



EnterpriseOne Tools 8.94

PeopleBook: ☐ Server ☐ and ☐ Workstation Administration

November 2004

EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration
SKU E1_TOOLS8.94TSW-B 1104

Copyright © 2004 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

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. The following open source software may be used in PeopleSoft products and the following disclaimers are provided.

Apache Software Foundation

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.

OpenSSL

Copyright (c) 1998-2003 The OpenSSL Project. All rights reserved.

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "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 OpenSSL PROJECT 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.

SSLey

Copyright (c) 1995-1998 Eric Young. All rights reserved.

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS 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 AUTHOR OR 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.

Loki Library

Copyright (c) 2001 by Andrei Alexandrescu. This code accompanies the book:

Alexandrescu, Andrei. "Modern C++ Design: Generic Programming and Design Patterns Applied". Copyright (c) 2001. Addison-Wesley. Permission to use, copy, modify, distribute and sell this software for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

Contents

General Preface

| | |
|--|------------|
| About This PeopleBook | xix |
| PeopleSoft Application Prerequisites..... | xix |
| PeopleSoft Application Fundamentals..... | xix |
| Documentation Updates and Printed Documentation..... | xx |
| Obtaining Documentation Updates..... | xx |
| Ordering Printed Documentation..... | xx |
| Additional Resources..... | xxi |
| Typographical Conventions and Visual Cues..... | xxii |
| Typographical Conventions..... | xxii |
| Visual Cues..... | xxiii |
| Country, Region, and Industry Identifiers..... | xxiii |
| Currency Codes..... | xxiv |
| Comments and Suggestions..... | xxiv |
| Common Elements Used in PeopleBooks..... | xxiv |

Preface

| | |
|--|--------------|
| PeopleSoft EnterpriseOne Preface..... | xxvii |
| Server and Workstation Administration..... | xxvii |

Chapter 1

| | |
|---|----------|
| Getting Started..... | 1 |
| Understanding the EnterpriseOne Configurable Network Computing Implementation PeopleBooks..... | 1 |
| Understanding Roles..... | 3 |
| Implementing Server and Workstation Administration..... | 3 |

Chapter 2

| | |
|--|----------|
| Understanding Server Administration for iSeries..... | 5 |
| Understanding Server Administration for iSeries..... | 5 |
| Understanding the EnterpriseOne iSeries Architecture and Process Flow for iSeries..... | 5 |
| Understanding EnterpriseOne Initialization for iSeries..... | 7 |
| Starting the Enterprise Server for iSeries..... | 8 |

| | |
|--|----|
| Prerequisites..... | 8 |
| Understanding the iSeries Library Structure for EnterpriseOne..... | 8 |
| Understanding Startup Options for the Enterprise Server for iSeries..... | 10 |
| Starting the Enterprise Server for iSeries Manually..... | 10 |
| Starting the Enterprise Server for iSeries Automatically..... | 11 |
| Shutting Down the Enterprise Server for iSeries..... | 12 |
| Prerequisite..... | 12 |
| Using iSeries Integrated File System Logging Support..... | 13 |
| Cleaning Up the Enterprise Server for iSeries..... | 13 |
| Prerequisite..... | 13 |
| Cleaning Up the Enterprise Server for iSeries..... | 14 |
| Cleaning Up the Enterprise Server for iSeries..... | 14 |
| Clearing the jde.log and jde.debug Files for iSeries..... | 14 |
| Setting Up a Printer for iSeries..... | 14 |
| Understanding Printer Setup for iSeries..... | 15 |
| Creating the OUTQ..... | 15 |
| Starting the OUTQ..... | 15 |
| Printing Multiple Copies to a Remote Printer..... | 15 |
| Administering Batch Processes for iSeries..... | 16 |
| Understanding Batch Process Administration for iSeries..... | 16 |
| Monitoring Batch Processes..... | 17 |
| Reviewing Batch Output Files..... | 18 |
| Encoding the Passwords of Users Who Submit Batch Jobs..... | 18 |
| Running Multiple Instances of EnterpriseOne on the iSeries..... | 18 |
| Prerequisites..... | 18 |
| Running Multiple Instances of EnterpriseOne..... | 19 |
| Understanding iSeries Database Security Parameters..... | 21 |
| Copying Libraries and Directories..... | 26 |
| Applying Security to Multiple Instances of EnterpriseOne on the iSeries..... | 27 |
| Creating an EnterpriseOne Subsystem on the iSeries..... | 28 |
| Administering Security EnterpriseOne Database Security for iSeries..... | 28 |
| Prerequisite..... | 29 |
| Understanding EnterpriseOne Database Security Administration..... | 29 |
| Setting Up iSeries Database Security for a Single EnterpriseOne Instance..... | 39 |
| Setting Up iSeries Database Security for Multiple EnterpriseOne Instances..... | 40 |
| Adding Administrators..... | 40 |
| Removing Administrative Authority from User Profiles..... | 41 |
| Displaying User Profile Information..... | 41 |

Chapter 3

| | |
|--|-----------|
| Understanding Server Administration for UNIX and Linux..... | 47 |
| Understanding Server Administration for UNIX and Linux..... | 47 |
| EnterpriseOne Directory Structure for UNIX and Linux..... | 47 |
| EnterpriseOne Architecture and Process Flow for UNIX and Linux..... | 48 |
| EnterpriseOne Initialization for UNIX and Linux..... | 50 |
| Starting the Enterprise Server for UNIX or Linux..... | 51 |
| Understanding Enterprise Server Startup for UNIX or Linux..... | 51 |
| Starting the Enterprise Server for UNIX or Linux Manually..... | 52 |
| Starting the Enterprise Server for HP-UX Automatically..... | 52 |
| Starting the Enterprise Server for AIX and Solaris Automatically..... | 53 |
| Starting the Enterprise Server for Linux Automatically..... | 54 |
| Verifying the EnterpriseOne Installation..... | 54 |
| Shutting Down the Enterprise Server for UNIX or Linux..... | 54 |
| Shutting Down the Enterprise Server for UNIX or Linux..... | 54 |
| Setting Up a Printer for UNIX or Linux..... | 55 |
| Administering Batch Processes for UNIX or Linux..... | 55 |
| Understanding Batch Process Administration for UNIX or Linux..... | 55 |
| Monitoring Batch Processes..... | 57 |
| Listing Batch Output Files..... | 57 |
| Running Reports from the Command Line for UNIX or Linux..... | 58 |
| Scheduling Reports from the Command Line for UNIX or Linux..... | 58 |
| Maintaining File Security for UNIX and Linux..... | 60 |
| File Security Maintenance for UNIX and Linux..... | 60 |
| HP-UX and Solaris Kernel Parameter Settings..... | 61 |
| Working with Linux Kernel Parameter Settings..... | 65 |
| Understanding Linux Kernel Parameter Settings..... | 66 |
| Setting Specification File Security..... | 67 |
| Setting Business Function File Security..... | 67 |
| Setting Executables Security..... | 67 |
| Setting jde.ini File Security..... | 68 |
| Working with AIX Kernel Parameter Settings for EnterpriseOne..... | 68 |
| Understanding AIX Kernel Parameter Settings for EnterpriseOne..... | 69 |
| Setting the Value of maxuproc..... | 69 |
| Viewing the System Parameters..... | 70 |
| Setting Tune Parameters..... | 70 |
| Running Multiple Instances of the EnterpriseOne Enterprise Server..... | 71 |
| Prerequisite..... | 71 |
| Running Multiple Instances of the EnterpriseOne Enterprise Server..... | 71 |
| Running Multiple Instances of the EnterpriseOne Enterprise Server..... | 71 |

Chapter 4

| | |
|--|-----------|
| Understanding Server Administration for Windows..... | 75 |
| Understanding Server Administration for Windows..... | 75 |
| Understanding the EnterpriseOne Directory Structure for Windows..... | 75 |
| Understanding EnterpriseOne Architecture and Process Flow for Windows..... | 76 |
| Understanding EnterpriseOne Initialization for Windows..... | 79 |
| Understanding JDE.INI Settings for Starting Batch Queues on Windows..... | 80 |
| Understanding Active Directory..... | 81 |
| Setting Up a Printer for Windows..... | 83 |
| Understanding Printer Setup for Windows..... | 83 |
| Understanding Windows Services, Accounts, and Permissions..... | 83 |
| Adding a Printer..... | 85 |
| Determining or Changing Printer Ownership..... | 85 |
| Setting Up User Accounts on an Enterprise Server..... | 85 |
| Changing the Domain..... | 85 |
| Adding a Local Account..... | 86 |
| Adding a User to the Administrators Group..... | 86 |
| Working with Network Services..... | 87 |
| Understanding Network Services..... | 87 |
| Setting Up the Network Service..... | 88 |
| Starting the Network Service..... | 88 |
| Stopping the Network Services..... | 88 |
| Cleaning Up the Enterprise Server for Windows..... | 89 |
| Uninstalling the Network Service..... | 89 |
| Starting the Enterprise Server for Windows Manually..... | 89 |
| Verifying the EnterpriseOne Installation..... | 90 |
| Administering Batch Processes for Windows..... | 90 |
| Understanding Batch Process Administration for Windows..... | 90 |
| Monitoring Batch Processes..... | 91 |
| Reviewing Batch Output Files..... | 91 |
| Running Reports from the Command Line for Windows..... | 92 |
| Scheduling Reports from the Command Line for Windows..... | 93 |
| Maintaining File Security for Windows..... | 94 |
| Running Multiple Instances of EnterpriseOne on Windows..... | 95 |
| Prerequisites..... | 95 |
| Running Multiple Instances of EnterpriseOne on Windows..... | 95 |
| Generating a Unique Identifier..... | 96 |
| Modifying the Server JDE.INI Files..... | 97 |
| Modifying the Workstation JDE.INI File..... | 98 |
| Uninstalling EnterpriseOne Services..... | 98 |

| | |
|---|----|
| Moving or Changing an EnterpriseOne Directory Tree..... | 98 |
|---|----|

Chapter 5

| | |
|--|------------|
| Server Administration Workbench..... | 101 |
| Understanding the Server Administration Workbench..... | 101 |
| Setting up Security Access for SAW..... | 102 |
| Setting Up Security Access for SAW..... | 102 |
| Monitoring EnterpriseOne with SAW on Windows..... | 102 |
| Understanding the SAW Interface..... | 103 |
| Adding to the JDE.INI File..... | 104 |
| Accessing SAW..... | 104 |
| Adding an EnterpriseOne Server for Querying..... | 104 |
| Checking EnterpriseOne Server Configurations..... | 104 |
| Checking EnterpriseOne Server Connectivity..... | 105 |
| Removing an EnterpriseOne Server from the Query List..... | 105 |
| Configuring EnterpriseOne Server Monitoring Settings..... | 105 |
| Understanding EnterpriseOne Server Monitoring Settings..... | 105 |
| Selecting the Server, Port, and Processes to Monitor..... | 107 |
| Setting up Event Notification..... | 107 |
| Setting up SAW Run-time Parameters..... | 108 |
| Monitoring EnterpriseOne Server Processes..... | 108 |
| Prerequisite..... | 109 |
| Monitoring EnterpriseOne Server Processes..... | 109 |
| Monitoring Enterprise Server Statistics..... | 110 |
| Understanding Enterprise Server Statistics..... | 110 |
| Monitoring Process Statistics for Enterprise Servers..... | 112 |
| Changing Process Settings..... | 112 |
| Viewing Log and Debug Log Files for Enterprise Server Processes..... | 112 |
| Printing Log and Debug Log Files for Enterprise Server Processes..... | 113 |
| Turning Logging on and off for Enterprise Server Processes..... | 113 |
| Viewing Process User Statistics for Enterprise Servers..... | 113 |
| Viewing Statistics for Batch Processes..... | 114 |
| Viewing Statistics for Server Disks..... | 114 |
| Monitoring EnterpriseOne Web Server Statistics..... | 114 |
| Understanding EnterpriseOne Web Server Statistics..... | 115 |
| Monitoring Data on Enterprise Servers Connected to the Web Server..... | 120 |
| Monitoring Web Server CallObjects..... | 120 |
| Monitoring Web Server Connection Statistics..... | 120 |
| Monitoring Web Server User Statistics..... | 120 |

| | |
|--|-----|
| Monitoring Web Server User Connection Statistics..... | 121 |
| Monitoring Web Server Logs..... | 121 |
| Querying EnterpriseOne Enterprise Servers..... | 121 |
| Understanding EnterpriseOne Enterprise Server Queries..... | 122 |
| Querying Enterprise Server Process Components..... | 124 |
| Querying Enterprise Server Disk Components..... | 124 |
| Querying Enterprise Server IPC Resources..... | 124 |
| Querying Enterprise Server Log Files..... | 125 |
| Querying the Enterprise Server Users..... | 125 |
| Querying the Enterprise Server Job Queue..... | 125 |
| Configuring SAW and Monitoring EnterpriseOne Servers from the Web..... | 126 |
| Understanding SAW and Server Monitoring from the Web..... | 126 |
| Creating an SMC File..... | 130 |
| Regenerating an SMC File..... | 131 |
| Modifying Advanced SMC Settings..... | 132 |
| Adding EnterpriseOne Servers..... | 132 |
| Adding Ports for Multiple Servers..... | 133 |
| Adding Events to Multiple Server Ports..... | 134 |
| Adding Statistics to Multiple Server Ports..... | 134 |
| Adding Profiles to Events..... | 135 |
| Customizing Server Ports..... | 135 |
| Changing Server Ports..... | 136 |
| Changing Statistics on a Server Port..... | 136 |
| Changing Events on a Server Port..... | 136 |
| Changing the Profiles of a Single Event..... | 137 |
| Running the SAW Agent in Windows..... | 137 |
| Running the SAW Agent in UNIX..... | 137 |
| Starting the SAW Agent on the AS/400..... | 138 |
| Monitoring Enterprise Servers from the Web..... | 138 |
| Understanding Enterprise Server Monitoring from the Web..... | 138 |
| Understanding the Server Summary View..... | 139 |
| Understanding the All Log Files View..... | 141 |
| Understanding the Active Log Files View..... | 141 |
| Understanding the PrintQueue Log Files View..... | 141 |
| Understanding the Disk Usage View..... | 141 |
| Understanding the INI Settings View..... | 142 |
| Understanding the Package Build Files View..... | 142 |
| Understanding the Environment Variables View..... | 142 |
| Troubleshooting Call-Object Processes..... | 142 |
| Monitoring the JAS Servers from the Web..... | 142 |

| | |
|--|-----|
| Understanding JAS Server Monitoring from the Web..... | 143 |
| Monitoring the Web Server System..... | 148 |
| Monitoring Virtual Clients..... | 150 |
| Monitoring Web Server Log Files..... | 151 |
| Using the Java Logging System Menu..... | 152 |
| Troubleshooting with the Web Server Monitor..... | 153 |
| Accessing the Web Server Monitor..... | 154 |
| Configuring the Web Server Monitor..... | 154 |
| Work with Connected Users..... | 154 |
| View Application Failure information..... | 155 |
| Monitoring Realtime Events..... | 156 |
| Monitoring JDENET Connection Pools..... | 158 |
| Monitoring Business Functions Running on Connected Enterprise Servers..... | 158 |
| Monitoring Database Connection Pools..... | 158 |
| Running SAW Web..... | 158 |
| Monitoring Web Server Log Files..... | 159 |
| Bouncing or Stopping the Web Server..... | 159 |
| Enabling Debug.log and Net Trace..... | 162 |
| Deleting Log Files..... | 163 |
| Monitoring System XPIe Servers from the Web..... | 163 |
| Monitoring EnterpriseOne with SAW on the iSeries..... | 167 |
| Prerequisite..... | 167 |
| Forms Used to Monitor EnterpriseOne with SAW on the iSeries..... | 167 |
| Understanding EnterpriseOne Monitoring with SAW on the iSeries..... | 168 |
| Accessing SAW..... | 169 |
| Affecting Server Processes..... | 169 |
| Displaying Disk Resources..... | 170 |
| Displaying, Printing, and Deleting Server Log Files..... | 170 |
| Modifying the Server jde.ini File..... | 170 |
| Monitoring EnterpriseOne with SAW on UNIX..... | 170 |
| Forms Used to Monitor EnterpriseOne with SAW on UNIX..... | 171 |
| Understanding EnterpriseOne Monitoring with SAW on UNIX..... | 171 |
| Accessing SAW..... | 172 |
| Displaying, Printing, and Deleting Server Log Files..... | 172 |
| Displaying and Deleting Log Files for Server Processes..... | 173 |
| Stopping Individual Kernel Processes..... | 174 |
| Displaying, Printing, and Deleting Server Log Files..... | 174 |
| Modifying the Server jde.ini File..... | 174 |
| Editing the Server JDE.INI File..... | 175 |
| To Activate or Deactivate Debugging..... | 175 |

| | |
|---|-----|
| To Change the Port Number of the Server..... | 175 |
| Monitoring EnterpriseOne with the Knowledge Module..... | 175 |
| Prerequisites..... | 176 |
| Forms Used to Monitor EnterpriseOne with the Knowledge Module..... | 176 |
| Understanding EnterpriseOne Monitoring with the Knowledge Module..... | 176 |
| Adding OWKM Files to the PATROL Environment..... | 180 |
| Loading OWKM Files..... | 180 |
| Adding a Host..... | 180 |
| Adding an EnterpriseOne Server..... | 181 |
| Monitoring Enterprise Servers Using KM Commands..... | 181 |

Chapter 6

| | |
|---|------------|
| Working with EnterpriseOne on Windows Terminal Server Edition..... | 185 |
| Understanding Windows Terminal Server Edition..... | 185 |
| Incorporating Citrix MetaFrame with TSE..... | 186 |
| Understanding TSE Restrictions in Multi-user Mode..... | 187 |
| Understanding Network Considerations..... | 187 |
| Understanding Performance Considerations..... | 188 |
| Setting Up EnterpriseOne on the Terminal Server..... | 188 |
| Setting Up EnterpriseOne on the Terminal Server..... | 188 |
| Troubleshooting EnterpriseOne on Windows Terminal Server Edition..... | 189 |
| Troubleshooting UBE Output Security on TSE..... | 189 |
| Submitting a UBE Locally and Running it on the TSE..... | 190 |
| Troubleshooting Import/Export with Microsoft Excel..... | 191 |
| Troubleshooting Specification Files are Locked..... | 191 |
| Reducing JITI Frequency..... | 191 |
| Troubleshooting User Cannot Restart EnterpriseOne..... | 191 |
| Troubleshooting Logging Off Versus Disconnecting..... | 191 |
| Troubleshooting Shortcuts Do Not Work in Email Messages..... | 192 |
| Troubleshooting Data Selection and Sequencing Criteria Lost..... | 192 |
| Troubleshooting Run-Time Error Occurs During Server Connection Test..... | 192 |
| Troubleshooting EnterpriseOne Development Tools Are Disabled..... | 192 |
| Troubleshooting Users Experience Problems Accessing EnterpriseOne..... | 192 |
| Troubleshooting Log Path is Incorrect..... | 192 |
| Troubleshooting Shortcut Path is Incorrect..... | 193 |
| Troubleshooting Only One User Can Log in to EnterpriseOne..... | 193 |

Chapter 7

| | |
|---|------------|
| Administering EnterpriseOne on a Unix Cluster..... | 195 |
| Understanding Clustering..... | 195 |
| Maintaining Multiple Instances of EnterpriseOne in a Clustered Environment..... | 196 |
| Setting Up Clustering..... | 196 |
| Configuring Oracle Parallel Server (OPS)..... | 196 |
| Setting Up an Oracle Package for MC/ServiceGuard..... | 197 |
| Setting Up an EnterpriseOne Package..... | 198 |
| Setting up HACMP for AIX Clustering..... | 200 |
| Understanding HACMP for AIX Clustering..... | 200 |
| Creating Group and User Accounts..... | 201 |
| Setting Up EnterpriseOne for HACMP..... | 202 |
| Editing the owenv Script..... | 202 |
| Editing the Start Resource Control Script..... | 203 |
| Editing the Stop Resource Control Script..... | 203 |
| Creating an Application Server..... | 204 |
| Prerequisite..... | 204 |
| Moving the Control Scripts..... | 204 |
| Defining an Application Sever..... | 205 |
| Defining Cluster Resources..... | 205 |
| Setting Up Sun Solaris Clustering..... | 206 |
| Understanding Sun Solaris Clustering..... | 206 |
| Modifying the SunStartResource.sh Script..... | 207 |
| Modifying the SunStopResource.sh Script..... | 207 |
| Modifying the owenv Script..... | 207 |
| Modifying the SunOracleMgr.sh Script..... | 207 |
| Registering EnterpriseOne with SUNClustering..... | 207 |
| Troubleshooting HP-UX Clustering..... | 208 |

Chapter 8

| | |
|---|------------|
| Administering EnterpriseOne on a Windows Cluster..... | 211 |
| Prerequisites..... | 211 |
| Upgrading EnterpriseOne in a Windows Cluster Environment..... | 212 |
| Setting Up EnterpriseOne on a Windows Cluster..... | 212 |
| Adding EnterpriseOne Network Resources to the Group..... | 214 |
| Starting and Stopping the EnterpriseOne Resources on the Cluster..... | 215 |
| Testing EnterpriseOne Connections on the Cluster..... | 215 |
| Setting Advanced Options for EnterpriseOne Resources..... | 216 |
| Setting Failover and Failback Parameters..... | 216 |

| | |
|---|-----|
| Troubleshooting Windows Clustering..... | 217 |
| Creating a Static IP Address for the PeopleSoft Network Name..... | 217 |
| Testing EnterpriseOne Connections on the Cluster..... | 218 |
| Setting Advanced Options for EnterpriseOne Resources..... | 218 |
| Setting Failover and Failback Parameters..... | 219 |
| Troubleshooting a Windows Cluster..... | 219 |
| Changing an IP Address..... | 219 |
| Reinstalling MSCS..... | 220 |

Chapter 9

| | |
|---|------------|
| Administering EnterpriseOne on an iSeries Cluster..... | 221 |
| Understanding iSeries Clustering..... | 221 |
| EnterpriseOne Objects Used with iSeries Clustering..... | 222 |
| Cluster Exit Program..... | 224 |
| Technical Considerations..... | 224 |
| Understanding iSeries-EnterpriseOne Architecture with Clustering..... | 225 |
| Minimum Setup Requirements for iSeries Server Nodes..... | 227 |
| Identifying the Cluster Name..... | 228 |
| Setting up the Enterprise Servers..... | 228 |
| Setting up the Client for Clustering..... | 228 |
| Setting up the Deployment Server..... | 229 |
| Setting up logical Data Sources..... | 229 |
| Setting up the logical data source for the system map..... | 230 |
| Setting up the logical data source for the server map..... | 230 |
| Setting up Database Data Sources..... | 231 |
| Setting up the server map database data sources..... | 231 |
| Setting up Object Configuration Manager for Clustering..... | 231 |
| Configuring OCM for logical data sources for the server map..... | 231 |
| Configuring OCM for logical data sources for the system map..... | 232 |
| Configuring OCM for database data sources..... | 233 |
| Configuring ODBC connections..... | 233 |
| Distributing the ODBC Setup from the Deployment Server..... | 234 |
| Identifying the Cluster Name on the Deployment Server..... | 234 |

Chapter 10

| | |
|--|------------|
| Backing Up EnterpriseOne Tables..... | 235 |
| Understanding Backup Requirements for Servers..... | 235 |
| Backing Up a Deployment Server..... | 235 |

| | |
|--|-----|
| Backing Up an Enterprise Server..... | 236 |
| Understanding EnterpriseOne Tables and Object Owner IDs..... | 238 |
| Backing Up EnterpriseOne Tables on Servers..... | 242 |
| Prerequisites..... | 243 |
| Creating a Backup for iSeries..... | 243 |
| Creating a Backup for Oracle on UNIX or Windows..... | 245 |
| Creating a Backup for SQL Server..... | 245 |
| Restoring a Backup File for Oracle on UNIX or Windows..... | 246 |
| Restoring a Backup File for iSeries..... | 246 |
| Restoring a Backup File for SQL Server..... | 247 |
| Restoring a Backup File for SQL Server on Windows..... | 247 |

Chapter 11

| | |
|--|------------|
| Working with SnapShot..... | 249 |
| Using SnapShot..... | 249 |
| Prerequisite..... | 249 |
| Saving EnterpriseOne in a SnapShot File..... | 249 |
| Restoring a SnapShot..... | 250 |
| Deleting a SnapShot..... | 250 |
| Running Snapshot from the Command Line..... | 250 |

Chapter 12

| | |
|--|------------|
| Generating Serialized Objects for the EnterpriseOne Web Server..... | 253 |
| Understanding Serialized Object Generation..... | 253 |
| Prerequisites..... | 254 |
| Installing eGenerator..... | 254 |
| Prerequisites..... | 254 |
| Setting the Default Storage Parameter (Oracle Only)..... | 254 |
| Setting the Default Storage Parameter (SQL Only)..... | 255 |
| Setting Up the Win32 Client to Generate Serialized Objects..... | 255 |
| Copying the genapp.ini File to the Default Windows Directory..... | 255 |
| Copying the jas.ini and tnsnames Files to the Generation Machine..... | 255 |
| Populating the Data Dictionary..... | 256 |
| Populating the Global Specs..... | 256 |
| Modifying gen.bat..... | 257 |
| Generating Serialized Objects..... | 258 |
| Generating Serialized Objects..... | 258 |
| Logging In..... | 259 |

| | |
|---|-----|
| Prerequisites..... | 259 |
| Understanding Logging In..... | 260 |
| Generating to the Web Server..... | 262 |
| Bypassing the Web Server and Generate Directly to a Database..... | 263 |
| Configuring eGenerator..... | 263 |
| Setting eGenerator Options..... | 263 |
| Choosing the HTMLOverrides File..... | 264 |
| Choosing Languages..... | 264 |
| Configuring the genapp.ini File..... | 264 |
| Generating All Standard Serialized Objects..... | 267 |
| Running the TAM Analyzer..... | 267 |
| Generating All EnterpriseOne Objects..... | 267 |

Chapter 13

| | |
|---|------------|
| Understanding Executable Files on the Workstation..... | 269 |
| EnterpriseOne Linked Executable Files..... | 269 |
| Standalone Executable Files..... | 272 |

Chapter 14

| | |
|--|------------|
| Troubleshooting the Workstation..... | 279 |
| Understanding Error Messages..... | 279 |
| Error Message Details..... | 279 |
| Error Messages Generated By Applications..... | 280 |
| Frequent Generic Error Messages..... | 280 |
| Memory Violations..... | 280 |
| Form and Grid Add Failures..... | 280 |
| Communication Failure..... | 281 |
| Troubleshooting the Workstation..... | 281 |
| Troubleshooting the Production Workstation..... | 282 |
| Understanding Production Workstation Troubleshooting..... | 282 |
| Performing Preliminary Troubleshooting..... | 283 |
| Troubleshooting Interactive Application Problems..... | 283 |
| Troubleshooting Batch Processes Resulting in No Data..... | 283 |
| Troubleshooting Batch Processes Displaying Errors on the Report..... | 284 |
| Troubleshooting Batch Processes Displaying Unexpected Data on the Report..... | 285 |
| Troubleshooting Batch Processes Ending in an Error when Submitted on the Server..... | 286 |
| Troubleshooting Local Data Availability Problems..... | 286 |
| Troubleshooting a Standalone Installation of EnterpriseOne..... | 286 |

| | |
|---|-----|
| Troubleshooting Enterprise Server Data Availability Problems..... | 286 |
| Troubleshooting Printing Problems..... | 287 |
| Troubleshooting .DLL Problems on a Production Workstation..... | 287 |
| Troubleshooting Data Source Setup Problems..... | 287 |
| Troubleshooting the Development Workstation..... | 288 |
| Understanding Development Workstation Troubleshooting..... | 288 |
| Troubleshooting .DLL Problems on a Development Workstation..... | 288 |
| Troubleshooting Event Rule Problems..... | 289 |
| Troubleshooting Business Function Problems..... | 289 |
| Working with the Workstation Log Files..... | 289 |
| Understanding the Workstation Log Files..... | 290 |
| Understanding Global Tables..... | 290 |
| Comprehending Logic Processing Logs..... | 290 |
| Understanding Application Development Logs..... | 291 |
| Understanding the Workstation jdedebug.log..... | 291 |
| Understanding the Batch Process Log..... | 292 |
| Understanding the sql.log..... | 292 |
| Using Log Files to Troubleshoot Strategies..... | 292 |
| Viewing Log Files..... | 293 |
| Setting Up the Workstation jde.log..... | 293 |
| Setting Up the Workstation jdedebug.log..... | 294 |
| Setting Up the Batch Process Log..... | 295 |
| Troubleshooting with the Compile Error Log..... | 296 |
| Troubleshooting with jdecpy.log..... | 297 |
| Troubleshooting with the sql.log..... | 297 |
| Activating sql.log..... | 297 |
| Troubleshooting ODBC Problems Using sql.log..... | 297 |
| Troubleshooting with the jdeinst.log..... | 298 |

Chapter 15

| | |
|--|------------|
| Troubleshooting the Enterprise Server..... | 299 |
| Understanding Enterprise Server Troubleshooting..... | 299 |
| Understanding the Enterprise Server jde.log File..... | 300 |
| Understanding the Enterprise Server jdedebug.log File..... | 302 |
| Understanding the Batch Process Log File..... | 305 |
| Viewing Enterprise Server Logs from the Workstation..... | 307 |
| Setting Up the Enterprise Server jde.log..... | 307 |
| Setting Up the Enterprise Server jdedebug.log..... | 308 |
| Setting Up the <batch_process>.log File..... | 308 |

| | |
|---|-----|
| Troubleshooting General Problems..... | 309 |
| Troubleshooting Communication Problems..... | 310 |
| Troubleshooting Deadlock Collision Problems..... | 311 |
| Troubleshooting Server Map Problems..... | 312 |
| Troubleshooting the iSeries Enterprise Server..... | 312 |
| Understanding iSeries Enterprise Server Troubleshooting..... | 312 |
| Troubleshooting iSeries Enterprise Server Installation..... | 313 |
| Troubleshooting Multiple Release Setup..... | 320 |
| Troubleshooting JDBNET..... | 320 |
| Troubleshooting Interprocess Communications..... | 321 |
| Troubleshooting the JDE.INI File..... | 322 |
| Troubleshooting the UNIX/Linux Enterprise Server..... | 323 |
| Understanding UNIX/Linux Enterprise Server Troubleshooting..... | 324 |
| Troubleshooting the JDE.INI File..... | 324 |
| Troubleshooting Copying EnterpriseOne Files to a Server..... | 325 |
| Troubleshooting Database Tables Configuration..... | 325 |
| Troubleshooting Printer Setup..... | 325 |
| Troubleshooting Email..... | 326 |
| Troubleshooting Multiple Release Setup..... | 326 |
| Troubleshooting Finding the Report Files..... | 326 |
| Troubleshooting JDBNET Server Not Found..... | 326 |
| Troubleshooting EnterpriseOne Testing..... | 327 |
| Troubleshooting the Windows Enterprise Server..... | 327 |
| Understanding Windows Enterprise Server Troubleshooting..... | 327 |
| Troubleshooting EnterpriseOne Account Setup..... | 328 |
| Troubleshooting Copying EnterpriseOne Files to a Server..... | 328 |
| Troubleshooting Database Table Configuration..... | 328 |
| Troubleshooting Printer Setup..... | 329 |
| Troubleshooting JDE.INI File Setup..... | 329 |
| Troubleshooting Finding the Log Files..... | 329 |
| Troubleshooting Testing with PORTTEST..... | 330 |
| Troubleshooting Running EnterpriseOne Manually..... | 331 |
| Troubleshooting Finding the Report Files..... | 332 |
| Troubleshooting Testing EnterpriseOne by Submitting a Report..... | 332 |
| Taking Ownership of a Printer..... | 334 |
| Stopping All EnterpriseOne Processes..... | 334 |
| Stopping EnterpriseOne Processes without Rights..... | 335 |
| Troubleshooting Email..... | 335 |
| Troubleshooting Web Servers..... | 336 |
| Understanding Web Server Troubleshooting..... | 336 |

Troubleshooting IIS and IBM HTTP Web Servers.....336

Troubleshooting JAS.....336

Troubleshooting Serialized Database and Generation Issues.....336

Troubleshooting SQL Server Issues.....337

Troubleshooting Problems Using Log Files.....337

Glossary of PeopleSoft Terms.....339

Index359

About This PeopleBook

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

This preface discusses:

- PeopleSoft application prerequisites.
- PeopleSoft application fundamentals.
- Documentation updates and printed documentation.
- Additional resources.
- Typographical conventions and visual cues.
- Comments and suggestions.
- Common elements in PeopleBooks.

Note. PeopleBooks document only page elements, such as fields and check boxes, that require additional explanation. If a page element 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. Elements that are common to all PeopleSoft applications are defined in this preface.

PeopleSoft Application Prerequisites

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

You might also want to complete at least one PeopleSoft introductory training course, if applicable.

You should be familiar with navigating the system and adding, updating, and deleting information by using PeopleSoft menus, and pages, forms, or windows. 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 PeopleSoft applications most effectively.

PeopleSoft Application Fundamentals

Each application PeopleBook provides implementation and processing information for your PeopleSoft applications. For some applications, additional, essential information describing the setup and design of your system appears in a companion volume of documentation called the application fundamentals PeopleBook. Most PeopleSoft product lines have a version of the application fundamentals PeopleBook. The preface of each PeopleBook identifies the application fundamentals PeopleBooks that are associated with that PeopleBook.

The application fundamentals PeopleBook consists of important topics that apply to many or all PeopleSoft applications across one or more product lines. Whether you are implementing a single application, some combination of applications within the product line, or the entire product line, you should be familiar with the contents of the appropriate application fundamentals PeopleBooks. They provide the starting points for fundamental implementation tasks.

Documentation Updates and Printed Documentation

This section discusses how to:

- Obtain documentation updates.
- Order printed documentation.

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'll find a variety of useful and timely materials, including updates to the full PeopleSoft documentation that is delivered on your PeopleBooks CD-ROM.

Important! 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, <https://www.peoplesoft.com/corp/en/login.jsp>

Ordering Printed Documentation

You can order printed, bound volumes of the complete PeopleSoft documentation that is delivered on your PeopleBooks CD-ROM. PeopleSoft makes printed documentation available for each major release shortly after the software is shipped. Customers and partners can order printed PeopleSoft documentation by using any of these methods:

- Web
- Telephone
- Email

Web

From the Documentation section of the PeopleSoft Customer Connection website, access the PeopleBooks Press website under the Ordering PeopleBooks topic. The PeopleBooks Press website is a joint venture between PeopleSoft and MMA Partners, the book print vendor. Use a credit card, money order, cashier's check, or purchase order to place your order.

Telephone

Contact MMA Partners at 877 588 2525.

Email

Send email to MMA Partners at peoplesoftpress@mmapartner.com.

See Also

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

Additional Resources

The following resources are located on the PeopleSoft Customer Connection website:

| Resource | Navigation |
|--|--|
| Application maintenance information | Updates + Fixes |
| Business process diagrams | Support, Documentation, Business Process Maps |
| Interactive Services Repository | Interactive Services Repository |
| Hardware and software requirements | Implement, Optimize + Upgrade, Implementation Guide, Implementation Documentation & Software, Hardware and Software Requirements |
| Installation guides | Implement, Optimize + Upgrade, Implementation Guide, Implementation Documentation & Software, Installation Guides and Notes |
| Integration information | Implement, Optimize + Upgrade, Implementation Guide, Implementation Documentation and Software, Pre-built Integrations for PeopleSoft Enterprise and PeopleSoft EnterpriseOne Applications |
| Minimum technical requirements (MTRs) (EnterpriseOne only) | Implement, Optimize + Upgrade, Implementation Guide, Supported Platforms |
| PeopleBook documentation updates | Support, Documentation, Documentation Updates |
| PeopleSoft support policy | Support, Support Policy |
| Prerelease notes | Support, Documentation, Documentation Updates, Category, Prerelease Notes |
| Product release roadmap | Support, Roadmaps + Schedules |
| Release notes | Support, Documentation, Documentation Updates, Category, Release Notes |
| Release value proposition | Support, Documentation, Documentation Updates, Category, Release Value Proposition |
| Statement of direction | Support, Documentation, Documentation Updates, Category, Statement of Direction |

| Resource | Navigation |
|-----------------------------|---|
| Troubleshooting information | Support, Troubleshooting |
| Upgrade documentation | Support, Documentation, Upgrade Documentation and Scripts |

Typographical Conventions and Visual Cues

This section discusses:

- Typographical conventions.
- Visual cues.
- Country, region, and industry identifiers.
- Currency codes.

Typographical Conventions

This table contains the typographical conventions that are used in PeopleBooks:

| Typographical Convention or Visual Cue | Description |
|--|---|
| Bold | Indicates PeopleCode function names, business function names, event names, system function names, method names, language constructs, and PeopleCode reserved words that must be included literally in the function call. |
| <i>Italics</i> | Indicates field values, emphasis, and PeopleSoft or other book-length publication titles. In PeopleCode syntax, italic items are placeholders for arguments that your program must supply. We also use italics when we refer to words as words or letters as letters, as in the following: Enter the letter <i>O</i> . |
| KEY+KEY | Indicates a key combination action. For example, a plus sign (+) between keys means that you must hold down the first key while you press the second key. For ALT+W, hold down the ALT key while you press the W key. |
| Monospace font | Indicates a PeopleCode program or other code example. |
| “ ” (quotation marks) | Indicate chapter titles in cross-references and words that are used differently from their intended meanings. |

| Typographical Convention or Visual Cue | Description |
|--|--|
| ... (ellipses) | Indicate that the preceding item or series can be repeated any number of times in PeopleCode syntax. |
| { } (curly braces) | Indicate a choice between two options in PeopleCode syntax. Options are separated by a pipe (). |
| [] (square brackets) | Indicate optional items in PeopleCode syntax. |
| & (ampersand) | When placed before a parameter in PeopleCode syntax, an ampersand indicates that the parameter is an already instantiated object. Ampersands also precede all PeopleCode variables. |

Visual Cues

PeopleBooks contain the following visual cues.

Notes

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

Note. Example of a note.

If the note is preceded by *Important!*, the note is crucial and includes information that concerns what you must do for the system to function properly.

Important! Example of an important note.

Warnings

Warnings indicate crucial configuration considerations. Pay close attention to warning messages.

Warning! Example of a warning.

Cross-References

PeopleBooks provide cross-references either under the heading “See Also” or on a separate line preceded by the word *See*. Cross-references lead to other documentation that is pertinent to the immediately preceding documentation.

Country, Region, and Industry Identifiers

Information that applies only to a specific country, region, or industry is preceded by a standard identifier in parentheses. This identifier typically appears at the beginning of a section heading, but it may also appear at the beginning of a note or other text.

Example of a country-specific heading: “(FRA) Hiring an Employee”

Example of a region-specific heading: “(Latin America) Setting Up Depreciation”

Country Identifiers

Countries are identified with the International Organization for Standardization (ISO) country code.

