivers, such as Client Access 400 and open database connectivity (ODBC).
JDECallObject	An application programming interface that is used by business functions to invoke other business functions.
JDEIPC	Communications programming tools that are used by server code to regulate access to the same data in multiprocess environments, communicate and coordinate between processes, and create new processes.
JDENET	PeopleSoft proprietary middleware software. JDENET is a messaging software package.
JDENET communications middleware	PeopleSoft proprietary communications middleware package for EnterpriseOne. It is a peer-to-peer, message-based, socket-based, multiprocess communications middleware solution. It handles client-to-server and server-to-server communications for all EnterpriseOne supported platforms.
just in time installation (JITI)	EnterpriseOne's method of dynamically replicating objects from the central object location to a workstation.
just in time replication (JITR)	EnterpriseOne's method of replicating data to individual workstations. EnterpriseOne replicates new records (inserts) only at the time that the user needs the data. Changes, deletes, and updates must be replicated using Pull Replication.
Kagami	In Japan, summarized invoices that are created monthly (in most cases) to reduce the number of payment transactions.
latitude	The X coordinate of the location of an item in the warehouse. The system can use latitude, longitude, and height when suggesting locations for putaway, replenishment, and picking.
laytime (or layhours)	The amount of time that is allotted to a tanker at berth to complete loading or discharging cargo. This time is usually expressed in running hours, and is fixed by prior agreement between the vessel owner and the company that is chartering the vessel. Laytime is stipulated in the charter, which states exactly the total of number of hours that are granted at both loading and unloading ports, and indicates whether such time is reversible. A statement of "Seventy-Two Hours, Reversible" means that a total of 72 hours is granted overall at both ports, and any time saved at one port can be applied as a credit at the other port. For example, if the vessel uses only 32 hours instead of 36 hours to load cargo, it can apply an additional four hours to the 36 hours allotted at the discharge port.

	Such considerations are important for purposes of computing demurrage.
leading zeros	A series of zeros that certain facilities in PeopleSoft systems place in front of a value that is entered. This situation normally occurs when you enter a value that is smaller than the specified length of the field. For example, if you enter 4567 in a field that accommodates eight numbers, the facility places four zeros in front of the four numbers that you enter. The result appears as 00004567.
ledger type	A code that designates a ledger which is used by the system for a particular purpose. For example, all transactions are recorded in the AA (actual amounts) ledger type in their domestic currency. The same transactions can also be stored in the CA (foreign currency) ledger type.
level break	The position in a report or text where a group of similar types of information ends and another one begins.
libro IVA	Monthly VAT report. In Italy, the term for the report that contains the detail of invoices and vouchers that were registered during each month.
line of business	A description of the nature of a company's work; also a tool to control the relationship with that customer, including product pricing.
linked service type	A service type that is associated with a primary service type. Linked service types can be cancelled, and the maintenance tasks are performed when the primary service type to which they are linked comes due. You can specify whether the system generates work orders for linked service types, as well as the status that the system assigns to work orders that have already been generated. Sometimes referred to as associated service types. See also primary service type and service type.
livro razao	In Brazil, a general ledger report.
load balancing	The act of distributing the number of processes proportionally to all servers in a group to maximize overall performance.
location workbench	During the Installation Workbench process, Location Workbench copies all locations that are defined in the installation plan from the Location Master table in the Planner data source to the System data source.
log files	Files that track operations for a process or application. Reviewing log files is helpful for troubleshooting problems. The file extension for log files is .LOG.
logic data source	Any code that provides data during runtime.
logical compartment	One of two ways that is identified in the transportation constants to display compartments on vehicles. Logical display numbers the compartments sequentially. For example, if two vehicles are on a trip and each vehicle has three compartments, the logical display is 1,2,3,4,5,6.
logical file	A set of keys or indices that is used for direct access or ordered access to the records in a physical file. Several logical files can have different accesses to a physical.
logical shelf	A logical, not physical, location for inventory that is used to track inventory transactions in loan/borrow, or exchange agreements with other companies. See also logical warehouse.

logical warehouse	Not a physical warehouse containing actual inventory, but a means for storing and tracking information for inventory transactions in loan/borrow, or exchange agreements with other companies.
longitude	The Y coordinate of the location of an item in the warehouse. The system can use latitude, longitude, and height when suggesting locations for putaway, replenishment, and picking.
LSV	Lastschriftverfahren. A Swiss auto debit format that is required by Telekurs (Payserv).
mail merge	A mass-mail facility that takes names, addresses, and (sometimes) pertinent facts about recipients and merges the information into a form letter or a similarly basic document.
mailmerge workbench	[In EnterpriseOne] An application that merges Microsoft Word 6.0 (or higher) word-processing documents with EnterpriseOne records to automatically print business documents.
main fuels	Usually refers to bulk fuel products, but sometimes includes packaged products.
maintenance loop	See maintenance route.
maintenance route	A method of performing PMs for multiple pieces of equipment from a single preventive maintenance work order. A maintenance route includes pieces of equipment that share one or more identical maintenance tasks which can be performed at the same time for each piece of equipment. Sometimes referred to as maintenance loop.
maintenance work order	In PeopleSoft EnterpriseOne systems, a term that is used to distinguish work orders created for the performance of equipment and plant maintenance from other work orders, such as manufacturing work orders, utility work orders, and engineering change orders.
manufacturing and distribution planning	Planning that includes resource and capacity planning, and material planning operations. Resource and capacity planning allows you to prepare a feasible production schedule that reflects your demand forecasts and production capability. Material Planning Operations provides a short-range plan to cover material requirements that are needed to make a product.
mapping	A set of instructions that describes how one data structure passes data to another.
master business function	An interactive master file that serves as a central location for adding, changing, and updating information in a database.
master business function	A central system location for standard business rules about entering documents, such as vouchers, invoices, and journal entries. Master business functions ensure uniform processing according to guidelines that you establish.
master table	A database table that is used to store data and information that is permanent and necessary to the system's operation. Master tables might contain data such as paid tax amounts, supplier names, addresses, employee information, and job information.
matching document	A document that is associated with an original document to complete or change a transaction. For example, a receipt is the matching document of an invoice.

media object	An electronic or digital representation of an object.
media storage objects	Files that use one of the following naming conventions that are not organized into table format: Gxxx, xxxGT, or GTxxx.
memory violation	An error that occurs as the result of a memory leak.
menu selection	An option on a menu that initiates a software function directly.
message center	A central location for sending and receiving all EnterpriseOne messages (system- and user-generated), regardless of the originating application or user.
messaging application programming interface (MAPI)	An architecture that defines the components of a messaging system and how they behave. It also defines the interface between the messaging system and the components.
metal content	A series of properties of a blended product that help to determine its suitability for a prescribed purpose.
metals management	The process of maintaining information about the location and status of durable product containers such as liquid petroleum gas (LPG) cylinders.
mobile inventory	Inventory that is transferred from a depot to a barge or truck for milk-run deliveries.
modal	A restrictive or limiting interaction that is created by a given condition of operation. Modal often describes a secondary window that restricts a user's interaction with other windows. A secondary window can be modal with respect to its primary window or to the entire system. A modal dialog box must be closed by the user before the application continues.
model work order	For scheduled preventive maintenance or for a condition-based alert, a model work order functions as a template for the creation of other work orders. You can assign model work orders to service types and condition-based alerts. When the service type comes due or the alert is generated, the system automatically generates a work order that is based on information from the model work order.
modeless	Not restricting or limiting interaction. Modeless often describes a secondary window that does not restrict a user's interaction with other windows. A modeless dialog box stays on the screen and is available for use at any time, but also permits other user activities.
multiple stocking locations	Authorized storage locations for the same item number at locations, in addition to the primary stocking location.
multitier architecture	A client/server architecture that allows multiple levels of processing. A tier defines the number of computers that can be used to complete some defined task.
named event rules (NER)	Also called business function event rules. Encapsulated, re-usable business logic that is created by using event rules, rather than C programming.
national language support (NLS)	Mechanisms that are provided to facilitate internationalization of both system and application user interfaces.
natureza da operação	Transaction nature. In Brazil, a code that classifies the type of commercial transaction to conform to the fiscal legislation.
negative pay item	An entry in an account that indicates a prepayment. For example, you might

	prepay a supplier before goods are sent or prepay an employee's forecasted expenses for a business trip. The system stores these pending entries, assigning them a minus quantity as debit amounts in a designated expense account. After the prepaid goods are received or the employee submits an expense report, entering the actual voucher clears all of the negative pay items by processing them as regular pay items. Note that a negative pay item can also result from entering a debit memo (A/P) or a credit memo (A/R).
net added cost	The cost to manufacture an item at the current level in the bill of material. Thus, for manufactured parts, the net added cost includes labor, outside operations, and cost extras applicable to this level in the bill of material, but not materials (lower-level items). For purchased parts, the net added cost also includes the cost of materials.
next status	The next step in the payment process for payment control groups. The next status can be either WRT (write) or UPD (update).
node	A termination point for two or more communications links. A node can serve as the control location for forwarding data among the elements of a network or multiple networks, as well as performing other networking and, in some cases, local processing.
non-inventory items	See non-stock items.
non-list price	A price for bulk products that is determined by its own algorithms, such as a rolling average or commodity price plus.
non-prime product	A manufactured product with revenue potential that is less than the product planned for, or scheduled to be produced.
non-stock items	Items that the system does not account for as part of the inventory. For example, office supplies, or packaging materials can be non-stock items.
nota fiscal	In Brazil, a legal document that must accompany all commercial transactions.
nota fiscal fatura	In Brazil, a nota fiscal and invoice information.
notula	In Italy, the process whereby a business does not recognize value added tax until the payment of a voucher.
object configuration manager (OCM)	EnterpriseOne's object request broker and the control center for the runtime environment. It keeps track of the runtime locations for business functions, data, and batch applications. When one of these objects is called, the Object Configuration Manager directs access to it by using defaults and overrides for a given environment and user.
object embedding	When an object is embedded in another document, an association is maintained between the object and the application that created it; however, any changes made to the object are also only kept in the compound document. See also object linking.
object librarian	A repository of all versions, applications, and business functions that are reusable in building applications.
object linking	When an object is linked to another document, a reference is created with the file in which the object is stored, as well as with the application that created it. When the object is modified, either from the compound document or directly through the file in which it is saved, the change is reflected in that application as well as