Region Identifiers

Regions are identified by the region name. The following region identifiers may appear in PeopleBooks:

- Asia Pacific
- Europe
- Latin America
- North America

Industry Identifiers

Industries are identified by the industry name or by an abbreviation for that industry. The following industry identifiers may appear in PeopleBooks:

- USF (U.S. Federal)
- E&G (Education and Government)

Currency Codes

Monetary amounts are identified by the ISO currency code.

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 send email comments to doc@peoplesoft.com.

While we cannot guarantee to answer every email message, we will pay careful attention to your comments and suggestions.

Common Elements Used in PeopleBooks

Address Book Number

Enter a unique number that identifies the master record for the entity. An address book number can be the identifier for a customer, supplier, company, employee, applicant, participant, tenant, location, and so on. Depending on the application, the field on the form might refer to the address book number as the customer number, supplier number, or company number, employee or applicant id, participant number, and so on.

| | |
|----------------------------|---|
| As If Currency Code | Enter the three-character code to specify the currency that you want to use to view transaction amounts. This code allows you to view the transaction amounts as if they were entered in the specified currency rather than the foreign or domestic currency that was used when the transaction was originally entered. |
| Batch Number | Displays a number that identifies a group of transactions to be processed by the system. On entry forms, you can assign the batch number or the system can assign it through the Next Numbers program (P0002). |
| Batch Date | Enter the date in which a batch is created. If you leave this field blank, the system supplies the system date as the batch date. |
| Batch Status | <p>Displays a code from user-defined code (UDC) table 98/IC that indicates the posting status of a batch. Values are:</p> <p><i>Blank:</i> Batch is unposted and pending approval.</p> <p><i>A:</i> The batch is approved for posting, has no errors and is in balance, but it has not yet been posted.</p> <p><i>D:</i> The batch posted successfully.</p> <p><i>E:</i> The batch is in error. You must correct the batch before it can post.</p> <p><i>P:</i> The system is in the process of posting the batch. The batch is unavailable until the posting process is complete. If errors occur during the post, the batch status changes to E.</p> <p><i>U:</i> The batch is temporarily unavailable because someone is working with it, or the batch appears to be in use because a power failure occurred while the batch was open.</p> |
| Branch/Plant | Enter a code that identifies a separate entity as a warehouse location, job, project, work center, branch, or plant in which distribution and manufacturing activities occur. In some systems, this is called a business unit. |
| Business Unit | Enter the alphanumeric code that identifies a separate entity within a business for which you want to track costs. In some systems, this is called a branch/plant. |
| Category Code | Enter the code that represents a specific category code. Category codes are user-defined codes that you customize to handle the tracking and reporting requirements of your organization. |
| Company | Enter a code that identifies a specific organization, fund, or other reporting entity. The company code must already exist in the F0010 table and must identify a reporting entity that has a complete balance sheet. |
| Currency Code | Enter the three-character code that represents the currency of the transaction. PeopleSoft EnterpriseOne provides currency codes that are recognized by the International Organization for Standardization (ISO). The system stores currency codes in the F0013 table. |
| Document Company | <p>Enter the company number associated with the document. This number, used in conjunction with the document number, document type, and general ledger date, uniquely identifies an original document.</p> <p>If you assign next numbers by company and fiscal year, the system uses the document company to retrieve the correct next number for that company.</p> |

If two or more original documents have the same document number and document type, you can use the document company to display the document that you want.

Document Number

Displays a number that identifies the original document, which can be a voucher, invoice, journal entry, or time sheet, and so on. On entry forms, you can assign the original document number or the system can assign it through the Next Numbers program.

Document Type

Enter the two-character UDC, from UDC table 00/DT, that identifies the origin and purpose of the transaction, such as a voucher, invoice, journal entry, or time sheet. PeopleSoft EnterpriseOne reserves these prefixes for the document types indicated:

P: Accounts payable documents.

R: Accounts receivable documents.

T: Time and pay documents.

I: Inventory documents.

O: Purchase order documents.

S: Sales order documents.

Effective Date

Enter the date on which an address, item, transaction, or record becomes active. The meaning of this field differs, depending on the program. For example, the effective date can represent any of these dates:

- The date on which a change of address becomes effective.
- The date on which a lease becomes effective
- The date on which a price becomes effective.
- The date on which the currency exchange rate becomes effective.
- The date on which a tax rate becomes effective.

Fiscal Period and Fiscal Year

Enter a number that identifies the general ledger period and year. For many programs, you can leave these fields blank to use the current fiscal period and year defined in the Company Names & Number program (P0010)

G/L Date (general ledger date)

Enter the date that identifies the financial period to which a transaction will be posted. The system compares the date that you enter on the transaction to the fiscal date pattern assigned to the company to retrieve the appropriate fiscal period number and year, as well as to perform date validations.

PeopleSoft EnterpriseOne Preface

This guide describes tasks required by EnterpriseOne network administrators to configure and optimize an EnterpriseOne network. Although this guide addresses the needs of EnterpriseOne network administrators, employees with other job functions might find the information useful or essential to their positions as well.

Server and Workstation Administration

The EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration describes how to perform these types of tasks:

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

CHAPTER 1

Getting Started

This chapter provides an overview of the Configurable Network Computing Implementation PeopleBooks and administrator roles and discusses how to get started.

Understanding the EnterpriseOne Configurable Network Computing Implementation PeopleBooks

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 these guides:

- EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation

Intended primarily for CNC specialists, this guide contains information about these topics:

- Middleware.
- Data sources.
- Path codes and environments.
- Object Configuration Manager.
- Modes of processing.
- A typical EnterpriseOne customer configuration.

- EnterpriseOne Tools 8.94 PeopleBook: System Administration

Intended for system administrators, this guide contains information about these 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.
- EnterpriseOne Tools 8.94 PeopleBook: Package Management

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

 - Package management planning and setup.
 - Modification rules.
 - Object management.
 - Building packages.
 - Deploying packages.
 - Server packages.
 - Multitier deployment.
- EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration

Written primarily for EnterpriseOne network administrators, this guide contains information about these 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 EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation) might also be responsible for building and deploying packages (described in the EnterpriseOne Tools 8.94 PeopleBook: Package Management).

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

- 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. The PeopleSoft company does not provide documentation but does provide these 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 these 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.

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.

Implementing Server and Workstation Administration

Some of the functionality to Administer EnterpriseOne servers and workstations is delivered as part of the standard PeopleSoft tools that are provided with all EnterpriseOne installations. The rest is provided by the particular servers (Windows, Unix, and iSeries) that you have configured as part of the EnterpriseOne network.

Before you begin to administer servers and workstations in your environment, you should complete these tasks:

- Install EnterpriseOne on the Deployment Server, Enterprise Servers, and Workstations.

See EnterpriseOne PeopleTools 8.11 Installation guide for the platform and database of the enterprise server.

- Set up security and user roles for EnterpriseOne users.

See *EnterpriseOne Tools 8.94 PeopleBook: System Administration*.

CHAPTER 2

Understanding Server Administration for iSeries

This chapter provides an overview of server administration for iSeries and describes how to:

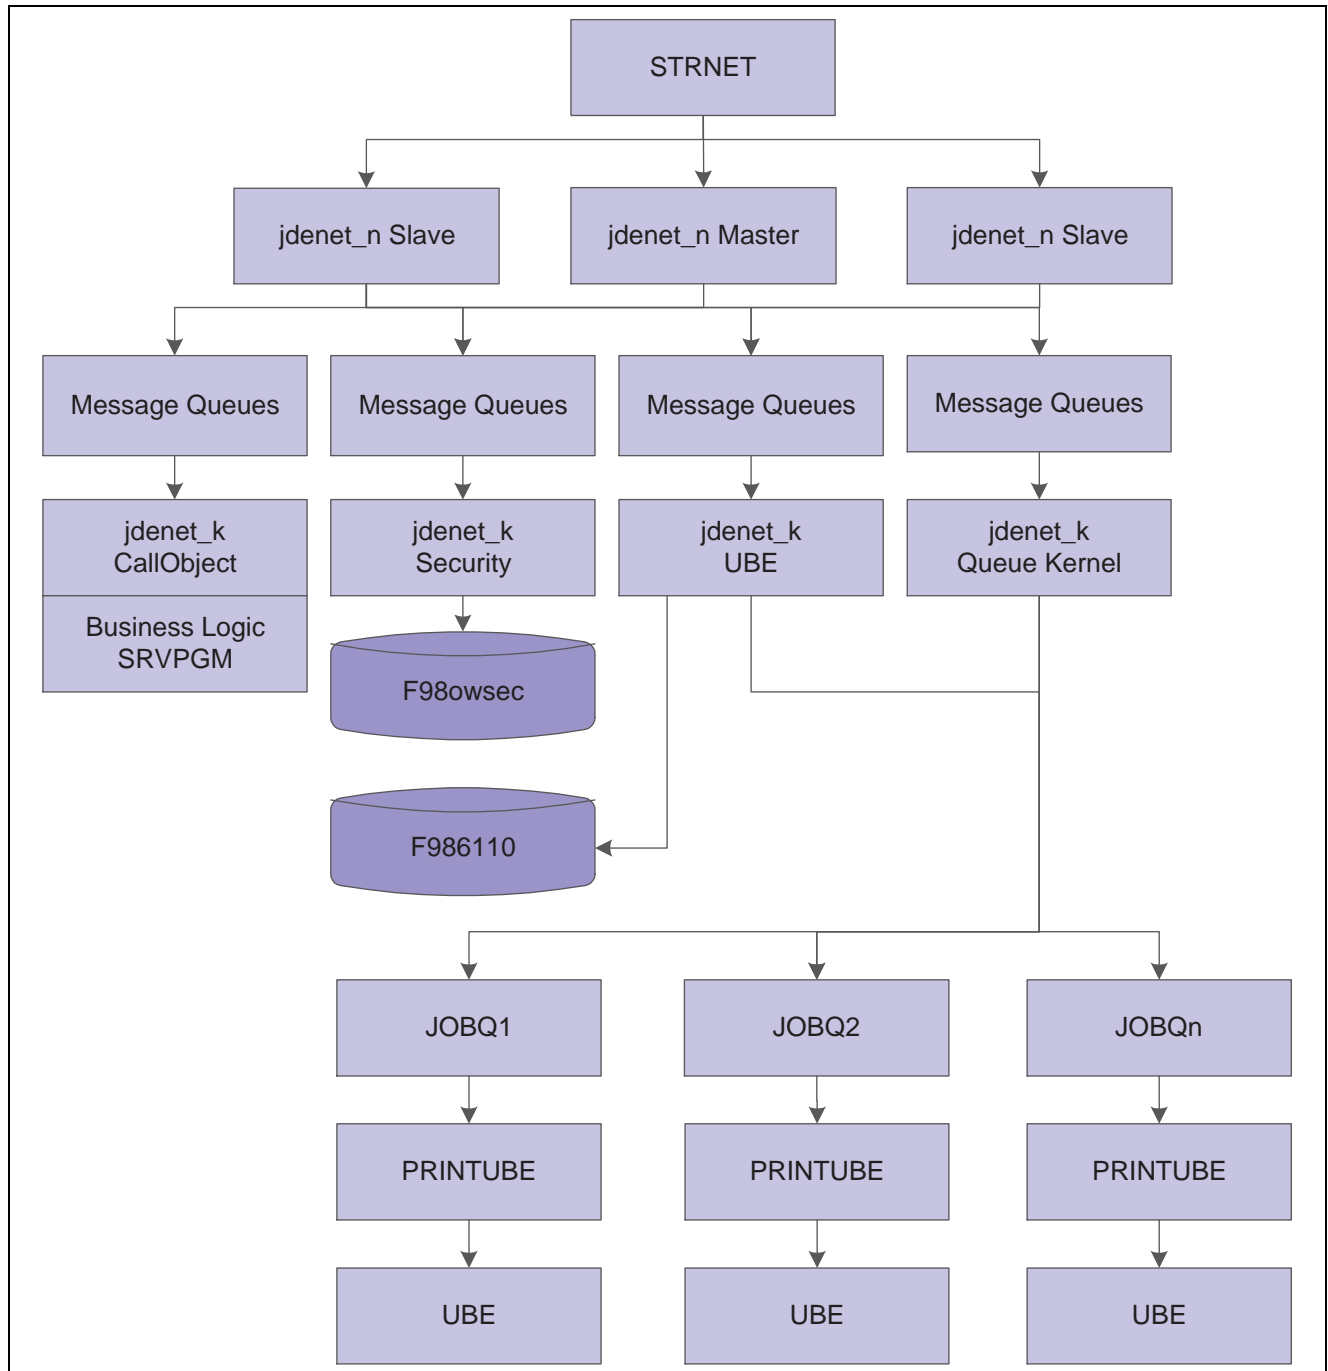
- Start the enterprise server for iSeries.
- Shut down the enterprise server for iSeries.
- Use iSeries integrated file system logging support.
- Clean up the enterprise server for iSeries.
- Set up a printer for iSeries.
- Administer batch processes for iSeries.
- Run multiple instances for EnterpriseOne on the iSeries.
- Administer security EnterpriseOne database security for iSeries.

Understanding Server Administration for iSeries

The PeopleSoft company 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 the enterprise server to ensure that EnterpriseOne runs properly.

Understanding the EnterpriseOne iSeries Architecture and Process Flow for iSeries

This flowchart illustrates the actions that the host server processes perform:



iSeries Process Flow

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.

The process flow is as follows:

1. The STRNET command runs the master NETWORK (JDENET_N) job in a newly started subsystem. 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 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, and this number must be the same for every client who wants 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 purpose of these slave processes is 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.

3. JDENET_K processes (kernel processes) 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 assigns each kernel process to a certain type.
4. 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 Service Program (SRVPGM). SRVPGM is loaded onto the JDENET_K processes and then called directly through a C function call.
5. The JDENET_K process that becomes a batch process kernel waits for requests to run batch processes from the client. When a request to run a batch process is submitted, these events occur:
 - JDENET_K (UBE kernel) adds a record to the F986110 database table with a status of W for waiting.
 - JDENET_K (UBE kernel) 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.

6. The PRINTUBE process runs the batch application, and changes the status of the record in the F986110 table to P for processing.
7. 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

This 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 this 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).

Note. We 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 display 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

This section provides an overview of the EnterpriseOne library structure and startup options for iSeries and describes how to:

- Start the enterprise server for iSeries manually.
- Start the enterprise server for iSeries automatically.

Prerequisites

Before you complete the tasks in this section:

- 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 these library to the top of the library list before you start EnterpriseOne on the enterprise server: *releaseSYS* (or the system library name). The variable *release* is the EnterpriseOne release level, such as E811SYS.

The *releaseSYS* library contains these objects:

| Object | Description |
|------------------|--|
| 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. |

| Object | Description |
|--------------------|---|
| 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 display 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 PSFTrelease, where release represents the EnterpriseOne release, such as /PSFT811. |
| 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 (*CMD) | The Convert Process ID to Job command, which returns the job information when the system 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. |
| 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 (*JOBQ) | The job description used by EnterpriseOne iSeries network jobs. |
| JDENET (*CLS) | The class used to create the routing entry for the JDENET subsystem. |

| Object | Description |
|---------------------|--|
| 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. |
| PSFTrelease (*SBSD) | The subsystem description under which the EnterpriseOne network jobs run. The variable release is the EnterpriseOne release level, such as PSFT811. |

Understanding Startup Options for the Enterprise Server for iSeries

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

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.

If you start the server automatically, we recommend that you separate the EnterpriseOne add library list entry (ADDLIB) 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. This recommendation also ensures that the 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.

Starting the Enterprise Server for iSeries Manually

To start the enterprise server for iSeries manually:

1. Sign on to the iSeries as *ONEWORLD*.
2. Start JDENet using this command:

```
STRNET
```

3. Start the PORTTEST program using this 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 in this 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 this command:

```
WRKACTJOB SBS(PSFTrelease)
```

The variable *release* is the EnterpriseOne release level that the site is using, such as PSFT811.

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).

Starting the Enterprise Server for iSeries Automatically

To start the enterprise server for iSeries automatically:

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(E811SYS 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 this command:

```
DSPSYSVAL SYSVAL(QSTRUPPGM)
```

4. Determine where the source of the program is located by executing this 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.

Shutting Down the Enterprise Server for iSeries

You can manually shut down the enterprise server for the iSeries with the command, ENDNET. This command is in the system library. For example, ENDNET causes EnterpriseOne to end the JDENet jobs and clean up all JDENet runtime structures.

Prerequisite

Ensure that the library is set correctly before performing this command.

See Also

[Chapter 2, “Understanding Server Administration for iSeries,” Understanding the iSeries Library Structure for EnterpriseOne, page 8](#)

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

These examples illustrate how to modify the JDE.INI file to enable easier access to log files from the workstation:

```
[DEBUG]
DebugFile=jdedebug
JobFile=jde.log
```

EnterpriseOne generates log files in the IFS root directory.

```
[DEBUG]
DebugFile=/psft811_a/jdedebug
JobFile=/psft811_a/jde.log
EnterpriseOne generates log files in the IFS directory called /psft811_a.
```

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

Cleaning Up the Enterprise Server for iSeries

This section provides an overview of cleaning up the enterprise server for iSeries and describes how to:

- Clean up the enterpriser server for iSeries.
- Clear the jde.log and jde.debug files for iSeries.

Prerequisite

Ensure that the 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.

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 after 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 enable 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 jdedebug files. This action keeps the files from becoming too large and removes old messages from it.

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

Cleaning Up the Enterprise Server for iSeries

To clean up the Enterprise server for iSeries:

From an iSeries command line, enter these IPCS commands:

```
DSPIPC
CLRIPC
```

Clearing the jde.log and jde.debug Files for iSeries

To clear the jde.log and jdedebug files for iSeries:

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

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

Where *release* is the EnterpriseOne release, such as psft811.

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

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

Where *release* is the EnterpriseOne release, such as psft811.

Setting Up a Printer for iSeries

This section provides an overview of setting up printers for iSeries and describes how to:

- Create the OUTQ.
- Start the OUTQ.
- Print multiple copies to a remote printer.

Understanding Printer Setup 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

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Setting Up EnterpriseOne Printing”

Creating the OUTQ

To create the OUTQ:

Enter this 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 the printer for additional information. For example, you must set this parameter to PASS for the IBM Network Printer 4317.

Starting the OUTQ

To start the OUTQ:

1. Enter this command:

```
STRRMTWTR outqname
```

For example:

```
STRRMTWTR QGPL/JDE_HP4PSB
```

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

```
RLSOUTQ outqname
```

For example, enter DEL '/PSFT*release*, where *release* is the EnterpriseOne release, as in PSFT811.

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.

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

Administering Batch Processes for iSeries

This section provides an overview of batch process administration for iSeries and describes how to:

- Monitor batch processes.
- Review batch output files.
- Encode the passwords of users who submit batch jobs.

Understanding Batch Process Administration 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, PSFT811, 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.

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 E811SYS/PRINTQUEUE file.

When users submit batch reports to run on the iSeries, a corresponding Portable Document Format (PDF) file is created on the enterprise server.

The default location for the PDF files is under the PRINTQUEUE folder of the installation directory in IFS, for example, /E811SYS/PRINTQUEUE. Users can access the PDF files directly on the enterprise server, or go to the submitted jobs on the client and view the PDF file.

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 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 this manner:

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



```
PrintImmediate=TRUE
```

Setting SaveOutput to TRUE causes the enterprise server to save the PDF files in E811SYS/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.

Finally, if you have the proper authority, you can run batch process reports from the server command line with this command:

```
RUNUBE USER(USER) PASSWORD(PASSWORD) ENVIRON(ENVIRONMENT)
REPORT(REPORTNAME) VERSION(VERSION)
```

Example: Running Reports from the Command Line for iSeries

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

```
RUNUBE USER(SF5488324) PASSWORD(PASSWORD) ENVIRON(PD811)
REPORT(R0006P) VERSION(XJDE0001)
```

This command begins processing version XJDE0001 of the report in the PD811 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 these example:

```
SBMJOB CMD(RUNUBE USER(SF5488324) PASSWORD(PASSWORD) ENVIRON(PD811)
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.

Monitoring Batch Processes

To monitor batch processes:

1. Sign on to the iSeries enterprise server using an administrative account.
2. Enter this command, substituting Subsystem with the appropriate subsystem name:

```
WRKACTJOB SBS (Subsystem)
```

Reviewing Batch Output Files

Complete this task to review the PDF output files.

To review batch output files:

1. From Windows Explorer, use this 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 /E811SYS/PrintQueue), and view the PDF files.

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

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

Running Multiple Instances of EnterpriseOne on the iSeries

This section provides an overview of running multiple instances of EnterpriseOne and database security parameters on the iSeries and describes how to:

- Copy libraries and directories.
- Apply security to multiple instances of EnterpriseOne on the iSeries.
- Create an EnterpriseOne subsystem on the iSeries.

Prerequisites

Before you complete the tasks in this section:

- 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.
- Set up security for multiple instances of EnterpriseOne.

Running Multiple Instances of EnterpriseOne

You might want to run multiple instances of EnterpriseOne on an iSeries server for these 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.

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

- 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 this 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. |

| Section in server JDE.INI file | Parameter | Purpose |
|--------------------------------|-------------------|--|
| [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. |

Similarly, to apply EnterpriseOne security throughout multiple instances, you use these 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 this 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] | JDTSTFile | 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. |

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

- Copy needed libraries and directories and modify the values of selected parameters in the .ini library.
To create an instance of EnterpriseOne on the iSeries, you copy these objects:
 - System library
 - System directory
 - Path code library
 - Path code directory
- Apply security to multiple instances of EnterpriseOne, if you desire to do so.

If you want to apply security to multiple instances of EnterpriseOne, complete the steps in these 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.

- Create a new EnterpriseOne subsystem identification.

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 we ship 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.

See Also

Chapter 2, “Understanding Server Administration for iSeries,” Administering Security EnterpriseOne Database Security for iSeries, page 28

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Environment Setup,” Working with an Environment

Chapter 2, “Understanding Server Administration for iSeries,” Setting Up iSeries Database Security for Multiple EnterpriseOne Instances, page 40

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Working with Servers,” Managing EnterpriseOne Subsystems

Understanding 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 Set Up 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 enable 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 these:

- 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 enables you to specify that the JDE.INI file is either not needed or not available.

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

Use a library name if all of these 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)

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 these 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.OMDATB2) that meet these criteria:

- F986101.OMDATP = F98611.OMDATP
- OMUGRP = *PUBLIC, OMSTSO = 'AV'
- OMSRVR = the host name

Modify System Profile

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 these:

- Modify or create the system profile that has not yet been modified. For example, you might enter this 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

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 these:

- Modify or create the JDE profile that has not been modified. For example, you might enter these 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

You can enter up to 10 security profiles at a time in this field to modify using the SETOWAUT program.

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.

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 these:

- Create or modify the profile by entering information such as these:
 - *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.

Sample Results for SETOWAUT in the EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration.

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, this 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 enable 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))

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.

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

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 these 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 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 we recommend 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.

Copying Libraries and Directories

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 this command:

```
CPYLIB E811SYS E811CST
```

Where E811CST 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 this command to create a temporary library:

```
CRTLIB TEMPLIB
```

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

```
CRTSAVF FILE (TEMPLIB/E811SYS)
```

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

```
SAV DEV ('/QSYS.LIB/TEMPLIB/E811SYS.FILE') OBJ ((' /E811SYS)) USEOPTBLK(*NO) DTACPR(*YES)
```

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

```
RST DEV ('/QSYS.LIB/TEMPLIB/E811SYS.FILE') OBJ ((' /E811sys/*' *INCLUDE/E811cst'))
```

Where E811cst 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 this command:

```
CPYLIB PRD811 CST811
```

Where CST811 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 this command to create a save file in the temporary library:

```
CRTSAVF FILE(TEMPLIB/PRD811)
```

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 this command:

```
SAV DEV ('QSYS.LIB/TEMPLIB/PRD811.FILE') OBJ ((' /prd811/*')) USEOPTBLK(*NO) DTACPR(*YES)
```

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

```
RST DEV('QSYS.LIB/TEMPLIB/PRD811.FILE') OBJ('/prd811/*' INCLUDE '/cst811'))
```

Where cst811 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 this 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 these 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

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

1. Copy the OCM library.
2. Copy the database libraries, such as SYS811, 811MAP, 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 this command:

```
CRTDIR DIR('/811CSTLOG')
```

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

5. Modify these 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 F98OWSEC library>
```

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

Creating an EnterpriseOne Subsystem on the iSeries

To create an EnterpriseOne subsystem on the iSeries:

1. Stop EnterpriseOne services.
2. From the iSeries main menu, type this command, and then press ENTER or press the F4 key:

```
CRTOWSBS
```

3. On the CREATE New EnterpriseOne Subsystem form, enter character values for these parameters, and then press ENTER:

- SUBSYSTEM
- SYSLIB

Note. 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.

4. To delete the old subsystem description from the system library, type this 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.

5. In the Work with Objects form, type 4 for Delete, and then press ENTER.
6. From the iSeries main menu, clear IPC memory by typing this command:

```
CLRIPC
```

7. From the iSeries main menu, start EnterpriseOne iSeries services by typing this command:

```
STRNET
```

Administering Security EnterpriseOne Database Security for iSeries

This section provides an overview of EnterpriseOne data base security administration and describes how to:

- Set up iSeries database security for a single EnterpriseOne instance.

- Set up iSeries database security for multiple EnterpriseOne instances.
- Add administrators.
- Remove administrative authority from user profiles.
- Display user profile information.

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.

Understanding EnterpriseOne Database Security Administration

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 enables you to enter specific security information for the 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 enables 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.

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 these:

- 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 Set Up 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 these 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. |

This 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. |

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.

You can set up security for separate instances of EnterpriseOne as well. 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.

Note. Use caution when you use EnterpriseOne security to lock a library that contains third-party software. We do not support the iSeries EnterpriseOne database security with third-party software.

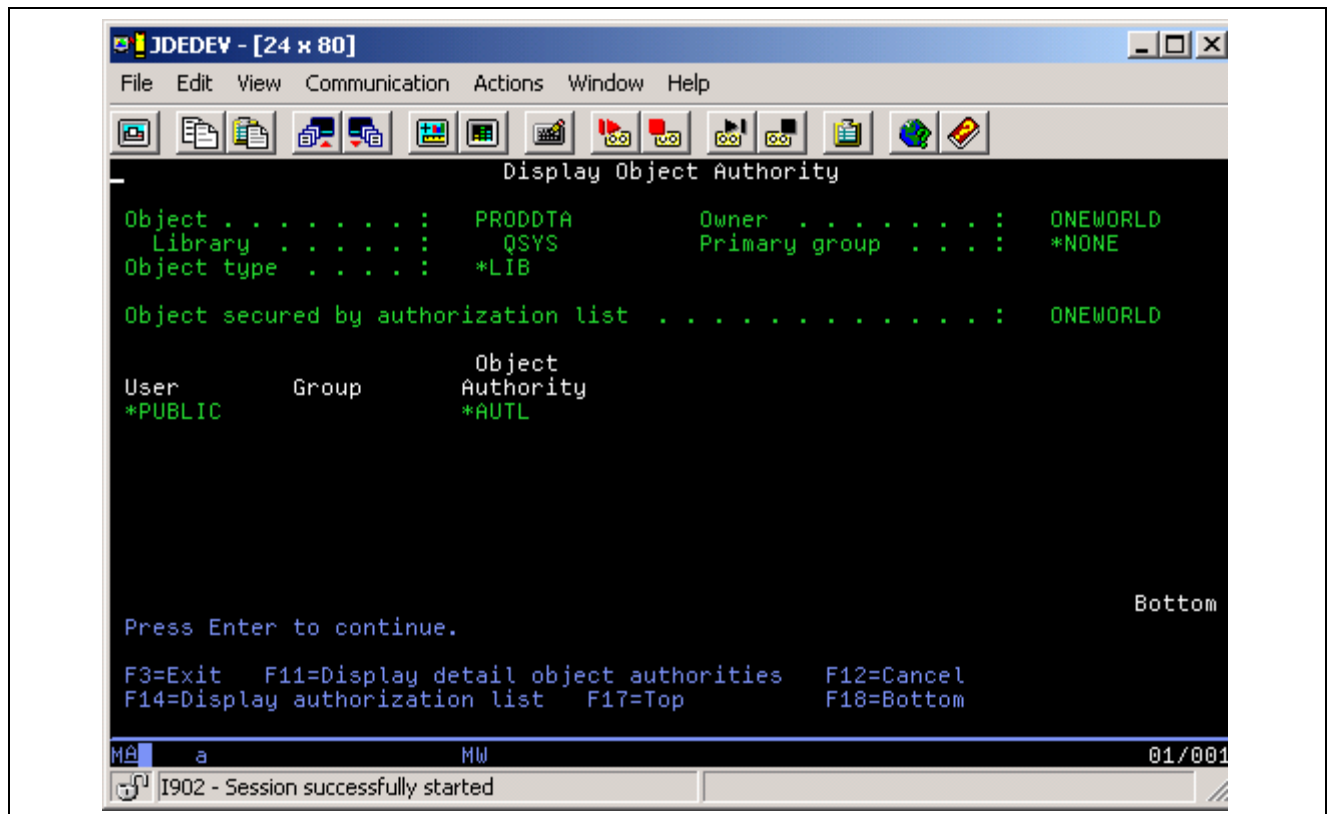
Sample Results for SETOWAUT

You can expect these 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.

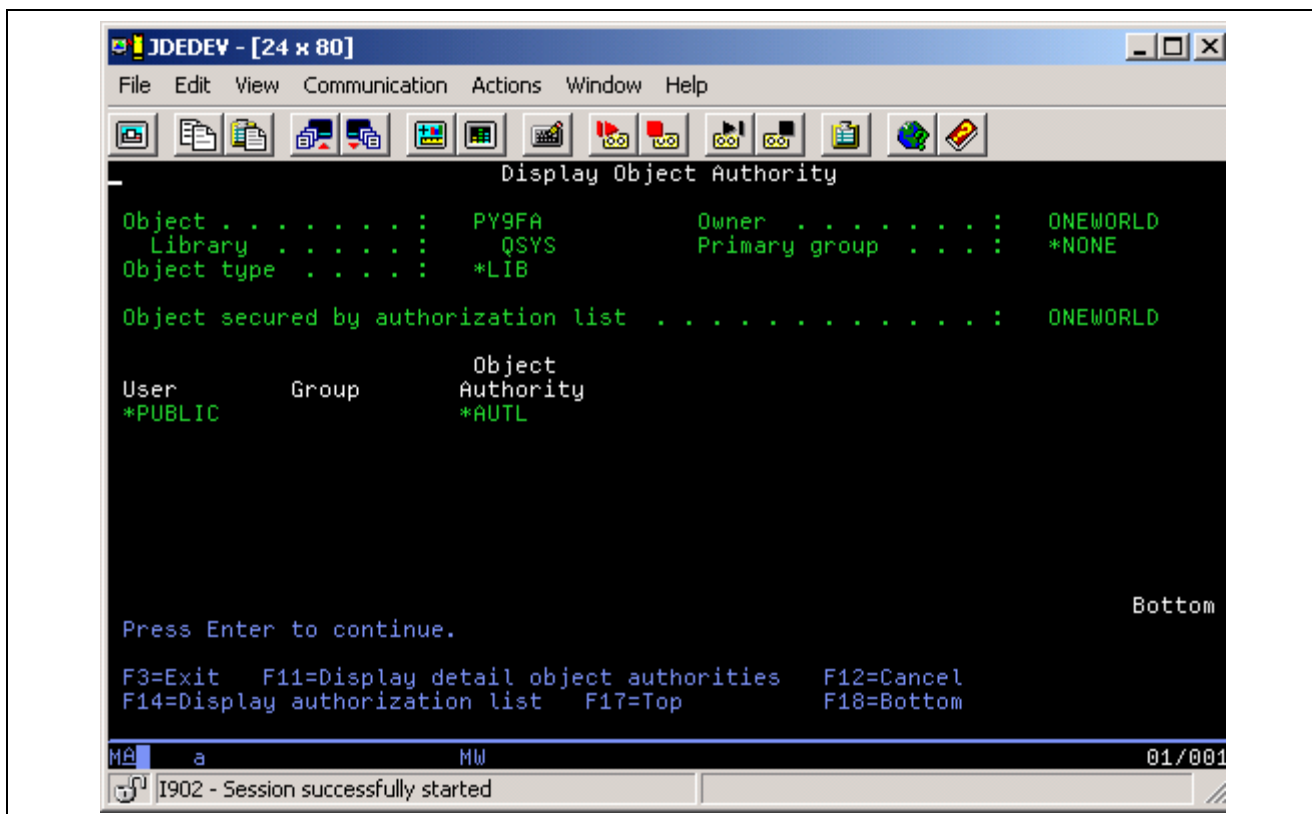
Similarly, if you set up multiple instances of EnterpriseOne and you display object authority, the value that appears is the name of the user profile for all objects except the SYSTEM library. The object authority for the SYSTEM library when you run multiple instances of EnterpriseOne is the name of 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.

This screen capture displays data source DSPOBJAUT:




Data source DSPOBJAUT

This screen capture displays pathcode DSPOBJAUT:




Pathcode DSPOBJAUT

This screen capture displays data source DSPLIBD:



Data source DSPLIBD

This screen capture displays pathcode DISLIBD:




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 these illustrations should correspond to the information that appears on the form. The required parameters are the object name, object type (*PGM), and the name of the library in which these objects reside.

This screen capture displays system library DSPOBJAUT:



System library DSPOBJAUT

This screen capture displays system library DSPLIBD:



System library DSPLIBD


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.

Sample Results for Authorization Lists

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

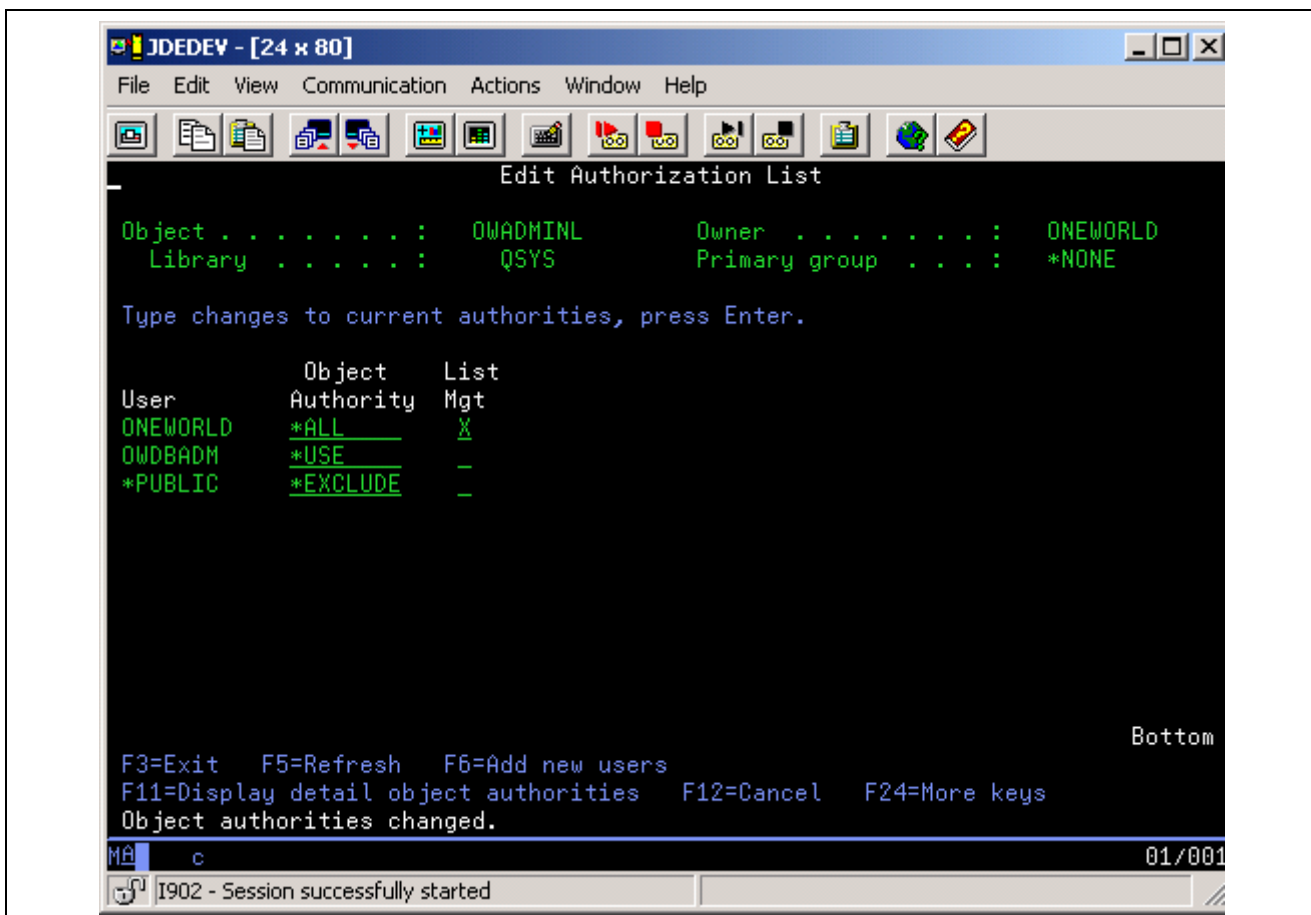
- IFS directories (specification files)
- WRKLNK - option 9 Work with authority

This screen capture displays DSPAUTL:



Display Authorization List

This screen capture displays DSPAUTL:



Edit Authorization List

Setting Up iSeries Database Security for a Single EnterpriseOne Instance

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

1. In the SETOWAUT library, on the command line, type this 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 ADDLIB command.

The Set Up OneWorld Authority (SETOWAUT) form appears.

2. On Set Up OneWorld Authority (SETOWAUT), complete the USPRF field with OneWorld, and then press ENTER:

The form displays additional security parameters. You can specify various security settings, including library access.

3. Complete these required fields, and then press ENTER:

- USRAUTL

Enter OWADMINL.

- TYPE

- INILIB
4. Complete any additional fields, if necessary.
 5. Press ENTER.

Setting Up iSeries Database Security for Multiple EnterpriseOne Instances

To set up iSeries database security for multiple EnterpriseOne instances:

1. In the SETOWAUT library, on the command line, type this command and press F4:

```
SETOWAUT
```

2. On Set Up OneWorld Authority (SETOWAUT), complete the USRPRF field, and then press ENTER:
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 these 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.

Adding Administrators

You can add administrators to the administrative authorization list by running the CRTOWADPRF command. The command also enables 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 this command and press F4:

```
CRTOWADPRF USRPRF
```

2. On Set Up OW User Profile (CRTOWADPRF), complete these 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 | Type the EnterpriseOne user profile name that you entered in the USRPRF field during setup. |
| ADMIN Authorization List | Type 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, CLRLCK, 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 this command and press F4:

```
RMVOWADPRF
```

2. On Remove OW Profile Authority (RMVOWADPRF), complete these fields and press ENTER:

| | |
|-----------------------------------|---|
| User Profile | Enter the name of the user from whom you want to remove authority. |
| Admin. Authorization List | Type the Admin. Authorization List name that you entered in the USRAUTL field during setup. |
| EnterpriseOne User Profile | Type the EnterpriseOne user profile name that you entered in the USRPRF field during setup. |

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 this command:

```
DSPUSRPRF
```

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 this example appears:

```
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 . . . . . : JDFOBJ7R2
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
```

CHAPTER 3

Understanding Server Administration for UNIX and Linux

This chapter provides an overview and discusses how to:

- Start the Enterprise Server for UNIX or Linux
- Shut Down the Enterprise Server for UNIX or Linux
- Set Up a Printer for UNIX or Linux
- Administer Batch Processes for UNIX or Linux
- Maintain File Security for UNIX and Linux
- Work with Linux Kernel Parameter Settings
- Work with AIX Kernel Parameter Settings for EnterpriseOne
- Run Multiple Instances of the EnterpriseOne Enterprise Server

Understanding Server Administration for UNIX and Linux

The PeopleSoft company 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 the UNIX or Linux enterprise server, you need to perform administrative procedures on the 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.

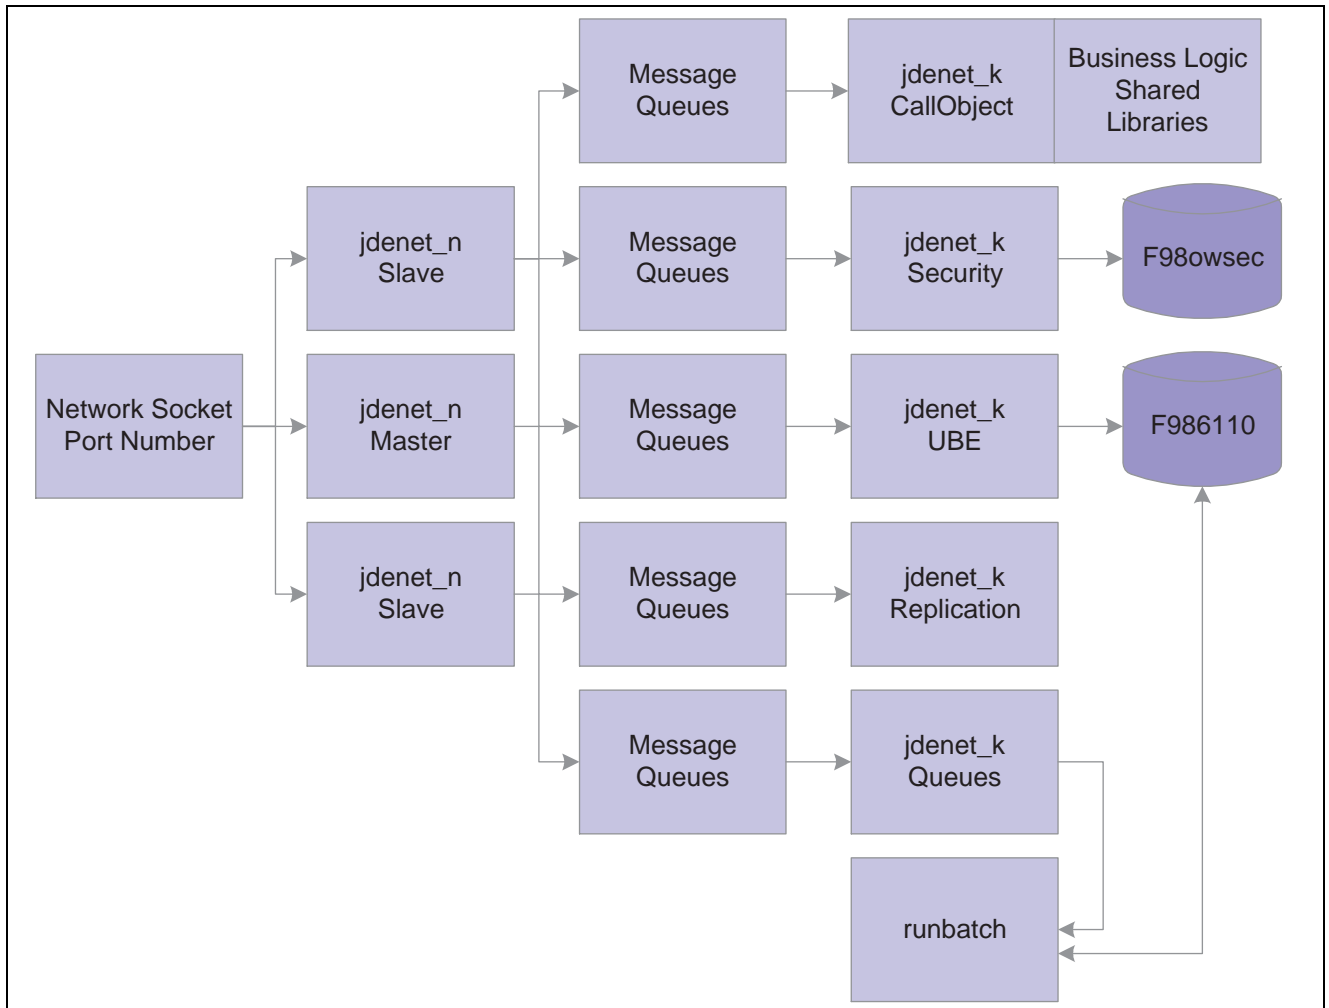
EnterpriseOne Directory Structure for UNIX and Linux

This 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/E811. Indented names indicate subdirectories of the directories, which are not indented.

| Directory | Description |
|------------|---|
| pathcode | <p>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 PD811 and DV811. These subdirectories are included:</p> <ul style="list-style-type: none"> • bin32 - Business function shared libraries. • Spec files - Binary data files in a PeopleSoft proprietary format. |
| system | <p>The main directory for the system-level executables, shared libraries, C header files, libraries, and localization files. These subdirectories are included:</p> <ul style="list-style-type: none"> • 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 | <p>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 these directories:</p> <ul style="list-style-type: none"> • 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 host server processes in this flowchart perform the indicated actions.



UNIX host server process

This 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 the `JDE.INI` file for a complete list.

EnterpriseOne Initialization for UNIX and Linux

This 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).

- These 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 this 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).

We strongly recommend 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 display 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

This section provides an overview of the enterprise server startup for UNIX or Linux and discusses how to:

- Start the enterprise server of UNIX or Linux manually.
- Start the enterprise server for HP-UX automatically.
- Start the enterprise server for AIX and Solaris automatically.
- Start the enterprise server for Linux automatically.
- Verify the EnterpriseOne installation.

Understanding Enterprise Server Startup 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 the 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 is the script that starts the EnterpriseOne system on the enterprise server. This script does this:

- 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.

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. These path is an example of where you might typically find this file:

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

Starting the Enterprise Server for UNIX or Linux Manually

To start the enterprise server for UNIX or Linux manually:

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 these 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.

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

- `RunOneWorld.sh`

This script starts the EnterpriseOne system.

3. Sign off the system.

Starting the Enterprise Server for HP-UX Automatically

To start the enterprise server for HP-UX automatically:

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

The script should contain only these:

```
#!/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.

- Using this 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
```

- Verify that these 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.

- Set `ORACLE_TERM` to *hp*.
- Set `ORAENV_ASK` to *NO*.
- If this command is in the profile, delete it:

```
unset ORAENV_ASK
```

Starting the Enterprise Server for AIX and Solaris Automatically

To start the enterprise server for AIX and Solaris automatically:

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

The script should contain only these:

```
#!/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.

- Add these line at the end of the text file named `inittab` in `/etc`:

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

- Verify that these 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.

- Set `ORACLE_TERM` to *hp*.
- Set `ORAENV_ASK` to *NO*.

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

- If this command is in the profile, you must delete it:

```
unset ORAENV_ASK
```

Starting the Enterprise Server for Linux Automatically

To start the enterprise server for Linux automatically:

Add these 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.

Verifying the EnterpriseOne Installation

To verify the EnterpriseOne installation:

After you start EnterpriseOne, execute these 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.

Shutting Down the Enterprise Server for UNIX or Linux

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

`EndOneWorld.sh` is the script that stops the EnterpriseOne system on the enterprise server. This script completes these 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.

Shutting Down the Enterprise Server for UNIX or Linux

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 these 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

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Setting Up EnterpriseOne Printing,” Printing Reports

Administering Batch Processes for UNIX or Linux

This section provides an overview of batch process administration for UNIX or Linux and discusses how to:

- Monitor batch processes.
- List batch output files.
- Run reports from the command line for UNIX or Linux.
- Schedule reports from the command line for UNIX or Linux.

Understanding Batch Process Administration 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, these processes start and run under the environment and security of the user who started them:

| Process | Description |
|----------|---|
| 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. |

Use the jdejobs command to monitor current batch processes. This example is a sample output:

```
pfst811 (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.

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.

You can list output files. The returned data looks similar to this:

```

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:

| Segment | Description |
|---------------|--|
| 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).

Monitoring Batch Processes

To monitor batch processes:

From the operating system prompt, enter this 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.

Listing Batch Output Files

To list batch output files:

1. From the operating system prompt, enter this 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 this command to list the files:

```
ls
```

Running Reports from the Command Line for UNIX or Linux

You can initiate batch process reports from the server command line by issuing this 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. These options apply to the runube command:

| Parameter | Description |
|-------------|--|
| 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

This example displays 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 enables 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. This table describes the CCYYMMDDHHMMSS variable:

| Segment | Description |
|---------|---|
| 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 this 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
```

This table describes the variables for this command:

| Segment | Description |
|---------|---|
| 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

This example displays 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 displays 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

This section describes file security maintenance for UNIX and Linux and kernel parameter settings for HP-UX and Solaris.

File Security Maintenance 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.

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.

Business function security should be similar to SPEC file security. This enables 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.

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 enable them to shut down EnterpriseOne.

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.

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 a text editor. 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 these parameters to the `/etc/system` file:

- `forceload: sys/msgsys`
- `forceload: sys/semsys`
- `forceload: 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.

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.

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

| Parameter | Definition |
|-----------|--|
| 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. |

This screen capture 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 this line: set shmyns shminfo_shmmax=4294967295:

Message Queues

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

| Parameter | Description |
|-----------|--|
| mesg | This value must be 1. System-V style message queues are valid. |
| msgmni | <p>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 the EnterpriseOne installation (per EnterpriseOne instance). You can use these equation to estimate the number of message queues necessary for EnterpriseOne:</p> $1 + \text{jdenet_n} + 2 \times \text{jdenet_k} + (\text{max number of concurrent runbatch, runube, and runprint processes})$ |
| 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 x msgmni for the requirements of EnterpriseOne. |
| msgmap | <p>The value for msgmap represents the number of entries in the map of free message segments. The default value of msgtql + 2 should be used. If the value of msgmap is less than the value of msgtql + 2, attempts to create a message queue or to send a message might fail.</p> <p>Note. This parameter is no longer used in Solaris 8.</p> |
| 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 msgseg x msgssz. For EnterpriseOne, a value of 32768 is reasonable. You can set a larger value as long as the product of msgseg x 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 msgmax = 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. These parameters, which do not apply to Solaris 8, determine the size and number of segments available throughout the system:

| Parameter | Description |
|-----------|---|
| 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 x the msgmni requirement for EnterpriseOne, or approximately 4096 per instance. |

Semaphores

These definitions apply to semaphores:

| Parameter | Description |
|-----------|--|
| sema | This value must be 1. System-V style message queues are valid. |
| semmni | <p>The value of semmni represents the maximum number of semaphore identifiers that can exist throughout the system.</p> <p>For EnterpriseOne, two identifiers exist for each instance of EnterpriseOne, so the default value supplied with the HP-UX and Solaris systems should suffice.</p> |
| semmap | <p>The value of semmap represents the number of entries in the map of free semaphores. The default value of semmni + 2 should suffice. If you decrease the value of semmap, attempts to create a semaphore set, which occurs during JDEIPC initialization, might fail.</p> <p>Note. This parameter is not used in Solaris 8.</p> |
| semmns | <p>The value of semmns 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.</p> <p>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.</p> |

| Parameter | Description |
|-----------|---|
| semmnu | <p>The value of semmnu 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, enable one for each EnterpriseOne process that can exist for all installations of EnterpriseOne on the system. Use these equation to determine this value:</p> $1 + \text{jdenet_n} + \text{jdenet_k} + \text{maximum number of runbatch processes} + \text{maximum number of runprint processes} + \text{maximum number of runube processes}$ <p>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.</p> <p>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.</p> <p>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.</p> |
| run | <p>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.</p> |
| semmsl | <p>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.</p> |

Shared Memory

These definitions apply to shared memory:

| Parameter | Description |
|-----------|--|
| 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, enable 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

These definitions apply to file descriptors:

| Parameter | Description |
|--------------|--|
| 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

This definition applies to processes:

| Parameter | Description |
|-----------|---|
| 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. |

Working with Linux Kernel Parameter Settings

This section provides an overview of Linux kernel parameter settings and describes how to:

- Set specification file security.
- Set business function file security.
- Set executables security.
- Set jde.ini file security.

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

These five entries in the /etc/sysctl.conf file affect EnterpriseOne IPC resources:

| Parameter | Description |
|---------------|---|
| kernel.sem | <p>This setting controls these four different semaphore limits:</p> <ul style="list-style-type: none"> • 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) <p>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.</p> |
| 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 | <p>This parameter defines the maximum number of message queues (identifiers) in the system. You can use these equation to estimate the number of message queues that are necessary for EnterpriseOne:</p> $1 + jdenet_n + 2 \times jdenet_k + (\text{max number of concurrent runbatch, runube, and runprint processes})$ |

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 the current values, review the file /proc/sys/fs/file-nr. This read-only file contains these 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 `fs.file-max` setting. This setting 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`

These lines are 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
```

Setting Specification File Security

To set specification file security:

Add this line to the `.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 the SPEC files should look similar to these example:

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

Setting Business Function File Security

To set business function file security:

1. Enter this command in the BSFN Source directory:

```
chmod 644 *.c
```

2. Enter this command in the BSFN Include directory:

```
chmod 644 *.h
```

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

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

Setting Executables Security

To set executables security:

UNRECOGNIZED STYLE ->class=singlestep>Enter this command:

```
cd $SYSTEM/bin32
chmod 540 *..sh
```

The access granted by this command gives 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 the EnterpriseOne executables should look similar to these example:

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

Setting jde.ini File Security

To set JDE.INI file security:

1. Enter this 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 the JDE.INI file should look similar to this:

```
-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 this 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.

Working with AIX Kernel Parameter Settings for EnterpriseOne

This section provides an overview of AIX kernel parameter settings for EnterpriseOne and describes how to:

- Set the value of maxuproc.
- View the system parameters.
- Set tune parameters.

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.

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.

Setting the Value of maxuproc

To set the value of maxuproc:

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

`smit`
3. In SMIT, select the System Environments item and then select the Change/Show Characteristics of Operating System item.
4. Change the value of Maximum number of processes to enable 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. This table lists these system parameters for general reference:

| Parameter | Description |
|-------------|--|
| 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 |

| Parameter | Description |
|-----------|--------------------------------|
| realmem | Size of usable physical memory |
| primary | Primary dump device |
| conslogin | System console login |

Viewing the System Parameters

To view the system parameters:

Enter this command:

```
lsattr-E-lsys0
```

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

Setting Tune Parameters

Setting the tune parameters requires you to run these 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 these tune parameters are set:

| Parameter | Value |
|--------------|----------------------------|
| 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

This section provides an overview of running multiple instances of the EnterpriseOne enterprise server and describes how to do so.

Prerequisite

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

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. These 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.

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.

See Also

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Environment Setup,” Adding an Environment

Running Multiple Instances of the EnterpriseOne Enterprise Server

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/E811
```

New .psft file:

```
export EVRHOME=/u02/PeopleSoft/E811
```

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

For example, type these:

```
mkdir -p /u02/PeopleSoft/E811
```

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.

These sample commands accomplish this:

```
cp -R /u01/PeopleSoft/E811/system /u02/PeopleSoft/E811
```

```
cp -R /u01/PeopleSoft/E811/ini /u02/PeopleSoft/E811
```

```
cp -R /u01/PeopleSoft/E811/DV811 /u02/PeopleSoft/E811
```

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 this:

```
mkdir -p /u02/PeopleSoft/E811/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/E811/log/ jdedebug.log
JobFile=/u02/PeopleSoft/E811/log/jde.log
```

```
[INSTALL]
B9=/u02/PeopleSoft/E811
```

```
[BSFN BUILD]
BuildArea=/u02/PeopleSoft/E811/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.

The values are defined by these parameters; but the numbers are only examples:

```
[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.

CHAPTER 4

Understanding Server Administration for Windows

This chapter provides an overview of Windows server administration and discusses how to:

- Set up a printer for Windows.
- Work with network services.
- Administer batch processes for Windows.
- Maintain file security for Windows.
- Run multiple instances of EnterpriseOne on Windows.

Understanding Server Administration for Windows

The PeopleSoft company 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 the enterprise server to ensure that the software runs properly.

Understanding the EnterpriseOne Directory Structure for Windows


This 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\E811\ddp). Indented names indicate subdirectories of the directories.

| Directory | Description |
|-----------|--|
| pathcode | <p>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 PD811 and DV811. These subdirectories are included:</p> <ul style="list-style-type: none">• bin32, which includes business function shared libraries.• spec, which includes specification files. These binary data files are in a PeopleSoft proprietary format. |

| Directory | Description |
|------------|--|
| system | <p>The main directory for the system-level executables, shared libraries, C header files, libraries, and localization files. These subdirectories are included:</p> <ul style="list-style-type: none"> • 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 | <p>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 these subdirectories:</p> <ul style="list-style-type: none"> • 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. |

Understanding EnterpriseOne Architecture and Process Flow for Windows

These host server processes perform the indicated actions:



Windows server processes

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).

This 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.
- This 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 CalloObj 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.
- This 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.
- When the UBE Jdenet_k (the kernel) writes to the database (batch application only), this occurs:
 1. The program is system\bin32\jdenet_k.exe.
 2. 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, this 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 these actions occur:
 - Cannot start the service name service on the enterprise server
 - Error 1069: The service did not start due to a logon failure

Understanding EnterpriseOne Initialization for Windows

This 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).
- These 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). We strongly recommend 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, select Programs, then Administrative Tools, and then Data Sources (ODBC).

- This information pertains to the setup of SQL Server ODBC drivers, using the ODBC Data Source Administrator applet:
 - 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.
- This information pertains to the setup of Client Access ODBC drivers, using the ODBC Data Source Administrator applet:
 - 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.

Understanding JDE.INI Settings for Starting Batch Queues on Windows

These 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

This table describes each setting:

| Setting | Description |
|-------------------------------|--|
| 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. |

Understanding 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 this 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 selects 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.

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.

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

```
SCPToPublish      PEOPLESOFT_ENTERPRISEONE_811_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.

These 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_811_SP1.

This 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_811_SP1 |
| SecurityServerSCP | Security server | PEOPLESOFT_ENTERPRISEONE_811_SP1 |
| LockManagerSCP | Lock manager | PEOPLESOFT_ENTERPRISEONE_811_SP1 |
| UnifiedLogonServerSCP | Unified logon server | PEOPLESOFT_ENTERPRISEONE_811_SP1 |

See Also

Windows-Based Systems in the *EnterpriseOne PeopleTools 8.11 Installation for Windows*

Setting Up a Printer for Windows

This section provides an overview of printer setup, services, accounts, and permissions for Windows and describes how to:

- Add a printer.
- Determine or change printer ownership.
- Set up user accounts on an enterprise server.
- Change the domain.
- Add a local account.
- Add a user to the administrators group.

Understanding Printer Setup 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 these:

- 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.

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.

When you create a Windows user account, you must associate that account with one of these 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.

In this table, you must define two types of service accounts and printer ownerships for the two types of printers:

| Printer Type | Account and Owner |
|--------------|--|
| 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 enable 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 these 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 these:

- 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.

Note. We strongly recommend that you use an account for a user who belongs to the local Administrators group.

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

Adding a Printer

To add a printer:

1. Click the Windows Start button.
2. Select Settings, and then select Printers.
3. Select 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

To determine or change printer ownership:

1. From Control Panel, select Printers.
2. Select a printer, right-click, and select 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, select 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, select Start, Settings, Control Panel, Administrative Tools, then Computer Management.
2. On the Tree tab, select Local Users and Groups, and then click the Users folder.

Changing the Domain

To change the domain:

1. From the main menu of User Manager, select User.
2. Select 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 select a network domain, all listed names represent network accounts. Likewise, if you select 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, select System Tools, and then select Local User and Groups.
3. From the Action menu, select New User.
4. On New User, complete these fields:
 - User name
 - Full Name
 - Description
 - Password
 - Confirm Password
5. Complete these options, as appropriate for the 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 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, select the domain of the user whom you want to add to the Administrators group.
4. Select 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

This section provides an overview of network services and describes how to:

- Set up the network service.
- Start the network service.
- Stop the network services.
- Clean up the enterpriser server for Windows.
- Uninstall the network service.
- Start the enterprise server for Windows manually.
- Verify the EnterpriseOne installation.

Understanding 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 these 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.

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

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. The process to uninstall this service removes these entries from the Windows registry:

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

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.

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

After you have installed and set up the Network service, you must start the service before EnterpriseOne can use it. Later, if you need to stop services, you must do so in the proper order.

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.

Setting Up the Network Service

To set up the network service:

1. From the Start menu, select Programs, Administrative Tools, and then Services.
2. Select 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 Release E811 is JDE 811 Network.

3. Click Action, 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

To start the Network service:

1. From the Services window, select 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 E811 is JDE 811 Network.

2. From the Action menu, click Start.
3. Use the Windows Task Manager to ensure that these 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, select 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 811 Network.

2. Use the Windows Task Manager to ensure that all EnterpriseOne processes are terminated.

This might take several minutes. These 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

Cleaning Up the Enterprise Server for Windows

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. Select 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 the process and select debug.
6. When the Visual C++ main window appears, select 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

To uninstall the Network services:

Run this program from the \system\bin32 directory:

```
jdesnet -u
```

Starting the Enterprise Server for Windows Manually

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.

To start the enterprise server for Windows manually:

1. On the enterprise server for Windows, sign on with administrator privileges.
If you used the user ID that we recommend, the value is *PSFT*.
2. On the Windows toolbar, from the Start menu, select Run, and then enter these 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 the EnterpriseOne installation with the PORTTEST program.

Note. When you run PORTTEST, make sure that one of this 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 these commands:

```
cd \PeopleSoft\E811\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

This section provides an overview of batch process administration for Windows and describes how to:

- Monitor batch processes.
- Review batch output files.
- Run reports from the command line for Windows.
- Schedule reports from the command line for Windows.

Understanding Batch Process Administration 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:

| Process | Description |
|--------------------------|--|
| <code>jdesnet.exe</code> | The network listener that listens for connection requests. |

| Process | Description |
|--------------|--|
| 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. Here are examples of both of these workstation settings:

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

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 the 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:

| Segment | Description |
|---------|--|
| 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. |

| Segment | Description |
|----------|--|
| 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 these 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. These options apply to the runube command:

| Parameter | Description |
|-------------|--|
| 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 default printer. |

Example: Running Reports from the Command Line for Windows

This example lists commands for executing a batch process report:

```
cd \PeopleSoft\E811\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 enables 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 these options apply:

| Parameter | Description |
|----------------------|---|
| \\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 | enables the program to interact with the Windows operating system desktop. |
| /EVERY:date | Specifies the days on which to run the job. 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 <code>runube</code> command with any of its parameters. |

Example: Scheduling Reports from the Command Line for Windows

This example lists a sample 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
```

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 these 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 enables 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

This section provides an overview of running multiple instances of EnterpriseOne on Windows and describes how to:

- Generate a unique identifier.
- Modify the server jde.ini files.
- Modify the workstation jde.ini file.
- Uninstall EnterpriseOne services.
- Move or change an EnterpriseOne directory tree.

Prerequisites

Before you complete the tasks in this section:

- Verify that you have enough disk space to create copies of the 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\E811 directory tree.

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.

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.

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. This 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\E811_2\log\jdedebug.log |
| [DEBUG] | JobFile= | Name of the log file that contains log data. | z:\PeopleSoft\E811_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\E811_2 |

| Section of server JDE.INI file | Parameter | Purpose | Example Value |
|--------------------------------|---------------------|---|--------------------------------------|
| [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 workstations. | 6005 |
| [JDENET] | serviceNameConnect= | The TCP/IP port number used by the server to send communications packets to servers. | 6005 |

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.

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 (ipcserv.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.

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

- serviceNameListen=
- serviceNameConnect=

See Also

[Chapter 4, “Understanding Server Administration for Windows,” Setting Up the Network Service, page 88](#)

Generating a Unique Identifier

To generate a unique identifier:

1. From the Start menu on the Windows taskbar, select Run, and then enter this 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 `uuidgen-?` command:

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

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 these parameters:

DebugFile= Type the name of the log file that will contain debugging information.

JobFile= Type the name of the file that will contain log information.

3. In the [INSTALL] section of the JDE.INI file, modify the values of these parameters:

StartServicePrefix= Type the value to be 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 811. The default value produces the service names JDE 811 Network and JDE 811 Queue.

B9= Type the name of 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 these parameters:

StartIPCKeyValue Type a number for the 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.

Note. 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= Type the unique string that is generated by the NT guidgen program.

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

serviceNameListen= Type the port number for 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= Type the port number for 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.

Modifying the Workstation JDE.INI File

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 these 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 this command:

```
C:\> d:\E811\system\bin32
```

3. To uninstall network services, enter this 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 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 this command:

```
C: \> d:\E811\system\bin32
```

2. To uninstall network services, enter this 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 using interprocess communications.

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

CHAPTER 5

Server Administration Workbench

This chapter provides an overview of the Server Administration Workbench (SAW) and describes how to:

- Set up security access for SAW.
- Monitor EnterpriseOne with SAW on Windows.
- Configure EnterpriseOne server monitoring settings.
- Monitor EnterpriseOne server processes.
- Monitor enterprise server statistics.
- Monitor EnterpriseOne Web server statistics.
- Query EnterpriseOne enterprise servers.
- Configure SAW and monitoring EnterpriseOne servers from the Web.
- Monitor enterprise servers from the Web.
- Monitor the JAS servers from the Web.
- Monitoring System XPIe Servers from the Web
- Monitor EnterpriseOne with SAW on the iSeries.
- Monitor EnterpriseOne with SAW on UNIX.
- Monitor EnterpriseOne with the Knowledge Module.

Understanding the 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.)

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.

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

See Also

EnterpriseOne Tools 8.94 PeopleBook: System Administration, "Understanding PeopleSoft EnterpriseOne Security"

Setting Up Security Access for SAW

To set up security access for SAW:

1. From the System Administration Tools menu (GH9011), select Security Workbench (P00950).
2. On Security Workbench, then select Setup Security and External Calls from the Form menu.
3. 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.
4. Enter `SERVERADMINISTRATIONWORKBENCH.EXE` in the Display Secured Item Executable field, and then click Find.
5. Expand the Unsecured node in the tree at the bottom of the form.
6. Select the Run Security option.
7. Drag the `SERVERADMINISTRATIONWORKBENCH.` item to the Secured node in the tree and click Find.
The `SERVERADMINISTRATIONWORKBENCH.` item should appear under the Secured node.
8. Click Close.

Monitoring EnterpriseOne with SAW on Windows

This section provides an overview of the SAW interface and describes how to:

- Add to the `jde.ini` file.
- Access SAW.

- Add an EnterpriseOne server for querying.
- Check EnterpriseOne server configurations.
- Check EnterpriseOne server connectivity.
- Remove an EnterpriseOne server from the query list.

Understanding the SAW Interface

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 these two tabs:

- **Query.** This tab lets you add servers and review statistics about their components at a particular point in time.
- **Monitor.** This tab enables 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.

Server Administration Workbench (SAW) is a separate EnterpriseOne executable that resides by default in the \E811\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 these 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.

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.

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

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.

Server Administration Workbench (SAW) enables 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.

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 enables SAW to create, transfer, and remove temporary files that are larger than 5 MB. This 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

To access SAW:

1. Open the \E811\SYSTEM\Bin32 directory and then run SERVERADMINISTRATIONWORKBENCH.exe.
2. On Sign On, complete these fields, and then click OK:
 - User ID
 - Password

Adding an EnterpriseOne Server for Querying

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, select 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 select the name of an EnterpriseOne server that is running on the Windows platform.

The ellipsis button is disabled if you select 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

To check server configurations:

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

SAW displays the connection port number when you select a server name.

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

Checking EnterpriseOne Server Connectivity

To check EnterpriseOne server connectivity:

1. On Server Administration Workbench, click the Query tab.
2. From the File menu, select Servers, and then select Connectivity.
3. On Server Connectivity, select 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

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 button of the server that you want to remove.

Note. You must click the server button, not a server component, such as Processes.

4. Click the Remove button.

Configuring EnterpriseOne Server Monitoring Settings

This section provides an overview of EnterpriseOne server monitor settings and describes how to:

- Select the server, port, and processes to monitor.
- Set up event notification.
- Set up SAW run-time parameters.

Understanding 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 enables you to do these:

- Select the servers that you want to monitor
- Select the processes that Server Administration Workbench (SAW) monitors
- Select 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 the server monitor settings in SAW using the EnterpriseOne Server Monitor Settings form. 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.

You make choices from these three tabs to configure the monitoring settings:

- Port/Server Selection, which enables you to set up servers, port connection numbers, and server processes to monitor.
- Notification Configuration, which enables you to select the server events that trigger administrator notification, the ways in which you are notified, and the frequency of notification.
- General, which enables 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.

From the General tab, you specify these 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 backs 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 before administrator notification is triggered
- UBE run-time threshold, which specifies the number of minutes for which SAW enables a batch process to run before administrator notification is triggered
- Queue threshold, which specifies the number of jobs in a queue that SAW enables before administrator notification is triggered

These threshold settings are recommended:

| Threshold Setting | Value |
|-----------------------------------|-------|
| 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 |

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, select the "This is a WebServer" option.
7. To set up server processes to monitor, select any of these options, if applicable:
 - Off line
 - UBE
 - Lock Manager
 - Replication
 - Security

Note. If you select 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, select from these 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 email notification, click the + button in the Email address field, enter an email 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 these options to indicate how you want to be notified:
 - Select 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

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 these 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:



SAW Server Monitor window

The information in the Server Monitor window enables 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 stoplight icon indicates that a process is running.

This table summarizes the server process parameters that appear in the Server Monitor window:

| Parameter | Description |
|--------------|--|
| 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.

Monitoring EnterpriseOne Server Processes

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 button on the menu bar.

Monitoring Enterprise Server Statistics

This section provides an overview of enterprise server statistics and describes how to:

- Monitor process statistics for enterprise servers.
- Change process settings.
- View log and debug log files for enterprise server processes.
- Print log and debug log files for enterprise server processes.
- Turn logging on and off for enterprise server processes.
- View process user statistics for enterprise servers.
- View statistics for batch processes.
- View statistics for server disks.

Understanding 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 these:

- 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 enables 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.

To display the processing statistics, you click a server button. Statistics about each process appear in the Process Statistics window.

This table describes the parameters that appear in the Process Statistics window:

| Parameter | Description |
|------------------------------------|--|
| 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 these additional tasks:

- View the debug log
- View the JDE log
- Print the debug log
- Print the JDE log
- Process users
- Select server logging
- Clear 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.

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.

The 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 these formula:

`maxKernelProcesses` minus the sum of individual kernel type `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 ($50-34=16$, $10+16=26$).

Monitoring Process Statistics for Enterprise Servers

To monitor process statistics for enterprise servers:

1. On SAW, select 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. Select 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, select the column heading.

Changing Process Settings

To change process settings:

1. On SAW, select the Monitor tab.
2. Select 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. Select a kernel or net process and right-click.
4. To change the number of net or kernel processes, select Change max number of this process type, enter a new value, and then click OK.
5. To disable a kernel process, select Disable Kernels.

Viewing Log and Debug Log Files for Enterprise Server Processes

Server Administration Workbench (SAW) enables 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, select an enterprise server icon.
2. On Process Statistics, right-click a network or kernel process and select View Debug Log or View JDE Log. The log or debug log file appears in a text editor.

Note. You can also select 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 the workstation. Server Administration Workbench (SAW) prints a text editor document to the 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, select an enterprise server icon.
2. On Process Statistics, right-click a network or kernel process and select Print Debug Log or Print JDE Log.

SAW prints the file to the 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 activate logging so that Server Administration Workbench (SAW) adds information to the log file for that process. When you have enough information, you can deactivate 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, select an enterprise server icon.
2. On Process Statistics, right-click a network or kernel process and select 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 select.

To view process user statistics for enterprise servers:

1. On Statistics, select an enterprise server icon.
2. On Process Statistics, right-click a CallObject kernel process and select 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 select the UBE icon, you can review these 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. Select 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 select the disk icon, you can review these 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. Select the Disk icon.

Statistics for each server disk appear on the Monitor Disk form.

Monitoring EnterpriseOne Web Server Statistics

This section provides an overview of EnterpriseOne Web server statistics and describes how to:

- Monitor data on enterprise servers connected to the Web server.
- Monitor Web server callobjects.
- Monitor Web server connection statistics.
- Monitor Web server user statistics.

- Monitor Web server user connection statistics.
- Monitor Web server logs.

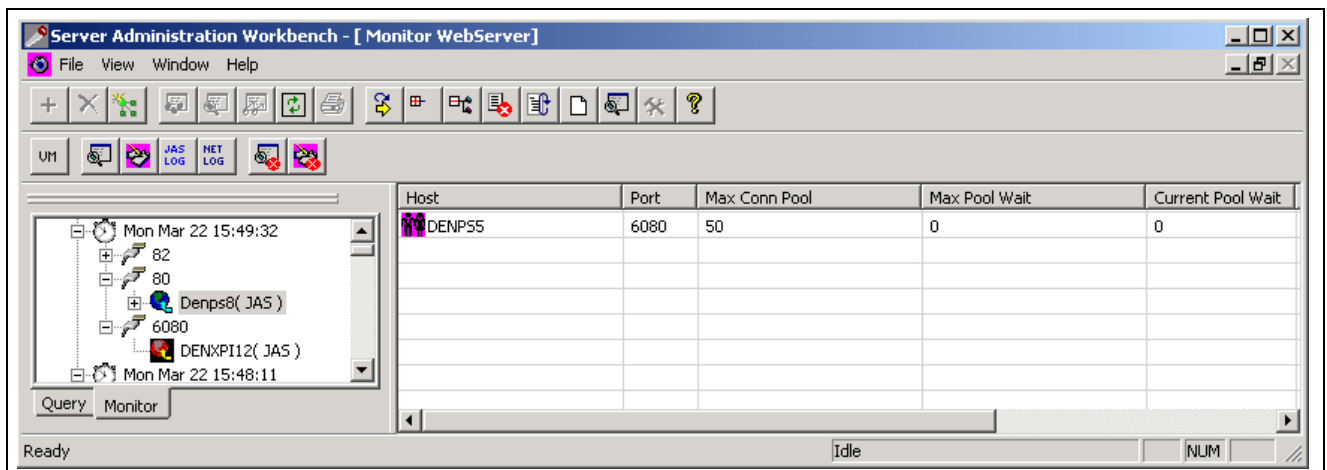
Understanding EnterpriseOne Web Server Statistics

You can also use Server Administration Workbench (SAW) to monitor statistics about EnterpriseOne Web servers from the 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 the 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 enable you to set up notification triggers.

When you select 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:



Monitor WebServer window

This window contains important connection information, as well as information about business functions running on the enterprise servers.

This table summarizes the data SAW displays on each enterprise server that is connected to the Web server:

| Display Element | Description |
|-----------------|--|
| 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. |

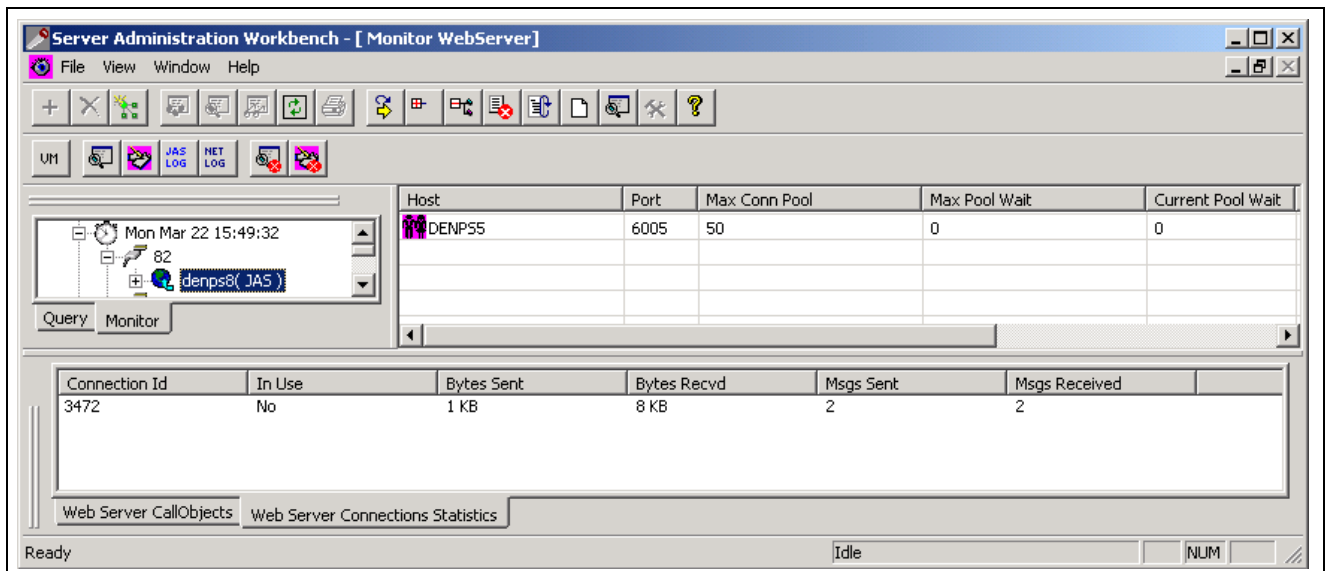
| Display Element | Description |
|---------------------|---|
| Max Pool Wait | The number of users waiting for a connection when the Max Conn Pool value has been exceeded. For example, if the Max Conn Pool value is 30, and five users are waiting to connect, the Max Pool Wait value is 5. As soon as the 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. |

Server Administration Workbench (SAW) enables 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.

This table summarizes the data that SAW displays about business functions that run on an enterprise server that is connected to the Web server:

| Display Element | Description |
|--------------------|---|
| 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. |

You select the Web Server Connections Statistics tab to review data on the connections made from the Web server to the enterprise server:



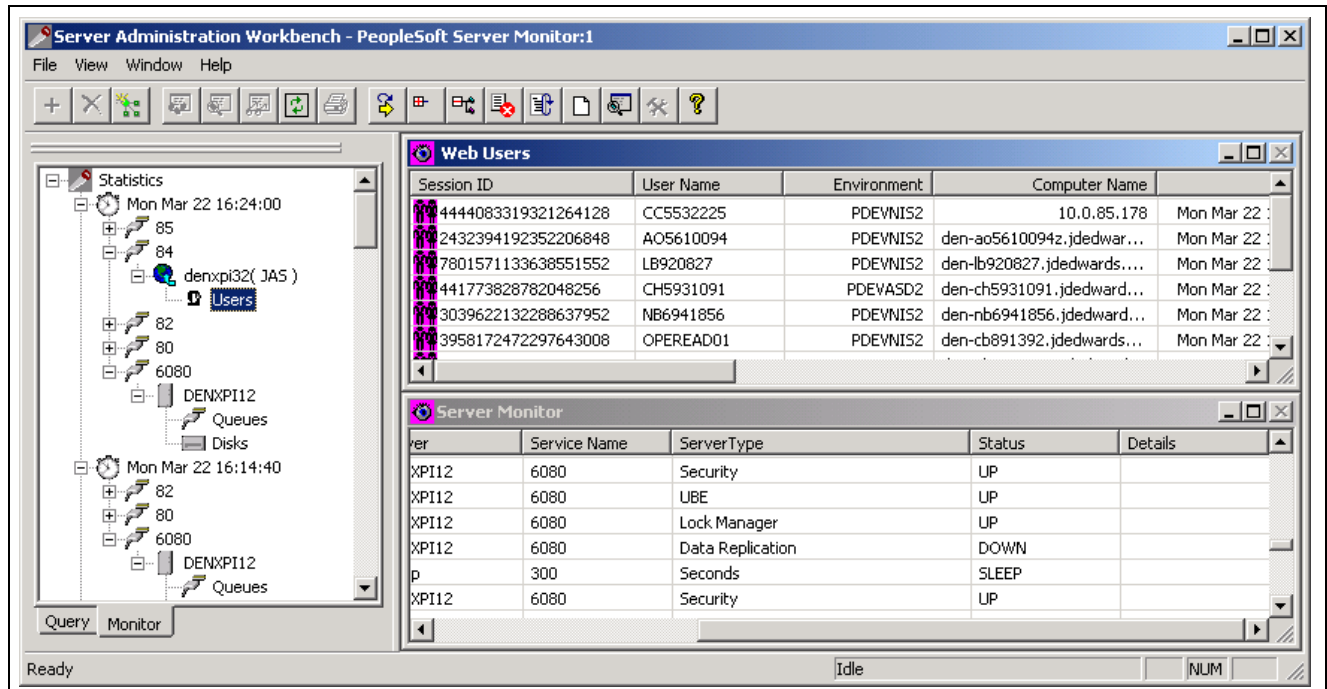
Web Server Connections Statistics tab

This table summarizes the data that Server Administration Workbench (SAW) displays on connections between the Web server and the enterprise server:

| Display Element | Description |
|-------------------|--|
| 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. |

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:



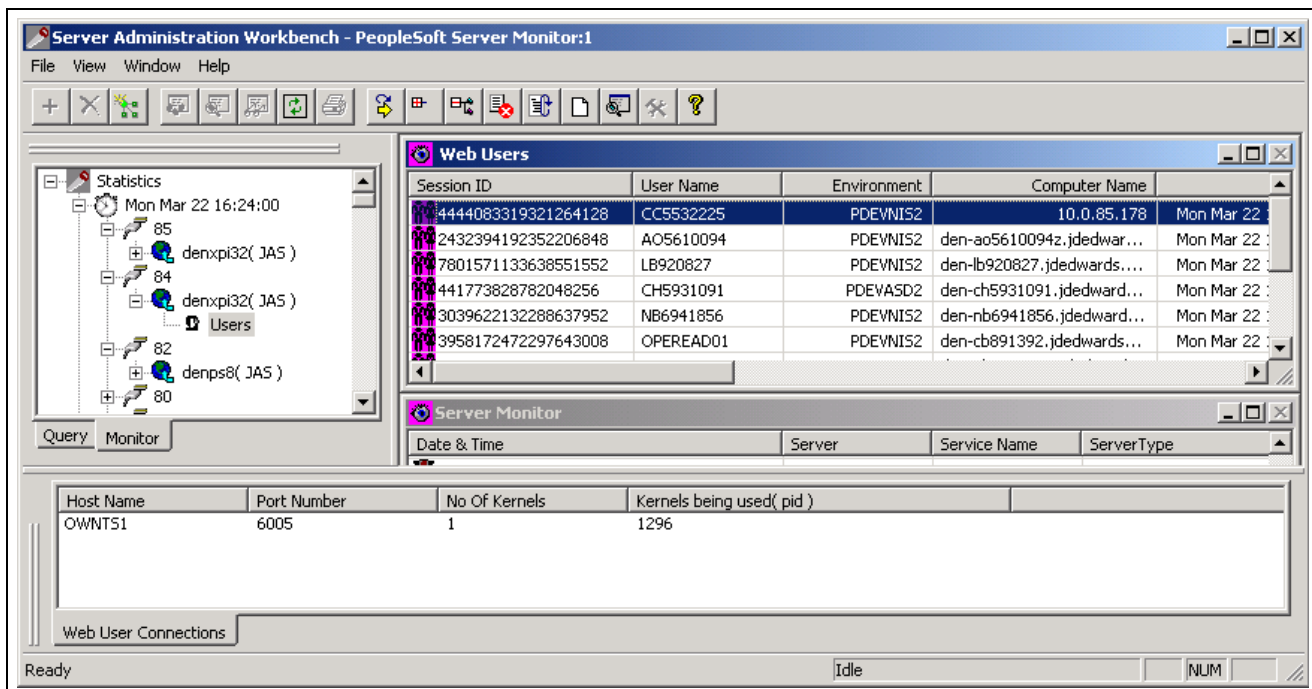
Expanding a node in SAW

The Web Users window displays information about each user who is connected to the Web server. This table summarizes the data that SAW displays for each Web user:

| Display Element | Description |
|---------------------------|--|
| 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. |

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.

Selecting Show Web User Connections opens a window that displays each Web server user's connection to the enterprise server:



Show Web User Connections window

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.

This table summarizes the connection data that appears for each Web user:

| Display Element | Description |
|--------------------------|---|
| 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. |

See Also

Chapter 5, "Server Administration Workbench," Monitoring the Web Server System, page 148

Monitoring Data on Enterprise Servers Connected to the Web Server

To monitor data on enterprise servers connected to the Web server:

1. On SAW, select the Monitor tab.
2. In the toolbar, select 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. Select 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

To monitor Web server CallObjects:

1. On Monitor WebServer, right-click a line containing the name of an enterprise server and then select Display CallObjects and Connections.

SAW opens a new window with these two tabs:

- Web Server CallObjects
- Web Server Connections Statistics

2. Select 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

To monitor Web server connection statistics:

1. On Monitor WebServer, right-click a line containing the name of an enterprise server, and then select Display CallObjects and Connections.
2. Select 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

To monitor Web server user statistics:

1. On SAW, select the Monitor tab.
2. In the toolbar, select 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. Select 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

To monitor Web server user connection statistics:

1. On Web Users, right-click a record that contains a user ID, and then select 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 these:

- 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 activate or deactivate JAS logging. JAS log gathers information about processes running on the Java Application Server.
- Get the NET log, which enables you to activate or deactivate 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 selecting 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

This section provides an overview of EnterpriseOne server queries and describes how to:

- Query enterprise server process components.
- Query enterprise server disk components.
- Query enterprise server IPC resources.
- Query enterprise server log files.
- Query the enterprise server users.
- Query the enterprise server job queue.

Understanding EnterpriseOne Enterprise Server Queries

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 these 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 select a component or subcomponent icon, SAW displays data about the component in a separate window.

For example, if you select 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 selecting a server icon, and then selecting 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.

Querying a server enables you to monitor kernel, network, and queue processes. When you select 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 these 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 select Information after right-clicking a kernel process, the Detailed Info - Kernel Process form appears.

If you select Information after right-clicking a network process, the Detailed Info - Network Process form appears.

A separate window on the Detailed Info - Network Process form displays these connection information:

- Connection number
- Connection type
- Internet Protocol (IP) address
- Long address

Server Administration Workbench (SAW) enables you to query interprocess communications, which enable applications to share data and memory and enable one application to control another.

SAW monitors these 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 enable 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 enable 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 these 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.

Server Administration Workbench (SAW) enables you to query for data about users connected to the enterprise server. When you select the View Users icon, SAW displays-in a separate window-information about each user. The information is displayed in these 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.

See Also

Chapter 5, “Server Administration Workbench,” Querying EnterpriseOne Enterprise Servers, page 121

Querying Enterprise Server Process Components

To query enterprise server process components:

1. On SAW, select the Query tab.
2. Expand the Processes node for an enterprise server you want to monitor and select one of these icons:
 - Kernel
 - Network
 - Queue
 - AllSAW populates a form window.
3. Select a column heading to sort the data in a column.
4. For additional information on a process, right-click and select Information.
5. To turn logging on or off for the process or to view and print logs, right-click and select the appropriate option.

Querying Enterprise Server Disk Components

Querying the server disk components produces the same information that is available from the Monitor tab.

To query enterprise server disk components:

1. On SAW, select the Query tab.
2. Expand the Disks node for an enterprise server you want to monitor and select one of these icons:
 - EnterpriseOne Disks
 - All Disks

Querying Enterprise Server IPC Resources

To query enterprise server IPC resources:

1. On SAW, select the Query tab.
2. Expand the IPC node for an enterprise server you want to monitor and select one of these icons:
 - Mutexes
 - SWMR Locks
 - Message Queues
 - Shared Memory
 - All

Querying Enterprise Server Log Files

Server Administration Workbench (SAW) enables you to query for summary server log and debug log information, as well as summary information for print queue logs. You can also select and display any log or debug log from the 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, select the Query tab.
2. Expand the Logs node for an enterprise server you want to monitor and select one of these 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 select View Selected.

SAW displays the log file or debug log file.

4. To delete the file, select 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.

Querying the Enterprise Server Users

To query the enterprise server users:

1. On SAW, select the Query tab.
2. Expand the node for an enterprise server you want to monitor and select 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) enables you to query for data about job queues on the enterprise server. When you select the Queue icon, SAW displays-in a separate window-information about each job in the queue. The information is displayed in these columns:

- Queue Name
- Queue Type
- Queue Status
- Queue Max Jobs

To query the enterprise server job queue:

1. On SAW, select the Query tab.
2. Expand the node for an enterprise server you want to monitor and select Queues.

SAW populates a separate window with information on jobs in the queue on the enterprise server.

Configuring SAW and Monitoring EnterpriseOne Servers from the Web

This section provides an overview of SAW and monitoring EnterpriseOne servers from the Web and describes how to:

- Create an SMC file.
- Regenerate an SMC file.
- Modify advanced SMC settings.
- Add EnterpriseOne servers.
- Add ports for multiple servers.
- Add events to multiple server ports.
- Add statistics to multiple server ports.
- Add profiles to events.
- Customize server ports.
- Change server ports.
- Change statistics on a server port.
- Change events on a server port.
- Change the profiles of a single event.
- Run the SAW agent in Windows.
- Run the SAW agent in UNIX.
- Start the SAW agent on the AS/400.

Understanding SAW and Server Monitoring 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 using email 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 these high-level tasks:

- 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 (email and pager addresses) to each event.
- Add, change, or delete ports, processes, and email addresses for selected server ports.

After 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 these 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 percent), 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.

After 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 the 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 emails 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 email 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.

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. This 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. |
| 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. |

Server Monitoring Configuration (SMC) includes these views that enable 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 selecting the appropriate link on the right side of the screen.

Each view gives administrators a different perspective on the data and enables 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 enables users to view the result of their requests in a separate frame in the same window. | <ul style="list-style-type: none"> • 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 |
| 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. | <ul style="list-style-type: none"> • 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 |

| View | Description | Available Functions |
|-----------------|--|--|
| Port View* | Displays an array of server ports by server name and port number. Graphical interface displays 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. | <ul style="list-style-type: none"> • 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 enables users to modify some of the settings. | <ul style="list-style-type: none"> • Force agent cycle • Modify events and email 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 display the same information in different formats.

From the Port View, Table View, and Frame View, you can access a monitoring interface for each server by selecting the link for that server. Each interface enables 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.

Creating an SMC File

Before you can monitor server ports from 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 these 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 the list of favorites.
3. On the Server Administration Workbench (SAW) WEB Version page, complete these fields:

Agent Web Server Name Type the host name of the web server.

Agent WEB Server port number Type the 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. Select OK to create the file.
6. On the SAW Home Page, complete these fields:

JAS Install Directory Type the 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 the 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 Type 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 Type 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 the company mail server, for example, mail.company_name.com. The SAW agent uses this server to send email/pager notifications when it detects faulty processes.

To access more information about any of the fields, select the Help button at the top of the screen.

7. Select Submit.

Regenerating an SMC File

To regenerate an SMC file:

If you update SAW to a new version for which the SMC file is no longer compatible, or if the SMC file is corrupted for any reason, you can regenerate the SMC file from a text file.

Each time you change or save the SMC file, SAW creates a text file that contains the same information as the SMC file, but in a text file format. If the SMC file is called saw.smc, then each time you change or save the 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. Select 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 specific requirements.

To modify advanced SMC settings:

1. On the SAW Home Page, select Advanced Settings from the Action field.
2. On SMC - Advanced Settings, modify any of these fields to suite the SMC configuration:

| | |
|------------------------------------|---|
| Agent cycle | Enter time (in seconds) between the start of one agent monitoring cycle and the next. |
| 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 email address | Change this value to false if you do not want the system to verify that the email 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. |
| 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 favorite link. |

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 these 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 these steps 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, select the Global View link to view the current configuration. This screen displays 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 select 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, perform one of these options:
 - To add a port number for several servers at the same time, complete the task Adding Ports for Multiple Servers in the EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration.
 - To add ports to a single server, complete the task Customizing Server Ports in the EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration.

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.

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 these sub steps to configure a port number:

- On SMC - Adding new Port, complete these fields:

Port Number Enter the port number you want to monitor (for example, 6004).

Description Enter a description of the port (for example, E1_SP4).

Server Type Select the down arrow and select the type of server from the list.

- Check the box to the left of each server that you want to assign this port number.

- 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, select the Global View link.

Adding Events to Multiple Server Ports

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 customize the 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.
2. On the SMC - Work with Events screen, select Activate next to the event you want to configure (in this example, Disk Percent Used).
3. For each of these events, enter a Warning Value and Alarm Value, or accept the default values:

| Event | Parameter |
|----------------------------|--|
| Disk Percent Used | Number (in percent) 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. Select the Global View to see the current configuration.

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 displaying 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, select Activate next to User Connection Graph.
3. On SMC - Activate Ports for Statistics, select each port number for which you want to include statistics.

4. Click the Action button.
5. Select 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 email 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 email 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 the 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 email addresses of a person you want to include on the notification list.

You must enter an email address in the First Email field. The 2nd and 3rd Email fields are optional. However, if you do not enter an email address in the 2nd and 3rd Email 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 select 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 these 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, select Remove Server next to the server you want to delete.
7. To modify a server configuration, select Modify Server next to the server you want to modify.
8. Change the fields as needed and click the Action button.

Changing Server Ports

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, select 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 these 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, select Remove Port next to the port number you want to delete.
7. To modify a port, select Modify Port next to the port number you want to modify.
8. Change the fields as needed and click the Action button.

Changing Statistics on a Server Port

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, select 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.

Changing Events on a Server Port

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 the 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.

Changing the Profiles of a Single Event

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 email 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 email address as needed, then click the Action button.

Note. You can add new email addresses, modify or delete an existing address, or change the person's name.

Running the SAW Agent in Windows

To run the SAW agent in Windows:

The script file for the SMC file that you create is created under the SAW database directory.

From Windows Explorer, go to the SAW database directory and double-click the SAW Agent script file.

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).

Running the SAW Agent in UNIX

To run the SAW agent in UNIX:

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.

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).

Starting the SAW Agent on the AS/400

To start the SAW Agent on the AS/400:

1. Log in to the AS/400.
2. Enter this command:

```
strqsh
```

This will start a Unix-like shell session.

3. Change directory to the SAW database directory.
4. Run the SAW agent by entering the name of the SAW Agent script. If the 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
```

5. In order to start the SAW Agent, enter this command:

```
SawAgent_saw_xe.smc.sh
```

Monitoring Enterprise Servers from the Web

This section provides an overview of the views you can use to monitor enterprise servers from the Web and describes how to troubleshoot call-object processes.

Understanding Enterprise Server Monitoring from the Web

The Enterprise Server Monitor enables 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. Select a view from the dialog box, and then click the Monitor button to view data. Each view contains a header with this information on the enterprise server:

- Enterprise server name and port number, such as sundev/6012.
- EnterpriseOne release, such as 811 - 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.

Understanding the Server Summary View

The Server Summary view enables you to conveniently monitor the status of the Enterprise server and the processes that are running on it.

The Server Summary view displays in a table these 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 button to remove the zombie process from the table.

The Server Summary view displays these 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 to activate or deactivate the debugging for the process.

Note. To view further details on server processes, click the process name in a row of the table.

Process Name

You can click the URL that identifies an enterprise server process name to view detailed information about network, kernel, and queue processes.

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 these 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 button 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 select, the Process information view displays additional information. For example, if you select a network process, the Process information view displays the values for these 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 select a kernel process, the Process information view displays these 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 select a queue process, the Process information view displays these 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 button, a web page containing the log file appears.

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.

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 trash can button. If you have SAW administrative rights, you can click the button to delete the dead process from the server.

See Also

Chapter 5, "Server Administration Workbench," Configuring SAW and Monitoring EnterpriseOne Servers from the Web, page 126

Understanding the All Log Files View

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 button in the Delete column of the table in the Log File Summary view.

See Also

Chapter 5, "Server Administration Workbench," Setting up Security Access for SAW, page 102

Understanding the 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.

Understanding the 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.

Understanding the Disk Usage View

If you click Disk Usage in the dialog box of the Enterprise Server Monitor, the Disk Usage Summary view appears.

The Disk Usage Summary view displays this 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.

Understanding the 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 enables 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 (>>).

Understanding the Package Build Files View

The Package Build Files view displays the Package Build directory and enables you to navigate through all subdirectories, navigate 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.

Understanding the 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.

Monitoring the JAS Servers from the Web

This section provides an overview of monitoring the JAS servers from the Web and describes how to:

- Monitor system XPIe servers from the Web.
- Monitor the Web server system.
- Monitor virtual clients.
- Monitor Web server log files.

- Use the Java logging system menu.
- Troubleshoot with the Web server monitor.
- Access the Web server monitor.
- Configure the Web server monitor.
- Monitor Web server users.
- Monitor JDENET connection pools.
- Monitor business functions running on connected enterprise servers.
- Monitor database connection pools.
- Run SAW Web.
- Monitor Web server Log files.
- Enable debug.log and net trace.
- Delete log files.

Understanding JAS Server Monitoring 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 these, in realtime:

- 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).

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 monitoring from one Web server to another by typing in the name of a different server.

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 stores 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.

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 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, the PeopleSoft company 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 enables an administrator to monitor all the virtual clients, and groups them by user. You can view virtual clients by clicking "View Virtual Clients" on the User List view. The virtual clients appear as records under each user session.

The User List view displays parameters and values for each user session. This table summarizes the information displayed 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. |

| Parameter in Table of User List View | Description |
|--------------------------------------|--|
| 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.

The Web Server Monitor enables 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 in which the 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 these 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

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.

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 these 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. This 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.

The JDBC Connection Pool Information view enables 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 these 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.

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.

This 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 |

| Web Server Monitor View | Parameter | EnterpriseOne Property | EnterpriseOne File Where Property Is Defined |
|-----------------------------|-----------------|---|--|
| 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 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 these 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. This 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 enabled 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. |

| System Summary Table Parameter | Description |
|--------------------------------|--|
| 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 in the Web Server Monitor, click the scroll button in the upper left-hand corner of the form. This drop-down menu displays each available view in the Web Server Monitor:

| 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. |
| realtime Events | Events currently being run on the Web Server. |
| 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

[Chapter 5, “Server Administration Workbench,” Monitoring JDENET Connection Pools, page 158](#)

[Chapter 5, “Server Administration Workbench,” Monitoring Business Functions Running on Connected Enterprise Servers, page 158](#)

Monitoring Virtual Clients

Multiple Application Framework (MAF) enables 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 enables 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 View Detail to view more specific information about the virtual client.

From the Virtual Client Detail form, you can view these:

- 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 list of available options. This view displays all of the virtual clients that are currently active:

PeopleSoft. Server Administration Workbench (SAW) for JAS Servers

Host: denxpi32 Port: 85

Views: Virtual Clients Refresh sawJas

Virtual Clients Fri Mar 19 11:53:35 MST 2004 [New Window](#) ?

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 |

Example of active virtual client display

You can view and modify certain JAS.INI settings by clicking the link 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.

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.

For more information about these settings, refer to the Web Server Installation guide for the platform you are using.

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 Log File view provides these menus:

- Log File buttons
- Java Logging System menu
- Web Server Files menu

The yellow buttons at the top of the screen enable 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.

Note. 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 the workstation, and they cannot be retrieved. Do not delete the files unless you have administrative privileges.

Using the Java Logging System Menu

The Java Logging System pull-down menu (located on the left side of the Log File view) includes these 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.

To view or change a property, click the Key to display its value. To change the value, type a new value and click the Submit (>>) button.

Work with Components

Select Work with Components to display the list of all the components of a Property Object in the jdelog.properties file. This option enables you to add, modify, or delete a component from the file:

- To add a component

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 the 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.

Web Server Files Menu

The Web Server Files Menu enables you to view any of these 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 enables you to navigate through all subdirectories, navigate 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. This table lists Web Server Monitor parameters that you can monitor for troubleshooting purposes:

| Parameter | Web Server Monitor Views | 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. |

| Parameter | Web Server Monitor Views | Parameter Meaning | Possible Problem Indication |
|------------------|--------------------------|-----------------------------|---|
| 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. |

Accessing the Web Server Monitor

To access the Web Server Monitor:

Open the web browser and enter a URL with these:

- 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

To configure the Web Server Monitor:

1. On Web Server Monitor, select Config from the pull-down menu in the Views field.
2. select or enter values in these 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, this sentence appears: The OCM cache has 0 entries.

Work with Connected Users

Complete options in this task to work with all the users connected to the selected JAS Server. In this task you can view user information, send messages to users, save user data, or log users off the system.

To work with connected users:

1. On SAW for JAS Servers, select the User List view.
2. On the User List view, click the Refresh button to get the latest user information.
The Web Server Monitor displays information on users connected to the web server.
3. Perform any of these actions to modify the view:

- Click any table column heading to sort the data by that column.
 - Click the link "New Window" on the upper right corner of the screen to start a new browser window with the content of the current window.
 - Click "View With Virtual Clients" on the upper left corner of the user list to view the list of all virtual clients for each user.
4. To send messages to users, save user data, or log users off the system, select the appropriate option from the drop-down menu.
- Broadcast a message to all virtual clients for all users

Note. Virtual client users will not see the broadcast message until they refresh their browser.

- Broadcast a message to all virtual clients for selected users
- Save data for all users
(This option saves data currently being entered by all users accessing the Web Server.)
- Save data for selected users
- Log out all virtual clients for all users

Note. The virtual client browser will not go away but their connection to the JAS will be terminated.

- Log out all virtual clients for selected users

You can perform any of these operations by selecting the appropriate action from the list of available options. To perform the operations for a select group of users, first select the check box next to each user you want to include, then select the appropriate option from the available list. You can perform any of these actions for all the virtual clients for all the users without manually selecting all of them.

View Application Failure information

The Application Failure Statistics screen provides detailed information for a selected failed application. You can list this information either by the failed application or by the user who accessed the application.

To view Application Failure information:

1. On the User List view, select the Application Failure Statistics link on the right side of the screen.
The program list application failures according to the name of the application.
2. To view application failures according to the name of the user, click View List By User on the left side of the screen.

Regardless of the option you select to view the failed application (user list or application list) the same detail information is available:

| Field | Definition |
|-------|--|
| Name | If viewed by Application, this is the name of the failed application. If viewed by User Name, this is the name of the user working with the failed application. |
| Label | Label of the failure. |

| Field | Definition |
|---------------|--|
| Reason | Reason for the failure. |
| Description | Description of the failure |
| Machine Name | Name of the machine on which the failure occurred. |
| Date and Time | Date and time of the failure. |

- To return to the User List view, click Back to User List on the right side of the screen.

Monitoring Realtime Events

The RealTime Events application in SAW enables administrator to monitor realtime event metrics and statistics, failed events, and event messages associated with subscribers. By accessing the Work with Real Time Events screen, administrators can perform these actions:

- View Metrics
- Work with Subscribers
- Work with Failed Events
- View Realtime Event Logs

To access the Work with Real Time Events view

Complete these steps to access the Work with Real Time Events view. Once you access the view, complete the subsequent tasks to view metrics and statistics and to purge messages or failed events.

- On SAW for JAS Servers, select the Work with Real Time Events view.
The Web Server Monitor displays the Work with Real Time Events view.
- Click the Refresh button to get the latest information on Real Time Events.
- Click any table column heading to sort the column data.

To view metrics

Access the Work with Real Time Events view.

On the Work with Real Time Events view, select View Metrics from the Select Action drop-down field.

The application displays realtime event metrics and statistics.

To work with subscribers

Access the Work with Real Time Events view.

On the Work with Real Time Events view, select Work with Subscribers from the Select Action drop-down field.

The application displays all the subscribers on the system.

Select Use this check box to select a subscriber for purging.

User Name Type a subscriber's name.

| | |
|------------------------|--|
| Description | Type a description of the Subscriber. |
| Queue messages | Number of messages in the subscriber's queue. |
| Routed messages | Number of events routed to this subscriber. |
| Action | Click this link to purge all the messages for this subscriber. |

To work with failed events

Access the Work with Real Time Events view.

- On the Work with Real Time Events view, select Work with Failed Events from the Select Action drop-down field.

| | |
|------------------------|--|
| Select | Use this check box to select a failed event to delete. |
| User Name | View the name of the user running the event when it failed. |
| ID | Unique Event ID. |
| Name | Name of the event. |
| BSFN | Business function run by the event when it failed. |
| Environment | Environment used by the Event and business function when it failed. |
| Application | Name of the application the event was running when it failed. |
| Host | Name of the machine where the event and business function were running when it failed. |
| Date & Time | Date and time of the failure. |
| Message | More detailed information about the failure. |
| Action | Click this link to delete this failed event. |

- From the Work with Failed Events view, you can delete one or more failed events by performing these actions:
 - To delete all failed events from the system:
Select Delete All Failed Events from the available list.
 - To delete only selected events:
In the Select column, select the check box for each event you want to delete, and select Delete Selected Failed Events from the available list.

To view Realtime log files

Access the Work with Real Time Events view.

- On the Work with Real Time Events view, select Work with failed events from the Select Action drop-down field.
- On Work with Real Time Log files, complete these fields and click the action (>>) button.

| | |
|-------------------------------|---|
| Enterprise Server Name | Type the name of the enterprise server for which you want to view to log files. |
| Enterprise Server Port | Type the port number of the server for which you want to obtain the log files. |

- Click the name of the log file to view it.

Monitoring JDENET Connection Pools

To monitor connection pools:

1. In the System Summary view of the Web Server Monitor, click the scroll button and select 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

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 select 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, select 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

To monitor database connection information:

1. In the System Summary view of the Web Server Monitor, click the scroll button and select 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.

Running SAW Web

To run SAW Web:

1. To run SAW Web, enter these in the web browser:
`http://WebServerName:port/jde/saw/SawWeb.html`
2. In the SAW Web window, select User List from the available list.

The application displays a list of all the users that are currently logged into this particular web server and port. There is a link at the top of the grid that displays 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 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. select Virtual Clients from the available 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 link.

This form enables 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

To monitor web server log files:

1. In the System Summary view of the Web Server Monitor, click the scroll button and select 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.

Bouncing or Stopping the Web Server

Use this option to bounce or stop an instance of the JAS Server running on either the WebSphere or WebLogic Application Server. When you run this option, SAW runs scripts on the Application Server machine (where JAS is running) in order to bounce or stop the JAS instance. These scripts are specific to the Application Server and version you are running. When you log into sawJas using sawJas.html, SAW automatically determines the type of Application Server the JAS is running on (WebSphere or WebLogic).

To configure the start and stop scripts (WebLogic only)

If you are running the Web Server on a WebLogic Application Server, before you can bounce the server, you must copy the contents of the start and stop scripts to the SAW scripts in the appropriate directory, and configure these files.

- For Windows machines:

1. From a command line, copy the contents of the startAS_JS_[PORT].cmd script (located in the domain directory of the JAS Server) to the sawstartAS_JS_[PORT].cmd script.

where [PORT] is the JAS port number.

For example, if the port is 87, enter this command:

```
copy startAS_JS_87.cmd sawstartAS_JS_87.cmd
```

2. Open the sawstartAS_JS_87.cmd script in a text editor and add this text before weblogic.Server at the end of the script:

```
-Dweblogic.management.username=%1 -Dweblogic.management.password=%2
```

For example:

```
"%JAVA_HOME%\bin\java" %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS%
-Dweblogic.Name=%SERVER_NAME% -Dweblogic.ProductionModeEnabled=
%PRODUCTION_MODE% -Djava.security.policy="%WL_HOME%\server\lib\
weblogic.policy" -Dweblogic.management.username=%1
-Dweblogic.management.password=%2 weblogic.Server
```

3. Copy the contents of the stopAS_JS_[PORT].cmd script into sawstopAS_JS_[PORT].cmd script.

For example, if the port is 87, enter this command:

```
copy stopAS_JS_87.cmd sawstopAS_JS_87.cmd
```

4. Open the sawstopAS_JS_87.cmd script in a text editor and remove these lines:

```
set /p WLSUSERNAME=Enter username for WebLogic server:
if /i [%WLSUSERNAME%] == [] (
echo Invalid username.
goto end )
```

5. In the sawstopAS_JS_87.cmd script, replace -username %WLSUSERNAME% with -username %1 -password %2 at the end of the script.

The last line of the script should read as follows:

```
"%JAVA_HOME%\bin\java" %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS%
weblogic.Admin -username %1 -password %2 -url %HOST%:%PORT%
SHUTDOWN -ignoreExistingSessions -timeout 5
```

- For Unix and iSeries machines:

1. From a command line, copy the contents of the startAS_JS_[PORT].sh script (located in the domain directory of the JAS Server) to the sawstartAS_JS_[PORT].sh script.

where [PORT] is the JAS port number.

For example, if the port is 87, enter this command:


```
cp startAS_JS_87.sh sawstartAS_JS_87.sh
```

2. Open the sawstartAS_JS_87.sh script in a text editor and add this text before weblogic.Server at the end of the script:

```
-Dweblogic.management.username=$1 -Dweblogic.management.password=$2
```

Here's an example of the last line of the script:

```
${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS}
-Dweblogic.Name=${SERVER_NAME} -Dweblogic.ProductionModeEnabled=
${PRODUCTION_MODE} -Djava.security.policy="${WL_HOME}/server/lib/
weblogic.policy" -Dweblogic.management.username=$1
-Dweblogic.management.password=$2 weblogic.Server
```

3. Copy the contents of the stopAS_JS_[PORT].sh script into sawstopAS_JS_[PORT].sh script.

For example, if the port is 87, enter this command:

```
cp stopAS_JS_87.sh sawstopAS_JS_87.sh
```

4. Open the sawstopAS_JS_87.sh script in a text editor and remove these lines:

```
read WLSUSERNAME
if [ "${WLSUSERNAME}" = "" ] ; then
echo "Invalid username."
exit 0
fi
```

5. In the sawstopAS_JS_87.cmd script, replace -username \${WLSUSERNAME} with -username \$1 -password \$2 at the end of the script.

The last line of the script should read as follows:

```
${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS}
weblogic.Admin -url ${HOST}:${PORT} -username $1 -password $2
SHUTDOWN -ignoreExistingSessions -timeout 5
```

To bounce or stop the Web Server

Complete this task to bounce or stop the Web Server:

1. On the Web Server Monitor, click the scroll button and select Bounce Server.
2. To ensure the system runs the appropriate scripts for the configuration, fill out the fields that are relevant to the Application Server:

| Field | Description |
|-------------------------|--|
| SAW script directory | The directory on the JAS Server where the specific scripts are located. Make sure that the default directory points to where the JAS instance is running. For example: C:\builds\ERP9_PROTOTYPE\Internet\dist\webclient\saw\scripts |
| Application Server Name | Required for WebSphere only. Note that the JAS Port Number is provided in this name. For example: AS_JAS_80 |

| Field | Description |
|-----------------------------|---|
| Node Name | Required for WebSphere 4.x only. The name of the node for the JAS instance. This value refers to the node name displayed in the WebSphere administrative tool. For example: AS_JAS_80 |
| User Name & Password | Required for WebSphere (Version 5.0 and higher) and WebLogic. For example: AS_JAS_80 |
| Broadcast message | If you select to send a broadcast message to all virtual clients before bouncing or stopping the server, type the message in this field. |
| Bounce or Stop | Select either to bounce the server or stop it. |
| WebSphere version | Required for WebSphere only. Select the version of the WebSphere Application Server you are running (4.x, 5.x, and so on). Note that each version runs a different script. |
| First perform these actions | Check any of the options listed to perform that action before you bounce or stop the server. <ul style="list-style-type: none"> • broadcast message to all users • save data for all users • logoff all users <p>Note. Before SAW stops or bounces the server, SAW performs each of the options you select in the order listed.</p> |

- Click the Submit button (>>) to perform the selected action.

Enabling Debug.log and Net Trace

To enable Debug.log and Net Trace:

- On the web server, locate the sawLogButtons.jsp file, which is typically in the \E811\internet\dist\jdewww\saw\ directory.
- Open the sawLogButtons.jsp file, locate these parameters and change the values to true, as follows:
 - boolean includeDeleteButtons = true
 - boolean includeTraceOnOffButtons = true
- 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.
- 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.
- If you want to enable Net Trace, select the Net Trace option.

You can enable Net Trace only if you have enabled the Debug Log option.
- 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.

Deleting Log Files

To delete log files:

1. In the System Summary view of the Web Server Monitor, click the scroll button and select 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.

Monitoring System XPIe Servers from the Web

The System XPIe Server Monitor enables 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 select from five views:

- Broker Summary
- Event Types
- Client Groups
- Client States
- Broker Logs

Broker Summary View

The Broker Summary view enables 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 select 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. This 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. |

| Parameter in Table of Broker Summary View | Parameter Meaning |
|---|--|
| 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 buttons 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 on which the broker is running by clicking the Server Statistics URL.

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 select 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. This 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. |

| Parameter in Table of Event Types View | Parameter Meaning |
|--|---|
| Storage Type | <p>An event attribute that determines how the event is stored in the broker. Storage types are:</p> <ul style="list-style-type: none"> • 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 button 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 select 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. This 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 select 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. This 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 owner's user name. |
| Sessions | Information about a client session. Click the magnifying glass button 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 the view to warnings, alerts, or information messages.

When you select 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

This section provides an overview of monitoring EnterpriseOne with SAW on the iSeries and describes how to:

- Access SAW.
- Display, print, and delete server log files.
- Modify the server jde.ini file.

Prerequisite

SAW must reside in the EnterpriseOne system library in the library list before you can access SAW from the command line.

Forms Used to Monitor EnterpriseOne with SAW on the iSeries

| Form Name | Form ID | Navigation | Usage |
|---------------------------------|---------|---|---|
| Work with Servers | N/A | SAW, select Work with Servers from the main menu. | Monitor a named server. |
| Work with Server Processes | N/A | SAW, select Work with Server Processes from the main menu. | Start or stop EnterpriseOne processes. Filter, sort, and view information about kernel, network, and queue processes. |
| Display EnterpriseOne Processes | N/A | On the Work with Server Processes form, click the Display EnterpriseOne Processes link. | Stop individual server process. View, debug, and delete process log files. |
| Work with Server Resources | N/A | SAW, select Work with Server Resources from the main menu. | Work with server resources. |
| Display EnterpriseOne Resources | N/A | SAW, select Display EnterpriseOne Resources from the main menu. | Display IP resources. |
| Display Disk Space Resources | N/A | On the Display EnterpriseOne Resources form, click the Display Disk Space Resources link. | Display disk resources. |
| Work with Server Log Files | N/A | SAW, select Work with Server Log Files from the main menu. | Display, print, and delete server log files. |

Understanding EnterpriseOne Monitoring 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 this information:

| Application | Description |
|----------------------------|--|
| Work with Servers | This option enables 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 enables 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 enables 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) which enables 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.

The Work with Server Processes form enables 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.

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 enables one process to pass packets to another process. |

| Resource | Description |
|-----------------------------------|---|
| Single Write Multiple Read (SWMR) | SWMR is a lock that enables 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 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.

The Work with Servers application enables 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.

Accessing SAW

You can access the Server Administration Workbench (SAW) from the command line.

To access SAW:

On the command line, enter SAW.

The SAW main menu appears.

Affecting Server Processes

Access the Display EnterpriseOne Processes form.

| Description | Glossary |
|-------------|---|
| Opt | <p>A field where you can enter a value to affect the disposition of the associated log file. Values:</p> <ul style="list-style-type: none"> • 4: Stop the process. • 7: View the log file. • 8: View the debug log file. • 9: 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.

Displaying Disk Resources

Access the Display Disk Resources form.

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

Access the Work with Server Log Files form.

| Description | Glossary |
|-------------|---|
| Opt | <p>A field where you can enter a value to affect the disposition of the associated log file. Values:</p> <ul style="list-style-type: none"> • 4: Delete the log file. • 5: View the log file. • 6: 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:

1. From the SAW main menu, select 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.
2. View and modify the jde.ini file as necessary and then press F3.
The Exit form for the SEU editor appears.
3. Save the 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

This section provides an overview of monitoring EnterpriseOne with SAW on UNIX and describes how to:

- Access SAW.
- Display, print, and delete server log files.
- Modify the server jde.ini file.

Forms Used to Monitor EnterpriseOne with SAW on UNIX

| Form Name | Form ID | Navigation | Usage |
|---------------------------------|---------|---|---|
| Work with Servers | N/A | SAW, select Work with Servers from the main menu. | Monitor a named server. |
| Work with Server Processes | N/A | SAW, select Work with Server Processes from the main menu. | Start or stop EnterpriseOne processes. Filter, sort, and view information about kernel, network, and queue processes. |
| Display EnterpriseOne Processes | N/A | On the Work with Server Processes form, click the Display EnterpriseOne Processes link. | Stop individual server process. View, debug, and delete process log files. |
| Work with Server Resources | N/A | SAW, select Work with Server Resources from the main menu. | Work with server resources. |
| Display EnterpriseOne Resources | N/A | SAW, select Display EnterpriseOne Resources from the main menu. | Display IP resources. |
| Display Disk Space Resources | N/A | On the Display EnterpriseOne Resources form, click the Display Disk Space Resources link. | Display disk resources. |
| Work with Server Log Files | N/A | SAW, select Work with Server Log Files from the main menu. | Display, print, and delete server log files. |
| Work with Server INI file | N/A | SAW, select Work with Server INI file from the main menu. | Modify the jde.ini file on the server. |

Understanding EnterpriseOne Monitoring with SAW on UNIX

The Server Administration Workbench (SAW) on UNIX provides access to this information:

- Work with Servers

This option enables 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.