	anywhere it has been linked. See also object embedding.
object linking and embedding (OLE)	A technology for transferring and sharing information among applications by allowing the integration of objects from diverse applications, such as graphics, charts, spreadsheets, text, or an audio clip from a sound program. OLE is a compound document standard that was developed by Microsoft Corporation. It enables you to create objects with one application, and then link or embed them in a second application. Embedded objects retain their original format and links to the application that created them. See also object embedding, object linking.
object management workbench (OMW)	The change management system that is used for EnterpriseOne development.
object-based technology (OBT)	A technology that supports some of the main principles of object-oriented technology: Classes. Polymorphism.I Inheritance. Encapsulation.
object-oriented technology (OOT)	Brings software development past procedural programming into a world of re-usable programming that simplifies development of applications. Object orientation is based on the following principles: Classes. Polymorphism.I Inheritance. Encapsulation.
offsetting account	An account that reduces the amount of another account to provide a net balance. For example, a credit of 200 to a cash account might have an offsetting entry of 200 to an A/P Trade (liability) account.
open database connectivity (ODBC)	Defines a standard interface for different technologies to process data between applications and different data sources. The ODBC interface comprises set of function calls, methods of connectivity, and representation of data types that define access to data sources.
open systems interconnection (OSI)	The OSI model was developed by the International Standards Organization (ISO) in the early 1980s. It defines protocols and standards for the interconnection of computers and network equipment.
order detail line	A part of an order that contains transaction information about a service or item being purchased or sold, such as quantity, cost, price, and so on.
order hold	A flag that stops the processing of an order because it has exceeded the credit or budget limit, or has another problem.
order-based pricing	Pricing strategy that grants reductions in price to a customer. It is based upon the contents and relative size (volume or value) of the order as a whole.
outbound document	A document that is sent to a trading partner using EDI. This term is also referred to as an outbound transaction.
outturn	The quantity of oil that is actually received into a buyer's storage tanks when a

	vessel is unloaded. For various reasons (vaporization, clingage to vessel tank walls, and so on), the amount of a product pumped into shore tankage at unloading is often less than the quantity originally loaded onto the vessel, as certified by the Bill of Lading. Under a delivered or CIF outturn transaction, the buyer pays only for the barrels actually “turned out” by the vessel into storage. When a buyer is paying CIF Bill of Lading figures, a loss of 0.5% of total cargo volume is considered normal. Losses in excess of 0.5%, however, are either chargeable to the seller or are covered by specialized insurance that covers partial, as well as total, loss of the cargo.
overhead	In the distillation process, that portion of the charge that leaves the top of the distillation column as vapor. This definition is strictly as it relates to ECS.
override conversion method	A method of calculating exchange rates that is set up between two specific currencies. For those specific currencies, this method overrides the conversion method in General Accounting Constants and does not allow inverse rates to be used when calculating currency amounts.
package / package build	A collection of software that is grouped into a single entity for modular installation. EnterpriseOne objects are installed to workstations in packages from the deployment server. A package can be compared to a bill of material or kit that indicates the necessary objects for that workstation and where the installation program can find them on the deployment server. It is a point-in-time “snapshot” of the central objects on the deployment server.
package location	The directory structure location for the package and its set of replicated objects. This location is usually \\deployment server\release\path_code\package\ package name. The replicated objects for the package are placed in the subdirectories under this path. This location is also where the package is built or stored.
package workbench	During the Installation Workbench process, Package Workbench transfers the package information tables from the Planner data source to the System - release number data source. It also updates the Package Plan detail record to reflect completion.
packaged products	Products that, by their nature, must be delivered to the customer in containers which are suitable for discrete consumption or resale.
pane/panel	A resizable subarea of a window that contains options, components, or other related information.
paper clip	An icon that is used to indicate that a media object is attached to a form or record.
parent/child form	A type of form that presents parent/child relationships in an application on one form: The left portion of the form presents a tree view that displays a visual representation of a parent/child relationship. The right portion of the form displays a detail area in browse mode. The detail area displays the records for the child item in the tree. The parent/child form supports drag and drop functionality.
parent/child relationship	See parent/component relationship.
parent/component relationship	1. In Capital Asset Management, the hierarchical relationship of a parent piece of equipment to its components. For example, a manufacturing line could be a parent

	and the machinery on the line could be components of the line. In addition, each piece of machinery could be a parent of still more components. 2. In Product Data Management, a hierarchical relationship of the components and subassemblies of a parent item to that parent item. For example, an automobile is a parent item; its components and subassemblies include: engine, frame, seats, and windows. Sometimes referred to as parent/child relationship.
partita IVA	In Italy, a company fiscal identification number.
pass-through	A process where data is accepted from a source and forwarded directly to a target without the system or application performing any data conversion, validation, and so on.
pay on consumption	The method of postponing financial liability for component materials until you issue that material to its consuming work order or rate schedule.
payment group	A system-generated group of payments with similar information, such as a bank account. The system processes all of the payments in a payment group at the same time.
PeopleSoft database	See JDEBASE Database Middleware.
performance tuning	The adjustments that are made for a more efficient, reliable, and fast program.
persistent object	An object that continues to exist and retains its data beyond the duration of the process that creates it.
pervasive device	A type of intelligent and portable device that provides a user with the ability to receive and gather information anytime, from anywhere.
planning family	A means of grouping end items that have similarity of design or manufacture.
plug-in	A small program that plugs into a larger application to provide added functionality or enhance the main application.
polymorphism	A principle of object-oriented technology in which a single mnemonic name can be used to perform similar operations on software objects of different types.
portal	A Web site or service that is a starting point and frequent gateway to a broad array of on-line resources and services.
Postfinance	A subsidiary of the Swiss postal service. Postfinance provides some banking services.
potency	Identifies the percent of an item in a given solution. For example, you can use an 80% potent solution in a work order that calls for 100% potent solution, but you would use 25% more, in terms of quantity, to meet the requirement ($100 / 80 = 1.25$).
preference profile	The ability to define default values for specified fields for a user defined hierarchy of items, item groups, customers, and customer groups. In Quality Management setup, this method links test and specification testing criteria to specific items, item groups, customers, or customer groups.
preflush	A work order inventory technique in which you deduct (relieve) materials from inventory when the parts list is attached to the work order or rate schedule.

preventive maintenance cycle	The sequence of events that make up a preventive maintenance task, from its definition to its completion. Because most preventive maintenance tasks are commonly performed at scheduled intervals, parts of the preventive maintenance cycle repeat, based on those intervals.
preventive maintenance schedule	The combination of service types that apply to a specific piece of equipment, as well as the intervals at which each service type is scheduled to be performed.
primary service type	A service type to which you can link related service types. For example, for a particular piece of equipment, you might set up a primary service type for a 1000-hour inspection and a linked service type for a 500-hour inspection. The 1000-hour inspection includes all of the tasks performed at 500 hours. When a primary service type is scheduled to be performed, the system schedules the linked service type. See also linked service type.
pristine environment	An EnterpriseOne environment that is used to test unaltered objects with PeopleSoft demonstration data or for training classes. You must have this environment so you can compare pristine objects that you modify.
processing option	A data structure that allows users to supply parameters that regulate the execution of a batch program or report.
product data management (PDM)	In PeopleSoft EnterpriseOne software, the system that enables a business to organize and maintain information about each item which it manufactures. Features of this system, such as bills of material, work centers, and routings, define the relationships among parents and components, and how they can be combined to manufacture an item. PDM also provides data for other manufacturing systems including Manufacturing Accounting, Shop Floor Management, and Manufacturing and Distribution Planning.
product line	A group of products with similarity in manufacturing procedures, marketing characteristics, or specifications that allow them to be aggregated for planning; marketing; and, occasionally, costing.
product/process definition	A combination of bill of material (recipe, formula, or both) and routing (process list). Organized into tasks with a statement of required consumed resources and produced resources.
production environment	An EnterpriseOne environment in which users operate EnterpriseOne software.
program temporary fix (PTF)	A representation of changes to PeopleSoft software that your organization receives on magnetic tapes or diskettes.
project	[In EnterpriseOne] A virtual container for objects being developed in Object Management Workbench.
projected cost	The target expenditure in added value for material, labor, and so on, during manufacture. See also standard cost.
promotion path	The designated path for advancing objects or projects in a workflow.
protocollo	See registration number.
PST	Provincial sales tax. A tax that is assessed by individual provinces in Canada.
published table	Also called a “Master” table, this is the central copy to be replicated to other machines and resides on the “publisher” machine. The Data Replication Publisher Table (F98DRPUB) identifies all of the published tables and their associated

	publishers in the enterprise.
publisher	The server that is responsible for the published table. The Data Replication Publisher Table (F98DRPUB) identifies all of the published tables and their associated publishers in the enterprise.
pull replication	One of the EnterpriseOne methods for replicating data to individual workstations. Such machines are set up as pull subscribers that use EnterpriseOne's data replication tools. The only time that pull subscribers are notified of changes, updates, and deletions is when they request such information. The request is in the form of a message that is sent, usually at startup, from the pull subscriber to the server machine that stores the Data Replication Pending Change Notification table (F98DRPCN).
query by example (QBE)	Located at the top of a detail area, this area is used to search for data to display in the detail area.
rate scheduling	A method of scheduling product or manufacturing families, or both. Also a technique to determine run times and quantities of each item within the family to produce enough of each individual product to satisfy demand until the family can be scheduled again.
rate type	For currency exchange transactions, the rate type distinguishes different types of exchange rates. For example, you can use both period average and period-end rates, distinguishing them by rate type.
real-time	Pertaining to information processing that returns a result so rapidly that the interaction appears to be instantaneous.
receipt routing	A series of steps that is used to track and move items within the receipt process. The steps might include in-transit, dock, staging area, inspection, and stock.
referential integrity	Ensures that a parent record cannot be deleted from the database when a child record for exists.
regenerable	Source code for EnterpriseOne business functions can be regenerated from specifications (business function names). Regeneration occurs whenever an application is recompiled, either for a new platform or when new functionality is added.
register types and classes	In Italian VAT Summary Reporting, the classification of VAT transactions.
relationship	Links tables together and facilitates joining business views for use in an application or report. Relationships that are created are based on indexes.
relevé d'identité bancaire (RIB)	In France, the term that indicates the bank transit code, account number, and check digit that are used to validate the bank transit code and account number. The bank transit code consists of the bank code and agency code. The account number is alphanumeric and can be as many as 11 characters. PeopleSoft supplies a validation routine to ensure RIB key correctness.
remessa	In Brazil, the remit process for A/R.
render	To include external data in displayed content through a linking mechanism.
repassé	In Brazil, a discount of the ICMS tax for interstate transactions. It is the adjustment between the interstate and the intrastate ICMS tax rates.