The Work with Server Processes form enables 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.

- Work with Server Resources

This option enables 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. You can select from these options:

- Message queues (MSG)

A message queue enables 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.

- 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), which enables 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 the 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 this command:

```
saw.sh
```

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:

Access the Display EnterpriseOne Processes form.

If you are viewing processes for the host server, press 5 to view detail information about the selected 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.

Displaying and Deleting Log Files for Server Processes

To display and delete log files for server processes:

Access the Display EnterpriseOne Processes form.

On Display EnterpriseOne Processes, select a process, then do one of these:

- Press 7 to view the log file.
Enables 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.
Enables you to view the jddebug.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 jddebug 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.

Stopping Individual Kernel Processes

To stop individual kernel processes:

Access the Display EnterpriseOne Processes form.

Select the process to stop and press 4.

Option 4 is valid only for the local server. This option enables 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, Printing, and Deleting Server Log Files

To display, print, and delete server log files:

Access the Work with Server Log Files form.

select a log file, and then do one of these:

- 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 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.

Modifying the Server jde.ini File

To modify the server JDE.INI file:

Access the Work with Server INI file form.

1. Select 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. Select 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.

Editing the Server JDE.INI File

To edit the server JDE.INI file:

Access the Work with Server INI file form.

1. 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 Activate or Deactivate Debugging

To activate or deactivate debugging:

Access the Work with Server INI file form.

1. 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

To change the port number of the server:

Access the Work with Server INI file form.

1. 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

This chapter provides an overview of monitoring EnterpriseOne with the Knowledge Module and describes how to:

- Add OWKM files to the PATROL environment.
- Load OWKM files.
- Add a host.
- Add an EnterpriseOne server.
- Monitor enterprise servers using KM commands.

Prerequisites

On the PATROL console, right-click the MainMap button, and then select these options:

- KM Commands
- Configure
- EnterpriseOne Monitored Servers
- EnterpriseOne KM (CLI)

Forms Used to Monitor EnterpriseOne with the Knowledge Module

| Form Name | Form ID | Navigation | Usage |
|------------------------------------|---------|---|------------------------------|
| Add Hosts | N/A | On the PATROL console, select Add from the Hosts menu. | Monitor a named server. |
| Add EnterpriseOne Monitored Server | N/A | On Configure EnterpriseOne Monitored Servers, select the Add option and then click Execute. | Add an EnterpriseOne server. |
| EnterpriseOne KM Configuration | N/A | On the PATROL console, right-click the OW_SYSTEM button. | Configure OWKM. |

Understanding EnterpriseOne Monitoring with the Knowledge Module

BMC Software has a product called PATROL, which enables 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 the system, you can use the EnterpriseOne Knowledge Module (OWKM) to monitor the processes on the 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.

This 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. |

| Component | Role |
|-----------|--|
| 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 this 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 buttons for each server component. You can double-click these buttons to monitor each component. In addition, you can select OWKM commands for each component to gather additional data about each one.

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.

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

This 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 the system. Default value is off. |
| Use Defaults | Retrieves and sets default values for each configuration parameter after you select the option and click OK. Default value is off. |

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.

This 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. |

| Operating System | Consideration |
|------------------|---|
| 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. |

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 button 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.

This 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 | Enables 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 a button 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.

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:

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:\E811\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 /u10/oneworld/system/owpatrol/owkm/psl
OW_install_km.sh <PATROL_DIRECTORY>
cd /u10/oneworld/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 the servers. This procedure needs to be done only once to configure the system.

To load OWKM files:

- On the PATROL console, select File, then Load KM.
- Select `OW_ALL.kml`.

The `SYSTEM_SETUP` icon should appear in the PATROL console. If you accidentally repeat this procedure at a later time, select 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, a button for the host appears in the PATROL console. You can monitor information about the host by double-clicking its button.

To add a host:

Access the Add Hosts form.

| Description | Glossary |
|----------------|---|
| Host Name | Enter the name of the local workstation. |
| Computer Class | From the list, select the name of the operating system for the local workstation. |

| Description | Glossary |
|-----------------|---|
| 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, select an unused port number. |

Adding an EnterpriseOne Server

Access the Add EnterpriseOne Monitored Server form.

| Description | Glossary |
|---------------|---|
| Host Name | Enter the name of the local workstation. |
| Port Number | Enter the server port number. |
| Web Server | Select this option if you are monitoring a web server. |
| JDE Base Path | Enter the path to the EnterpriseOne installation directory on the 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:\E811\system. For UNIX, you might enter /u15/E811/system. |

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 buttons in the PATROL console. KM Commands give you additional monitoring options for selected server components.

Note. You can select any available enterprise server component for monitoring by right-clicking the server button and choosing KM Commands and View.

This table lists and summarizes the KM Commands options for each enterprise server component in the PATROL console:

| Component | KM Commands Parameters | Comments |
|-----------|--|--|
| Server | <ul style="list-style-type: none"> • View • Refresh Parameters • Stop Monitoring • Detailed Status | <ul style="list-style-type: none"> • View enables user to select an individual server component and view information about it. • 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 Disks, Logs, Processes, and Resources components in one form. |
| Disks | <ul style="list-style-type: none"> • View • Detailed Status | View enables user to view a list of disks and the percentage used for each. Select a disk and click View Details for host, port, path, mount point, and usage data. |
| INI | <ul style="list-style-type: none"> • View • Edit | <ul style="list-style-type: none"> • View enables users to view only the server's jde.ini file. • Edit enables users with Server Administration Workbench (SAW) administrator privileges to edit the server's jde.ini file. |
| Logs | <ul style="list-style-type: none"> • View • Detailed Status | <ul style="list-style-type: none"> • View displays lists of log files, debug files, server logs, and print queue logs. Users can sort, select, and delete logs. Select an individual log and click Execute to view the log. • Detailed Status displays log type, file count, and disk usage information. |

| Component | KM Commands Parameters | Comments |
|-----------|---|---|
| Processes | <ul style="list-style-type: none"> • View • Detailed Status • Clear Missing Process Alarms | <ul style="list-style-type: none"> • 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 enables an administrator to delete blinking (alarm) icons for failed processes. |
| Resources | <ul style="list-style-type: none"> • View • Detailed Status | <ul style="list-style-type: none"> • View displays a list of resource names and types. Select 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. |

CHAPTER 6

Working with EnterpriseOne on Windows Terminal Server Edition

This chapter provides an overview of the Windows Terminal Server Edition (TSE) and describes how to:

- Set up EnterpriseOne on the terminal server.
- Troubleshoot EnterpriseOne on Windows Terminal Server Edition.

Understanding Windows Terminal Server Edition

Windows Terminal Server Edition (TSE) provides an excellent solution for EnterpriseOne in a WAN environment. TSE enables 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.

Windows Terminal Server Edition (TSE) is a multi-user extension to the Microsoft Windows family of operating systems. TSE enables 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.

This list 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 enables 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.

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 the 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 jddebug.log files.

This table lists the capabilities of TSE RDP and MetaFrame:

| Capability | TSE RDP | MetaFrame |
|--------------------------------|---|---|
| Client Platforms | <p>TSE client software runs on these platforms:</p> <ul style="list-style-type: none"> • Windows 16-bit • Windows 32-bit • Some RDP-equipped Windows terminals • Windows CE | <p>MetaFrame runs on these platforms:</p> <ul style="list-style-type: none"> • 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. | <p>MetaFrame supports these standards:</p> <ul style="list-style-type: none"> • 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. |
| CCPDD:cut/copy/paste/drag/drop | Windows 2000 supports cut, copy, paste, drag, and drop. | MetaFrame enables CCPDD between the session window and the underlying Windows desktop. |

| Capability | TSE RDP | MetaFrame |
|-------------------|--|--|
| Device Mapping | TSE RDP enables you to map local devices for printing through a work-around. | MetaFrame enables 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. |

Understanding 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, these 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 the 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.

Understanding 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.

Understanding Performance Considerations

When you add any ICA session, change the Window Colors display properties to 16 color mode or to the lowest setting that the software and hardware allow. You can access display properties from the Control Panel.

Also, for the Citrix setup, you need to select 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.

See Also

Hardware and Software Requirements in the *EnterpriseOne PeopleTools 8.11 Installation Guide*

Setting Up EnterpriseOne on the Terminal Server

To set up EnterpriseOne on the Terminal Server:

1. Install Windows Terminal Server Edition on the machine that you will use as the 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 the 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 select the All users begin with common application settings option on the Change User Option dialog box.

Install these software packages 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 select 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.

Troubleshooting EnterpriseOne on Windows Terminal Server Edition

This section describes how to:

- Troubleshoot UBE output security on TSE.
- Submit a UBE locally and run it on the TSE.
- Troubleshoot import/export with Microsoft Excel.
- Troubleshoot specification files are locked.
- Reduce JITI frequency
- Troubleshoot user cannot restart EnterpriseOne.
- Troubleshoot logging off versus disconnecting.
- Troubleshoot shortcuts do not work in email messages.
- Troubleshoot data selection and sequencing criteria lost.
- Troubleshoot run-time error occurs during server connection test.
- Troubleshoot EnterpriseOne development tools are disabled.
- Troubleshoot users experience problems accessing EnterpriseOne.
- Troubleshoot log path is incorrect.
- Troubleshoot shortcut path is incorrect.
- Troubleshoot only one user can log in to EnterpriseOne.

Troubleshooting UBE Output Security on TSE

An EnterpriseOne TSE user sends PDF files by default from the Enterprise server to the local \E811\PrintQueue directories using the Work With Servers (P986116) application. Users select 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 this 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 user profile 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.

Submitting a UBE Locally and Running it on the 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:

1. From the System Administration Tools menu (GH9011), select Logical Data Sources (P986115).
2. On Logical Data Sources, select the TSE machine name with System as the data source and click Select.
3. On Work with Data Sources, click Add.
4. On Data Source Revisions, complete these fields:
 - Data Source User
Enter DB to specify a Local Data Source.
 - Data Source Name
Enter TSE Local.
 - Data Source Type
Enter N to specify MSDE/ODBC.
 - DLL Name
Enter JDBODBC.DLL.
 - Database Name
Enter System Local.
 - Server Name
Enter LOCAL.
 - Platform
Enter LOCAL.
5. When you run the UBE on the TSE, on the Work with Batch Versions from, select the report you want to run and click Select.
6. On Version Prompting, select Advanced from the Form menu.
7. On Advanced Version Prompting, select the Override Location option and click OK.
8. On Version Prompting, select the Data Selection option and click Submit.
9. On JDE Data Sources, select 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.

We recommend using Microsoft Office 2000 on TSE. To install Microsoft Office 2000 in a TSE environment, access this 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 these 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 this task:

Reducing JITI Frequency

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:

```
[INTERACTIVE RUNTIME]
EXCEPTION_Enabled=False
```

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 Email Messages

Workflow provides the ability to send shortcuts to EnterpriseOne applications using email messages. This function does not work when the email application, such as Microsoft Outlook, is not currently active on the terminal server. When the email application invokes the shortcut, the operating system attempts to launch the shortcut on the local machine and not on the terminal server.

Run the email application and EnterpriseOne on the same terminal server machine.

Troubleshooting Data Selection and Sequencing Criteria Lost

This 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, we instruct 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 select 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 this 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.

CHAPTER 7

Administering EnterpriseOne on a Unix Cluster

This chapter provides an overview of clustering and describes how to:

- Maintain multiple instances of EnterpriseOne in a clustered environment.
- Set up clustering.
- Set up HACMP for AIX clustering.
- Set up EnterpriseOne for HACMP.
- Create an application server.
- Set up Sun Solaris clustering.
- Troubleshoot HP-UX clustering.

Understanding Clustering

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 the enterprise.

Clustering enables 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 the 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 the 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.

Maintaining Multiple Instances of EnterpriseOne in a Clustered Environment

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. These 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
```

where

- *Service Name or Port Number* is the actual port number or the name of a service that you enter into the `/etc/services` file.
- *Numeric Value* is the IPC key value of the EnterpriseOne instance.

The IPC key values should differ by at least 1000 between any two EnterpriseOne instances.

Setting Up Clustering

This section describes how to:

- Configure Oracle Parallel Server (OPS).
- Set up an Oracle package for MC/ServiceGuard.
- Set up an EnterpriseOne package.

Configuring Oracle Parallel Server (OPS)

Oracle Parallel Server (OPS) enables concurrent database access from multiple nodes in a cluster. If you use OPS, you must install MC/LockManager, not MC/ServiceGuard.

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.

To set up Oracle Parallel Server 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. This 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=jdel
export ORAENV_ASK=NO
. $ORACLE_HOME/bin/oraenv
su oracle -c '$ORACLE_HOME/bin/lsnrctl start'
su oracle -c '$ORACLE_HOME/bin/svrnmgsl' <<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 enables 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.
This 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/svrmgr1' <<EOF1
connect internal
startup
exit
EAOF1
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 these 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 these setting:

```

[CLUSTER]
PrimaryNode=Package IP name

```

where `Package IP name` 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 enables you to utilize the installation defaults and the start and end scripts provided by the PeopleSoft company.

Depending on the needs of the enterprise, you might want to install EnterpriseOne on a shared volume group. This setup enables multiple nodes in a cluster to access a single version of EnterpriseOne, but only one node at a time. This setup also enables 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.

This 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/owenvsu jde
<< EOF1mv $OWHOME/log/jde*.log
$OWHOME/log/oldlogscd
$SYSTEM/bin32RunOneWorld.sh
EOF1
test_return 51
}
```

This 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 the 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 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 $OWHOMELog/jde*.log $OWHOMELog/oldlogs
```

This line moves any logs in the log directory to a backup log directory, which you create. This command is particularly important if the 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 $OWHOMELog/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 we provide 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.

Setting up HACMP for AIX Clustering

This section provides an overview of HACMP for AIX clustering and describes how to create group and user accounts.

Understanding HACMP for AIX Clustering

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 the LAN, disk capacity, and performance requirements.

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) enables 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 enable the system administrator to add users, files, and security functions without stopping mission-critical jobs.

HACMP provides several configuration options, including these:

- 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 enables customers to select 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 DB2 Parallel Edition and Oracle 8 Parallel Server.

Several components make up the HACMP environment, including these:

- 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).

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 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 the 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 the selection by searching all password and group files on the node where the new user and group will be created.
4. Enter this command on the command line:


```
# smit hacmp
```
5. From the menu, select Cluster System Management, then Cluster Users & Groups, then Groups, and then Add a Group to the Cluster.
6. Select 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, select Users and then Add a User to the Cluster.
9. Select the same resource group that you chose for adding a group.
10. Add a user psft, assign it the pre-selected unique ID number, and select 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.

This section describes how to:

- Edit the owenv script.
- Edit the Start Resource Control script.
- Edit the Stop Resource Control script.

Editing the owenv Script

To edit the owenv script:

1. In Windows Explorer, go to the /\$SYSTEM/bin32 directory and open the owenv file.
2. Edit the emphasized lines:

```
#!/bin/ksh
## set OWHOME to point to the base install path for EnterpriseOne
export OWHOME=/ow2/PeopleSoft/E811
## 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_PATHexport PATH=$PATH:$SYSTEM/bin32
```

3. Save and close the file.

Editing the Start Resource Control Script

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 emphasized lines:

```
# _____
#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 emphasized 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 EnterpriseOne Tools 8.94 PeopleBook: Server and Workstation Administration.

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.

Editing the Stop Resource Control Script

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 lines:

```
# _____
#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.

This section describes how to:

- Move the control scripts.
- Define an application server.
- Define cluster resources.

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 to meet the requirements. Before you can create an application server, you must move the scripts to a non-shared disk directory.

Moving the Control Scripts

To move the control scripts:

1. Log in as root and enter these commands:

```
#export SYSTEM=< EnterpriseOne system directory path>
```

Where EnterpriseOne system directory path is the path to the system directory. An example of the path is /ow2/PeopleSoft/811/system.

```
# cd /usr/sbin/cluster
```

```
# mkdir scripts
# cd scripts
```

Be sure to include the periods, preceded by a single space, in these commands:

```
# cp $SYSTEM/bin32/StartResource.sh .
# cp $SYSTEM/bin32/owenv .
# chmod 755 *
```

2. Repeat this step on all failover nodes and FTP over the modified script files.

Defining an Application Server

To define an application server:

1. From the command line, enter this command:

```
# smit hacmp
```

2. From the menu, select Cluster Configuration, then Cluster Resources, then Define Application Servers, and then Add an Application Server.
3. Complete these 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.

Defining 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.

To define cluster resources:

1. From the command line, enter this command:

```
# smit hacmp
```

2. From the menu, select Cluster Configuration, then Cluster Resources, and then Change/Show Resource for a Resource Group.
3. Select 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.

Setting Up Sun Solaris Clustering

This section provides an overview of Sun Solaris clustering and describes how to:

- Modify the SunStartResource.sh script.
- Modify the SunStopResource.sh script.
- Modify the owenv script.
- Register EnterpriseOne with SUNClustering.

Understanding Sun Solaris Clustering

This software provides higher availability for the applications because it enables you to recover almost instantaneously from a power failure or hardware problem. It also enables 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.

You must have a disk that is accessible to all machines in the cluster, and this disk must be large enough to accommodate the 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.

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.

The SunStartResource.sh 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.

The SunStopResource.sh 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.

The owenv script sets various UNIX environment variables that are needed by EnterpriseOne. The script is called from within the SunStartResource.sh and SunStopResource.sh scripts.

The SunOracleMgr.sh 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.

See Also

[Chapter 7, “Administering EnterpriseOne on a Unix Cluster,” Registering EnterpriseOne with SUNClustering, page 207](#)

Modifying the SunStartResource.sh Script

To modify the SunStartResource.sh script:

1. Under Global Variables, there is a call to /suncldata/peoplesoft/E811_sp1/system/bin32/owenv. Change /suncldata/peoplesoft/E811_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.

Modifying the SunStopResource.sh Script

To modify the SunStopResource.sh script:

1. Under Global Variables, change the call to /suncldata/peoplesoft/E811_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.

Modifying the owenv Script

To modify the owenv script:

1. Set OWHOME to be the base directory of EnterpriseOne, for example, /suncldata/peoplesoft/E811_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.

Modifying the SunOracleMgr.sh Script

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 the 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 this 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 this command with no options:

```
/opt/SUNWcluster/bin/hareg
```

If the line containing the service name you assigned to EnterpriseOne contains off, then enter this command:

```
/opt/SUNWcluster/bin/hareg -y [service name]
```

This sets the data service to “on,” which enables the data service to switch physical hosts when appropriate.

3. To test whether the ccluster switches, enter this 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.

Troubleshooting HP-UX Clustering

These sections address specific problems with HP-UX clustering.

Problems with Oracle Parallel Server (OPS)

Complete these 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 these 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 this 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 this 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 PeopleSoft 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.

CHAPTER 8

Administering EnterpriseOne on a Windows Cluster

Configuration of MSCS has changed with EnterpriseOne 8.9. This information is only valid beginning with EnterpriseOne 8.9 and should not be used with previous releases. For more information about how to install and set up a Microsoft cluster, see the appropriate Microsoft documentation.

This section discusses how to:

- Upgrade EnterpriseOne in a Windows Cluster environment
- Set up EnterpriseOne on a Windows Cluster
- Create a Static IP Address for the PeopleSoft Network Name
- Test EnterpriseOne Connections on the Cluster
- Set Advanced Options for EnterpriseOne Resources
- Set Failover and Failback Parameters
- Troubleshoot a Windows Cluster

Prerequisites

Before you complete the tasks in this section:

- Partition the disk array to logically divide the software components. Typically, you will set up these 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 the 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 redundancy. One network card will communicate with the public network, and the second card will connect between nodes. This setup enables the 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.

See Also

Starting the Windows Enterprise Server in the *EnterpriseOne PeopleTools 8.11 Installation Guide*

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Working with Servers,” Managing Server Jobs

Upgrading EnterpriseOne in a Windows Cluster Environment

If you are already running PeopleSoft EnterpriseOne in a cluster but are upgrading to EnterpriseOne 8.9 or later, you must change the cluster configuration for the cluster to operate properly. EnterpriseOne 8.9 and later releases require use of an IP address, and a network name separate from the cluster IP address and network name. (In previous releases, we used the cluster network name and IP address for the PeopleSoft enterprise server name.)

Since the cluster name exists in PeopleSoft ini files and tables, the recommended solution is to rename the cluster and create new resources for the PeopleSoft group using the old cluster name and a different IP address. This requires changing the static IP address associated with the old cluster name to avoid an IP address conflict. Consult the Microsoft documentation on renaming a cluster. After the cluster is renamed, continue with these instructions.

Setting Up EnterpriseOne on a Windows Cluster

This task explains how to set up EnterpriseOne on a Windows cluster. This example is for a two node cluster.

To set up EnterpriseOne on a Windows cluster:

1. Consult the PeopleSoft MTRs to determine the appropriate Windows OS level. Install the OS on each node. Use member servers and not domain controllers.
2. Install Microsoft Cluster Server (MSCS) software on each node using an account that has administrator authority.

On Windows 2003 Advanced Edition MSCS is already installed. Refer to the appropriate Microsoft documentation for specific instructions on the installation of MSCS.

3. Access the cluster to verify the installation of MSCS.
4. Install EnterpriseOne on the disk designated for PeopleSoft EnterpriseOne Enterprise server. All nodes will share a single copy of EnterpriseOne. The name used for the Enterprise server will be the PeopleSoft network name and used by PeopleSoft EnterpriseOne to reference the cluster.
5. On the node which currently owns the PeopleSoft disk, install the network service, open a DOS command window. Change the directory location to the system/bin32 directory for EnterpriseOne.
6. Type *jdesnet -I*.
A message will be returned stating that the service was installed.
7. Exit the command prompt and open services.
8. Right click the PeopleSoft Network Service and select Properties.
9. Click This Account and change the account to a valid domain account that is in the local administrators group.
10. In the server jde.ini, verify that all references to the enterprise server are set to the PeopleSoft Network name.
11. Open cluster administrator.
12. Right click the cluster name, and select New, Group.
13. Type in the cluster name and a description.
14. Click Next.
15. Add the required nodes to the preferred owners list.
16. Click Finish.
17. When the cluster group is successfully created, click OK.
18. Right click the cluster administrator screen and select New, then Resource.
Three resources will need to be created to complete the setup. The first resource is the IP address resource for EnterpriseOne.
19. Type a name and description for the EnterpriseOne IP address. Resource type is the IP Address. The group name should be the group just created
20. Click Next.
21. Verify that all required nodes are preferred owners.
22. Click Next.
The next screen defines dependencies. The IP resource should not have any dependencies.
23. Click Next.
24. Add the IP address and subnet mask.
25. Select Public from the available list, and select Enable NetBios for this address.
26. Click Finish.
A pop-up should appear verifying successful configuration of the resource
27. To set up the Network Name resource, right click the cluster administrator screen and select New — Resource.
28. Fill in the name with the PeopleSoft EnterpriseOne Network name you have chosen. Add a description, select Network Name from the available list and select the group you created earlier

29. Click Next.

30. Verify that all required nodes are preferred owners.

31. Click Next.

The next screen defines dependencies. The Network Name should have the IP address as a dependency

32. Click Next.

Specify the Network Name in the Name field.

33. Select the DNS registration check box. Do not select the Enable Kerberos Authentication check box.

34. Click Finish. The program should display a message verifying successful configuration of the resource

35. Click OK.

36. Right click the cluster administrator screen and select New — Resource.

37. Fill in a name and a description that will help you identify the PeopleSoft network service. Select Generic Service as the resource type, and select the PeopleSoft group you created earlier.

38. Click Next.

39. Verify that all required nodes are preferred owners.

40. Click Next.

The next screen defines dependencies. The only required dependency that needs to be added is the PeopleSoft Network Name resource.

41. Click Next.

42. Type in the service name exactly as it appears in Microsoft services. Leave start parameters blank.

43. Select the check box for Use Network Name.

44. Click Next.

Do not change the Registry Replication screen.

45. Click Finish.

The program should display a message verifying successful configuration of the resource

The installation of the resources is complete.

EnterpriseOne can be started and stopped as necessary by using the cluster administrator, right click the network service resource, and selecting online or offline as needed. Because the network IP address and the network name are dependent on the PeopleSoft network service resource, all of the resources will start in the proper order when the network service resource is brought online.

Adding EnterpriseOne Network Resources to the Group

To add the EnterpriseOne network resources to the cluster group:

1. From the Cluster Administrator main menu, select New and then Resource from the File menu.
2. On New Resource, complete these fields for the JDE network resources, and then click Next.
 - Name
 - Description
 - Resource Type

Select Generic Service from the list.

- Group Name

Select the cluster group.

Note. Do not select the option to Run this resource in a separate Resource Monitor.

3. On Possible Owners, verify that the possible owners match the preferred owners, and then click Next.
4. On Dependencies, select 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.
5. On Generic Service Parameters, do these, and then click Next:
 - Type the JDE network service name as it appears in the NT Services applet.
 - Leave the Startup Parameters field blank.
6. On Registry Replication, click Finish.
7. On Cluster Administrator, click OK.

Starting and Stopping the EnterpriseOne Resources on the Cluster

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 select Online.
2. To stop network or queue services or both services, in Cluster Administrator, right-click the EnterpriseOne resource and select Offline.

Testing EnterpriseOne Connections on the Cluster

EnterpriseOne network and queue services must be set up before you can test 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. To verify that EnterpriseOne is installed correctly on the cluster server, enter these commands to run porttest:

```
cd \PeopleSoft\ddp\xxxx\System\Bin32
porttest userID password environment
```

Where *xxxx* is the specific release of EnterpriseOne (for example, 811), *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 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 by right-clicking the node name and choosing *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.

Setting Advanced Options for EnterpriseOne Resources

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. Select a resource, and then select Properties from the File menu.
3. Click the Advanced tab, and select one of these options:
 - Do Not Restart
 - RestartIf you want to disable failover, select 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

To set failover and failback parameters:

1. From the Cluster Administrator main menu, select the appropriate group, and then select Properties from the File menu.
2. Click the Failover tab, and review these 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 these 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.

Troubleshooting Windows Clustering

This section includes these troubleshooting tasks:

Changing an IP Address

After moving the cluster equipment to another area, we had to change the IP address for cluster nodes and any virtual machines connected to them. In order to perform this task, we 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.

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.

Creating a Static IP Address for the PeopleSoft Network Name

Starting with EnterpriseOne 8.9, you must create a network name for PeopleSoft. The network name is used for the enterprise cluster server name. If you are installing a cluster aware DBMS, it also requires one or more static IP addresses.

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 each node only requires one network card, Microsoft recommends using two network cards as a backup. One network card will communicate with the public network, and the second card will connect between nodes. This setup enables the 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 ensure that EnterpriseOne and the database do not share disk resources. EnterpriseOne resources should be in a group other than the cluster group. EnterpriseOne resources will have their own network name and IP address.

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 same cluster group.

If you use Oracle, SQLServer, or IBM UDB cluster services, the cluster might experience dependencies on database resources. Consult the appropriate Oracle, SQLServer, or IBM UDB documentation for more information.

Testing EnterpriseOne Connections on the Cluster

Complete this task to test the EnterpriseOne connection.

Note. EnterpriseOne network and queue services must be set up before you can test connections on the cluster.

To test an EnterpriseOne Connection:

1. Sign on to Windows.
2. Using Cluster Administrator, verify that all EnterpriseOne resources are online for the EnterpriseOne group.
3. To verify that EnterpriseOne is installed correctly on the cluster server, enter these commands to run porttest:

- `cd \jdedwardsoneworld\ddp\xxx\System\Bin32`

where *xxx* is the specific release of EnterpriseOne, for example, E811.

- `porttest userID password environment`

Where *userID* is a valid EnterpriseOne user, *password* is the password for that user, and *environment* is the environment you are verifying.

Note. These parameters are case-sensitive.

If EnterpriseOne is installed and configured correctly, the porttest program initializes a user and an environment. The program should display messages indicating the selection of records out of an EnterpriseOne table. If it does not, review the `jde_####.log` file that was specified in the `jde.ini` file on the enterprise server.

4. Sign on to an EnterpriseOne workstation, 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. Stop EnterpriseOne on the first node using the Cluster Administration tools.
6. To stop cluster services on the first node, right-click the node name and choosing Cluster Service. All groups will failover to the second node.
7. Verify that all resources are online on the second node and repeat these steps.

See Also

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Working with Servers,” Managing Server Jobs

Setting Advanced Options for EnterpriseOne Resources

This task explains how to set up advanced options for EnterpriseOne on a Windows NT cluster in a two-server configuration.

1. From the Cluster Administrator main menu, open the group that contains the resource that you want to modify.
2. Select a resource, and then select Properties from the File menu.

3. Click the Advanced tab, and select one of these options:
 - Do Not Restart
 - Restart

If you want to disable failover, select 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

This task explains how to set up fail parameters for EnterpriseOne on a Windows NT cluster in a two-server configuration.

1. From the Cluster Administrator main menu, select the appropriate group, then select Properties from the File menu.
2. Click the Failover tab, and review these fields:
 - Threshold: Determines the number of failovers attempted by MSCS.
 - Period: Determines the period of time during which MSCS attempts failovers.
3. Click the Failback tab and review these 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.

Troubleshooting a Windows Cluster

If you encounter problems with the Windows cluster, review these topics for possible solutions.

Changing an IP Address

If you move the cluster equipment to another area, you may have to change the IP address for cluster nodes and any virtual machines connected to them. To perform this task, complete these steps:

1. Remove all defined resources and groups.
2. Uninstall the DBMS cluster software.
3. Uninstalling MSCS.
4. Change the physical IP address in the nodes.
5. Reinstall MSCS using the new IP address for the virtual machine.
6. Reinstall SQL Server Cluster Manager and Oracle Failsafe.

Reinstalling MSCS

If you already installed a version of Microsoft Cluster software, before you reinstall, you must uninstall the software using Add/Remove programs.

CHAPTER 9

Administering EnterpriseOne on an iSeries Cluster

This section includes an overview of iSeries clustering and discusses how to:

- Identify the cluster name
- Set up the enterprise servers
- Set up the client for clustering
- Set up the deployment server
- Set up logical data sources
- Set up database data sources
- Set up Object Configuration Manager
- Distribute ODBC setup from the deployment server
- Identify the cluster name on the deployment server

Understanding iSeries Clustering

iSeries clustering is a platform-specific software solution that provides users with continuous access to EnterpriseOne even when the primary server becomes unavailable. You can switchover from a primary server to a backup server either automatically or manually.

An iSeries cluster consists of more than one node, although not necessarily more than one physical machine. For example, you can use logical partitioning to represent several nodes on a single machine, or you can maintain more than one iSeries machine, each of which represents a cluster node.

A cluster name is associated with a floating or takeover IP address. Each node in the cluster has an IP address associated with a TCP/IP interface. At any time, only one node in the cluster has the interface activated. That node is the primary node, on which EnterpriseOne services are running. All other nodes in the cluster are designated as backup nodes, and the TCP/IP interface is inactive.

iSeries nodes participating in a single cluster use the cluster software to do these tasks:

- Replicate object and data changes from the primary node to backup nodes so that any backup node can assume the primary server role when an interruption in service occurs
- Read a specifier file that identifies the objects and data that must be replicated and the locations of those specified objects and data
- Use a backup node to monitor the primary node for availability.
- If the primary node becomes unavailable, use an exit program to activate the floating IP address on the first backup node associated with the cluster name.

- Restart EnterpriseOne on the first backup node, thus making it the primary node
- Queue changes to objects and data so that an original primary node can be updated once it becomes available for service

You can use EnterpriseOne and iSeries clustering software to support two-tier configurations (iSeries functioning as data server only), virtual three-tier configurations (iSeries functioning as both logic server and data server), and three-tier server configurations (separate iSeries machines functioning as a logic server and a data server).

IBM recognizes three partners who market the iSeries clustering solution:

- DataMirror Corporation
- Lakeview Technologies
- Vision Solutions

For more details on these vendor solutions, consult the CNC specialist.

EnterpriseOne Objects Used with iSeries Clustering

The PeopleSoft company requires you to download the latest iSeries Clustering objects from the Update Center. This topic discusses each of the EnterpriseOne objects required for use with iSeries clustering software:

- Specifier file
- SETOWCLST command
- Application data areas
- Exit program

Consult the CNC specialist if you need to modify any of these files.

Specifier File

The specifier file identifies all EnterpriseOne objects and data that must be replicated from the primary server node to the backup node. Replication ensures that you have a backup of essential business information on the backup node should the primary node fail.

The information in the specifier file enables the creation of application and data cluster resource groups (CRGs). Application CRGs identify nodes in a cluster that can be used to run a particular program or group of programs on the enterprise server. Data CRGs represent the locations of data and objects in the cluster. With this information established, the clustering software is able to replicate the specific set of objects and data to a specific server node.

Note. Do not attempt to modify the specifier file directly. To customize the file, consult the CNC specialist.

This table summarizes the items that must be replicated for EnterpriseOne-iSeries clustering:

| Replication Item | Files and Directories for Replication |
|----------------------------------|---|
| All data used for an environment | Examples: PRODDTA/*PF PRODCTL/*PF |

| Replication Item | Files and Directories for Replication |
|------------------------------------|---------------------------------------|
| Object Librarian | OL811/*PF |
| Central Objects | COPD811/*PF |
| Data dictionary | DD811/*PF |
| PrintQueue Directory | /E811SYS/printqueue |
| EnterpriseOne jobs table (F986110) | SVM811/F986110 |

Note. Clustering vendor solutions shouldn't replicate EnterpriseOne specification files in the integrated file system (IFS).

SETOWCIST Command

You can use the SETOWCIST command to access the specifier file. This command is necessary to update the specifier file if you add, remove, or modify the name or location of a EnterpriseOne iSeries library or object.

Running the SETOWCLST Command

To run the SETOWCLST Command:

1. Sign onto the iSeries system and add the EnterpriseOne system library to the library list.
2. Enter this command:

```
setowclst
```

3. Enter one of these commands and press ENTER:

```
*ADD
```

```
*REMOVE
```

```
*CHANGE
```

```
*CLEARALL
```

Note. THE *CLEARALL command removes all replicated objects. If you select this command, no further steps are required.

4. In the IFS Object field, enter one of these:
 - *Y* if the clustering object is an IFS object
 - *N* if the clustering object is a QSYS object
 - If the clustering object is an IFS object, enter the name and the path of the IFS clustering object.

Note. Replicate all objects in a directory by typing * and the path, for example, */E811SYS/PRINTQUEUE/**.

- If the clustering object is a QSYS object, enter the name and the library of the QSYS clustering object.

Note. Special values of *ALL and *PF (physical file) can also be used as valid objects.

Application Data Areas

A data area is an object used to communicate data such as variable values between programs within a job and between jobs. Application data areas contain information about resilient resources in the EnterpriseOne clustering setup. Resilient resources are objects and other information located on more than one cluster node. Important resilient resources include data and programming objects critical to running EnterpriseOne.

The clustering software uses EnterpriseOne information to create and keep track of CRGs, which identify nodes in a cluster and the types and locations of EnterpriseOne resilient resources.

There are two application data areas in the EnterpriseOne-iSeries clustering environment:

- Input application data area
- Output application data area

Input Application Data Area

The input data area is used to communicate information about the EnterpriseOne application to the IBM cluster middleware business partner. The cluster middleware uses this information to create CRGs. The input application data area will contain information about EnterpriseOne, its resilience information, and information about required data.

Output Application Data Area

The output data area is used by the IBM cluster middleware business partner software to track the use of the CRG it created for use in EnterpriseOne. The CRG identifies the nodes in a cluster used to run programs on the enterprise server and the locations of data and objects in the cluster.

Cluster Exit Program

The iSeries clustering software invokes the exit program, which is called CLSTR_EXIT, when a failure on the primary node requires a failover to the backup node. The program stops EnterpriseOne services running on the backup node until replication of all EnterpriseOne objects and data identified in the specifier file has been completed. After replication has completed, EnterpriseOne services restart.

Technical Considerations

To aid in the implementation of iSeries clustering software with EnterpriseOne, we make these suggestions:

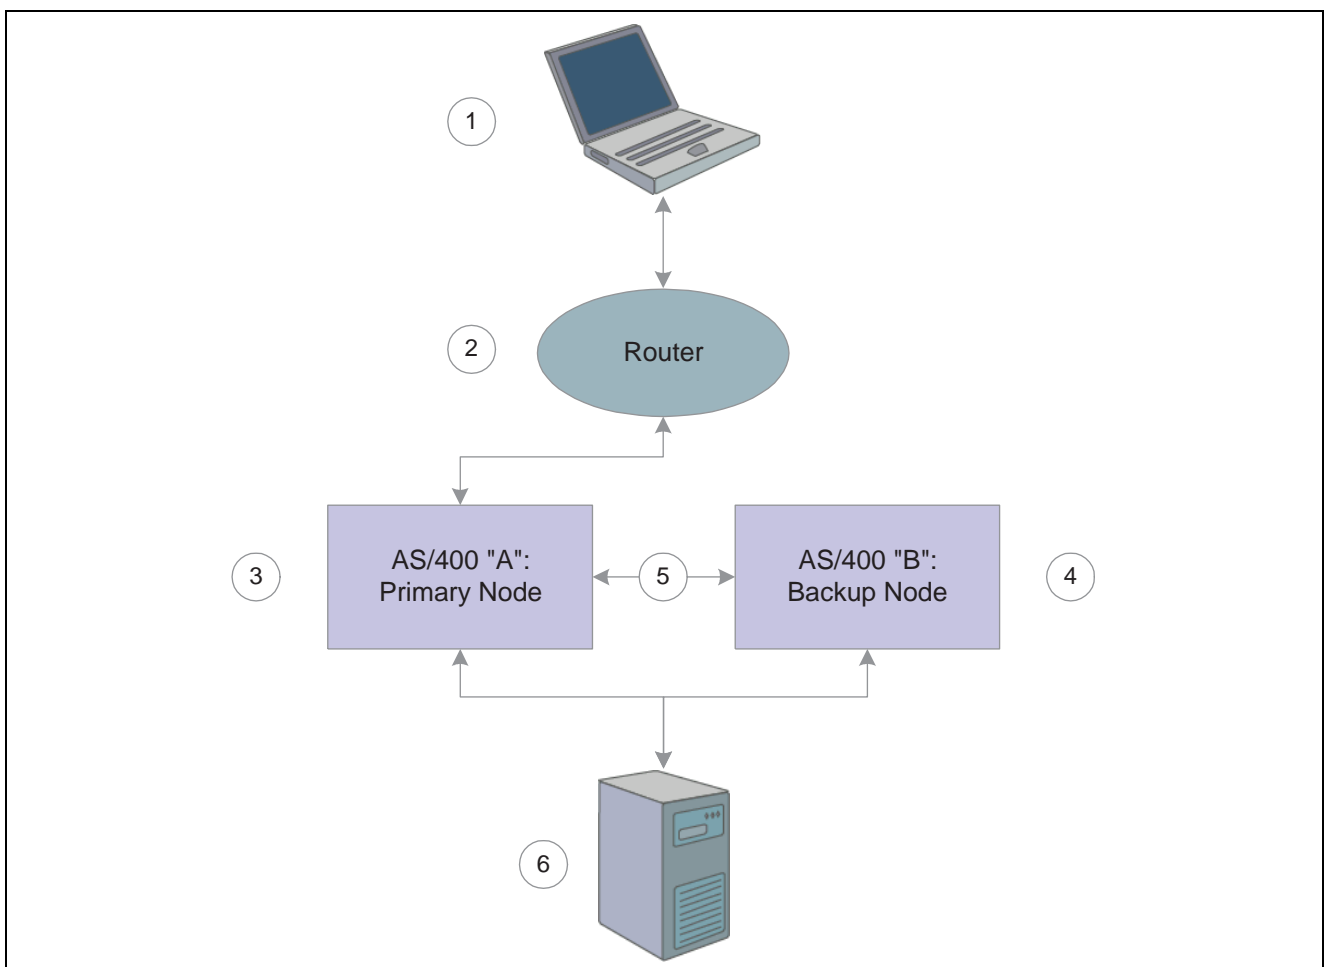
- Represent iSeries nodes either by a logical partition or by an entire machine, depending on the system.
- Use separate pipes (LAN cards) to minimize the effects of clustering software functions on end user activities. For example, dedicate one pipe to running applications and one pipe to object/data replication and node monitoring.
- Consult IBM documentation on iSeries clustering solutions for further details on handling LAN card setup.
- Use the EnterpriseOne deployment server to deliver EnterpriseOne packages to server cluster nodes. Do not attempt to use the clustering replication processes to deploy EnterpriseOne packages between cluster nodes.
- Use the clustering software to replicate object/data changes made by EnterpriseOne users.
- Define to the specifier file on each node any changes you make to EnterpriseOne library names or the locations of EnterpriseOne-supplied objects, if those libraries and objects were listed in the specifier file provided by the PeopleSoft company.

- Evaluate each EnterpriseOne cumulative release, electronic software update (ESU), service pack, or program temporary Fix (PTF) to determine if any changed objects are listed in the specifier file. If any are, be sure to define the changes to the file.

The setup and configuration of EnterpriseOne discussed in this chapter assumes a virtual three-tier setup for a two-node server cluster. The setup should be used as a general reference guide only. Many variations on the setup are possible and will be determined by the specific requirements of the organization. Consult the Cluster Middleware Vendor or CNC specialist for additional details on how to configure the iSeries system to use clustering.

Understanding iSeries-EnterpriseOne Architecture with Clustering

This graphic illustrates the main components of a virtual three-tier clustering setup:



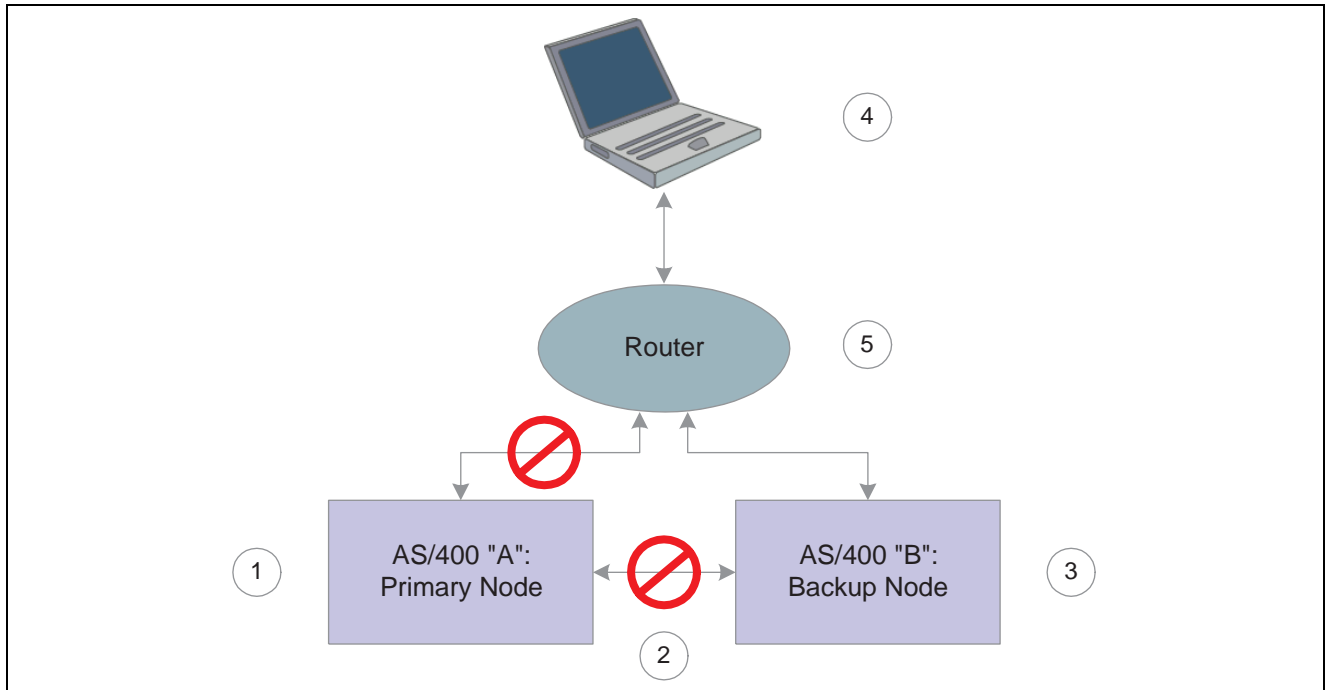
iSeries Clustering Architecture — normal operation

The roles of each component in the cluster are as follows:

- Workstation (1) attempts to access the primary node.
- Router (2) directs workstation access request to the primary node using the cluster name and the associated floating IP address.
- Primary node (3) provides services to the workstation and sends changes to objects specified in the specifier file to the backup node.

- Backup node (4) monitors primary node and applies changes to data and replicated objects using static IP address (5).
- Deployment server (6) updates both primary and backup nodes with changes to application objects using package deployment.

This graphic illustrates the role of each component in the clustering environment when a failover scenario occurs:

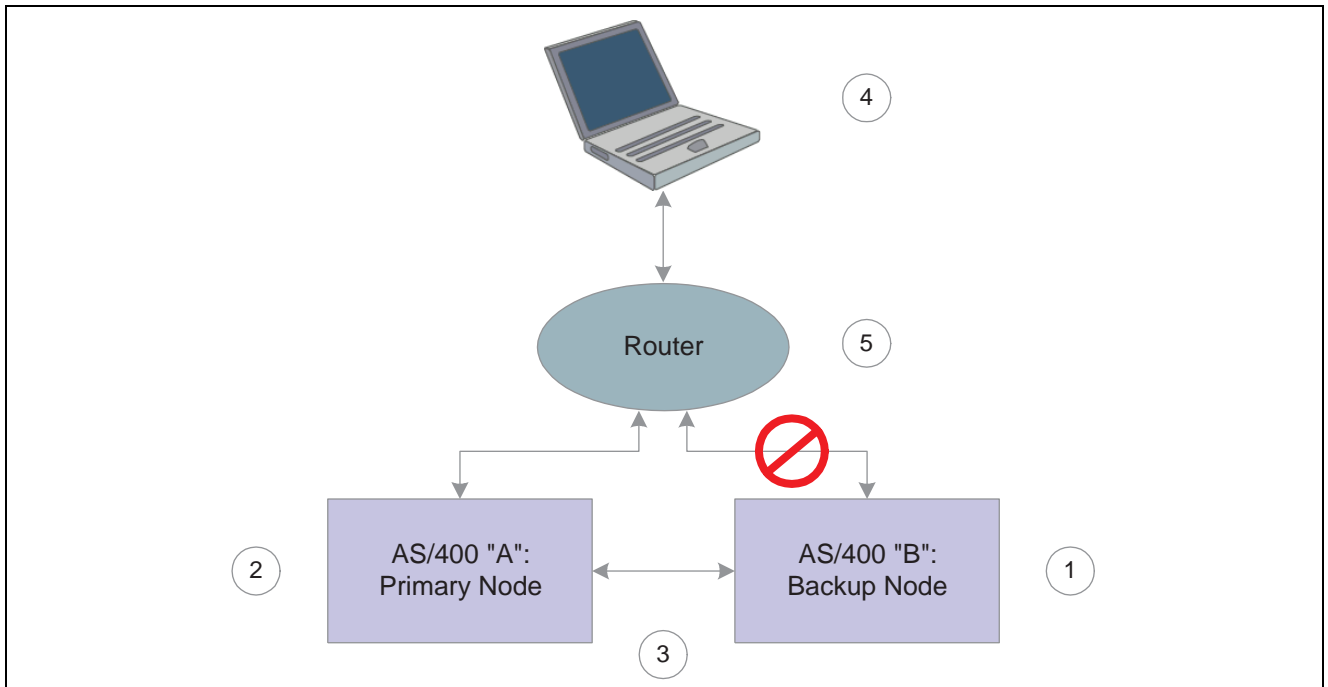


iSeries Clustering Architecture — failover scenario

The roles of each component in the cluster in the failover scenario are as follows:

- Ex-primary node (1) ends communication with the workstation and the backup node.
- Object and data changes using static IP address end, as does monitoring of nodes(2).
- Clustering software on iSeries B detects failure of iSeries A, activates failover TCP /IP interface, applies queued changes to objects listed in clustering specifier file, and makes EnterpriseOne services available on iSeries B (3).
- Workstation (4) requests server services and waits for a response for a length of time specified in its jde.ini file. After EnterpriseOne notifies the workstation of a lost server connection, the workstation attempts to reconnect.
- Router (5) receives reconnection attempt from workstation and directs request to the new primary node using the cluster name and the associated floating IP address.

After the failover, the ex-primary node should eventually become available once again for service. This graphic illustrates the role of each component in the clustering environment when the return to normal operations occurs:



iSeries Clustering Architecture — return to normal

The roles of each component in the cluster in the return to normal operations are as follows:

- The clustering software stops the failover interface on iSeries B (1), ending communication with the workstation.
- The clustering software starts the failover interface on iSeries A (2) and applies any queued changes to objects listed in the clustering specifier file.
- Replication and monitoring services restart (3).
- Workstation (4) requests server services, waits for a response for a length of time specified in its jde.ini file. After EnterpriseOne notifies the workstation of a lost server connection, the workstation attempts to reconnect.
- Router (5) receives the reconnection attempt from workstation and directs the request to the primary node using the cluster name and the associated floating IP address.

For a full discussion of platform, hardware, and LAN and WAN configurations, consult IBM clustering documentation as well as documentation provided by each of IBM's approved clustering software partners.

Minimum Setup Requirements for iSeries Server Nodes

Each node in a virtual three-tier cluster configuration must be set up with these components:

- A host table and domain names server (DNS) entry with the cluster name that is associated with a floating IP address.
- An IP interface with the same floating IP address as was entered in the host table. The floating IP address is used to find the backup node when the primary node is unavailable.
- A second IP interface with a unique address. This address is used for object and data replication.
- Identical operating system releases (such as OS/400 V5R2 or higher) Identical OS/400 PTF levels
- Identical versions of EnterpriseOne at identical service pack levels

- Identical clustering exit programs, object specifier files, and data areas for EnterpriseOne.
- Identical copies of the business data objects listed in the clustering object specifier file
- iSeries clustering software installed and configured. Download it from the Update Center.
- A server jde.ini file with a [CLUSTER] section that defines the cluster name

Note. This requirements list applies only to a virtual three-tier configuration. With additional nodes, you can have different clustering exit programs, object specifier files, data areas, and so on.

Identifying the Cluster Name

In an iSeries clustering environment, EnterpriseOne uses a logical cluster name to make services available to clients and to manage connections between nodes. You must define the name on the DNS and in the iSeries host names table. The cluster name is associated with a floating IP address.

For example, you might set up two iSeries machines that you name DEINS3 and DENIS4. The cluster name you set up in the DNS might be DENISZ. DENISZ is the cluster name associated with a floating IP address that is defined as a TCP /IP interface and listed in the host table of each node in the cluster.

The client references the cluster name when it requests EnterpriseOne services from the enterprise server. Therefore, services are not tied to a single physical machine. The node with the active TCP /IP interface associated with the cluster name provides the services to the client that requests them.

Setting up the Enterprise Servers

You must identify the cluster to each node in the cluster by adding a [CLUSTER] section to the server jde.ini file. If the cluster name is DENISZ, you add this entry to the jde.ini file:

```
[CLUSTER]
PrimaryNode=DENISZ
```

Note. If you referenced the logical cluster name or virtual host name (for example, DENISZ) in the installation plan, you only need to modify the Enterprise Server's jde.ini file as described previously. You don't need to complete any of these setups since all the data sources and ODBC setting are already pointing to the logical host name.

Setting up the Client for Clustering

You must set up the client jde.ini file for it to connect to the cluster name you identified. To set up the client jde.ini file, make changes to these parameters:

```
[DB system settings]
Server=DENISZ
[SECURITY]
SecurityServer=DENISZ
```

where DENISZ is the logical cluster name you specified in the DNS and iSeries host names table.

Setting up the Deployment Server

You must set up the deployment server to connect to each designated enterprise server node in the cluster. You configure the [DB SYSTEM SETTINGS] section and [SECURITY] section of the deployment server jde.ini so that it can connect to the logical data source of the enterprise server.

You should review these Deployment server jde.ini section entries:

```
[DB SYSTEM SETTINGS]
Base Datasource=System - 811
Server=DENIS3
Database=System - 811
Load Library=JDBODBC.DLL
Decimal Shift=Y
Julian Dates=Y
Use Owner=N
Secured=Y
Type=I
DatabaseName2=SY811

[SECURITY]
DataSource=System - 811
```

where System - 811 is the logical data source used to establish initial Object Configuration Manager (OCM) settings, and DENIS3 is the name of an enterprise server node in the cluster.

After you configure the deployment server jde.ini file for a server node in the cluster, you log onto EnterpriseOne on the deployment server as an administrative user and set up logical data sources, database data sources, and OCM for the node. You then configure the deployment server jde.ini for another cluster node. You repeat the sequence of tasks for each node in the cluster.

Setting up logical Data Sources

After you configure the deployment server jde.ini file for clustering, you must set up the logical data sources on each enterprise server node for both system map and server map.

Note. System map configures the client connection to logical data on the enterprise server. Server map configures the server connection to logical data on the enterprise server.

This topic discusses the steps required to complete these tasks:

- Setting up the logical data source for the system map
- Setting up the logical data source for the server map

See Also

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Data Sources,” Understanding Data Sources

Setting up the logical data source for the system map

To set up the logical data source for the system map:

1. On the Systems Administration Tools menu (GH9011), select Logical Data Sources (p986115).
2. Select the logical data source you identified in the Base Datasource= parameter of the [DB SYSTEM SETTINGS] section of the deployment server jde.ini file and click Select.

The default is System – 811

3. In the Work With Data Sources form, click Add.

The Data Source Revisions form appears.

4. Complete these fields to create the logical data source:

| | |
|------------------------------------|---|
| Data Source Name | Enter the name of the cluster. |
| Data Source Use | Type <i>SVR</i> . |
| Platform | Type <i>AS400</i> |
| Logical Server Name | Enter the name of the cluster. Make sure the name matches the logical cluster name you created in the DNS or in the iSeries host names table. |
| Server Map Data Source Name | Type <Server Name> - 811 Server Map |

5. Click OK.

EnterpriseOne launches a form prompting you to create a new ODBC data source. Because you are configuring a logical ODBC data source for the cluster, not a database data source, click Cancel.

Setting up the logical data source for the server map

To set up the logical data source for the server map:

1. On the Systems Administration Tools menu (GH9011), select Logical Data Sources (P986115).
2. Select the logical data source for the enterprise server node server map and click Select.

The default is <Server Name> - 811 Server Map.

3. In the Work With Data Sources form, click Add.
4. Complete these fields to complete the logical data source for the server:

| | |
|------------------------------------|---|
| Data Source Name | Enter the name of the cluster. |
| Data Source Use | Type <i>SVR</i> |
| Platform | Type <i>AS400</i> |
| Logical Server Name | Enter the name of the cluster. Make sure the name matches the logical cluster name you created in the DNS or in the iSeries host names table. |
| Server Map Data Source Name | Type <Server Name> - 811 Server Map |

5. Click OK.

For a full description of the fields in the Data Source Revisions form, see Adding or Modifying a Data Source in the EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation.

Setting up Database Data Sources

After you set up the logical data sources for the cluster, you must set up the database data sources for the enterprise server node. A database data source identifies to EnterpriseOne the database information EnterpriseOne needs to identify and connect to a database, including the type and location of the data.

See Also

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Data Sources,” Understanding Database Data Sources

Setting up the server map database data sources

To set up the server map database data sources:

1. On the Systems Administration Tools menu (GH9011), select Database Data Sources (P986115).
2. In the Machine Search and Select form, select the logical data source for the enterprise server node server map and click Select.
The default is <Server Name> - 811 Server Map.
3. In the Work With Data Sources form, click Find.
4. Select a data source, such as Business Data -PROD, for the enterprise server node and click Select.
5. In the Data Source Revisions form, change the value in the Database Server Name field to the name of the logical cluster.
6. Repeat steps 4 and 5 for each data source.

Setting up Object Configuration Manager for Clustering

After you have set up logical data sources and database data sources, use OCM to map objects in each clustering environment.

Note. Complete OCM configuration requires that you complete the steps discussed in this topic for each clustering environment.

This topic discusses the tasks you complete to set up OCM for clustering:

- Configuring OCM for logical data sources for the server map
- Configuring OCM for logical data sources for the system map
- Configuring OCM for database data sources
- Configuring ODBC connections

Configuring OCM for logical data sources for the server map

To configure OCM for logical data sources for the server map:

1. On the Systems Administration Tools menu (GH9011), select Object Configuration Manager (P986110).

2. In the Machine Search and Select form, select the logical data source for the enterprise server node server map and click Select.

The default is <Server Name> - 811 Server Map.

3. On the Work With Object Mappings form, click Add.
4. Complete these fields and click OK:

Environment Name Enter the name of an environment that will use the cluster.

Object Name Enter DEFAULT

Primary Data Source Enter the name of the logical cluster.

System Role Type **PUBLIC*

Object Type Enter *BSFN* for business functions.

Data Source Mode Type *P*

5. Click OK.
6. On the Work With Object Mappings form, make sure that the business function DEFAULT OCM mapping for the environment is active.
7. Click Close.
8. Repeat this task for each clustering environment.

Configuring OCM for logical data sources for the system map

To configure OCM for logical data sources for the system map:

1. On the Systems Administration Tools menu (GH9011), select Object Configuration Manager (P986110).
2. On the Machine Search and Select form, select the logical data source you identified in the Base Datasource parameter of the [DB SYSTEM SETTINGS] section of the deployment server jde.ini file and click Select.

The default is System - 811.

3. On the Work With Object Mappings form, click Add.
4. Complete these fields and click OK:

Environment Name Enter the name of an environment that will use the cluster.

Object Name Enter DEFAULT

Primary Data Source Enter the name of the logical cluster.

System Role Type **PUBLIC*

Object Type Type *BSFN*

Data Source Mode Type *P*

Note. If you want to run UBEs on the server by DEFAULT, then you must set up another logical data source for UBEs and use the logical cluster name.

5. Determine the business functions that should be run on client workstations, not on the enterprise server.

Note. You can run the batch application Create Server Business Function OCM Records (R986140) to accomplish this task. Be sure to create a new version in proof mode. To do so, select enter a value of O to Processing Option 1 when you submit the report.

6. For each business function that runs locally, launch the Work With Object Mappings form, complete these fields and click OK:

| | |
|----------------------------|---|
| Environment Name | Enter the name of an environment that will use the cluster. |
| Object Name | Enter the name of the business function that you want to run on the client workstation. |
| Primary Data Source | Type <i>LOCAL</i> |
| System Role | Type a system role |
| Object Type | Type <i>BSFN</i> |

7. In the Work With Object Mappings form, make sure that the business function DEFAULT and LOCAL OCM mappings for the environment are active. If they are not, change the status to active
8. Click Close.
9. Repeat this task for each clustering environment.

Configuring OCM for database data sources

To configure OCM for database data sources:

1. On the Systems Administration Tools menu (GH9011), select Database Data Sources (P986115).
2. On the Machine Search and Select form, select the logical data source for the system map and click Select.
3. On the Work With Data Sources form, click Find.
4. Select a data source used by the clustering environment and click Select.
5. On the Server Name field of the Data Source Revisions form, enter the name of the logical cluster and click OK.
6. Repeat steps 3 through 5 for each data source used by the clustering environment.
7. Repeat this task for each clustering environment.

Note. If more than one environment shares a data source, verify that you set up all the environments for clustering. If an environment sharing a data source is not a clustering environment, you might need to set up independent data sources for the environments.

Configuring ODBC connections

To configure ODBC connections:

1. In the ODBC Data Source Administrator form, click the System DSN tab.
2. Select a data source used by the environment that will use clustering and click Configure.
The Client Access Express ODBC Setup (32-bit) form appears.
3. Click the General tab.
4. From the iSeries combo box, select the name of the logical cluster and click OK.

5. Repeat these steps for each data source used in the cluster.

Distributing the ODBC Setup from the Deployment Server

After you have completed the configuration of all enterprise server nodes in the cluster, you must distribute the ODBC configuration you set up on the deployment server to all clients (workstations and servers) that will connect to enterprise server nodes in the cluster. Servers that might connect to enterprise server cluster nodes include Windows Terminal Servers and web servers. The deployment server handles the distribution of the ODBC configuration information.

To accomplish the ODBC setup distribution task, you create a .reg file on the deployment server. The .reg file is an executable that contains the ODBC registry settings you set up on the deployment server. After you create the .reg file, client machines must run it to get the saved ODBC settings and set up their ODBC connections.

Administrators can deploy the .reg file to clients that connect to enterprise servers in the cluster.

To distribute the ODBC setup from the deployment server:

1. From the Windows Start menu, run regedit.exe.
2. In the Registry Editor, browse to HKEY_LOCAL_MACHINE\SOFTWARE\ ODBC\ ODBC.INI.
3. Click Registry and select Export Registry File.
4. Save the .reg file.

Identifying the Cluster Name on the Deployment Server

To complete setup for clustering, you identify the cluster name in the deployment server jde.ini file. Doing so enables the deployment server to look for the cluster when it attempts to connect to an enterprise server node. The active node makes the connection.

Note. Remember to identify the cluster name in the jde.ini files of any servers, such as web servers or Windows Terminal Servers that need to connect to enterprise servers in the cluster.

To identify the cluster name on the deployment server:

1. Open the jde.ini file of the deployment server.
2. Replace any reference to the server (for example, Server=DENIS3) with the cluster name (for example, Server=DENISZ).

CHAPTER 10

Backing Up EnterpriseOne Tables

This section provides an overview of server backup requirements and describes how to backup EnterpriseOne tables on servers.

Understanding Backup Requirements for Servers

A well-planned backup strategy is essential to protect the enterprise information assets. Rigorously following the backup strategy will provide insurance against data lost by acts of nature, hardware or software failure, or human error. The backup strategy must balance the level of protection you need against the physical constraints of the system, such as information storage capacity.

We recommend that the backup strategy include these:

- 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.

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.

Note. You should outline and implement the backup strategy before you begin the Prototype phase of implementation.

Backing Up a Deployment Server

EnterpriseOne on the deployment server includes these 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 PD811, DV811, and PY811 directories. If you modify help files, create a backup of the HELPS directory. If the media objects reside on the deployment server, create a backup of the MEDIA OBJ directory.

If important data, such as system data, resides on the deployment server, create nightly backups of the EnterpriseOne data sources (Oracle or SQL Server). For example, if the central objects or Object Management Workbench resides on the deployment server, create a nightly backup.

Backing Up an 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

These EnterpriseOne iSeries libraries should be backed up:

Note. Shut down the database before you create any backups.

- All EnterpriseOne system libraries.
- JDEOW
- SYS811
- E811SYS
- SVM811
- EnterpriseOne data dictionary library: DD811.
- EnterpriseOne Object Management Workbench library: OL811.
- All EnterpriseOne production libraries (This example is for pristine and production):
- PD811
- PY811
- PRODDTA
- PRSTDTA
- All EnterpriseOne business data libraries:
- PRODDTA
- CRPDTA
- PRSTDTA
- TESTDTA

- All EnterpriseOne control libraries:
 - PRODCCTL
 - CRPCTL
 - TESTCTL
 - PRSTCTL
- All EnterpriseOne versions libraries:
 - PD811DNT
 - PY811DNT
 - PS811DNT
 - DV811DNT
- IFS (Integrated File System) libraries:
 - PD811
 - PY811
 - PS811
 - TS811
 - DV811
- IBM libraries that require backups:
 - QCPA
 - QGPL
- Central objects on the deployment server in Oracle or Microsoft SQL Server database.

UNIX

On an EnterpriseOne UNIX system, backup these database files:

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/E811 directory. For example, /u03/PeopleSoft/E811/*.
- 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

On an EnterpriseOne Windows system, back up these database files:

Note. Shut down the database before you create any backups.

- System files.

PeopleSoft\ddp\E811 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, PSFT811) to a backup database.

Note. We recommend that you use the backup tool provided by the RDBMS vendor.

Understanding EnterpriseOne Tables and Object Owner IDs

These 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 the database.

System Tables

The Object Owner for System tables is sys811.

- F00053
- F000531
- F000532
- F0092
- F00921
- F00924
- F0093
- F0094
- F00941
- F00942
- F00945
- F00946
- F00948
- F00950
- F00960
- F98111
- F986101
- F98611
- F986115

- F986116
- F98613
- F986150
- F986151
- F986152
- F98616
- F986161
- F986162
- F986163
- F986164
- F986165
- F98701
- F98800D
- F98811D
- 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 Management Workbench (OMW) Tables

The Object Owner for OMW tables is obj811.

- F00165
- F9860
- F9861
- F9862
- F9863
- F9865

Data Dictionary Tables

The Object Owner for the Data Dictionary tables is dd811.

- F00165
- F9200
- F9202
- F9203
- F9207
- F9210
- F9211

Server Map Tables

The Object Owner for Server Map tables is svm811.

- F986101
- F98611
- F986110
- F986111
- F986113
- F98DRPCN
- F98DRLOG

Control Tables

The Object Owners for the Control Tables are:

- Control Tables - PROD: prodtcl

- Control Tables - CRP: crpctl
- Control Tables - TEST: testctl
- Control Tables - PS811: prstctl

The Control Tables are listed:

- F0002
- F00021
- F0004
- F0004D
- F0005
- F0005D
- F0082
- F00821
- F00825
- F00826
- F0083
- F0084

Versions Tables

The Object Owners of the Versions tables are:

- Versions - PD811: PD811
- Versions - PY811: PY811
- Versions - DV811: DV811
- Versions - PS811: PS811

The Versions tables are listed:

- F983051
- F98306

Central Objects

The Object Owners of the Central Objects tables are:

- Central Objects - PD811: pd811
- Central Objects - PY811: py811
- Central Objects - DV811: dv811
- Central Objects - PS811: PS811

The Central Objects tables are listed:

- F980011
- F980021

- F983051
- F98306
- F98710
- F98711
- F98712
- F98713
- F98720
- F98740
- F98741
- F98743
- F98745
- F98750
- F98751
- F98752
- F98753
- F98760
- F98761
- F98762
- F98950

Business Data

The Object Owners of the Business Data tables are:

- Business Data - PROD: proddta
- Business Data - CRP: crpdta
- Business Data - TEST: testdta
- Business Data - PS811: prstdta

Backing Up EnterpriseOne Tables on Servers

This section describes how to:

- Create a backup for iSeries.
- Creating a Backup for Oracle on UNIX or Windows.
- Creating a Backup for SQL Server.
- Restoring a Backup File for Oracle on UNIX or Windows.
- Restoring a Backup File for iSeries.
- Restoring a Backup File for SQL Server.

- Restoring a Backup File for SQL Server on Windows.

Prerequisites

Before you complete the tasks in this section:

- If you are using SQL Server or Oracle, verify that you have enough disk space for the 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.

Creating a Backup for iSeries

To create a backup for iSeries:

1. On a tape drive, back up these libraries, depending on which path codes you have installed:

| Library name | Description |
|--------------|-----------------------------|
| SYS811 | System library |
| SVM811 | Server Map |
| OL811 | Object Librarian |
| DD811 | Data Dictionary |
| COPY811 | Central Objects - Prototype |
| COPS811 | Central Objects - PS811 |
| COPD811 | Central Objects - PROD |
| CODV811 | 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 |
| PY811DNT | Versions for CRP |

| Library name | Description |
|--------------|-----------------------------|
| PD811DNT | Versions for PROD |
| PS811DNT | Versions for PRST |
| DV811DNT | Versions for DEV |
| E811SYS | Server system library |
| JDEOW | PeopleSoft Installation |
| PY811 | Server modules - Prototype |
| PY811FA | Package Library - Prototype |
| PS811 | Server modules - PS811 |
| PS811FA | Package Library - PS811 |
| PD811 | Server modules - PROD |
| PD811FA | Package Library - PROD |
| DV811 | Server modules - DEV |
| DV811FA | Package Library - DEV |

2. Back up these IFS structure with the subdirectories:

| Library name | Description |
|--------------|---|
| PSFT811 | Logging directory |
| E811SYS | Kernel spec and XML |
| PY811 | Spec files for Prototype |
| PS811 | Spec files for PSFT |
| PD811 | Spec files for PROD |
| DV811 | Spec files for DEV |
| PeopleSoft | Contains the spec files for each path code. \PeopleSoft\PACKAGES\PY811FA\SPEC*.* \PeopleSoft\PACKAGES\PS811FA\SPEC*.* \PeopleSoft\PACKAGES\PD811FA\SPEC*.* \PeopleSoft\PACKAGES\DV811FA\SPEC*.* |

Creating a Backup for Oracle on UNIX or Windows

To create a backup for Oracle on UNIX or Windows:

1. From the Oracle Enterprise Manager Tool, open Data Manager and select Export from the Data menu.
2. Type the name for the export utility .dmp file.
Click the Browse button to select the directory where the .dmp file will reside.
3. Click Next.
4. On the Object Selection form, select 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 select the item to export. Use the arrows to move objects to and from the Selected Objects form.

5. On the Tuning form, select generate a log file, if needed.
6. Click Next.

Note. Select the Generate Log File option and enter a log file name or use Browse to select a log file.

7. On the Advanced Options form, take the default values or select 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.

When the export process is completed, you will receive these message: "Export terminated successfully without warnings."

10. If errors or warnings exist, check the log file to review the export process.

Creating a Backup for SQL Server

To create a backup for SQL Server:

1. From SQL Enterprise Manager, select Database/Object Transfer from the Tools menu.
2. On the Database/Object Transfer form, select a destination server and database on which to create backup copies of the 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 the objects.

4. Perform either of these tasks 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 the backup tables transferred to the new database with data.

Restoring a Backup File for Oracle on UNIX or Windows

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 select Import.
2. Type the name of the import utility .dmp file.
3. Click Next.
4. On the Object Selection form, select 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, select 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 displays 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.

Note. 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, select 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 generate a log file, if needed.
8. Click Next.

Note. Select the Generate Log File options and enter a log file name or use Browse to select a log file.

9. On the Advanced Options form, select the Increment Type. If you followed the instructions to create a backup, select None for Increment Type and click Next.
10. On the Summary form, verify that all chosen objects and options are correct.

Note. 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 these message: "Process terminated successfully with no warnings."

If errors or warnings exist, check the log file to review the export process.

13. Perform a SELECT statement to verify that the backup tables are populated with data.

Restoring a Backup File for iSeries

To restore a backup file for iSeries:

Restore the libraries and IFS directories that you backed up from tape.

Restoring a Backup File for SQL Server

To restore a backup file for SQL Server:

1. Verify that the Choose 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.
2. From SQL Enterprise Manager, select Database/Object Transfer from the Tools menu.
3. On the Database/Object Transfer form, select a destination server and database from which to transfer backup copies of the tables.

Note. The source server and the destination server can be the same, but the database must be different.

4. Deselect the Transfer All Objects option, but keep all of the other default settings.
5. Click the Choose Objects button, select the objects that you want to transfer, and then click OK to return to the Database/Object Transfer form.
6. Click the Start Transfer button.
The Database/Object Transfer tool moves the objects.
7. Perform either of these 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 the backup tables transferred to the new database with data.

Restoring a Backup File for SQL Server on Windows

To restore a backup file for SQL Server on Windows:

1. 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.
2. Generate scripts for the tables you want to restore and then drop the tables.
3. Use SQL to recreate the scripts for the tables.
4. From the command line, type this command:

```
bcp [[database_name.]owner.] table_name(in|out) datafile /n /u /p /s
```

5. Perform a SELECT statement to verify that data populates the backup tables.

CHAPTER 11

Working with SnapShot

This chapter describes Snapshot.

Using SnapShot

The SnapShot application enables you to install multiple versions of EnterpriseOne on a single workstation. Use SnapShot to rename the installation directory, \E811, and to move executable files, security tables, the jde.ini file, 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 enables you to switch between the versions without requiring you to perform any administrative tasks.

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.

Prerequisite

Before you run SnapShot, copy SnapShot.exe to a directory outside of the \E811 directory. If you try to run SnapShot from inside the \E811 directory with EnterpriseOne open or with an EnterpriseOne directory open, you will receive an error message.

Saving EnterpriseOne in a SnapShot File

To save EnterpriseOne in a SnapShot file:

1. Run SnapShot.exe.

The default directory for SnapShot is \E811\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 the SnapShot file and determine the folder name where the SnapShot file will reside.

3. Click OK to complete the SnapShot.

The 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.

Restoring a SnapShot

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, select a SnapShot to restore and then click Restore.

SnapShot restores the 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.

Deleting a SnapShot

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, select a SnapShot to delete and then click Delete.

SnapShot deletes the directory that stores the 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 these 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 811.

- 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.

If needed, you can also use these switches in the command line:

- /U

This switch disables the user interface, including warnings and errors. Use these 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 these 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 select 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 the 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 these syntax:

```
SnapShot [/R] [/Ttarget dir]
```


CHAPTER 12

Generating Serialized Objects for the EnterpriseOne Web Server

This chapter provides an overview of serialized object generation and describes how to:

- Install eGenerator.
- Generate serialized objects.
- Log in.
- Configure eGenerator.
- Generate all standard serialized objects.

Understanding Serialized Object Generation

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 these 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 the EnterpriseOne objects, you should first test the custom modifications and then generate serialized JAS objects from the upgraded path code.

We strongly recommend that you generate all objects so that JAS is working with the same specs as the 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 these tasks to install eGenerator and generate EnterpriseOne JAS objects from a set of EnterpriseOne objects.

To generate serialized objects for the Java server, the developers must install and use the eGenerator, a PeopleSoft tool that turns EnterpriseOne specifications into Java code, which enables 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.

See Also

Generating Serialized Objects for the EnterpriseOne Web Server in the *Web Server Installation Guide EnterpriseOne PeopleTools 8.11 Installation Guide for Windows NT-Based Systems*

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.

Installing eGenerator

This section describes how to:

- Set the default storage parameter (Oracle only).
- Set the default storage parameter (SQL only)
- Set up the Win32 client to generate serialized objects.
- Copy the genapp.ini file to the default Windows directly.
- Copy the jas.ini and tnsnames files to the generation machine.
- Populate the data dictionary.
- Populate the global specs.
- Modify gen.bat.

Prerequisites

Before you complete the tasks in this section:

- Verify that the paths indicated in the gen.bat file on the generation machine are valid for the configuration. If needed, change them to point to the correct local jar files on the 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 (Oracle Only)

To set the default storage parameter (Oracle only):

Log in to the Oracle database and enter this 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.

Setting the Default Storage Parameter (SQL Only)

To set the default storage parameter (SQL only):

Open the SQL database and set the database Maximum File Size parameter to Unrestricted filegrowth.

Complete this step for each environment (for example, PY811) to which you are installing the Java objects.

Setting Up the Win32 Client to Generate Serialized Objects

After 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 this parameter and value:

```
WebAdmin=1
```

Setting the value for this parameter to 1 enables 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 these customized settings for the Java generation 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.

Copying the genapp.ini File to the Default Windows Directory

To copy the genapp.ini file to the default Windows directory:

1. On the generation machine, navigate to the x:\E811\system\generator directory, and locate the genapp.ini file.
2. Copy the genapp.ini file to the default Windows directory on the generation machine.

The default Windows directory for Windows 2000 is x:\WINNT

Copying the jas.ini and tnsnames Files to the Generation Machine

To copy the jas.ini file and tnsnames file to the generation machine:

1. On the JAS server, navigate to this 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 these files to the x:\E811\system\Generator\WEB-INF directory on the generation machine:

- jas.ini
- tnsnames.ora (Oracle only)

Note. These directories are the default locations of the files, but the actual directory paths might be different.

(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 this setting: tns setting in the [JDBj-Oracle] section (Oracle only).

Populating the Data Dictionary

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. Select 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], select Override Location and click OK.
7. On Batch Applications - [Version Prompting], click Submit.
8. On Batch Applications - [PSFT Data Sources], select Local and click Submit.

The application starts a UBE that populates databases stored in the x:\E811 directory. This process runs for an hour or more.

9. When the process is complete, copy these files from the x:\E811 directory to the \spec directory under the path code of the environment you will use to generate the serialized Java objects (for example, x:\E811\PY811\spec):
 - dddict.xdb
 - ddtext.xdb
 - dddict.ddb
 - ddtext.ddb

Populating the Global Specs

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. Select 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], select Override Location and click OK.
6. On Batch Applications - [Version Prompting], click Submit.

7. On Batch Applications - [JDE Data Sources], select 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 the 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:\E811\system\Generator directory on the generation machine.
2. Right click gen.bat and select Edit.
3. In the gen.bat configuration files, change these 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:\E811\system\Generator\WEB-INF. |
| set DBDRIVERS= | <p>Paths and file names of the database JDBC drivers. (Required for Direct generation only.) Enter the path to the JDBC driver for the particular database:</p> <p>For Oracle: classes12.zip</p> <p>For SQL 2000: msbase.jar, mssqlserver.jar, msutil.jar</p> <p>For DB2/UDB: db2java.zip</p> <p>For DB2 (iSeries): jt400.jar</p> |

| Parameter | Setting |
|------------|---|
| Safe Mode | <p>Adding the safe mode switch in gen.bat makes parallel generation more stable and less prone to hanging when using IBM JDK 1.3.1. When the safe mode is activated, the Generator uses 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.</p> <p>To activate Safe Mode, add these bolded text to the command line in the gen.bat file:</p> <pre>%JDK%\bin\java %GENERATOR_OPT% -Dmode=safe -Ddefault_path=%JAS_ INI% -classpath %GEN_CLASSPATH% com.jdedwards.runtime.generator.Generator %1</pre> |
| JASLOGGER= | <p>WEB-INF\lib\log4j.jar</p> <p>Delete these comment and replace it with the path to the folder containing the jdelog.properties file:</p> <p><folder containing jdelog.properties>.</p> <p>For example, if the default location for the jdelog.properties file is the WEB-INF directory, then the JASLOGGER setting is:</p> <p>WEB-INF\lib\log4j.jar;WEB-INF</p> |

Note. If the 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 the Win32 client with a JAS build, can generate objects of compiled Java bytecode, and you have access to all EnterpriseOne objects when you run the Java server.

Generating Serialized Objects

To generate serialized objects:

1. From the generation machine, log off EnterpriseOne.
2. Under the /JAS directory on the 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 the generation machine before you run gen.bat.

3. Log on to a Java environment in EnterpriseOne (for example, JPY811).

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.

4. On Java & HTML Generator, complete the Web Server Name field by entering the name of the JAS Server.

Note. If you are running multiple Web server ports on the generation machine, type machine name:port, where machine name is the name of the generation machine and port is the port number.

5. Click Advanced Settings.
6. On Log Files, type the name and path to the Status Log and Error Log files and click OK.
7. Select the "Generate all objects" option.
8. Select these options and click OK:
 - Business Views
 - Tables

The system generates all EnterpriseOne business views and tables.

9. Verify that the DOS window indicates the generation process has completed successfully.
10. Select the "Generate all objects" option.
11. Select all options except Business Views and Tables, and click OK.
- The generation process takes five to ten hours, depending on the speed of the generation machine.
12. 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

This section provides an overview of logging in to eGenerator and describes how to:

- Generate to the Web server.
- Bypass the Web server and generating directly to a database.

Prerequisites

Before you complete the tasks in this section:

- Ensure that these files and folders reside on the workstation:
 - gen.bat
 - OWResource.jar

- xalan.jar
- xerces.jar
- images.jar
- classes folder
- For direct generation to a database, you also need these files:
 - jas.ini
 - jdbj.ini
 - Database drivers
 - tnsnames.ora (Oracle only)
- Verify that these sections in the jas.ini file are configured correctly. If necessary, change them to match the 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.
- Verify that these sections in the JDBj.ini file are configured correctly.
 - [JDBj-SPEC DATA SOURCE]

Change the settings to match the configuration.
 - [JDBj-ORACLE]

Even if you are not using Oracle as the database, a path to this database must be present in these setting:

```
tns=c:\E811\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 this path.