replenishment point	The location on or near the production line where additional components or subassemblies are to be delivered.
replication server	A server that is responsible for replicating central objects to client machines.
report design aid (RDA)	The EnterpriseOne GUI tool for operating, modifying, and copying report batch applications.
repost	In Sales, the process of clearing all commitments from locations and restoring commitments, based on quantities from the Sales Order Detail table (F4211).
resident	Pertaining to computer programs or data while they remain on a particular storage device.
retorno	In Brazil, the receipt process for A/R.
RIB	See relevé d'identité bancaire.
ricevute bancarie (RiBa)	In Italy, the term for accounts receivable drafts.
riepilogo IVA	Summary VAT monthly report. In Italy, the term for the report that shows the total amount of VAT credit and debit.
ritenuta d'acconto	In Italy, the term for standard withholding tax.
rollback	[In database management] A feature or command that undoes changes in database transactions of one or more records.
rollup	See cost rollup.
row exit	[In EnterpriseOne] An application shortcut, available as a button on the Row Exit bar or as a menu selection, that allows users to open a form that is related to the highlighted grid record.
runtime	The period of time when a program or process is running.
SAD	The German name for a Swiss payment format that is accepted by Postfinance.
SAR	See software action request.
scalability	The ability of software, architecture, hardware, or a network to support software as it grows in size or resource requirements.
scripts	A collection of SQL statements that perform a specific task.
scrub	To remove unnecessary or unwanted characters from a string.
search/select	A type of form that is used to search for a value and return it to the calling field.
selection	Found on PeopleSoft menus, selections represent functions that you can access from a menu. To make a selection, type the associated number in the Selection field and press Enter.
serialize	To convert a software object into a stream of bytes to store on a disk or transfer across a network.
server map	The server view of the object configuration mapping.
server workbench	During the Installation Workbench process, Server Workbench copies the server

	configuration files from the Planner data source to the System release number data source. It also updates the Server Plan detail record to reflect completion.
service interval	The frequency at which a service type is to be performed. Service intervals can be based on dates, periods, or statistical units that are user defined. Examples of statistical units are hours, miles, and fuel consumption.
service type	An individual preventive maintenance task or procedure, such as an inspection, lubrication, or overhaul. Service types can apply to a specific piece of equipment or to a class of equipment. You can specify that service types come due based on a predetermined service interval, or whenever the task that is represented by the service type becomes necessary.
servlet	A [small] program that extends the functionality of a Web server by generating dynamic content and interacting with Web clients by using a request-response paradigm.
share path	The network node under which one or more servers or objects reside.
shop floor management	A system that uses data from multiple system codes to help develop, execute, and manage work orders and rate schedules in the enterprise.
silent mode	A method for installing or running a program that does not require any user intervention.
silent post	A type of post that occurs in the background without the knowledge of the user.
simulated cost	After a cost rollup, the cost of an item, operation, or process according to the current cost scenario. This cost can be finalized by running the frozen update program. You can create simulated costs for a number of cost methods—for example, standard, future, and simulated current costs. See also cost rollup.
single-byte character set (SBCS)	An encoding scheme in which each alphabetic character is represented by one byte. Most Western languages, such as English, can be represented by using a single-byte character set.
single-level tracking	Finding all immediate parents where a specific lot has been used (consumed).
single-voyage (spot) charter	An agreement for a single voyage between two ports. The payment is made on the basis of tons of product delivered. The owner of the vessel is responsible for all expenses.
slimer	A script that changes data in a table directly without going through a regular database interface.
smart field	A data dictionary item with an attached business function for use in the Report Design Aid application.
SOC	The Italian term for a Swiss payment format that is accepted by Postfinance.
soft commitment	The number of items that is reserved for sales orders or work orders in the primary units of measure.
soft error	An error from which an operating system or program is able to recover.
software action request (SAR)	An entry in the AS/400 database that is used for requesting modifications to PeopleSoft software.

SOG	The French term for a Swiss payment format that is accepted by Postfinance.
source directory	The path code to the business function source files belonging to the shared library that is created on the enterprise server.
special period/year	The date that determines the source balances for an allocation.
specification merge	The Specification merge is comprised of three merges: Object Librarian merge (via the Object Management Workbench). Versions List merge. Central Objects merge. The merges blend customer modifications with data that accompanies a new release.
specification table merge workbench	During the Installation Workbench process, Specification Table Merge Workbench runs the batch applications that update the specification tables.
specifications	A complete description of an EnterpriseOne object. Each object has its own specification, or name, which is used to build applications.
spot charter	See single-voyage charter.
spot rates	An exchange rate that is entered at the transaction level. Spot rates are not used on transactions between two EMU member currencies because exchange rates are irrevocably fixed to the euro.
stamp tax	In Japan, a tax that is imposed on drafts payable, receipts over 30000 Japanese yen, and all contracts. The party that issues any of the above documents is responsible for this tax.
standalone	Operating or capable of operating independently of certain other components of a computer system.
standard cost	The expected, or target cost of an item, operation, or process. Standard costs represent only one cost method in the Product Costing system. You can also calculate, for example, future costs or current costs. However, the Manufacturing Accounting system uses only standard frozen costs.
standard costing	A costing method that uses cost units that are determined before production. For management control purposes, the system compares standard costs to actual costs and computes variances.
subprocess	A process that is triggered by and is part of a larger process, and that generally consists of activities.
subscriber table	The Subscriber table (F98DRSUB), which is stored on the Publisher Server with the Data Replication Publisher table (F98DRPUB), that identifies all of the subscriber machines for each published table.
summary	The presentation of data or information in a cumulative or totaled manner in which most of the details have been removed. Many systems offer forms and reports that summarize information which is stored in certain tables. Contrast with detail.
super backflush	To create backflush transactions for material, labor, or both, against a work order at predefined pay points in the routing. By doing so, you can relieve inventory

	and account for labor amounts at strategic points throughout the manufacturing process.
supersession	Specification that a new product is replacing an active product on a specified effective date.
supplemental data	Additional types of data for customers and suppliers. You can enter supplemental data for information such as notes, comments, plans, or other information that you want in a customer or supplier record. The system maintains this data in generic databases, separate from the standard master tables (Customer Master, Supplier Master, and Address Book Master).
supplying location	The location from which inventory is transferred once quantities of the item on the production line have been depleted. In kanban processing, the supplying location is the inventory location from which materials are transferred to the consuming location when the containers are replenished.
system code	A numeric or alphanumeric designation that identifies a specific system in EnterpriseOne software.
system function	[In EnterpriseOne] A named set of pre-packaged, re-usable instructions that can be called from event rules.
table access management (TAM)	The EnterpriseOne component that handles the storage and retrieval of user defined data. TAM stores information such as data dictionary definitions; application and report specifications; event rules; table definitions; business function input parameters and library information; and data structure definitions for running applications, reports, and business functions.
table conversion workbench	During the Installation Workbench process, Table Conversion Workbench runs the table conversions that change the technical and application tables to the format for the new release of EnterpriseOne. It also updates the Table Conversions and Controls detail records to reflect completion.
table design aid (TDA)	An EnterpriseOne GUI tool for creating, modifying, copying, and printing database tables.
table event rules	Use table event rules to attach database triggers (or programs) that automatically run whenever an action occurs against the table. An action against a table is referred to as an event. When you create an EnterpriseOne database trigger, you must first determine which event will activate the trigger. Then, use Event Rules Design to create the trigger. Although EnterpriseOne allows event rules to be attached to application events, this functionality is application-specific. Table event rules provide embedded logic at the table level.
table handle	A pointer into a table that indicates a particular row.
table space	[In relational database management systems] An abstract collection of containers in which database objects are stored.
task	[In Solution Explorer and EnterpriseOne Menu] A user defined object that can initiate an activity, process, or procedure.
task view	A group of tasks in Solution Explorer or EnterpriseOne Menu that are arranged in a tree structure.
termo de abertura	In Brazil, opening terms for the transaction journal.

termo de encerramento	In Brazil, closing terms for the transaction journal.
three-tier processing	The task of entering, reviewing, approving, and posting batches of transactions.
three-way voucher match	The process of comparing receipt information to supplier's invoices to create vouchers. In a three-way match, you use the receipt records, the purchase order, and the invoice to create vouchers.
threshold percentage	In Capital Asset Management, the percentage of a service interval that you define as the trigger for maintenance to be scheduled. For example, you might set up a service type to be scheduled every 100 hours with a threshold percentage of 90 percent. When the equipment accumulates 90 hours, the system schedules the maintenance.
throughput agreement	A service agreement in which a business partner agrees to store and manage product for another business partner for a specified time period. The second partner actually owns the stock that is stored in the first partner's depot, although the first partner monitors the stock level; suggests replenishments; and unloads, stores, and delivers product to the partner or its customers. The first partner charges a fee for storing and managing the product.
throughput reconciliation	Reconcile confirmed sales figures in a given period with the measured throughput, based on the meter readings. This process is designed to catch discrepancies that are due to transactions not being entered, theft, faulty meters, or some combination of these factors. This reconciliation is the first stage. See also operational reconciliation.
token	[In Object Management Workbench] A flag that is associated with each object which indicates whether you can check out the object.
tolerance range	The amount by which the taxes that you enter manually can vary from the tax that is calculated by the system.
TP monitor	Transaction Processing monitor. A monitor that controls data transfer between local and remote terminals and the applications that originated them. TP monitors also protect data integrity in the distributed environment and can include programs that validate data and format terminal screens.
tracing	The act of researching a lot by going backward, to discover its origin.
tracking	The act of researching a lot by going forward, to discover where it is used.
transaction set	An electronic business transaction (EDI Standard document) composed of segments.
transclude	To include the external data in the displayed content through a linking mechanism.
transfer order	An order that is used to ship inventory between branch/plants within your company and to maintain an accurate on-hand inventory amount. An interbranch transfer order creates a purchase order for the shipping location and a sales order for the receiving location.
translation adjustment account	An optional G/L account used in currency balance restatement to record the total adjustments at a company level.
translator software	The software that converts data from an application table format to an EDI Standard Format, and from EDI Standard Format to application table format. The

	data is exchanged in an EDI Standard, such as ANSI ASC X12, EDIFACT, UCS, or WINS.
tree structure	A type of graphical user interface that displays objects in a hierarchy.
trigger	Allows you to attach default processing to a data item in the data dictionary. When that data item is used on an application or report, the trigger is invoked by an event which is associated with the data item. EnterpriseOne also has three visual assist triggers: Calculator. Calendar. Search form.
two-way voucher match	The process of comparing purchase order detail lines to the suppliers' invoices to create vouchers. You do not record receipt information.
universal batch engine (UBE)	[In EnterpriseOne] A type of application that runs a noninteractive process.
unnormalized	Data that is a random collection of data elements with repeating record groups scattered throughout. Also see Normalized.
user overrides merge	The User Overrides merge adds new user override records into a customer's user override table.
user-defined code (UDC)	A value that a user has assigned as being a valid entry for a given or specific field.
utility	A small program that provides an addition to the capabilities which are provided by an operating system.
variable numerator allocations	A procedure that allocates or distributes expenses, budgets, adjustments, and so on, among business units, based on a variable.
variable quantity	A term that indicates the bill of material relationship between a parent item and its components or ingredients. When a bill of material component has a variable quantity relationship to its parent, the amount of the component changes when the software calculates parts list requirements for different work order quantities. Contrast with fixed quantity.
variance	1. In Product Costing and Manufacturing Accounting, the difference between the frozen standard cost, the current cost, the planned cost, and the actual cost. For example, the difference between the frozen standard cost and the current cost is an engineering variance. Frozen standard costs come from the Cost Components table, and the current costs are calculated by using the current bill of material, routing, and overhead rates. 2. In Capital Asset Management, the difference between revenue that is generated by a piece of equipment and costs that are incurred by the equipment.
versions list merge	The Versions List merge preserves any non-XJDE and non-ZJDE version specifications for objects that are valid in the new release as well as their processing options data.
VESR	Verfahren Einzahlungsschein mit Referenznummer. The processing of an ESR pay slip with reference line through accounts receivable and accounts payable.
visual assist	Forms that can be invoked from a control to assist the user in determining what

	data belongs in the control.
voucher logging	The process of entering vouchers without distributing amounts to specific G/L accounts. The system initially distributes the total amount of each voucher to a G/L suspense account, where it is held until you redistribute it to the correct G/L account.
wareki date format	In Japan, a calendar format, such as Showa or Heisei. When a new emperor begins to reign, the government chooses the title of the date format and the year starts over at one. For instance, January 1, 1998, is equal to Heisei 10, January 1st.
wash down	A minor cleanup between similar product runs. Sometimes used in reference to the sanitation process of a food plant.
wchar_t	An internal type of a wide character. Used for writing portable programs for international markets.
web server	A server that sends information as requested by a browser and uses the TCP/IP set of protocols.
work order life cycle	In Capital Asset Management, the sequence of events through which a work order must pass to accurately communicate the progress of the maintenance tasks that it represents.
workfile	A system-generated file that is used for temporary data processing.
workflow	According to the Workflow Management Coalition, workflow means “the automation of a business process, in whole or part, during which documents, information, or tasks are passed from one participant to another for action, according to a set of procedural rules.”
workgroup server	A network server usually containing subsets of data that are replicated from a master network server.
WorldSoftware architecture	The broad spectrum of application design and programming technology that PeopleSoft uses to achieve uniformity, consistency, and complete integration throughout its software.
write payment	A step in processing payments. Writing payments includes printing checks, drafts, and creating a bank tape table.
write-off	A method for getting rid of inconsequential differences between amounts. For example, you can apply a receipt to an invoice and write off the difference. You can write off both overpayments and underpayments.
Z file	For store and forward (network disconnected) user, EnterpriseOne store-and-forward applications perform edits on static data and other critical information that must be valid to process an order. After the initial edits are complete, EnterpriseOne stores the transactions in work tables on the workstation. These work table are called Z files. When a network connection is established, Z files are uploaded to the enterprise server; and the transactions are edited again by a master business function. The master business function then updates the records in your transaction files.
z-process	A process that converts inbound data from an external system into an EnterpriseOne software table or converts outbound data into an interface table for

	an external system to access.
zusammenfassende meldung	In Germany, the term for the EU Sales Listing.

Index

A

- Accessing log files
 - examples, 17
- Accessing Server Administration Workbench
 - iSeries, 192
 - UNIX, 197
 - Windows, 114
- Accessing the Web Server Monitor, 168
- Accounts
 - clusters for AIX users, 239
 - Windows
 - printer setup, 89, 92
 - troubleshooting EnterpriseOne setup, 364
- Add failures
 - forms and grids, 308
- Adding a host to the PATROL console, 205
- Adding a local account
 - printer setup
 - Windows, 92
- Adding a user to the administrators group
 - printer setup
 - Windows, 93
- Adding administrators
 - database security
 - iSeries, 40
- Adding an EnterpriseOne server for querying
 - using SAW for Windows, 115
- Adding an EnterpriseOne server to the PATROL console, 206
- Adding OWKM files to the PATROL environment, 204
- Adding the EnterpriseOne network and queue resources to the group, 251
 - Windows cluster, 251
- Administering batch processes
 - iSeries, 20
 - encoding passwords, 22
 - monitoring processes, 20
 - reviewing batch output files, 21
 - running reports, 22
 - scheduling reports, 23
 - Linux
 - scheduling reports, 67
 - UNIX, 63
 - monitoring processes, 63
 - reviewing batch output files, 64
 - running reports, 66
 - scheduling reports, 67
 - Windows, 99
 - monitoring processes, 99
 - reviewing batch output files, 99
 - running reports, 100
 - scheduling reports, 101
- Administrators
 - CNC, 6
 - hardware, 6
 - network, 6
 - third-party software, 6
 - Windows
 - printer setup, 93
- AIX
 - clusters
 - application server, 242
 - EnterpriseOne setup, 240
 - groups, 239
 - HACMP software, 237
 - multiple instances of EnterpriseOne, 244
 - troubleshooting, 244
 - user accounts, 239
 - HACMP software, 237
 - kernel parameter settings
 - CPU binding, 81
 - disk striping, 80
 - maxuproc, 78
 - system parameters, 79
 - tune parameters, 79
 - viewing system parameters, 79
- Application development logs
 - workstations, 318
- Applications
 - consultants, 6
 - developers, 6
 - error messages, 307
 - project leaders, 6
- Applying security on the iSeries, 25
- Architecture
 - enterprise servers
 - iSeries, 8
 - UNIX, 56
 - Windows, 85
- Authorization lists
 - iSeries
 - database security, 46
 - samples, 52
- Automatic start

- enterprise servers
 - iSeries, 15
 - UNIX (HP 9000), 60
 - UNIX (RS/6000), 61
 - UNIX (Sun Solaris), 61
 - Windows, 95

B

- Backing up EnterpriseOne tables, 255
 - creating a backup
 - iSeries, 264
 - Oracle on UNIX or Windows, 266
 - SQL Server, 267
 - restoring a backup file
 - iSeries, 269
 - Oracle on UNIX or Windows, 267
 - SQL Server, 269
 - SQL Server on Windows, 270
 - server requirements
 - deployment server, 255
 - enterprise server, 256
 - servers, 264
- Backing up servers
 - EnterpriseOne tables and object owner IDs, 258
 - requirements
 - deployment server, 255
 - enterprise server, 256
- Batch output files
 - iSeries
 - reviewing, 21
 - UNIX
 - reviewing, 64
 - Windows
 - reviewing, 99
- Batch processes
 - general event flow
 - log file, 325
 - iSeries
 - administration, 20
 - encoding passwords, 22
 - monitoring processes, 20
 - PDF output locations, 21
 - running reports, 22
 - scheduling reports, 23
 - Linux
 - scheduling reports, 67
 - log files
 - enterprise server, 331
 - Enterprise server, 340, 343
 - workstation event flow, 325
 - workstation setup, 325

- troubleshooting
 - error when run on a server, 312
 - errors on a report, 311
 - no data, 310
 - unexpected data on a report, 312

- UNIX
 - administration, 63
 - monitoring processes, 63
 - running reports, 66
 - scheduling reports, 67

- Windows
 - administration, 99
 - monitoring processes, 99
 - running reports, 100
 - scheduling reports, 101
- workstations
 - troubleshooting, 310

- Batch queues
 - settings for starting batch queues, 102

- BSFNLIB security parameter, 39

- Business functions
 - security files
 - UNIX, 69
 - Windows, 106
 - troubleshooting development workstations, 316

C

- Changing alarm ranges from the PATROL console, 213

- Changing printer ownership, 91

- Changing process settings using SAW for Windows, 124

- Checking EnterpriseOne server configurations using SAW for Windows, 116

- Checking EnterpriseOne server connectivity using SAW for Windows, 116

- Citrix Metaframe
 - incorporating with Terminal Server Edition, 221

- Cleaning up the enterprise server

- iSeries, 18
 - cleanup, 18
 - clearing the log files, 18
 - Windows, 96

- Clusters, 230

- AIX
 - application server, 242
 - EnterpriseOne setup, 240
 - groups, 239
 - HACMP software, 237
 - multiple instances of EnterpriseOne, 244

- troubleshooting, 244
- user accounts, 239
- HP-UX, 230
 - MC/ServiceGuard, 231
 - multiple instances of EnterpriseOne, 235
 - Oracle package, 231
 - troubleshooting, 236
- HP-UX, 232
- Solaris, 246
 - registering EnterpriseOne with SUNClustering, 248
- Windows, 249
 - advanced options for EnterpriseOne resources, 253
 - failover and failback parameters, 254
 - network and queue resources, 251
 - setting up EnterpriseOne, 250
 - starting and stopping EnterpriseOne resources, 252
 - testing EnterpriseOne connections, 252
 - troubleshooting, 250
- Communications
 - troubleshooting
 - Enterprise servers, 346
 - workstations, 308
- Compile error log
 - workstation, 318, 326
- Configurable Network Computing
 - administrator, 6
 - consultant, 6
- Configurable Network Computing Implementation Guide, 4
- Configuring database tables
 - troubleshooting
 - XE "Databases: iSeries:table configuration problems" iSeries, 349
 - UNIX, 362
- Configuring EnterpriseOne server monitoring
 - settings using SAW for Windows
 - selecting a server, port, and processes to monitor, 118
 - setting up event notification, 118
 - setting up run-time parameters, 119
- Configuring OWKM, 207
- Configuring SAW, 139
- Configuring server monitoring settings using SAW for Windows, 117
- Configuring the Web Server Monitor, 168
- Consultants
 - application, 6
 - CNC, 6
 - custom solution, 6

- hardware, network, and third-party software, 6
- Copying EnterpriseOne to a server
 - UNIX troubleshooting, 361
- CPU binding for AIX, 81
- Creating a backup
 - iSeries, 264
 - Oracle on UNIX or Windows, 266
 - SQL Server, 267
- Creating a separate instance of EnterpriseOne enterprise server, 81
- Creating an application server
 - AIX
 - cluster resources, 243
 - clusters, 242
 - defining a server, 243
 - moving control scripts, 242
- Creating an EnterpriseOne subsystem on the iSeries, 29
- Creating group and user accounts
 - AIX
 - clusters, 239
- Creating the OUTQ
 - iSeries printer, 19
- Custom solution consultant, 6

D

- Data sources
 - troubleshooting
 - Server Map, 347
 - setup for production workstation, 314
- Databases
 - iSeries
 - security, 31
 - setting up security, 33
 - table configuration problems, 349
 - Oracle
 - HACMP software, 239
 - production workstations
 - troubleshooting enterprise server data availability, 313
 - UNIX
 - increasing the maxprocess limit, 62
 - starting the enterprise server, 62
 - table configuration problems, 362
- Deadlock collisions
 - troubleshooting, 347
- Debugging
 - Enterprise server, 345
 - jde.ini file
 - turning on or off, 203
- Defining an application server

- AIX
 - clusters, 243
- Defining cluster resources
 - AIX, 243
- Deleting a SnapShot file, 274
- Deleting log files for an enterprise server
 - process from the PATROL console, 212
- Deleting log files for server processes
 - iSeries, 194
 - UNIX, 199
- Deleting server log files
 - iSeries, 195
 - UNIX, 198
- Deleting web server log files from the PATROL
 - console, 216
- Deployment servers
 - backup requirements, 255
- Determining printer ownership, 91
- Developers
 - application, 6
- Development workstation
 - troubleshooting
 - business functions, 316
 - DLL files, 315
 - event rules, 316
- Development workstations
 - troubleshooting, 315
- Directory structures
 - UNIX, 54
- Disabling jde.log, 321
- Disabling jdedebug.log, 323
- Disk space
 - displaying resources
 - iSeries, 195
 - UNIX, 201
- Disk striping
 - AIX
 - kernel parameter settings, 80
- Displaying disk space resources
 - iSeries, 195
 - UNIX, 201
- Displaying EnterpriseOne server processes
 - iSeries, 193
 - UNIX, 198
- Displaying IPC resources
 - iSeries, 195
 - UNIX, 201
- Displaying log files for server processes
 - iSeries, 194
 - UNIX, 199
- Displaying server log files
 - iSeries, 195

- UNIX, 198
- Displaying user profile information
 - database security
 - iSeries, 46
- DLL files
 - troubleshooting
 - development workstation, 315
 - production workstation, 314
- Domains
 - Windows
 - printer setup, 91
- DTAPATH datapath security parameter, 36