Understanding Logging In

Using eGenerator, you can generate EnterpriseOne objects in one of these ways:

- Generate to the Web server
- Bypass the Web server and generate directly to a database

When you generate directly to the database, 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. The advantages to this generation method are that you can generate objects before setting up the Java server, it is faster than generating to the Web server, and it is easier to maintain than generating to the Web server. The disadvantage is that generating directly to the database is more difficult to set up initially.

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.

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 prevents eGenerator from launching. If a non-fatal error is detected, the application displays an error message but enables 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 these 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 these 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 these key parameters in the file: <ul style="list-style-type: none"> • 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 |

Generating 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:

1. Access the generation machine.
2. Under the x:\E811\system\Generator directory on the 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.

Ensure that the genapp.ini file is located in the WINNT directory on the generation machine.

A sample genapp.ini file is included in the section called Configuring the genapp.ini File.

3. On EnterpriseOne Sign On, log on to a Java environment in EnterpriseOne (for example, JPY811).

Note. 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.

4. On eGenerator, complete the JAS Server Name field.

Type the name of the JAS server that you want to log in to.

The system gets the database path from the jas.ini file. eGenerator generates all objects to this web server.

Note. If you are running multiple web server ports on the generation machine, type machine name:port, where machine name is the name of the generation machine and port is the port number.

5. 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\DV811\SPEC.

Note. You can use this field to point to the latest specs on the deployment server and generate objects from the most recent, undeployed specs.

6. Click Connect.

The system connects to the web server that you indicated. A message appears on the status bar when login is complete.

Bypassing the Web Server and Generate Directly to a Database

Select direct generation by selecting the Direct Generation check box at the bottom of the Connect to Web Server screen of the eGenerator.

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.

2. On EnterpriseOne Sign On, log in to a Java environment (for example, JPY811).
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 each particular environment. The eGenerator has these modes for generating Web objects:

- Sequential Generation
- Parallel Generation

In sequential generation mode, web objects are generated one after the other when you select the Generate-All Objects option. It is the most stable mode of generation, requires the least system resources (for example, memory), and is safe to use when JITI (Just in Time Install) is occurring. However, it is also the slowest mode of generation.

Parallel generation 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. It is much faster than Sequential Generation, especially on multi-processor workstations. On the other hand, parallel generation 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.

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 the generation requirements:

Logging

Specify whether to log the generation process by selecting the check boxes 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. |
| Report Licensing | Select the check box to generate Report Licensing after every report. |
| User Option | Choose the appropriate option to generate a public version of the applications or a personal version. Personal versions are available only to the person who created the versions. |
| Application Entry Point | Select this option to generate Application Entry Point information after every application. |

Choosing the HTMLOverrides File

Select the HTML Overrides file to generate HTML objects with advanced features, such as multi-line edit.

To select the HTMLOverrides File:

1. Launch Gen.bat.
2. On EnterpriseOne Sign On, log in to EnterpriseOne.
3. On eGenerator, complete the HTMLOverrides File field:

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 select languages:

1. On eGenerator, select Options->Languages from the pull-down menu.
2. Select the check box 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.

These 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 enable you access to all EnterpriseOne objects when you run the 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.

This section describes how to:

- Run the TAM analyzer.
- Generate all EnterpriseOne objects.

Running the TAM Analyzer

The TAM Analyzer enables eGenerator to generate all objects without saving them to the database, which enables 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. Select 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

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.

5. 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.

6. Select one of these options:

- To resume the previous generation process, click RESUME OLD SESSION.

We recommend you select 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.

CHAPTER 13

Understanding Executable Files on the Workstation

This chapter describes linked and standalone executable files.

EnterpriseOne Linked Executable Files

These 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.

The executables in this 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. |

| Executable | Description | Call Details |
|---------------|--|--|
| 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 enable green screens to be displayed through EnterpriseOne. | Called by EnterpriseOne kernel. |
| 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 | We ship 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 email 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). |

| Executable | Description | Call Details |
|-------------------|---|--|
| 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 Autopilot or EventCapture. | Called by 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. |

| Executable | Description | Call Details |
|-----------------|---|---------------------------|
| 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. |
| 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 these 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. |

| Executable | Description | Run Instructions |
|---------------|---|--|
| LogViewer.exe | <p>This program employs a user friendly interface to view and modify plain ASCII EnterpriseOne files such as:</p> <ul style="list-style-type: none"> • jdedebug.log • jde.log • olt.log • jde.ini | Double-click the executable. |
| MOConv.exe | <p>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.</p> | Exit EnterpriseOne and then double-click the executable. |
| Nettest.exe | <p>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.</p> | Double-click. Required argument: Enterprise Server Name |
| Regdlls.exe | <p>This program adds these dynamic link libraries (DLLs) to the registry:</p> <p>jdetapitest.dll</p> <p>jdetapicomtek.dll</p> <p>Register these DLLs before using RunTAPI.exe.</p> | Run from a DOS window in the EnterpriseOne system/bin32 directory. |
| RunTAPI.exe | <p>This program controls interoperability between EnterpriseOne and telephone switching systems. It is a snap-in (harness) to ComTech CTI Server objects. It requires jde.ini file changes and Regdlls.exe before it can be run. Contact Customer Support for full documentation.</p> | Double-click. |
| SABridge.exe | <p>This Object Export Facility displays the names of the EnterpriseOne objects along with their descriptions and corresponding product codes.</p> | Double-click the executable. |

| Executable | Description | Run Instructions |
|---------------|---|---|
| SnapShot.exe | <p>This program manages multiple workstation installations on the same PC.</p> <p>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.</p> | 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. | <p>Run from a DOS window. Required arguments:</p> <ul style="list-style-type: none"> • 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 select Properties. | Open a DOS window, change the directory to the desired target, and double-click the executable. |
| 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 | <p>Creates CORBA wrappers around EnterpriseOne business functions. This is a command line utility that requires a script file as an input.</p> <p>GenCORBA generates COBRA interfaces for EnterpriseOne business functions.</p> | <p>Syntax:</p> <p>GenCORBA[options] [libraries]</p> <p>For example:</p> <p>GenCORBA /Cat /UserID Devuser1 /Password Denuser1 /Environment ADEVHPO2 CAEC</p> |
| GenJava.exe | GenJava provides access to EnterpriseOne business functions by generating pure Java interfaces to them. | <p>Run GenJava.</p> <p>Syntax:</p> <p>GenJava [options] [libraries]</p> <p>For example:</p> <p>GenJava /Cat /UserID Devuser1 /Password Denuser1 /Environment ADEVHPO2 CAEC</p> |
| 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 using the command line. |

| Executable | Description | Run Instructions |
|-------------------|---|--|
| 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. |
| 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. |

| Executable | Description | Run Instructions |
|---------------|--|---|
| vapplayer.exe | Virtual Autopilot Player enables 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 debugging information generated by an EnterpriseOne application run under Autopilot or EventCapture. | Create a desktop icon or Start Menu button. No command line arguments are used. |
| UTBrowse.exe | <p>UTB is a tool used for viewing the records in tables. We also use it to view local EnterpriseOne object specifications.</p> <p>UTBrowse.exe uses these two libraries in the bin32 directory:</p> <ul style="list-style-type: none"> • datautils.dll • envtool.dll. | Type UTB in the EnterpriseOne fast path or click the executable. |
| tda.exe | Use TDA is to modify EnterpriseOne tables. | <p>On the Command Line, type tda.exe -idtablename, where tablename is the name of the table you want to modify. For example, F0101.</p> <p>You do not need to run EnterpriseOne before running tda.</p> |
| 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. |

| Executable | Description | Run Instructions |
|-----------------------------------|---|---|
| tamtool.exe | <p>tamtool can perform these functions:</p> <ul style="list-style-type: none"> • 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 | <p>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.</p> <p>tampack.exe is included with the workstation and the deployment server. tampack.exe creates a translated copy of TAM files (RDASPEC.DDB, GBRSPEC.DDB, etc.) on the PC.</p> <p>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.</p> | You must run tampack.exe from a DOS shell and pass in parameters. |
| 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 for 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. |

CHAPTER 14

Troubleshooting the Workstation

This chapter provides an overview over workstation troubleshooting and error messages and describes how to:

- Troubleshoot the workstation.
- Troubleshoot the production workstation
- Troubleshoot the development workstation.
- Work with the workstation log files.

Understanding Error Messages

This section describes the nature of some of EnterpriseOne error messages to help you correct the errors or determine the steps that caused the error.

Error Message Details

When you encounter an error, right-click the error in the error message window and select 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 this example, the source file is c:\E811\MSTR811\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 the error includes the name of the 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 disable the constant, you avoid the error.

Note. If you see a source file description that begins with c:\E811\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 for conditional statements that determine whether to activate 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.

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.

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.

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 these techniques to troubleshoot these errors:

- Look at the `jddebug.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.

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.

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 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 these 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 displays no data. It displays only the report headers and the text No Data Selected.
- Batch process displays 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

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Object Configuration Manager,” Running the Object Configuration Copy Report

Troubleshooting the Production Workstation

This section provides an overview of troubleshooting on the production workstation and describes how to:

- Perform preliminary troubleshooting.
- Troubleshooting interactive application problems.
- Troubleshoot batch processes resulting in no data.
- Troubleshoot batch processes displays errors on the report.
- Troubleshoot batch processes displays unexpected data on the report
- Troubleshoot batch processes ending in an error when submitted on the server.
- Troubleshoot local data availability problems.
- Troubleshoot enterprise server data availability problems.
- Troubleshoot printing problems.
- Troubleshoot .DLL problems on a production workstation.
- Troubleshoot data source setup problems.

Understanding Production Workstation Troubleshooting

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.

As a system administrator, you can perform preliminary troubleshooting on the production workstation to verify the nature of the problem. You will also want to isolate problems to a user's particular workstation and environment.

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 select 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.

A user might encounter several problems when attempting to run a batch process. For example, the output might display 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.

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.

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

EnterpriseOne Tools 8.94 PeopleBook: Package Management, “Understanding the Package Build Process”

EnterpriseOne Tools 8.94 PeopleBook: Workflow Tools

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Using the Universal Table Browser”

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Object Configuration Manager,” Running the Object Configuration Copy Report

Performing Preliminary Troubleshooting

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

To troubleshoot interactive application problems:

1. Select one of these 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 select Full Description.

The system displays a full glossary of the error and includes information for resolving the issue.

Troubleshooting Batch Processes Resulting in No Data

This task provides a solution to previously discussed 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, activate the applicable logs and continue with these steps.
4. Run the batch process and locate the batch process log for the report you ran.

EnterpriseOne names this log with these convention:

```
report name_version_other identifiers.log
```

This log is located in the local directory under \E811\PrintQueue. If you ran report R04431, it would appear in the local directory, E811\PrintQueue, as follows:

```
R04431_XJDE0001_D960823_T104512.log
```

5. 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 the main source for debugging batch processes. However, you can look in the jde.log and jdedebug.log for errors or failures of any kind.
6. 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.

Troubleshooting Batch Processes Displaying Errors on the Report

Access Error Messages from the Data Dictionary Design menu (GH951).

Error Messages form

To troubleshoot batch processes displaying errors on the report:

1. On Error Messages, complete the Glossary Group field.
2. Complete these optional fields:
 - Language
 - Alpha Description
3. On the grid, enter values in the Data Item field and click Find to narrow the search to the 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, select Glossary from the Row menu.

Troubleshooting Batch Processes Displaying Unexpected Data on the Report

To troubleshoot batch processes displaying 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. Activate 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.

Troubleshooting Batch Processes Ending in an Error when Submitted on the Server

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.

Troubleshooting Local Data Availability Problems

Inspect the local database at `\E811\pathcode\data\JDELocal_PD811.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 these conventions: report_name_version_other identifiers.log. This log is located in the local directory, `\E811\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 the SQL statement, which means that data selection is not the cause of the problem.

Troubleshooting a 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 the system name for a standalone installation.

The solution is to disable the DNS name in Windows.

Troubleshooting Enterprise Server Data Availability Problems

If the 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 the 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 the 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 the report.

For example, if you leave the 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 the 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 display 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.

These printing problems can occur:

- The batch application produces error messages on the report, for example, Invalid Company Number.
- The report batch process displays 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.

- The 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 these:

```
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 \E811\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

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

This section provides an overview of troubleshooting the development workstation and describes how to:

- Troubleshoot .DLL problems on a development workstation.
- Troubleshoot event rules problems.
- Troubleshoot business function problems.

Understanding Development Workstation Troubleshooting

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.

Problems with workstation .DLL files are indicated if you receive an error message such as these:

```
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 E811\path code\bin32 and E811\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 E811\system\bin32\busbuild.exe.

Troubleshooting .DLL Problems on a Development Workstation

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, select 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, select 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 enables you to debug the event rules for an application or batch process.

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++ enables 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

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Object Configuration Manager,” Troubleshooting Business Function Processing Problems

Working with the Workstation Log Files

This section provides an overview of the various workstation logs and tables and describes how to:

- Use log files to troubleshoot strategies.
- View log files.
- Set up the workstation jde.log.
- Set up the workstation jdedebug.log.
- Set up the batch process log.
- Troubleshoot with the compile error log.
- Troubleshoot with jdecpy.log.
- Troubleshoot with the sql.log.
- Activate sql.log.
- Troubleshoot ODBC problems using sql.log.
- Troubleshoot with jdeinst.log.

Understanding 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.

If you do not use data replication in the configuration, ignore error messages that refer to these tables in the JDE.LOG and JDEDEBUG.LOG:

- F98DRPUB
- F98DRENV
- F98DRSUB
- F98DRLOG
- F98DRPCN

Understanding 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 the 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 the current development. You cannot edit the contents of the disk cache.

We recommend 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 E811/pathcode/spec directory.

In general, on EnterpriseOne workstations, logs are classified in these categories:

- Logic processing.
- Batch processing.
- Application development (compiling and generating).
- Object Management Workbench transactions.

Comprehending Logic Processing Logs

You use two major log file sources for troubleshooting processing faults on the workstation:

- jde.log

This log displays fatal errors. Jde.log can track any fault that might occur within EnterpriseOne.

- jdedebug.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.

Understanding Application Development Logs

For EnterpriseOne workstations in application development environments, you can use these 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 `jdeinst.log` to identify faults in EnterpriseOne silent installation:

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.

Understanding the Workstation `jddebug.log`

The workstation `jddebug.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 `jddebug.log` to track errors. Instead, it uses this log to track the timing of EnterpriseOne processes.

You can use `jddebug.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 the process failed and you have logging turned on, look in the `jddebug.log` for these 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.

Understanding 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 \E811\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 displays 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.

The batch process log file displays the process flow in batch processes. This flow is completed in these 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.

Understanding the sql.log

In sql.log, the important lines for you to search are:

- **SELECT * FROM** (bolded in these 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.

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 the 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

EnterpriseOne Tools 8.94 PeopleBook: Configurable Network Computing Implementation, “Object Configuration Manager,” Troubleshooting Business Function Processing Problems

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, select the View System Log option.
3. On Log Viewer, select Open from the File menu to locate and open an EnterpriseOne log file.

You can also use the View menu selection to select log files. If you have viewed log files previously, the File menu keeps a history of those files.

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 the 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 the settings for the job file variable:

```
[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.

Setting Up the Workstation jdedebug.log

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 the 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 these 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 these parameter in the [DEBUG] section:

```
[DEBUG]
Output=output parameter
```

Where the 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 `jddebug.log` file. The debug level is determined by these parameter in the [DEBUG] section:

```
[DEBUG]
Output=debug level
```

The level parameter controls the debug level. You can specify any combination of allowable values using commas as delimiters. The default setting is `LEVEL=BSFN,EVENTS`. 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. Values are listed:

- **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:

6. 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.
7. 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.
8. 3: Traces statements as the software enters and exits each event. Specifying this value is the equivalent of specifying the `EVENTS` value.
9. 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.
10. 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.
11. Save the changes and close the `jde.ini` file.

Setting Up the Batch Process Log

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 the 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 these 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 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. Values are:

The <batch_report>.log file is not saved.

The <batch_report>.log file is saved in the workstation EnterpriseOne print queue directory (E811\PrintQueue).

4. Save the changes and close the jde.ini file.

Troubleshooting 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 E811 directory tree; for example, c:\E811\PD811\LOG.

Troubleshooting 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.

Troubleshooting 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, activate 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 these processes are called:

- Process SQL Statements
- Receive Results

Activating sql.log

To activate sql.log:

1. From the Windows Control Panel, select 32bitODBC.
2. On Data Sources, click Options.

Note. Leave the Stop Tracing Automatically option turned on. Because this log grows rapidly, we recommend 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.

Troubleshooting ODBC Problems Using sql.log

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 these 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 the 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.

Troubleshooting with the jdeinst.log

Use jdeinst.log to view the status of the EnterpriseOne silent installation. The silent installation mode enables 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.

CHAPTER 15

Troubleshooting the Enterprise Server

This chapter provides an overview and describes how to:

- Troubleshoot the iSeries enterprise server.
- Troubleshoot the UNIX/Linux enterprise server.
- Troubleshoot the Windows enterprise server.
- Troubleshoot Web servers.

Understanding Enterprise Server Troubleshooting

Using these techniques, 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.

You might encounter these 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 these two major log file sources for troubleshooting processing faults on the Enterprise server:

- jde.log

This log displays fatal errors. It can track any fault that might occur within EnterpriseOne.

- jdedebug.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 jdedebug.log to track errors; instead, this log is used to track the timing of EnterpriseOne processes.

Batch Processing Logs

You can use the batch process log to identify faults in EnterpriseOne processing related to 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

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Working with Servers”

Chapter 13, “Understanding Executable Files on the Workstation,” page 269

Chapter 5, “Server Administration Workbench,” Monitoring EnterpriseOne with SAW on Windows, page 102

Understanding 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 jdedebug.log. This example has combinations for the jde.ini parameter setting for enabling or disabling server logs.

Enable jde.log

Below is an example of the jde.log file with debug logging enabled.

```
[DEBUG]
Output=NONE
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Enable jde.log and jdedebug.log

Below is an example of the jde.log file with debug logging enabled and output to a file.

```
[DEBUG]
Output=FILE
LogErrors=1
```

```
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Disable jde.log

Below is an example of the jde.log file with debug logging disabled.

```
[DEBUG]
Output=NONE
JobFile=blank/invalid location/name (1)
DebugFile= blank/invalid location/name (2)
```

Files and members generated by the jde.log will be located in JobFile. EnterpriseOne uses these 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 jdedebug.log will be located in DebugFile. EnterpriseOne uses these naming conventions:

```
jdedebug_process_ID.log
```

Where jdedebug is the file name prefix, process_ID is a uniquely named process ID, and log is the file suffix or extension.

Note. 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 jdedebug will be located in DebugFile. EnterpriseOne uses these naming conventions:

```
jdedebug_process_ID
```

Where jdedebug 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 the 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 this default setting:

```
[DEBUG]
JobFile=jde.log
```

Files generated by the `jde.log` are located in JobFile. EnterpriseOne uses this 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 these commands and if logging is enabled, the system places the log files in the `/u13/PeopleSoft/E811/system/bin32` directory:

```
cd /u13/PeopleSoft/E811/system/bin32
RunOneWorld.sh
```

If you start EnterpriseOne with these 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/E811/system/bin32/RunOneWorld.sh
```

If you set up the 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`.

Example: Enterprise Server `jde.log`

This example of the `jde.log` from the Enterprise server displays 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.

See Also

[Chapter 2, “Understanding Server Administration for `iSeries`,” page 5](#)

Understanding the Enterprise Server `jdedebug.log` File

You can use the Enterprise server `jdedebug.log` 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. The log contains API calls and SQL statements as well as other messages.

You can use `jdedebug.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 the process failed and you have logging turned on, look in the `jddebug.log` for these 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`. This example has valid setting combinations for enabling or disabling server `jddebug.log`.

Enable `jddebug.log`

```
[DEBUG]
Output=FILE
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Enable `jde.log` and `jddebug.log`

```
[DEBUG]
Output=BOTH
LogErrors=1
JobFile=valid location/name (1)
DebugFile=valid location/name (2)
```

Disable `jddebug.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 these 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 these 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 jdedebug.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

We recommend that you create a separate directory on the Enterprise server for logs. You should set up the jde.ini file to explicitly direct log files to that directory. For jdedebug.log, these setting controls the location:

```

[DEBUG]
DebugFile=jdedebug.log

```

For Enterprise servers, the files generated by the jdedebug.log will be located in DebugFile. EnterpriseOne uses these naming conventions:

```
jdedebug_process_ID.log (jdedebug_JobNumber.log for iSeries)
```

Where jdedebug 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 these commands:

```
cd /u13/PeopleSoft/E811/system/bin32 RunOneWorld.sh
```

and assuming that logging is enabled, the system places the log files in the /u13/PeopleSoft/E811/system/bin32 directory. Similarly, on a UNIX machine, if you start EnterpriseOne with these commands:

```
cd /usr/PeopleSoft /u13/PeopleSoft/E811/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 the 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 jdedebug.log on the Enterprise Server

EnterpriseOne processes create logs as jdedebug_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 jdedebug_123456.log.

Understanding 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 (E811\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 displays 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 displays the process flow in batch processes. This 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.MCADDS, 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, these 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. These are values for UBEDebugLevel:

| Setting | Description |
|---------|--|
| 0 | No error messages |
| 3 | Object-level messages |
| 4 | Event rule messages and SQL statements (plus levels 1-3) |

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
```

Setting Up the Enterprise Server jde.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, E811SYS).
 - 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 this 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/E811. In this case, the directory with the JDE.INI file (the file name is case-sensitive) is JDE_BASE=/u04/PeopleSoft/E811/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 the 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 these 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. 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 the changes and close the file.

Setting Up the Enterprise Server jddebug.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 the settings for the debug file variable:

```
[DEBUG]
DebugFile=debug file
```

Where debug file specifies the name of the jddebug.log file (JDEDEBUG member for iSeries). For non-iSeries Enterprise servers, the default value is jddebug.log. For iSeries Enterprise servers, the default value is JDEDEBUG.

4. Enable or disable the logging of events to the jddebug.log file (JDEDEBUG member for iSeries) by setting these 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 jddebug.log) is enabled. 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).

And where Output=output parameter has these values:

NONE: No trace information is written to jddebug.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.

EXCFIL: 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 jddebug.log (JDEDEBUG member for iSeries).

5. Save the changes and close the jde.ini file.

Setting Up the <batch_process>.log File

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 the 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 these 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. 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. 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 (E811\PrintQueue).

4. Save the changes and close the jde.ini file.

Troubleshooting 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 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 these 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 PD811 or DV811.

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 these 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 documentation for Debugging C++.

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 the 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 the workstation, exit EnterpriseOne and turn logging on in the `jde.ini`.
3. Run the application on the server again, and then check the `jde.log` file to see if any of these 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 JDENET is down.
4. 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 workstation EnterpriseOne software, although you can run a standalone version on Windows server platforms.
5. Check the `services` file on the workstation (located in the 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.
6. If you receive a Communication Failure message, try resubmitting the application. A time-out may have occurred.

```
[JDENET]
netTrace=0/1 (disabled/enabled)
```

7. Look in the log file for these 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

Chapter 5, "Server Administration Workbench," Monitoring Enterprise Server Statistics, page 110

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 the environments.

See the section specific to the platform type for more information about PORTTEST.

Troubleshooting the iSeries Enterprise Server

This section provides an overview of troubleshooting the iSeries enterprise server and describes how to:

- Troubleshoot iSeries enterprise server installation.
- Troubleshoot multiple release setup.
- Troubleshoot JDBNET.
- Troubleshoot interprocess communications.
- Troubleshoot jde.ini file.

Understanding iSeries Enterprise Server Troubleshooting

This subsection 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 the 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 these in the JDE.LOG files:

```
JDB3100011 - Failed to get location of table F983051 for environment PD811
Look for messages similar to these 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.

Note. To complete the resolutions provided for this issues, you must sign on to the Enterprise server using an account that has administrative privileges.

See Also

Chapter 5, “Server Administration Workbench,” Monitoring Enterprise Server Statistics, page 110

Troubleshooting iSeries Enterprise Server Installation

This subsection 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. |

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 these message in the JDE.LOG file: Databases: iSeries:table configuration problems iSeries JDB3300011 - Failed to get location of table F983051 for environment PD811 | <ul style="list-style-type: none"> • 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, /E811sys/ini/JDE.INI. |
| You need more information on using the iSeries .INI file | Review the notes and descriptions of .INI settings. |

Troubleshooting: 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 /PSFT811/LogFiles, enter this command:

```
MKDIR DIR('/PSFT811') DTAAUT(*RWX) OBJAUT(*ALL)
```

Followed by:

```
MKDIR DIR('/PSFT811/LogFiles') DTAAUT(*RWX) OBJAUT(*ALL)
```

Logging might be turned off in the .INI. Activate logging in the .INI using these settings in the [DEBUG] section:

```
[DEBUG]
```

```
LogErrors=1
```

```
Output=FILE
```

Where variable names and descriptions are as follows:

LogErrors values are:

0 = Do not generate logs.

1 = Create logs.

Output values (always in upper case) are:

NONE = Do not write debug messages to any output device.

FILE = Write messages to log files.

Troubleshooting: Not enough relevant information is written to the log files

Additional logging information may need to be turned on in the .INI. Set these keys in the .INI for additional information to be output to the log files:

```
[JDENET]
```

```
netTrace=1
```

```
[JDEIPC]
```

```
ipcTrace=1
```

```
[DEBUG]
```

```
TAMTraceLevel=1
```

```
[UBE]
```

```
UBEDebugLevel=6
```

```
[TCEngine]
```

```
TraceLevel=10
```

Where variable names and descriptions are as follows:

netTrace values are:

0 = Do not generate JDENet error messages (that is, communication between platforms).

1 = Generate JDENet error messages.

ipcTrace 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 values are:

0 = Do not generate Table Access Management (TAM) error messages (that is, regarding specification files).

1 = Generate TAM error messages.

UBEDebugLevel 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 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, activate 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 the EnterpriseOne iSeries configuration and setup.

| Issue | Resolution |
|--|---|
| An error with OCM occurred. | Verify that the OCM is correct for the environment. Disable the security server 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. | <p>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</p> <p>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.</p> |
| An error with the security server occurred | <p>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.</p> <p>If the library list is correct, the problem could be because the activation group under which the job is running on the iSeries may retain some of the information from previous attempts. Log off, log on, and run PORTTEST again.</p> |
| An error with the security server occurred | <p>The supplied user name or password might not match any names or passwords in the EnterpriseOne security table. Try one of these:</p> <ul style="list-style-type: none"> • Run PORTTEST with a valid user name and password. • Add the given user name and password to the EnterpriseOne security table. |
| You get these message on the screen: Invalid parms PORTTEST <USER> <PWD> <ENV> | <p>You might not have included the correct number of arguments to PORTTEST. Use these arguments:</p> <p>User - A valid EnterpriseOne user ID.</p> <p>Password - Password for the EnterpriseOne user ID.</p> <p>Environment - A valid EnterpriseOne environment.</p> |

| Issue | Resolution |
|---|---|
| Fewer than 99 F0902 records are written to the screen by PORTTEST | <p>A possible PORTTEST failure. Examine the log files.</p> <ul style="list-style-type: none"> 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 occurs in the log file | The environment might not be set up correctly. Any errors in the affected .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 JDENET

| Issue | Resolution |
|---|---|
| NETWORK dies immediately. | <p>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.</p> <p>The startIPCKeyValue in the .INI file could be used by another version of EnterpriseOne. Try one of these:</p> <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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. Select 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 select 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. | <ul style="list-style-type: none"> 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 |
|--|---|
| <p>You get these message: Communication Failure with <server name></p> | <p>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.</p> <p>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.</p> <p>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 jde.ini must match the serviceNameListen in the jde.ini on the server.</p> <p>Other communications problems might exist. Run SERVERADMINISTRATIONWORKBENCH.exe (found in the system\bin32 directory on the workstation). This program displays only the machines on the specified port (also known as service) that are running EnterpriseOne (either client or server). Use this information to track down the problem:</p> <ul style="list-style-type: none"> • 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. |
| <p>The report does not display any data.</p> | <p>No data might exist in the database for the report that you are running, or you do not have access to the data. Try these:</p> <ul style="list-style-type: none"> • 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 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. <p>If no data is found, it could be because:</p> <ul style="list-style-type: none"> • 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. |

| Issue | Resolution |
|---|--|
| The report does not display 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. | <p>The environment might not be set up correctly. Try these:</p> <ul style="list-style-type: none"> • 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. |
| <p>You get these message:</p> <p>Communication Failure with <server name></p> <p>This error occurs sometimes on the workstation</p> <p>Restarting JDENET_N sometimes gets rid of the error</p> <p>SAW (running on the workstation) cannot always see the server</p> <p>You can ping the server from the workstation</p> | <p>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.</p> <p>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. JDENET then uses the IP address associated with the given network card.</p> |

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: Email and PPAT

| Issue | Resolution |
|--|--|
| <p>The batch application, server package installation, or table conversion log file (in the PrintQueue directory) displays the message:</p> <p>PPAT:troubleshooting iSeries</p> <p>XE Email:troubleshooting: iSeries</p> <p>DoSendMessage Error: User 5600427 does not exist in the address book file (F0101).</p> | <p>The particular user may not be found in the Address Book table (F0101). Add the user to the Address Book table (F0101).</p> |

See Also

[Chapter 2, “Understanding Server Administration for iSeries,” Understanding the iSeries Library Structure for EnterpriseOne, page 8](#)

[Chapter 15, “Troubleshooting the Enterprise Server,” Troubleshooting the JDE.INI File, page 322](#)

Troubleshooting Multiple Release Setup

This subsection explains how to troubleshoot problems that can occur with multiple releases on the iSeries.

| Issue | Resolution |
|---|---|
| <p>When you try to run multiple releases of EnterpriseOne at the same time, conflicts seem to occur between each release.</p> | <p>Each installed release of EnterpriseOne may not have its own unique set of keys in the .INI. Change these keys in one or both .INI files:</p> <p>[JDEIPC]</p> <p>startIPCKeyValue</p> <p>[JDENET]</p> <p>serviceNameListen</p> <p>serviceNameConnect</p> <p>Variable names and descriptions:</p> <p>startIPCKeyValue</p> <p>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.</p> <p>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.</p> <p>serviceNameListen</p> <p>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.</p> <p>serviceNameConnect</p> <p>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.</p> <p>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.</p> |

Troubleshooting JDBNET

Enterprise servers: iSeries:JDBNET

This subsection explains how to troubleshoot problems that can occur with JDBNET.

| Issue | Resolution |
|---|---|
| You do not know how JDBNET is used | <p>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.</p> <p>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.</p> |
| 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 | <p>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 this information:</p> <ul style="list-style-type: none"> • 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 (using 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 subsection explains how to troubleshoot problems that can occur with Interprocess Communication (IPC).

| Issue | Resolution |
|--|--|
| <p>EnterpriseOne jobs cannot communicate with one another with these symptoms:</p> <p>PORTTEST fails</p> <p>The security server on the iSeries fails</p> <p>UBE submission fails</p> <p>If you activated ipcTrace in the [JDEIPC] section of the server jde.ini file, an error similar to these should appear in the JDEDEBUG.log:</p> <p>IPC2100017 createIPC Msgq (name Port6005) failed, errno=3484: A damaged object was encountered</p> | <p>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.</p> <ul style="list-style-type: none"> • Use the *CNTRLD option to end an iSeries job. <hr/> <p>Note. You might still have damaged IPC message queues if the iSeries-controlled ending times out.</p> <hr/> <ul style="list-style-type: none"> • Run these program to verify whether a damaged message queue exists. You must have V4R1 PTF# SF45946. <p>CALL QPOZIPCS PARM('-aqE')</p> <p>This program generates a spool file called IPCS that contains information about message queues on the system. Look for these output:</p> <pre>KEY MODE 0x00000000 ----- 0x00000000 --RW----- 0x00000000 --RW----- 0x00000000 --RW----- 0x00000000 --RW----- 0x00001234 D-RW----RW-</pre> <p>In this example, the message queue 0x00001234 is damaged. To fix, stop and restart JDENET using these commands:</p> <pre>ENDNET CLRIPC STRNET</pre> <p>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.</p> <p>Note. Some of the message queues might be damaged even if the JDEDEBUG.log file does not indicate that any damage exists.</p> |

Troubleshooting the JDE.INI File

This subsection explains how to troubleshoot problems that can occur with the JDE.INI file. These notes apply to the .INI file in the E811SYS 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.
- We recommend that you place any incidental comments on a separate line adjacent to the key to which the comment applies. Be sure to include a preceding semicolon. Comments can be included at the end of the key's 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, we recommend 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 these 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 values, contact PeopleSoft Worldwide Customer Support Services.
- YES
- ON
- TRUE
- 1

Likewise, these 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.

Troubleshooting the UNIX/Linux Enterprise Server

This section provides an overview for troubleshooting the UNIX/Linux enterprise server and describes how to:

- Troubleshoot the jde.ini file.
- Troubleshoot copying EnterpriseOne files to a server.
- Troubleshoot database table configuration.
- Troubleshoot printer setup.
- Troubleshoot email.
- Troubleshoot multiple release setup.
- Troubleshoot finding the report files.

- Troubleshoot JDBNET server not found.
- Troubleshoot EnterpriseOne testing.

Understanding UNIX/Linux Enterprise Server Troubleshooting

This section describes some typical problems that you might encounter and their solutions. When troubleshooting, follow these guidelines:

- Check the 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 the 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.

Note. To complete the resolutions provided for this issues, you must sign on to the UNIX enterprise server using an account that has administrative privileges.

Troubleshooting the JDE.INI File

To locate the JDE.INI file, search in the system/bin32 subdirectory. For example, /u01/PeopleSoft/E811/ini/JDE.INI. These 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.
- We recommend that you place incidental comments on a separate line adjacent to 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 these may mean to activate 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 values, contact PeopleSoft Worldwide Customer Support Services.
- YES
- ON
- TRUE
- 1

These 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 Files 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 the system administrator.

Troubleshooting Database Tables Configuration

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 Printer Setup

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 these, then the printer is not configured on the server or is not online:

```
"lp: destination aPrinter non-existent"
```

Contact the system administrator for assistance.

For Linux, do not set up a print queue that translates files to postscript. The 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 Email

If the report, server package installation, or table conversion log file (in the PrintQueue directory) displays 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 these:

- 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/E811SYS/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 this 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 provides an overview of Windows enterprise server troubleshooting and describes how to:

- Troubleshoot EnterpriseOne account setup.
- Troubleshoot copying EnterpriseOne files to a server.
- Troubleshoot database table configuration.
- Troubleshoot printer setup.
- Troubleshoot jde.ini file setup.
- Troubleshoot finding the log files.
- Troubleshoot testing with PORTTEST.
- Troubleshoot running EnterpriseOne manually.
- Troubleshoot finding report files.
- Troubleshoot testing EnterpriseOne by submitting a report.
- Take ownership of a printer.
- Stop all EnterpriseOne processes.
- Stop all EnterpriseOne processes without rights.
- Troubleshoot email.

Understanding Windows Enterprise Server Troubleshooting

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 the 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.

Note. To complete the resolutions provided for this issues, you must sign on to the Windows Enterprise server using an account that has administrative privileges.

Troubleshooting EnterpriseOne Account Setup

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 Files 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 Database Table Configuration

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 Printer Setup

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 JDE.INI File Setup

If you cannot find the JDE.INI file:

- Search in the system\bin32 subdirectory in the EnterpriseOne tree. For example, z:\PeopleSoft\E811\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 there are no paths, the logs are in the system\bin32 directory. The log files are named according to these 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 jdidebug_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 jdidebug.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 jdidebug_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 .jdidebug.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 these 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 these 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 these:

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, select the service in Windows Control Panel, then go to Services and click 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 these 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 the EnterpriseOne network is not running, start the EnterpriseOne network service.
- Verify the EnterpriseOne network is running by doing these:
 - 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, select the EnterpriseOne network service in Windows Control Panel, select Services, and then select Startup. Note the account name. If you are using Windows 2000, select the EnterpriseOne network service in the Windows Control Panel, select Services, and then select 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 there is no location, listed, the files are in the PrintQueue directory of the EnterpriseOne base directory. For example, z:\PeopleSoft\E811\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 these:
 - 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 lists only the machines on the specified port (also known as service) that are running EnterpriseOne (either client or server). Troubleshoot using these:
 - 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 these 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).
- These 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 EnterpriseOne Tools 8.94 PeopleBook: System.
 - 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.

- All jobs sent to this printer using the current server will conform 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 list 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 these:

- Select a different report.
 - Add data to the database.
 - Change the processing options for the report.
 - Change the OCM and 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 jddebug.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 specific database 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 these 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. Stop EnterpriseOne and determine if PORTTEST runs correctly. If not, correct those problems and then run EnterpriseOne manually.