E

- Editing the enterprise server JDE.INI file
 - AIX
 - clusters, 245
- Editing the enterprise server jde.ini file from the
 - PATROL console, 211
- Editing the owenv script
 - AIX
 - clusters, 240
- Editing the server jde.ini file
 - UNIX, 202
- Editing the start resource control script
 - AIX
 - clusters, 240
- Editing the stop resource control script
 - AIX
 - clusters, 241
- E-mail
 - troubleshooting
 - iSeries, 355
 - UNIX, 362
 - Windows, 373
- Encoding the passwords of users who submit
 - batch jobs
 - iSeries, 22
- Ending iSeries server processes, 194
- Enterprise server components
 - monitored by OWKM, 208
- enterprise server initialization
 - iSeries, 11
- Enterprise servers
 - backup requirements, 256
 - communication problems, 346
 - creating a separate instance, 81
 - debugging, 345
 - general problems, 344
 - iSeries
 - architecture, 8
 - automatic start, 15

- cleaning up the server, 18
 - e-mail, 355
 - initialization, 11
 - installation problems, 349
 - interprocess communication (IPC), 357
 - JDBNET, 357
 - jde.ini file, 358
 - manual start, 14
 - multiple release setup, 356
 - PPAT, 355
 - setting up a printer, 19
 - shutting down the server, 16
 - starting the server, 12
 - iSeries
 - backup requirements, 256
 - Linux
 - architecture, 56
 - automatic start, 61
 - shutting down the server, 62
 - log files
 - batch processes, 340
 - batch processing, 331
 - jde.log, 332
 - jddebug.log, 337
 - logic processing, 331
 - viewing files, 332
 - monitoring using KM commands, 209
 - running multiple instances, 81
 - Server Map data source
 - problems, 347
 - troubleshooting
 - data availability, 313
 - types of problems, 331
 - UNIX
 - architecture, 56
 - automatic start (HP 9000), 60
 - automatic start (RS/6000), 61
 - automatic start (Sun Solaris), 61
 - backup requirements, 257
 - directory structure, 54
 - initialization, 58
 - installation problems, 360
 - JDBNET, 363
 - jde.ini file, 360
 - jde.ini file security, 70
 - manual start, 59
 - multiple-release setup problems, 362
 - setting up a printer, 63
 - shutting down the server, 62
 - starting the server, 59
 - Windows
 - automatic start, 95
 - backup requirements, 258
 - cleanup, 96
 - directory structure, 84
 - installation problems, 364
 - JDE.INI file security, 106
 - manual start, 97
 - Network services, 94
 - setting up a printer, 89
 - user accounts, 91
 - EnterpriseOne
 - UNIX
 - directory structure, 54
 - Windows
 - testing by submitting report, 369
 - EnterpriseOne
 - backing up tables, 255
 - EnterpriseOne DB admin profile security
 - parameter, 38
 - Error messages
 - workstations, 306
 - communication failures, 308
 - details, 307
 - form and grid add failures, 308
 - frequent and generic, 307
 - generated by applications, 307
 - memory violations, 308
 - Event rules
 - workstations
 - troubleshooting, 316
 - Executable file security
 - UNIX, 70
 - Windows, 106
 - Executable files on the workstation, 296
- ## F
- Failover and failback parameters
 - Windows cluster, 254
 - Failures
 - communication, 308
 - form and grid add, 308
 - File descriptors
 - kernel parameter settings
 - HP-UX, 76
 - Sun Solaris, 76
 - File security
 - UNIX
 - business function files, 69
 - EnterpriseOne executables, 70
 - jde.ini file (enterprise server), 70
 - specification files, 69
 - Windows
 - business function files, 106

JDE.INI file (enterprise server), 106
specification files, 106

Files

restoring a backup file
iSeries, 269
Oracle on UNIX or Windows, 267
SQL Server, 269
SQL Server on Windows, 270
viewing batch output files, 21

Finding the log files

iSeries, 350

Finding the report files

UNIX, 362
Windows, 369

Forms

add failures, 308

G

Generating serialized objects for the web server

setting up the Win32 client to generate
serialized objects, 277

Generating serialized web objects, 282

Global tables

workstations, 317

Grids

add failures, 308

H

HACMP software

AIX

clusters, 237

application server

cluster, 242

cluster resources, 243

control scripts, 242

definition, 243

EnterpriseOne setup, 240

jde.ini file, 245

owenv script, 240

start resource control script, 240

stop resource control script, 241

groups, 239

how it works, 237

installation

Oracle database, 239

user login accounts, 238

user accounts, 239

Hardware

administrators, 6

consultants, 6

HP 9000

enterprise server

automatic start, 60

HP-UX

clusters, 230

EnterpriseOne package, 232

MC/ServiceGuard, 231

multiple instances of EnterpriseOne, 235

Oracle package, 231

troubleshooting, 236

kernel parameter settings, 71

file descriptors, 76

message queues, 73

processes, 76

semaphores, 74

shared memory, 75

I

Identifying an EnterpriseOne instance, 25

on the iSeries, 24

Implementation methodology

roles, 5

INI files

iSeries

setup problems, 349

INILIB security parameter, 36

Initialization

troubleshooting

iSeries, 11

UNIX, 58

Windows, 87

Installation

EnterpriseOne

verification for UNIX, 62

HACMP software

Oracle database, 239

user login accounts, 238

troubleshooting

standalone installation, 313

verification for Windows, 98

Installation considerations, 238

Interactive applications

production workstation

troubleshooting, 309

Interprocess communication

troubleshooting iSeries, 357

IPC resources

Linux, 77

IPC resources for

iSeries, 195

IPC resources for EnterpriseOne

UNIX, 201

IPCS

- setup for database security
 - iSeries, 40
- iSeries
 - enterprise server
 - architecture, 8
 - automatic start, 15
 - initialization, 11
 - installation problems, 349
 - jde.ini file, 358
 - manual start, 14
 - multiple release setup, 356
 - process flow, 8
 - shutting down the server, 16
 - starting the server, 12
 - integrated file system logging support, 17
 - library structure
 - EnterpriseOne, 12
 - restoring a backup file, 269
 - Server Administration Workbench
 - accessing SAW, 192
 - disk space resources, 195
 - displaying processes, 193
 - ending a server process, 194
 - IPC resources, 195
 - log files, 194
 - modifying the jde.ini file, 196
 - monitoring EnterpriseOne, 191
 - remote servers, 192
 - server log files, 195
 - starting the EnterpriseOne server, 193
 - stopping EnterpriseOne server, 193
- iSeries
 - administering servers, 8
 - authorization lists
 - database security, 46
 - samples, 52
 - backing up tables, 264
 - batch output files
 - reviewing, 21
 - batch processes
 - administration, 20
 - encoding passwords, 22
 - monitoring processes, 20
 - PDF output location, 21
 - running reports, 22
 - scheduling reports, 23
 - database security
 - administrators, 40
 - removing administrative authority, 41
 - user profile information, 46
 - enterprise server
 - backup requirements, 256

- cleaning up the server, 18
 - e-mail, 355
 - interprocess communication (IPC), 357
 - JDBNET, 357
 - PPAT, 355
 - setting up a printer, 19
- EnterpriseOne database security, 31
- jde.log
 - clearing the file, 18
- jddebug.log
 - clearing the file, 18
- log files, 17
- setting up EnterpriseOne database security, 33
- iSeries database security parameters
 - BSFNLIB, 39
 - DTAPATH datapath, 36
 - EnterpriseOne DB admin profile, 38
 - INILIB, 36
 - modify JDE profile, 37
 - modify security profile, 38
 - modify system profile, 37
 - secure log path, 39
 - type, 35

J

- JDBNET
 - troubleshooting
 - iSeries, 357
- jde.ini file
 - modifying the file
 - iSeries, 196
 - troubleshooting
 - iSeries, 358
 - UNIX, 360
- UNIX
 - changing the server port number, 203
 - editing the file, 202
 - modifying the file, 202
 - overview, 202
 - security, 70
 - turning on or off debugging, 203
- JDE.INI file
 - UNIX
 - [NETWORK QUEUE SETTINGS], 64
 - Windows
 - [NETWORK QUEUE SETTINGS], 100
 - security, 106
 - settings for starting batch queues, 102
- jde.log
 - disabling the file, 321
 - enterprise server

- definition, 331
- Enterprise server
 - overview, 332
 - setup, 335
- iSeries
 - clearing the file, 18
 - locating the file, 325
 - workstation, 318
 - setup, 320
- jdecpy.log
 - examples
 - successful copy, 327
 - unsuccessful copy, 327
 - workstation, 318, 327
- jddebug.log
 - clearing the file, 18
 - definition, 318
 - disabling the file, 323
 - enterprise server
 - overview, 331
 - Enterprise server
 - disabling the log file, 338
 - enabling the log file, 338
 - naming conventions, 339
 - overview, 337
 - reading the log file, 337
 - recommendations, 338
 - server locations, 339
 - setting up, 339
 - location on a workstation, 322
 - workstation
 - setup, 322
- jdeinst.log
 - definition, 318
 - purpose, 330
- JDENET
 - iSeries
 - shutting down, 355
 - troubleshooting, 353
- JITI
 - reducing the frequency
 - B732, 226
 - B733, 226
 - B7331, 227
 - B7332, 227
 - B7334, 227

K

- Kernels
 - AIX
 - CPU binding, 81
 - disk striping, 80

- maxuproc, 78
- parameter settings, 78
- system parameters, 78
- tune parameters, 79
- viewing system parameters, 79
- HP-UX
 - file descriptors, 76
 - message queues, 73
 - parameter settings, 71
 - processes, 76
 - semaphores, 74
 - shared memory, 75
- Linux
 - parameter settings, 76
- Sun Solaris
 - file descriptors, 76
 - message queues, 73
 - parameter settings, 71
 - processes, 76
 - semaphores, 74
 - shared memory, 75
- KM commands
 - monitoring web servers, 215

L

- Libraries
 - iSeries
 - installation problems, 349
 - library structure, 12
 - Linux
 - administering servers, 54
 - batch processes
 - scheduling reports, 67
 - enterprise server
 - architecture, 56
 - shutting down the server, 62
 - starting automatically, 61
 - file limits, 77
 - file security, 68
 - IPC resources, 77
 - kernel parameter settings, 76
- Listing batch output files
 - UNIX, 64
- Loading OWKM files, 205
- Locating the jde.log, 325
- Locating the workstation jddebug.log, 322
- Log files
 - clearing jde.log, 19
 - clearing jddebug.log, 19
 - enterprise server
 - batch processing logs, 331
 - Enterprise server

- batch processes, 340
 - jde.log, 332
 - jddebug.log, 337
 - viewing logs, 332
 - enterprise server batch process log, 331
 - enterprise servers, 331
 - EnterpriseOne workstations
 - categories, 317
 - iSeries, 17
 - location problems, 350
 - jde.log
 - locating the file, 325
 - jddebug.log
 - location on a workstation, 322
 - workstation
 - application development, 318
 - batch processes, 325
 - compile errors, 326
 - global tables, 317
 - jde.log, 320
 - jdecpy.log, 327
 - jddebug.log, 322
 - jdeinst.log, 330
 - logic processing logs, 318, 331
 - sql.log, 328, 329
 - troubleshooting strategies, 319
 - troubleshooting workstation processes, 317
 - viewing log files, 319
 - viewing server logs, 332
 - workstation batch process log, 325, 340, 343
 - Logging
 - integrated file system logging support
 - iSeries, 17
 - server log files
 - iSeries, 194, 195
 - UNIX, 198, 199
 - Logic processing logs
 - enterprise server, 331
 - workstations, 318
- ## M
- Maintaining file security
 - executable files, 106
 - Linux
 - overview, 68
 - UNIX
 - business function files, 69
 - jde.ini (enterprise server), 70
 - overview, 68
 - specification files, 69
 - Windows
 - business function files, 106
 - JDE.INI (enterprise server), 106
 - overview, 106
 - specification files, 106
 - Maintaining multiple instances of EnterpriseOne
 - AIX, 244
 - HP-UX, 235
 - Manual start
 - enterprise servers
 - iSeries, 14
 - UNIX, 59
 - Windows, 97
 - Maxuproc
 - AIX
 - system parameters, 78
 - MC/ServiceGuard
 - setting up an Oracle package, 231
 - Memory
 - kernel parameter settings
 - HP-UX, 75
 - Sun Solaris, 75
 - Memory violations
 - workstations, 308
 - Message queues
 - kernel parameter settings
 - HP-UX, 73
 - Sun Solaris, 73
 - Modify JDE profile security parameter, 37
 - Modify security profile security parameter, 38
 - Modify system profile security parameter, 37
 - Modifying the server jde.ini file
 - iSeries, 196
 - UNIX, 202
 - Monitoring EnterpriseOne
 - Web Server Monitor
 - configuring, 168
 - Monitoring EnterpriseOne enterprise server
 - statistics using SAW for Windows
 - monitoring server process statistics, 123
 - viewing statistics for server disks, 127
 - Monitoring EnterpriseOne with the Web Server Monitor
 - configuring, 168
 - Monitoring batch processes
 - iSeries, 20
 - UNIX, 63
 - Windows, 99
 - Monitoring business functions running on
 - connected enterprise servers using the Web Server Monitor, 174
 - Monitoring data on enterprise servers connected to the web server using SAW for Windows, 128

- Monitoring database connection pools using the Web Server Monitor, 177
- Monitoring enterprise server log files from the PATROL console, 211
- Monitoring enterprise server statistics using SAW for Windows, 122
 - printing log and debug log files, 126
 - turning logging on and off, 126
 - viewing log and debug log files, 125
 - viewing process user statistics, 126
 - viewing statistics for batch processes, 126
- Monitoring enterprise servers using KM commands, 209
- Monitoring EnterpriseOne Knowledge Module, 203, 206
 - adding a host to the PATROL console, 205
 - adding OWKM files to the PATROL environment, 204
 - changing alarm ranges, 213
 - deleting log files for an enterprise server process, 212
 - deleting web server log files, 216
 - editing the enterprise server jde.ini file, 211
 - loading OWKM files, 205
 - monitoring, 210
 - monitoring web server components, 214
 - monitoring enterprise server log files, 211
 - monitoring web server components, 214
 - monitoring web server components from the PATROL console, 215
 - monitoring web server connection pools, 217
 - monitoring web server log files, 216
 - monitoring web server users, 217
 - operating system-specific consideration, 204
 - setting logging flags for the web server, 217
 - viewing enterprise server disk information, 213
 - viewing log files for an enterprise server process, 212
- Web Server Monitor, 167
 - monitoring business functions running on connected enterprise servers, 174
 - monitoring database connection pools, 177
 - monitoring JDENET connection pools, 173
 - monitoring the Web server system, 169
 - monitoring Web server log files, 181
 - monitoring Web server users, 171
 - troubleshooting the web server, 186
- Monitoring EnterpriseOne enterprise server statistics using SAW for Windows
 - changing process settings, 124
- Monitoring EnterpriseOne server processes XE "Server Administration Workbench:Windows:monitoring EnterpriseOne server processes" using SAW for Windows, 120
- Monitoring EnterpriseOne Servers from the web, 139
- Monitoring EnterpriseOne web server log files from the PATROL console, 216
- Monitoring EnterpriseOne web server statistics using SAW for Windows, 127
- Monitoring EnterpriseOne with Server Administration Workbench iSeries
 - accessing SAW, 192
 - IPC and disk space resources, 194
 - list of applications, 191
 - modifying the jde.ini file, 196
 - server log files, 195
 - server processes, 193
- UNIX, 196
 - accessing SAW, 197
 - IPC and disk space resources, 200
 - jde.ini file, 202
 - remote servers, 197
 - server log files, 198
 - server processes, 198
- Windows, 114
 - accessing SAW, 114
- Monitoring EnterpriseOne with the Web Server Monitor, 167
 - monitoring business functions running on connected enterprise servers, 174
 - monitoring database connection pools, 177
 - monitoring JDENET connection pools, 173
 - monitoring the Web server system, 169
 - monitoring Web server log files, 181
 - monitoring Web server users, 171
 - troubleshooting the web server, 186
- Monitoring JDENET connection pools using the Web Server Monitor, 173
- Monitoring remote servers iSeries, 192
 - UNIX, 197
- Monitoring server process statistics using SAW for Windows, 123
- Monitoring the Web server system using the Web Server Monitor, 169

- Monitoring the Web server users using the Web Server Monitor, 171
- Monitoring virtual clients, 178
- Monitoring web server CallObjects XE
 - "Monitoring web server statistics using SAW for Windows:monitoring Web server CallObjects"using SAW for Windows, 128
- Monitoring web server connection pools from the PATROL console, 217
- Monitoring web server connection statistics using SAW for Windows, 130
- Monitoring Web server log files using the Web Server Monitor, 181
- Monitoring web server statistics using SAW for Windows
 - monitoring data on enterprise servers connected to the web server, 128
 - monitoring Web server CallObjects, 128
 - monitoring web server connection statistics, 130
 - monitoring web server user logs, 133
- Monitoring Web server statistics using SAW for Windows
 - monitoring Web server user connection statistics, 132
 - monitoring Web server user statistics, 130
- Monitoring Web server user connection statistics using SAW for Windows, 132
- Monitoring web server user logs using SAW for Windows, 133
- Monitoring Web server user statistics XE
 - "Monitoring Web server statistics using SAW for Windows:monitoring Web server user statistics"using SAW for Windows, 130
- Monitoring web server users from the PATROL console, 217
- Monitoring web servers
 - using KM commands, 215
- Moving the control scripts
 - AIX clusters, 242
- Multiple release setup
 - troubleshooting
 - iSeries, 356
 - UNIX, 362
- Multi-user mode
 - restrictions with Terminal Server Edition, 222

N

- Network
 - administrators, 6
 - consultants, 6
 - resources

- Windows cluster, 251
- services for Windows, 94
 - setup, 95
 - starting services, 95
 - stopping services, 96
 - uninstalling services, 97
- Terminal Server Edition, 223

O

- Object owner IDs
 - EnterpriseOne tables, 258
- ODBC
 - troubleshooting
 - sql.log, 329
- Operating system-specific considerations for Knowledge Module, 204
- Oracle
 - backing up tables on UNIX or Windows, 266
 - database
 - HACMP software, 239
 - restoring a backup file, 267
 - setting up a package
 - MC/ServiceGuard, 231
- Output files
 - viewing batch output files, 21
- OUTQ
 - creating for iSeries printer, 19
 - starting for iSeries printer, 19
- ownev, 247

P

- Package Management Guide, 5
- Packages
 - EnterpriseOne for a cluster, 232
 - Oracle for MC/ServiceGuard, 231
- Parameter settings
 - AIX
 - CPU binding, 81
 - disk striping, 80
 - maxuproc, 78
 - tune parameters, 79
 - viewing system parameters, 79
 - HP-UX memory, 75
 - Sun Solaris memory, 75
- Passwords
 - encoding passwords for iSeries users who submit batch jobs, 22
- Performance
 - AIX
 - CPU binding, 81
 - diskstriping, 80

- kernel parameter settings, 79
- Terminal Server Edition, 223
- Performing preliminary troubleshooting
 - production workstation, 309
- Permissions
 - Windows
 - printer setup, 89
- Port number
 - UNIX
 - changing the jde.ini file, 203
- Portable Document Format
 - iSeries output location, 21
- PORTTEST
 - iSeries, 351
 - Windows
 - verifying the EnterpriseOne installation, 252
- PPAT
 - troubleshooting iSeries, 355
- Printers
 - iSeries
 - creating the OUTQ, 19
 - printing multiple copies to a remote printer, 20
 - setup, 19
 - starting the OUTQ, 19
 - UNIX
 - setup, 63
 - Windows
 - accounts, 89
 - adding a printer, 90
 - administrators, 93
 - changing domains, 91
 - local account, 92
 - permissions, 89
 - printer ownership, 91
 - services, 89
 - setup, 89
 - setup problems, 365
 - user accounts, 91
 - workstations
 - troubleshooting, 313
- Printing log and debug log files using SAW for Windows, 126
- Printing multiple copies to a remote printer, 20
- Printing server log files
 - iSeries, 195
 - UNIX, 198
- Process flow
 - enterprise servers
 - iSeries, 8
- Processes

- iSeries
 - monitoring batch processes, 20
- kernel parameter settings
 - HP-UX, 76
 - Sun Solaris, 76
- UNIX
 - monitoring batch processes, 63
- Windows
 - monitoring batch processes, 99
- Production workstation
 - troubleshooting, 309
 - batch processes, 310
 - data source setup, 314
 - DLL files, 314
 - enterprise server data availability, 313
 - interactive applications, 309
 - preliminary, 309
 - printing, 313
 - standalone installation, 313
- Project leaders
 - application, 6

Q

- Querying enterprise server disk components
 - using SAW for Windows, 137
- Querying enterprise server IPC resources using SAW for Windows, 137
- Querying enterprise server job queue using SAW for Windows, 139
- Querying enterprise server log files using SAW for Windows, 138
- Querying enterprise server process components
 - using SAW for Windows, 135
- Querying enterprise server users using SAW for Windows, 139
- Querying enterprise servers using SAW for Windows
 - querying enterprise server disk components, 137
 - querying enterprise server IPC resources, 137
 - querying enterprise server log files, 138
 - querying enterprise server process components, 135
 - querying enterprise server users, 139
 - querying the enterprise server job queue, 139
- Querying EnterpriseOne enterprise servers using SAW for Windows, 134
- Queue resources
 - Windows cluster, 251
- Queue services
 - Windows, 94
 - setup, 95

- starting services, 95
- stopping services, 96
- uninstalling services, 97

R

- Read
 - sql.log, 328
- Reducing JITI frequency
 - B732, 226
 - B733, 226
 - B7331, 227
 - B7332, 227
 - B7334, 227
- Remote printers
 - printing multiple copies on the iSeries, 20
- Removing a server a server from the query list
 - using SAW for Windows, 117
- Removing administrative authority from user profiles
 - iSeries, 41
- Reports
 - batch processes
 - errors, 311
 - iSeries, 22, 23
 - unexpected data, 312
 - UNIX, 66, 67
 - Windows, 100, 101
 - UNIX
 - scheduling, using at and batch commands, 67
 - Windows
 - file location problems, 369
- Requirements
 - backing up servers, 255
- Restoring a backup file
 - iSeries, 269
 - Oracle on UNIX or Windows, 267
 - SQL Server, 269
 - SQL Server on Windows, 270
- Restoring a SnapShot file, 274
- Reviewing batch output files
 - iSeries, 21
 - UNIX, 64
 - Windows, 99
- RS/6000
 - enterprise server
 - starting automatically, 61
- Running batch process reports from the server
 - command line
 - iSeries, 22
 - UNIX, 66
 - Windows, 100

- Running EnterpriseOne JDENET iSeries

- troubleshooting, 353

- Running EnterpriseOne manually

- Windows

- troubleshooting, 368

- Running multiple instances of EnterpriseOne

- on the iSeries, 23

- Running SnapShot from the command line, 274

S

- Saving EnterpriseOne in a SnapShot file, 273

- SAW for Windows

- configuring EnterpriseOne server monitoring settings

- selecting a server, port, and processes to monitor, 118

- setting up event notification, 118

- setting up run-time parameters, 119

- monitoring enterprise server statistics

- printing log and debug log files, 126

- turning logging on and off, 126

- viewing process user statistics, 126

- viewing statistics for batch processes, 126

- monitoring EnterpriseOne enterprise server statistics

- changing process settings, 124

- monitoring server process statistics, 123

- viewing log and debug log files, 125

- viewing statistics for server disks, 127

- monitoring web server statistics

- monitoring data on enterprise servers

- connected to the web server, 128

- monitoring Web server CallObjects, 128

- monitoring web server connection

- statistics, 130

- monitoring web server user logs, 133

- monitoring Web server statistics

- monitoring Web server user connection

- statistics, 132

- monitoring Web server user statistics, 130

- querying enterprise servers

- querying enterprise server disk

- components, 137

- querying enterprise server IPC resources, 137

- querying enterprise server job queue, 139

- querying enterprise server log files, 138

- querying enterprise server process

- components, 135

- querying enterprise server users, 139

- Scheduling reports

- iSeries
 - command line, 23
- UNIX
 - at and batch commands, 67
 - command line, 67
- Windows
 - command line, 101
- Scripts
 - Solaris clustering
 - owenv, 247
 - SunOracleMgr.sh, 247
 - SunStartResource.sh, 246
 - SunStopResource.sh, 247
- Secure log path security parameter, 39
- Security
 - iSeries
 - EnterpriseOne databases, 31
 - setup for EnterpriseOne databases, 33
 - Linux
 - file security, 68
 - UNIX
 - business function files, 69
 - file security, 68
 - jde.ini file (enterprise server), 70
 - specification files, 69
 - Windows
 - business function files, 106
 - executable files, 106
 - file security, 106
 - JDE.INI file (enterprise server), 106
 - specification files, 106
- Selecting a server, port, and processes for monitoring using SAW for Windows, 118
- Semaphores
 - kernel parameter settings
 - HP-UX, 74
 - Sun Solaris, 74
- Server Administration Workbench
 - Windows, 115
 - checking EnterpriseOne server configurations, 116
 - checking EnterpriseOne server connectivity, 116
 - configuring server monitoring settings, 117
 - jde.ini setting for creation of temporary log files, 114
 - monitoring enterprise server statistics, 122
 - monitoring EnterpriseOne server processes, 120
 - monitoring EnterpriseOne web server statistics, 127
 - querying EnterpriseOne enterprise servers, 134
 - removing a server from the query list, 117
- Server Administration Workbench (SAW), 113
 - iSeries
 - accessing SAW, 192
 - disk space resources, 195
 - displaying processes, 193
 - ending server processes, 194
 - IPC resources, 195
 - log files, 194
 - modifying the jde.ini file, 196
 - monitoring EnterpriseOne, 191
 - remote servers, 192
 - server log files, 195
 - starting the EnterpriseOne server, 193
 - stopping the EnterpriseOne server, 193
 - UNIX
 - accessing SAW, 197
 - disk space resources, 201
 - displaying processes, 198
 - editing the jde.ini file, 202
 - IPC resources, 201
 - log files, 199
 - modifying the jde.ini file, 202
 - monitoring EnterpriseOne, 196
 - remote servers, 197
 - server log files, 198
 - server port number, 203
 - stopping server processes, 200
 - stopping the EnterpriseOne server, 198
 - turning on or off debugging, 203
 - Windows
 - accessing SAW, 114
 - monitoring EnterpriseOne, 114
- Server and Workstation Administration Guide, 5
- Server Map data source
 - problems, 347
- Server monitoring interface
 - accessing, 157
- Servers
 - AIX
 - application server for cluster, 242
 - setting value of maxuproc, 78
 - backing up tables, 264
 - backup requirements
 - deployment server, 255
 - batch processes
 - error when run on a server, 312
 - command line
 - batch process reports for iSeries, 22, 23
 - batch process reports for UNIX, 66, 67

- batch process reports for Windows, 100, 101
- deployment servers
 - backup requirements, 255
- enterprise servers
 - backup requirements, 256
- iSeries
 - cleaning up the enterprise server, 18
 - IPC and disk space resources, 194
 - log files, 195
 - logging support, 17
 - modifying the jde.ini file, 196
 - monitoring remote servers, 192
 - server processes, 193
 - setting up a printer, 19
 - starting the enterprise server, 12
- iSeries
 - administration, 8
- Linux
 - administration, 54
- UNIX
 - administration, 54
 - IPC and disk space resources, 200
 - jde.ini file, 202
 - jde.ini file security, 70
 - log files, 198
 - monitoring remote servers, 197
 - problems copying EnterpriseOne, 361
 - server processes, 198
 - setting up a printer, 63
 - shutting down the enterprise server, 62
 - starting the enterprise server, 59
- Windows
 - administration, 84
 - cleanup, 96
 - JDE.INI file security, 106
 - Network services, 94
 - setting up a printer, 89
- Services
 - Windows
 - Network services, 94
 - printer setup, 89
- Setting advanced options for EnterpriseOne
 - resources, 253
 - Windows cluster, 253
- Setting business function file security
 - UNIX, 69
 - Windows, 106
- Setting EnterpriseOne web server logging flags
 - from the PATROL console, 217
- Setting failover and failback parameters, 253
 - Windows cluster, 254
- Setting specification file security
 - UNIX, 69
 - Windows, 106
- Setting the value of maxuproc on AIX, 78
- Setting up a EnterpriseOne package, 232
- Setting up a printer
 - iSeries
 - creating the OUTQ, 19
 - printing multiple copies to a remote printer, 20
 - starting the OUTQ, 19
 - Windows
 - accounts, 89
 - adding a printer, 90
 - administrators, 93
 - changing domains, 91
 - local account, 92
 - permissions, 89
 - printer ownership, 91
 - services, 89
 - troubleshooting, 365
 - user accounts, 91
- Setting up an Oracle package for MC/ServiceGuard, 231
- Setting up EnterpriseOne accounts
 - Windows
 - troubleshooting, 364
- Setting up EnterpriseOne for HACMP, 240
 - control script
 - start resource, 240
 - stop resource, 241
 - jde.ini file
 - enterprise server, 245
 - owenv script, 240
- Setting up EnterpriseOne on a Windows cluster, 250
 - advanced options for EnterpriseOne
 - resources, 253
 - failover and failback parameters, 254
 - network and queue resources, 251
 - starting and stopping EnterpriseOne
 - resources, 252
 - testing EnterpriseOne connections, 252
- Setting up EnterpriseOne on the terminal server, 223
- Setting up event notification using SAW for
 - Windows, 118
- Setting up iSeries EnterpriseOne database
 - security, 31, 33
 - administrators, 40
 - displaying user profile information, 46
 - removing administrative authority

- user profiles, 41
- Setting up run-time parameters using SAW for Windows, 119
- Setting up the batch process log
 - Enterprise server, 343
 - workstation, 325
- Setting up the Enterprise server jde.log, 335
- Setting up the Enterprise server jdedebug.log, 339
- Setting up the EnterpriseOne iSeries .INI file
 - troubleshooting, 349
- Setting up the Network and Queue services
 - Windows, 95
- Setting up the Win32 client tp generate
 - serialized web objects, 277
- Setting up the workstation jde.log, 320
- Setting up the workstation jdedebug.log, 322
- Setting up user accounts on an enterprise server
 - Windows
 - printers, 91
- Shared memory
 - kernel parameter settings
 - HP-UX, 75
 - Sun Solaris, 75
- Shutting down JDENET
 - iSeries
 - troubleshooting, 355
- Shutting down the enterprise server
 - iSeries, 16
 - Linux, 62
 - UNIX, 62
- Single occurrence report
 - scheduling for UNIX, 67
- SMC views
 - working with, 154
- SMIT application
 - setting the value of maxuproc, 78
- SnapShot, 271
 - deleting a file, 274
 - restoring a file, 274
 - running SnapShot from the command line, 274
 - saving EnterpriseOne in a file, 273
 - the process, 271
- Solaris
 - clusters
 - modifying scripts, 246
 - registering EnterpriseOne with SUNClustering, 248
 - requirements for Solaris, 246
 - setting up a cluster, 246
 - enterprise server
 - automatic start, 61
 - kernel parameter settings, 71
 - file descriptors, 76
 - message queues, 73
 - processes, 76
 - semaphores, 74
 - shared memory, 75
 - Specification file security
 - UNIX, 69
 - Windows, 106
- SQL Server
 - backing up tables, 267
 - restoring a backup file, 269
 - Windows
 - restoring a backup file, 270
- sql.log
 - read, 328
 - workstation, 318, 328, 329
 - troubleshooting ODBC, 329
 - turning on sql.log, 329
 - turning on the log, 329
- Standalone executable files, 299
- Starting EnterpriseOne resources on the cluster
 - Windows, 252
- Starting the enterprise server
 - iSeries
 - automatic, 15
 - library structure, 12
 - manual, 14
 - iSeries, 12
 - Linux
 - automatic, 61
 - UNIX, 59
 - automatic (HP 9000), 60
 - automatic (RS/6000), 61
 - automatic (Sun Solaris), 61
 - database settings, 62
 - manual, 59
 - Windows
 - manual, 97
- Starting the EnterpriseOne server
 - iSeries, 193
- Starting the Network and Queue services
 - Windows, 95
- Starting the OUTQ for iSeries printer setup, 19
- Stopping all EnterpriseOne processes
 - Windows
 - troubleshooting, 373
- Stopping EnterpriseOne
 - Windows
 - stopping processes using Visual C++, 373
- Stopping EnterpriseOne as run manually

- Windows
 - stopping all processes, 373
 - troubleshooting, 372
- Stopping EnterpriseOne resources on the cluster
 - Windows, 252
- Stopping server processes
 - UNIX, 200
- Stopping the EnterpriseOne server
 - iSeries, 193
 - UNIX, 198
- Stopping the Network and Queue services
 - Windows, 96
- SUNClustering
 - registering EnterpriseOne, 248
- SunOracleMgr.sh, 247
- SunStartResource.sh, 246
- SunStopResource.sh, 247
- System Administration Guide, 4
- System functionality
 - AIX
 - kernel parameter settings, 78

T

- Tables
 - backing up EnterpriseOne tables, 255
 - backing up servers, 264
 - object owner IDs, 258
- Terminal Server Edition
 - EnterpriseOne on Windows, 219
 - incorporating Citrix Metaframe, 221
 - network considerations, 223
 - performance considerations, 223
 - restrictions in multi-user mode, 222
 - setting up EnterpriseOne, 223
 - troubleshooting, 224
 - EnterpriseOne UBE output security, 225
- Testing
 - iSeries
 - PORTTEST, 351
 - submitting a report, 354
- Testing EnterpriseOne
 - Windows
 - connections on a cluster, 252
- Testing EnterpriseOne by submitting a report
 - Windows
 - default printer orientation, 371
 - ownership of a printer, 371
- Third-party software
 - administrators, 6
 - consultants, 6
- Troubleshooting
 - AIX on a cluster, 244

- batch processes
 - error when run on a server, 312
 - errors on a report, 311
 - no data, 310
 - unexpected data on a report, 312
- data selection and sequencing criteria lost, 228
- development workstation, 315
- Enterprise server, 348
- enterprise servers, 331
 - batch process log, 331
 - types of problems, 331
 - UNIX, 359
 - Windows, 363
- Enterprise servers, 344
 - batch process log, 340
 - communication problems, 346
 - general problems, 344
 - jde.log, 332
 - jddebug.log, 337
 - Server Map data source, 347
 - viewing logs, 332
- Enterprise servers:, 348
- EnterpriseOne development tools are disabled, 228
- HP-UX on a cluster, 236
- import/export with Microsoft Excel, 226
- iSeries enterprise server
 - .INI file, 349
 - architecture, 8
 - database table configuration, 349
 - e-mail, 355
 - initialization, 11
 - installation, 349
 - interprocess communication (IPC), 357
 - JDBNET, 357
 - jde.ini file, 358
 - library installation verification, 349
 - log file location, 350
 - PPAT, 355
 - running JDENET, 353
 - shutting down JDENET, 355
 - testing by submitting a report, 354
 - testing with PORTTEST, 351
- iSeries enterprise server
 - multiple release setup, 356
- log path is incorrect, 229
- logging off versus disconnecting, 227
- ODBC
 - sql.log, 329
- only one user can log in to EnterpriseOne, 229

- run-time error occurs during server
 - connection test, 228
 - shortcut path is incorrect, 229
 - shortcuts do not work in e-mail messages, 228
 - specification files are locked, 226
 - submit UBE locally to TSE, 225
 - Terminal Server Edition, 224
 - EnterpriseOne UBE output security, 225
 - UBE output security on TSE, 224
 - UNIX enterprise server, 359
 - architecture, 56
 - copying EnterpriseOne to a server, 361
 - database table configuration, 362
 - directory structure, 54
 - e-mail, 362
 - finding the report file, 362
 - initialization, 58
 - installation, 360
 - JDBNET, 363
 - jde.ini file, 360
 - multiple release setup, 362
 - user cannot restart EnterpriseOne, 227
 - users experience problems accessing EnterpriseOne, 228
 - Windows enterprise server, 363
 - architecture, 85
 - clusters, 250
 - directory structure, 84
 - e-mail, 373
 - finding data, 372
 - initialization, 87
 - installation, 364
 - log file location, 365
 - report file location, 369
 - running EnterpriseOne manually, 368
 - setting up EnterpriseOne accounts, 364
 - stopping all EnterpriseOne processes, 373
 - stopping EnterpriseOne as run manually, 372
 - testing by submitting a report, 369
 - using Visual C++ to stop processes, 373
 - Windows Enterprise server
 - setting up a printer, 365
 - workstations, 306
 - batch process log, 325
 - batch processes, 310
 - business functions, 316
 - compile error log, 326
 - data source setup, 314
 - development, 315
 - DLL files (development), 315
 - DLL files (production), 314
 - enterprise server data availability, 313
 - error messages, 306
 - event rules, 316
 - interactive applications, 309
 - jdecpy.log, 327
 - jdeinst.log, 330
 - log files, 317
 - preliminary troubleshooting, 309
 - printing, 313
 - production, 309
 - setting up the jde.log, 320
 - setting up the jdedebug.log, 322
 - sql.log, 328, 329
 - standalone installation, 313
 - strategies using log files, 319
 - viewing log files, 319
 - viewing server logs, 332
 - Troubleshooting EnterpriseOne servers
 - web servers, 374
 - IIS and IBM HTTP web servers, 374
 - JAS, 374
 - serialized database and generation issues, 374
 - Troubleshooting EnterpriseOne web servers, 374
 - Troubleshooting general problems on the Enterprise server
 - deadlock collision problems, 347
 - Troubleshooting IIS and IBM HTTP web servers, 374
 - Troubleshooting JAS, 374
 - Troubleshooting serialized database and generation issues, 374
 - Troubleshooting using the Web Server Monitor, 186
 - Tune parameters
 - AIX
 - CPU binding, 81
 - disk striping, 80
 - kernel parameter settings, 79
 - Turning logging on and off using SAW for Windows, 126
 - Turning on or off debugging
 - jde.ini file
 - UNIX, 203
 - Turning on sql.log, 329
 - Type security parameter, 35
- ## U
- Uninstalling the Network and Queue services
 - Windows, 97
 - UNIX

- administering servers, 54
- AIX
 - kernel parameter settings, 78
- backing up Oracle tables, 266
- batch output files
 - reviewing, 64
- batch processes
 - administration, 63
 - monitoring processes, 63
 - running reports, 66
 - scheduling reports, 67
- business function files
 - security, 69
- creating a separate instance of the enterprise server, 81
- enterprise server
 - architecture, 56
 - backup requirements, 257
 - directory structure, 54
 - initialization, 58
 - installation problems, 360
 - JDBNET, 363
 - jde.ini file, 360
 - manual start, 59
 - multiple-release setup problems, 362
 - setting up a printer, 63
 - shutting down the server, 62
 - starting the server, 59
- EnterpriseOne executable files
 - security, 70
- file security, 68
- jde.ini file
 - security, 70
- kernel parameter settings
 - HP-UX, 71
 - Linux, 76
 - Sun Solaris, 71
- restoring an Oracle backup file, 267
- running multiple EnterpriseOne enterprise servers, 81
- Server Administration Workbench
 - accessing SAW, 197
 - disk space resources, 201
 - displaying processes, 198
 - editing the jde.ini file, 202
 - IPC resources, 201
 - log files, 199
 - modifying the jde.ini file, 202
 - monitoring EnterpriseOne, 196
 - remote servers, 197
 - server log files, 198
 - server port number, 203

- stopping a server process, 200
 - stopping the EnterpriseOne server, 198
 - turning on or off debugging, 203
- Solaris
 - setting up a cluster, 246
- specification files
 - security, 69
- verifying the EnterpriseOne installation, 62
- User accounts
 - Windows
 - printer setup, 91
- User login accounts
 - HACMP software, 238
- User profiles
 - iSeries database security
 - displaying information, 46
 - removing administrative authority, 41
- Using iSeries integrated file system logging support, 17

V

- Verify XE "Libraries: iSeries:installation problems"ing the library installation troubleshooting the iSeries, 349
- Viewing batch output files, 21
- Viewing enterprise server disk information from the PATROL console, 213
- Viewing log and debug log files using SAW for Windows, 125
- Viewing log files
 - iSeries, 17
 - workstation, 319
- Viewing log files for an enterprise server process from the PATROL console, 212
- Viewing process user statistics using SAW for Windows, 126
- Viewing server logs from the workstation, 332
- Viewing statistics for batch processes using SAW for Windows, 126
- Viewing statistics for server disks using SAW for Windows, 127
- Viewing system parameters
 - AIX, 79
- Virtual clients
 - monitoring, 178

W

- Windows
 - administering the server, 84
 - backing up Oracle tables, 266

- batch XE "Administrating batch processes:Windows"processes administration, 99
- batch output files
 - reviewing, 99
- batch processes
 - monitoring processes, 99
 - running reports, 100
 - scheduling reports, 101
- business function files
 - security, 106
- clusters, 249
 - advanced options for EnterpriseOne resources, 253
 - failover and failback parameters, 254
 - network and queue resources, 251
 - setting up EnterpriseOne, 250
 - starting and stopping EnterpriseOne resources, 252
 - testing EnterpriseOne connections, 252
 - troubleshooting, 250
- enterprise server
 - automatic start, 95
 - backup requirements, 258
 - cleanup, 96
 - directory structure, 84
 - initialization, 87
 - installation problems, 364
 - manual start, 97
 - Network services, 94
 - process flow, 85
 - Queue services, 94
 - setting up a printer, 89
- EnterpriseOne on Terminal Server Edition, 219
- file security, 106
- JDE.INI file
 - security, 106
 - settings for starting batch queues, 102
- printer setup
 - ac XE "Printers:Windows:accounts"counts, 89
 - permissions, 89
 - services, 89
- restoring an Oracle backup file, 267
- Server Administration Workbench
 - accessing SAW, 114
 - jde.ini file setting for creation of temporary log files, 114
 - monitoring EnterpriseOne, 114
- specification files
 - security, 106
- SQL Server
 - restoring a backup file, 270
 - verifying the EnterpriseOne installation, 98
- Working with Enterprise server log files
 - batch process log
 - overview, 340
 - setup, 343
 - jde.log
 - overview, 332
 - reading the log file, 337
 - setup, 335
 - jdedebug.log
 - disabling the log file, 338
 - enabling the log file, 338
 - naming conventions, 339
 - overview, 337
 - recommendations, 338
 - server locations, 339
 - setting up, 339
- Working with EnterpriseOne Network and Queue services
 - Windows
 - cleanup, 96
 - manually starting the enterprise server, 97
 - setup, 95
 - stopping services, 96
 - uninstalling services, 97
- Working with jdecpy.log, 327
 - examples
 - successful copy, 327
 - unsuccessful copy, 327
- Working with jdeinst.log, 330
- Working with Network and Queue services
 - Windows
 - starting services, 95
- Working with SMC views, 154
- Working with sql.log, 328, 329
 - troubleshooting ODBC, 329
 - turning on the log, 329
- Working with the batch process log
 - Enterprise server
 - overview, 340
 - setup, 343
- Working with the compile error log, 326
- Working with the Enterprise server jde.log
 - overview, 332
 - setup, 335
- Working with the Enterprise server jdedebug.log
 - disabling the log file, 338
 - enabling the log file, 338
 - naming conventions, 339
 - overview, 337

- reading the log file, 337
- recommendations, 338
- server locations
 - setup, 339
- setting up, 339
- Working with the workstation log files
 - application development, 318
 - batch process log
 - setup, 325
 - compile error log, 326
 - global tables, 317
 - jde.log
 - setup, 320
 - jdecpy.log, 327
 - example of successful copy, 327
 - example of unsuccessful copy, 327
 - jddebug.log
 - setup, 322
 - jdeinst.log, 330
 - logic processing logs, 318
 - overview, 317
 - sql.log, 328, 329
 - troubleshooting ODBC, 329
 - turning on the log, 329
 - troubleshooting strategies, 319
 - viewing log files, 319
- Workstations
 - development (troubleshooting)
 - business functions, 316
 - DLL files, 315
 - event rules, 316
 - error messages, 306
 - details, 307
 - global tables, 317
 - log files
 - application development, 318
 - batch processes, 325
 - compile errors, 326
 - jde.log, 320
 - jdecpy.log, 327
 - jddebug.log, 322
 - jdeinst.log, 330
 - sql.log, 328, 329
 - troubleshooting strategies, 319
 - viewing files, 319
 - logic processing logs, 318
 - production (troubleshooting)
 - batch processes, 310
 - data source setup, 314
 - DLL files, 314
 - enterprise server data availability, 313
 - interactive applications, 309
 - preliminary troubleshooting, 309
 - printing, 313
 - standalone installation, 313