Taking Ownership of a Printer

To take ownership of a printer:

1. From the Windows Start menu, select Settings, and then Printers.
2. Right-click the desired printer.
3. Select Properties and then the Privileges tab.
4. 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 these task:

5. From the Windows Start menu, select Settings, and then Printers.
6. Right-click the desired printer.
7. Select Document Defaults.
8. Select the desired default orientation.
9. Click OK.

Stopping All EnterpriseOne Processes

If you need to stop the EnterpriseOne processes that you started from the command prompt, for example, jdenet_n, stop any of these 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 EnterpriseOne processes::

1. Run the Windows Task Manager.
2. Select the Processes tab.
3. Select one of the running processes.
4. Click End Process.
5. Repeat for each process to be stopped.

Stopping EnterpriseOne Processes without Rights

Use this task if you do not have the rights to stop the processes.

To stop EnterpriseOne processes without rights:

1. Log on to Windows in an account that has rights to stop processes.
2. Stop processes using Visual C++:
3. Run the Windows Task Manager.
4. Select the Processes tab.
5. Select one of the running processes.
6. Click Debug Process.
Visual C++ will start.
7. Click the X in the upper right-hand corner to close Visual C++. Do not save the project workspace.
This should kill the runaway process.
8. Repeat these steps for each runaway process. If they still do not end, reboot the machine.

Troubleshooting Email

If the report, server package installation, or table conversion log file in the PrintQueue directory displays 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

EnterpriseOne Tools 8.94 PeopleBook: System Administration, “Understanding the Jde.ini File Settings”

Troubleshooting Web Servers

This section provides an overview of Web server troubleshooting and describes how to:

- Troubleshoot IIS and IBM HTTP Web servers.
- Troubleshoot JAS.
- Troubleshoot serialized database and generation issues.
- Troubleshoot SQL server issues.
- Troubleshoot problems using log files.

Understanding Web Server Troubleshooting

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:\E811\internet\jas.log or ;debuglog=d:\E811\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 E811\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 these:

- [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 the 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 these commands in SQL *Plus:

```
ANALYZE TABLE owner.F989999 COMPUTE STATISTICS
```

- For SQL Server, issue these 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 the 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, select Tools, and then select Internet Options. In Internet Options, click the Advanced tab, scroll down to the section titled Java VM, and select these 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.
- 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.

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.

The 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.

See Also

Chapter 5, “Server Administration Workbench,” Monitoring the JAS Servers from the Web, page 142

Chapter 5, “Server Administration Workbench,” Configuring EnterpriseOne Server Monitoring Settings, page 105

Glossary of PeopleSoft Terms

| | |
|------------------------------|---|
| absence entitlement | This element defines rules for granting paid time off for valid absences, such as sick time, vacation, and maternity leave. An absence entitlement element defines the entitlement amount, frequency, and entitlement period. |
| absence take | This element defines the conditions that must be met before a payee is entitled to take paid time off. |
| academic career | In PeopleSoft Enterprise Campus Solutions, all course work that a student undertakes at an academic institution and that is grouped in a single student record. For example, a university that has an undergraduate school, a graduate school, and various professional schools might define several academic careers—an undergraduate career, a graduate career, and separate careers for each professional school (law school, medical school, dental school, and so on). |
| academic institution | In PeopleSoft Enterprise Campus Solutions, an entity (such as a university or college) that is independent of other similar entities and that has its own set of rules and business processes. |
| academic organization | In PeopleSoft Enterprise Campus Solutions, an entity that is part of the administrative structure within an academic institution. At the lowest level, an academic organization might be an academic department. At the highest level, an academic organization can represent a division. |
| academic plan | In PeopleSoft Enterprise Campus Solutions, an area of study—such as a major, minor, or specialization—that exists within an academic program or academic career. |
| academic program | In PeopleSoft Enterprise Campus Solutions, the entity to which a student applies and is admitted and from which the student graduates. |
| accounting class | In PeopleSoft Enterprise Performance Management, the accounting class defines how a resource is treated for generally accepted accounting practices. The Inventory class indicates whether a resource becomes part of a balance sheet account, such as inventory or fixed assets, while the Non-inventory class indicates that the resource is treated as an expense of the period during which it occurs. |
| accounting date | The accounting date indicates when a transaction is recognized, as opposed to the date the transaction actually occurred. The accounting date and transaction date can be the same. The accounting date determines the period in the general ledger to which the transaction is to be posted. You can only select an accounting date that falls within an open period in the ledger to which you are posting. The accounting date for an item is normally the invoice date. |
| accounting split | The accounting split method indicates how expenses are allocated or divided among one or more sets of accounting ChartFields. |
| accumulator | You use an accumulator to store cumulative values of defined items as they are processed. You can accumulate a single value over time or multiple values over time. For example, an accumulator could consist of all voluntary deductions, or all company deductions, enabling you to accumulate amounts. It allows total flexibility for time periods and values accumulated. |
| action reason | The reason an employee's job or employment information is updated. The action reason is entered in two parts: a personnel action, such as a promotion, termination, or change from one pay group to another—and a reason for that action. Action reasons are used by PeopleSoft Human Resources, PeopleSoft Benefits Administration, |

| | |
|--------------------------------|---|
| | PeopleSoft Stock Administration, and the COBRA Administration feature of the Base Benefits business process. |
| action template | In PeopleSoft Receivables, outlines a set of escalating actions that the system or user performs based on the period of time that a customer or item has been in an action plan for a specific condition. |
| activity | <p>In PeopleSoft Enterprise Learning Management, an instance of a catalog item (sometimes called a class) that is available for enrollment. The activity defines such things as the costs that are associated with the offering, enrollment limits and deadlines, and waitlisting capacities.</p> <p>In PeopleSoft Enterprise Performance Management, the work of an organization and the aggregation of actions that are used for activity-based costing.</p> <p>In PeopleSoft Project Costing, the unit of work that provides a further breakdown of projects—usually into specific tasks.</p> <p>In PeopleSoft Workflow, a specific transaction that you might need to perform in a business process. Because it consists of the steps that are used to perform a transaction, it is also known as a step map.</p> |
| address usage | In PeopleSoft Enterprise Campus Solutions, a grouping of address types defining the order in which the address types are used. For example, you might define an address usage code to process addresses in the following order: billing address, dormitory address, home address, and then work address. |
| adjustment calendar | In PeopleSoft Enterprise Campus Solutions, the adjustment calendar controls how a particular charge is adjusted on a student's account when the student drops classes or withdraws from a term. The charge adjustment is based on how much time has elapsed from a predetermined date, and it is determined as a percentage of the original charge amount. |
| administrative function | In PeopleSoft Enterprise Campus Solutions, a particular functional area that processes checklists, communication, and comments. The administrative function identifies which variable data is added to a person's checklist or communication record when a specific checklist code, communication category, or comment is assigned to the student. This key data enables you to trace that checklist, communication, or comment back to a specific processing event in a functional area. |
| admit type | In PeopleSoft Enterprise Campus Solutions, a designation used to distinguish first-year applications from transfer applications. |
| agreement | In PeopleSoft eSettlements, provides a way to group and specify processing options, such as payment terms, pay from a bank, and notifications by a buyer and supplier location combination. |
| allocation rule | In PeopleSoft Enterprise Incentive Management, an expression within compensation plans that enables the system to assign transactions to nodes and participants. During transaction allocation, the allocation engine traverses the compensation structure from the current node to the root node, checking each node for plans that contain allocation rules. |
| alternate account | A feature in PeopleSoft General Ledger that enables you to create a statutory chart of accounts and enter statutory account transactions at the detail transaction level, as required for recording and reporting by some national governments. |
| analysis database | In PeopleSoft Enterprise Campus Solutions, database tables that store large amounts of student information that may not appear in standard report formats. The analysis database tables contain keys for all objects in a report that an application program can use to reference other student-record objects that are not contained in the printed report. For instance, the analysis database contains data on courses that are considered for satisfying a requirement but that are rejected. It also contains information on |

| | |
|-------------------------------|--|
| | courses captured by global limits. An analysis database is used in PeopleSoft Enterprise Academic Advisement. |
| AR specialist | Abbreviation for <i>receivables specialist</i> . In PeopleSoft Receivables, an individual in who tracks and resolves deductions and disputed items. |
| arbitration plan | In PeopleSoft Enterprise Pricer, defines how price rules are to be applied to the base price when the transaction is priced. |
| assessment rule | In PeopleSoft Receivables, a user-defined rule that the system uses to evaluate the condition of a customer's account or of individual items to determine whether to generate a follow-up action. |
| asset class | An asset group used for reporting purposes. It can be used in conjunction with the asset category to refine asset classification. |
| attribute/value pair | In PeopleSoft Directory Interface, relates the data that makes up an entry in the directory information tree. |
| audience | In PeopleSoft Enterprise Campus Solutions, a segment of the database that relates to an initiative, or a membership organization that is based on constituent attributes rather than a dues-paying structure. Examples of audiences include the Class of '65 and Undergraduate Arts & Sciences. |
| authentication server | A server that is set up to verify users of the system. |
| base time period | In PeopleSoft Business Planning, the lowest level time period in a calendar. |
| benchmark job | In PeopleSoft Workforce Analytics, a benchmark job is a job code for which there is corresponding salary survey data from published, third-party sources. |
| billing career | In PeopleSoft Enterprise Campus Solutions, the one career under which other careers are grouped for billing purposes if a student is active simultaneously in multiple careers. |
| bio bit or bio brief | In PeopleSoft Enterprise Campus Solutions, a report that summarizes information stored in the system about a particular constituent. You can generate standard or specialized reports. |
| book | In PeopleSoft Asset Management, used for storing financial and tax information, such as costs, depreciation attributes, and retirement information on assets. |
| branch | A tree node that rolls up to nodes above it in the hierarchy, as defined in PeopleSoft Tree Manager. |
| budgetary account only | An account used by the system only and not by users; this type of account does not accept transactions. You can only budget with this account. Formerly called "system-maintained account." |
| budget check | In commitment control, the processing of source transactions against control budget ledgers, to see if they pass, fail, or pass with a warning. |
| budget control | In commitment control, budget control ensures that commitments and expenditures don't exceed budgets. It enables you to track transactions against corresponding budgets and terminate a document's cycle if the defined budget conditions are not met. For example, you can prevent a purchase order from being dispatched to a vendor if there are insufficient funds in the related budget to support it. |
| budget period | The interval of time (such as 12 months or 4 quarters) into which a period is divided for budgetary and reporting purposes. The ChartField allows maximum flexibility to define operational accounting time periods without restriction to only one calendar. |

| | |
|------------------------------------|--|
| business event | <p>In PeopleSoft Receivables, defines the processing characteristics for the Receivable Update process for a draft activity.</p> <p>In PeopleSoft Sales Incentive Management, an original business transaction or activity that may justify the creation of a PeopleSoft Enterprise Incentive Management event (a sale, for example).</p> |
| business unit | A corporation or a subset of a corporation that is independent with regard to one or more operational or accounting functions. |
| buyer | In PeopleSoft eSettlements, an organization (or business unit, as opposed to an individual) that transacts with suppliers (vendors) within the system. A buyer creates payments for purchases that are made in the system. |
| campus | In PeopleSoft Enterprise Campus Solutions, an entity that is usually associated with a distinct physical administrative unit, that belongs to a single academic institution, that uses a unique course catalog, and that produces a common transcript for students within the same academic career. |
| catalog item | In PeopleSoft Enterprise Learning Management, a specific topic that a learner can study and have tracked. For example, "Introduction to Microsoft Word." A catalog item contains general information about the topic and includes a course code, description, categorization, keywords, and delivery methods. A catalog item can have one or more learning activities. |
| catalog map | In PeopleSoft Catalog Management, translates values from the catalog source data to the format of the company's catalog. |
| catalog partner | In PeopleSoft Catalog Management, shares responsibility with the enterprise catalog manager for maintaining catalog content. |
| categorization | Associates partner offerings with catalog offerings and groups them into enterprise catalog categories. |
| category | In PeopleSoft Enterprise Campus Solutions, a broad grouping to which specific comments or communications (contexts) are assigned. Category codes are also linked to 3C access groups so that you can assign data-entry or view-only privileges across functions. |
| channel | In PeopleSoft MultiChannel Framework, email, chat, voice (computer telephone integration [CTI]), or a generic event. |
| ChartField | A field that stores a chart of accounts, resources, and so on, depending on the PeopleSoft application. ChartField values represent individual account numbers, department codes, and so forth. |
| ChartField balancing | You can require specific ChartFields to match up (balance) on the debit and the credit side of a transaction. |
| ChartField combination edit | The process of editing journal lines for valid ChartField combinations based on user-defined rules. |
| ChartKey | One or more fields that uniquely identify each row in a table. Some tables contain only one field as the key, while others require a combination. |
| checkbook | In PeopleSoft Promotions Management, enables you to view financial data (such as planned, incurred, and actual amounts) that is related to funds and trade promotions. |
| checklist code | In PeopleSoft Enterprise Campus Solutions, a code that represents a list of planned or completed action items that can be assigned to a staff member, volunteer, or unit. Checklists enable you to view all action assignments on one page. |

| | |
|--|---|
| class | <p>In PeopleSoft Enterprise Campus Solutions, a specific offering of a course component within an academic term.</p> <p>See also <i>course</i>.</p> |
| Class ChartField | <p>A ChartField value that identifies a unique appropriation budget key when you combine it with a fund, department ID, and program code, as well as a budget period. Formerly called <i>sub-classification</i>.</p> |
| clearance | <p>In PeopleSoft Enterprise Campus Solutions, the period of time during which a constituent in PeopleSoft Contributor Relations is approved for involvement in an initiative or an action. Clearances are used to prevent development officers from making multiple requests to a constituent during the same time period.</p> |
| clone | <p>In PeopleCode, to make a unique copy. In contrast, to <i>copy</i> may mean making a new reference to an object, so if the underlying object is changed, both the copy and the original change.</p> |
| cohort | <p>In PeopleSoft Enterprise Campus Solutions, the highest level of the three-level classification structure that you define for enrollment management. You can define a cohort level, link it to other levels, and set enrollment target numbers for it.</p> <p>See also <i>population</i> and <i>division</i>.</p> |
| collection | <p>To make a set of documents available for searching in Verity, you must first create at least one collection. A collection is set of directories and files that allow search application users to use the Verity search engine to quickly find and display source documents that match search criteria. A collection is a set of statistics and pointers to the source documents, stored in a proprietary format on a file server. Because a collection can only store information for a single location, PeopleSoft maintains a set of collections (one per language code) for each search index object.</p> |
| collection rule | <p>In PeopleSoft Receivables, a user-defined rule that defines actions to take for a customer based on both the amount and the number of days past due for outstanding balances.</p> |
| comm key | <p>See <i>communication key</i>.</p> |
| communication key | <p>In PeopleSoft Enterprise Campus Solutions, a single code for entering a combination of communication category, communication context, communication method, communication direction, and standard letter code. Communication keys (also called <i>comm keys</i> or <i>speed keys</i>) can be created for background processes as well as for specific users.</p> |
| compensation object | <p>In PeopleSoft Enterprise Incentive Management, a node within a compensation structure. Compensation objects are the building blocks that make up a compensation structure's hierarchical representation.</p> |
| compensation structure | <p>In PeopleSoft Enterprise Incentive Management, a hierarchical relationship of compensation objects that represents the compensation-related relationship between the objects.</p> |
| condition | <p>In PeopleSoft Receivables, occurs when there is a change of status for a customer's account, such as reaching a credit limit or exceeding a user-defined balance due.</p> |
| configuration parameter catalog | <p>Used to configure an external system with PeopleSoft. For example, a configuration parameter catalog might set up configuration and communication parameters for an external server.</p> |
| configuration plan | <p>In PeopleSoft Enterprise Incentive Management, configuration plans hold allocation information for common variables (not incentive rules) and are attached to a node without a participant. Configuration plans are not processed by transactions.</p> |

| | |
|--------------------------|--|
| constituents | In PeopleSoft Enterprise Campus Solutions, friends, alumni, organizations, foundations, or other entities affiliated with the institution, and about which the institution maintains information. The constituent types delivered with PeopleSoft Enterprise Contributor Relations Solutions are based on those defined by the Council for the Advancement and Support of Education (CASE). |
| content reference | Content references are pointers to content registered in the portal registry. These are typically either URLs or iScripts. Content references fall into three categories: target content, templates, and template pagelets. |
| context | <p>In PeopleCode, determines which buffer fields can be contextually referenced and which is the current row of data on each scroll level when a PeopleCode program is running.</p> <p>In PeopleSoft Enterprise Campus Solutions, a specific instance of a comment or communication. One or more contexts are assigned to a category, which you link to 3C access groups so that you can assign data-entry or view-only privileges across functions.</p> <p>In PeopleSoft Enterprise Incentive Management, a mechanism that is used to determine the scope of a processing run. PeopleSoft Enterprise Incentive Management uses three types of context: plan, period, and run-level.</p> |
| control table | Stores information that controls the processing of an application. This type of processing might be consistent throughout an organization, or it might be used only by portions of the organization for more limited sharing of data. |
| cost profile | A combination of a receipt cost method, a cost flow, and a deplete cost method. A profile is associated with a cost book and determines how items in that book are valued, as well as how the material movement of the item is valued for the book. |
| cost row | A cost transaction and amount for a set of ChartFields. |
| course | <p>In PeopleSoft Enterprise Campus Solutions, a course that is offered by a school and that is typically described in a course catalog. A course has a standard syllabus and credit level; however, these may be modified at the class level. Courses can contain multiple components such as lecture, discussion, and lab.</p> <p>See also <i>class</i>.</p> |
| course share set | In PeopleSoft Enterprise Campus Solutions, a tag that defines a set of requirement groups that can share courses. Course share sets are used in PeopleSoft Enterprise Academic Advisement. |
| current learning | In PeopleSoft Enterprise Learning Management, a self-service repository for all of a learner's in-progress learning activities and programs. |
| data acquisition | In PeopleSoft Enterprise Incentive Management, the process during which raw business transactions are acquired from external source systems and fed into the operational data store (ODS). |
| data elements | <p>Data elements, at their simplest level, define a subset of data and the rules by which to group them.</p> <p>For Workforce Analytics, data elements are rules that tell the system what measures to retrieve about your workforce groups.</p> |
| dataset | A data grouping that enables role-based filtering and distribution of data. You can limit the range and quantity of data that is displayed for a user by associating dataset rules with user roles. The result of dataset rules is a set of data that is appropriate for the user's roles. |
| delivery method | In PeopleSoft Enterprise Learning Management, identifies the primary type of delivery method in which a particular learning activity is offered. Also provides |

default values for the learning activity, such as cost and language. This is primarily used to help learners search the catalog for the type of delivery from which they learn best. Because PeopleSoft Enterprise Learning Management is a blended learning system, it does not enforce the delivery method.

In PeopleSoft Supply Chain Management, identifies the method by which goods are shipped to their destinations (such as truck, air, rail, and so on). The delivery method is specified when creating shipment schedules.

| | |
|-----------------------------------|---|
| delivery method type | In PeopleSoft Enterprise Learning Management, identifies how learning activities can be delivered—for example, through online learning, classroom instruction, seminars, books, and so forth—in an organization. The type determines whether the delivery method includes scheduled components. |
| directory information tree | In PeopleSoft Directory Interface, the representation of a directory's hierarchical structure. |
| division | In PeopleSoft Enterprise Campus Solutions, the lowest level of the three-level classification structure that you define in PeopleSoft Enterprise Recruiting and Admissions for enrollment management. You can define a division level, link it to other levels, and set enrollment target numbers for it. See also <i>population</i> and <i>cohort</i> . |
| document sequencing | A flexible method that sequentially numbers the financial transactions (for example, bills, purchase orders, invoices, and payments) in the system for statutory reporting and for tracking commercial transaction activity. |
| dynamic detail tree | A tree that takes its detail values—dynamic details—directly from a table in the database, rather than from a range of values that are entered by the user. |
| edit table | A table in the database that has its own record definition, such as the Department table. As fields are entered into a PeopleSoft application, they can be validated against an edit table to ensure data integrity throughout the system. |
| effective date | A method of dating information in PeopleSoft applications. You can predate information to add historical data to your system, or postdate information in order to enter it before it actually goes into effect. By using effective dates, you don't delete values; you enter a new value with a current effective date. |
| EIM ledger | Abbreviation for <i>Enterprise Incentive Management ledger</i> . In PeopleSoft Enterprise Incentive Management, an object to handle incremental result gathering within the scope of a participant. The ledger captures a result set with all of the appropriate traces to the data origin and to the processing steps of which it is a result. |
| elimination set | In PeopleSoft General Ledger, a related group of intercompany accounts that is processed during consolidations. |
| entry event | In PeopleSoft General Ledger, Receivables, Payables, Purchasing, and Billing, a business process that generates multiple debits and credits resulting from single transactions to produce standard, supplemental accounting entries. |
| equitization | In PeopleSoft General Ledger, a business process that enables parent companies to calculate the net income of subsidiaries on a monthly basis and adjust that amount to increase the investment amount and equity income amount before performing consolidations. |
| equity item limit | In PeopleSoft Enterprise Campus Solutions, the amounts of funds set by the institution to be awarded with discretionary or gift funds. The limit could be reduced by amounts equal to such things as expected family contribution (EFC) or parent contribution. Students are packaged by Equity Item Type Groups and Related Equity Item Types. This limit can be used to assure that similar student populations are packaged equally. |

| | |
|----------------------------------|---|
| event | <p>A predefined point either in the Component Processor flow or in the program flow. As each point is encountered, the event activates each component, triggering any PeopleCode program that is associated with that component and that event. Examples of events are FieldChange, SavePreChange, and RowDelete.</p> <p>In PeopleSoft Human Resources, also refers to an incident that affects benefits eligibility.</p> |
| event propagation process | <p>In PeopleSoft Sales Incentive Management, a process that determines, through logic, the propagation of an original PeopleSoft Enterprise Incentive Management event and creates a derivative (duplicate) of the original event to be processed by other objects. Sales Incentive Management uses this mechanism to implement splits, roll-ups, and so on. Event propagation determines who receives the credit.</p> |
| exception | <p>In PeopleSoft Receivables, an item that either is a deduction or is in dispute.</p> |
| exclusive pricing | <p>In PeopleSoft Order Management, a type of arbitration plan that is associated with a price rule. Exclusive pricing is used to price sales order transactions.</p> |
| fact | <p>In PeopleSoft applications, facts are numeric data values from fields from a source database as well as an analytic application. A fact can be anything you want to measure your business by, for example, revenue, actual, budget data, or sales numbers. A fact is stored on a fact table.</p> |
| financial aid term | <p>In PeopleSoft Enterprise Campus Solutions, a combination of a period of time that the school determines as an instructional accounting period and an academic career. It is created and defined during the setup process. Only terms eligible for financial aid are set up for each financial aid career.</p> |
| forecast item | <p>A logical entity with a unique set of descriptive demand and forecast data that is used as the basis to forecast demand. You create forecast items for a wide range of uses, but they ultimately represent things that you buy, sell, or use in your organization and for which you require a predictable usage.</p> |
| fund | <p>In PeopleSoft Promotions Management, a budget that can be used to fund promotional activity. There are four funding methods: top down, fixed accrual, rolling accrual, and zero-based accrual.</p> |
| gap | <p>In PeopleSoft Enterprise Campus Solutions, an artificial figure that sets aside an amount of unmet financial aid need that is not funded with Title IV funds. A gap can be used to prevent fully funding any student to conserve funds, or it can be used to preserve unmet financial aid need so that institutional funds can be awarded.</p> |
| generic process type | <p>In PeopleSoft Process Scheduler, process types are identified by a generic process type. For example, the generic process type SQR includes all SQR process types, such as SQR process and SQR report.</p> |
| gift table | <p>In PeopleSoft Enterprise Campus Solutions, a table or so-called <i>donor pyramid</i> describing the number and size of gifts that you expect will be needed to successfully complete the campaign in PeopleSoft Contributor Relations. The gift table enables you to estimate the number of donors and prospects that you need at each gift level to reach the campaign goal.</p> |
| GL business unit | <p>Abbreviation for <i>general ledger business unit</i>. A unit in an organization that is an independent entity for accounting purposes. It maintains its own set of accounting books.</p> <p>See also <i>business unit</i>.</p> |
| GL entry template | <p>Abbreviation for <i>general ledger entry template</i>. In PeopleSoft Enterprise Campus Solutions, a template that defines how a particular item is sent to the general ledger. An item-type maps to the general ledger, and the GL entry template can involve multiple general ledger accounts. The entry to the general ledger is further controlled</p> |

by high-level flags that control the summarization and the type of accounting—that is, accrual or cash.

GL Interface process

Abbreviation for *General Ledger Interface process*. In PeopleSoft Enterprise Campus Solutions, a process that is used to send transactions from PeopleSoft Enterprise Student Financials to the general ledger. Item types are mapped to specific general ledger accounts, enabling transactions to move to the general ledger when the GL Interface process is run.

group

In PeopleSoft Billing and Receivables, a posting entity that comprises one or more transactions (items, deposits, payments, transfers, matches, or write-offs).

In PeopleSoft Human Resources Management and Supply Chain Management, any set of records that are associated under a single name or variable to run calculations in PeopleSoft business processes. In PeopleSoft Time and Labor, for example, employees are placed in groups for time reporting purposes.

incentive object

In PeopleSoft Enterprise Incentive Management, the incentive-related objects that define and support the PeopleSoft Enterprise Incentive Management calculation process and results, such as plan templates, plans, results data, user interaction objects, and so on.

incentive rule

In PeopleSoft Sales Incentive Management, the commands that act on transactions and turn them into compensation. A rule is one part in the process of turning a transaction into compensation.

incur

In PeopleSoft Promotions Management, to become liable for a promotional payment. In other words, you owe that amount to a customer for promotional activities.

initiative

In PeopleSoft Enterprise Campus Solutions, the basis from which all advancement plans are executed. It is an organized effort targeting a specific constituency, and it can occur over a specified period of time with specific purposes and goals. An initiative can be a campaign, an event, an organized volunteer effort, a membership drive, or any other type of effort defined by the institution. Initiatives can be multipart, and they can be related to other initiatives. This enables you to track individual parts of an initiative, as well as entire initiatives.

inquiry access

In PeopleSoft Enterprise Campus Solutions, a type of security access that permits the user only to view data.

See also *update access*.

institution

In PeopleSoft Enterprise Campus Solutions, an entity (such as a university or college) that is independent of other similar entities and that has its own set of rules and business processes.

item

In PeopleSoft Inventory, a tangible commodity that is stored in a business unit (shipped from a warehouse).

In PeopleSoft Demand Planning, Inventory Policy Planning, and Supply Planning, a noninventory item that is designated as being used for planning purposes only. It can represent a family or group of inventory items. It can have a planning bill of material (BOM) or planning routing, and it can exist as a component on a planning BOM. A planning item cannot be specified on a production or engineering BOM or routing, and it cannot be used as a component in a production. The quantity on hand will never be maintained.

In PeopleSoft Receivables, an individual receivable. An item can be an invoice, a credit memo, a debit memo, a write-off, or an adjustment.

item shuffle

In PeopleSoft Enterprise Campus Solutions, a process that enables you to change a payment allocation without having to reverse the payment.

| | |
|-----------------------------|---|
| joint communication | In PeopleSoft Enterprise Campus Solutions, one letter that is addressed jointly to two people. For example, a letter might be addressed to both Mr. Sudhir Awat and Ms. Samantha Mortelli. A relationship must be established between the two individuals in the database, and at least one of the individuals must have an ID in the database. |
| keyword | In PeopleSoft Enterprise Campus Solutions, a term that you link to particular elements within PeopleSoft Student Financials, Financial Aid, and Contributor Relations. You can use keywords as search criteria that enable you to locate specific records in a search dialog box. |
| KPI | An abbreviation for <i>key performance indicator</i> . A high-level measurement of how well an organization is doing in achieving critical success factors. This defines the data value or calculation upon which an assessment is determined. |
| LDIF file | Abbreviation for <i>Lightweight Directory Access Protocol (LDAP) Data Interchange Format file</i> . Contains discrepancies between PeopleSoft data and directory data. |
| learner group | In PeopleSoft Enterprise Learning Management, a group of learners who are linked to the same learning environment. Members of the learner group can share the same attributes, such as the same department or job code. Learner groups are used to control access to and enrollment in learning activities and programs. They are also used to perform group enrollments and mass enrollments in the back office. |
| learning components | In PeopleSoft Enterprise Learning Management, the foundational building blocks of learning activities. PeopleSoft Enterprise Learning Management supports six basic types of learning components: web-based, session, webcast, test, survey, and assignment. One or more of these learning component types compose a single learning activity. |
| learning environment | In PeopleSoft Enterprise Learning Management, identifies a set of categories and catalog items that can be made available to learner groups. Also defines the default values that are assigned to the learning activities and programs that are created within a particular learning environment. Learning environments provide a way to partition the catalog so that learners see only those items that are relevant to them. |
| learning history | In PeopleSoft Enterprise Learning Management, a self-service repository for all of a learner's completed learning activities and programs. |
| ledger mapping | You use ledger mapping to relate expense data from general ledger accounts to resource objects. Multiple ledger line items can be mapped to one or more resource IDs. You can also use ledger mapping to map dollar amounts (referred to as <i>rates</i>) to business units. You can map the amounts in two different ways: an actual amount that represents actual costs of the accounting period, or a budgeted amount that can be used to calculate the capacity rates as well as budgeted model results. In PeopleSoft Enterprise Warehouse, you can map general ledger accounts to the EW Ledger table. |
| library section | In PeopleSoft Enterprise Incentive Management, a section that is defined in a plan (or template) and that is available for other plans to share. Changes to a library section are reflected in all plans that use it. |
| linked section | In PeopleSoft Enterprise Incentive Management, a section that is defined in a plan template but appears in a plan. Changes to linked sections propagate to plans using that section. |
| linked variable | In PeopleSoft Enterprise Incentive Management, a variable that is defined and maintained in a plan template and that also appears in a plan. Changes to linked variables propagate to plans using that variable. |
| LMS | Abbreviation for <i>learning management system</i> . In PeopleSoft Enterprise Campus Solutions, LMS is a PeopleSoft Student Records feature that provides a common set of interoperability standards that enable the sharing of instructional content and data between learning and administrative environments. |

| | |
|-------------------------------|--|
| load | In PeopleSoft Inventory, identifies a group of goods that are shipped together. Load management is a feature of PeopleSoft Inventory that is used to track the weight, the volume, and the destination of a shipment. |
| local functionality | In PeopleSoft HRMS, the set of information that is available for a specific country. You can access this information when you click the appropriate country flag in the global window, or when you access it by a local country menu. |
| location | Locations enable you to indicate the different types of addresses—for a company, for example, one address to receive bills, another for shipping, a third for postal deliveries, and a separate street address. Each address has a different location number. The primary location—indicated by a <i>1</i> —is the address you use most often and may be different from the main address. |
| logistical task | In PeopleSoft Services Procurement, an administrative task that is related to hiring a service provider. Logistical tasks are linked to the service type on the work order so that different types of services can have different logistical tasks. Logistical tasks include both preapproval tasks (such as assigning a new badge or ordering a new laptop) and postapproval tasks (such as scheduling orientation or setting up the service provider email). The logistical tasks can be mandatory or optional. Mandatory preapproval tasks must be completed before the work order is approved. Mandatory postapproval tasks, on the other hand, must be completed before a work order is released to a service provider. |
| market template | In PeopleSoft Enterprise Incentive Management, additional functionality that is specific to a given market or industry and is built on top of a product category. |
| mass change | In PeopleSoft Enterprise Campus Solutions, mass change is a SQL generator that can be used to create specialized functionality. Using mass change, you can set up a series of Insert, Update, or Delete SQL statements to perform business functions that are specific to the institution. See also <i>3C engine</i> . |
| match group | In PeopleSoft Receivables, a group of receivables items and matching offset items. The system creates match groups by using user-defined matching criteria for selected field values. |
| MCF server | Abbreviation for <i>PeopleSoft MultiChannel Framework server</i> . Comprises the universal queue server and the MCF log server. Both processes are started when <i>MCF Servers</i> is selected in an application server domain configuration. |
| merchandising activity | In PeopleSoft Promotions Management, a specific discount type that is associated with a trade promotion (such as off-invoice, billback or rebate, or lump-sum payment) that defines the performance that is required to receive the discount. In the industry, you may know this as an offer, a discount, a merchandising event, an event, or a tactic. |
| meta-SQL | Meta-SQL constructs expand into platform-specific Structured Query Language (SQL) substrings. They are used in functions that pass SQL strings, such as in SQL objects, the SQLExec function, and PeopleSoft Application Engine programs. |
| metastring | Metastrings are special expressions included in SQL string literals. The metastrings, prefixed with a percent (%) symbol, are included directly in the string literals. They expand at run time into an appropriate substring for the current database platform. |
| multibook | In PeopleSoft General Ledger, multiple ledgers having multiple-base currencies that are defined for a business unit, with the option to post a single transaction to all base currencies (all ledgers) or to only one of those base currencies (ledgers). |
| multicurrency | The ability to process transactions in a currency other than the business unit's base currency. |

| | |
|---|---|
| national allowance | In PeopleSoft Promotions Management, a promotion at the corporate level that is funded by nondiscretionary dollars. In the industry, you may know this as a national promotion, a corporate promotion, or a corporate discount. |
| need | In PeopleSoft Enterprise Campus Solutions, the difference between the cost of attendance (COA) and the expected family contribution (EFC). It is the gap between the cost of attending the school and the student's resources. The financial aid package is based on the amount of financial need. The process of determining a student's need is called <i>need analysis</i> . |
| node-oriented tree | A tree that is based on a detail structure, but the detail values are not used. |
| pagelet | Each block of content on the home page is called a pagelet. These pagelets display summary information within a small rectangular area on the page. The pagelet provide users with a snapshot of their most relevant PeopleSoft and non-PeopleSoft content. |
| participant | In PeopleSoft Enterprise Incentive Management, participants are recipients of the incentive compensation calculation process. |
| participant object | Each participant object may be related to one or more compensation objects. See also <i>compensation object</i> . |
| partner | A company that supplies products or services that are resold or purchased by the enterprise. |
| pay cycle | In PeopleSoft Payables, a set of rules that define the criteria by which it should select scheduled payments for payment creation. |
| payment shuffle | In PeopleSoft Enterprise Campus Solutions, a process allowing payments that have been previously posted to a student's account to be automatically reapplied when a higher priority payment is posted or the payment allocation definition is changed. |
| pending item | In PeopleSoft Receivables, an individual receivable (such as an invoice, a credit memo, or a write-off) that has been entered in or created by the system, but hasn't been posted. |
| PeopleCode | PeopleCode is a proprietary language, executed by the PeopleSoft application processor. PeopleCode generates results based upon existing data or user actions. By using business interlink objects, external services are available to all PeopleSoft applications wherever PeopleCode can be executed. |
| PeopleCode event | An action that a user takes upon an object, usually a record field, that is referenced within a PeopleSoft page. |
| PeopleSoft Internet Architecture | The fundamental architecture on which PeopleSoft 8 applications are constructed, consisting of a relational database management system (RDBMS), an application server, a web server, and a browser. |
| performance measurement | In PeopleSoft Enterprise Incentive Management, a variable used to store data (similar to an aggregator, but without a predefined formula) within the scope of an incentive plan. Performance measures are associated with a plan calendar, territory, and participant. Performance measurements are used for quota calculation and reporting. |
| period context | In PeopleSoft Enterprise Incentive Management, because a participant typically uses the same compensation plan for multiple periods, the period context associates a plan context with a specific calendar period and fiscal year. The period context references the associated plan context, thus forming a chain. Each plan context has a corresponding set of period contexts. |
| person of interest | A person about whom the organization maintains information but who is not part of the workforce. |

| | |
|-----------------------------|--|
| personal portfolio | In PeopleSoft Enterprise Campus Solutions, the user-accessible menu item that contains an individual's name, address, telephone number, and other personal information. |
| plan | In PeopleSoft Sales Incentive Management, a collection of allocation rules, variables, steps, sections, and incentive rules that instruct the PeopleSoft Enterprise Incentive Management engine in how to process transactions. |
| plan context | In PeopleSoft Enterprise Incentive Management, correlates a participant with the compensation plan and node to which the participant is assigned, enabling the PeopleSoft Enterprise Incentive Management system to find anything that is associated with the node and that is required to perform compensation processing. Each participant, node, and plan combination represents a unique plan context—if three participants are on a compensation structure, each has a different plan context. Configuration plans are identified by plan contexts and are associated with the participants that refer to them. |
| plan template | In PeopleSoft Enterprise Incentive Management, the base from which a plan is created. A plan template contains common sections and variables that are inherited by all plans that are created from the template. A template may contain steps and sections that are not visible in the plan definition. |
| planned learning | In PeopleSoft Enterprise Learning Management, a self-service repository for all of a learner's planned learning activities and programs. |
| planning instance | In PeopleSoft Supply Planning, a set of data (business units, items, supplies, and demands) constituting the inputs and outputs of a supply plan. |
| population | In PeopleSoft Enterprise Campus Solutions, the middle level of the three-level classification structure that you define in PeopleSoft Enterprise Recruiting and Admissions for enrollment management. You can define a population level, link it to other levels, and set enrollment target numbers for it. See also <i>division</i> and <i>cohort</i> . |
| portal registry | In PeopleSoft applications, the portal registry is a tree-like structure in which content references are organized, classified, and registered. It is a central repository that defines both the structure and content of a portal through a hierarchical, tree-like structure of folders useful for organizing and securing content references. |
| price list | In PeopleSoft Enterprise Pricer, enables you to select products and conditions for which the price list applies to a transaction. During a transaction, the system either determines the product price based on the predefined search hierarchy for the transaction or uses the product's lowest price on any associated, active price lists. This price is used as the basis for any further discounts and surcharges. |
| price rule | In PeopleSoft Enterprise Pricer, defines the conditions that must be met for adjustments to be applied to the base price. Multiple rules can apply when conditions of each rule are met. |
| price rule condition | In PeopleSoft Enterprise Pricer, selects the price-by fields, the values for the price-by fields, and the operator that determines how the price-by fields are related to the transaction. |
| price rule key | In PeopleSoft Enterprise Pricer, defines the fields that are available to define price rule conditions (which are used to match a transaction) on the price rule. |
| primacy number | In PeopleSoft Enterprise Campus Solutions, a number that the system uses to prioritize financial aid applications when students are enrolled in multiple academic careers and academic programs at the same time. The Consolidate Academic Statistics process uses the primacy number indicated for both the career and program at the institutional level to determine a student's primary career and program. The system also uses the |

| | |
|----------------------------|---|
| | number to determine the primary student attribute value that is used when you extract data to report on cohorts. The lowest number takes precedence. |
| primary name type | In PeopleSoft Enterprise Campus Solutions, the name type that is used to link the name stored at the highest level within the system to the lower-level set of names that an individual provides. |
| process category | In PeopleSoft Process Scheduler, processes that are grouped for server load balancing and prioritization. |
| process group | In PeopleSoft Financials, a group of application processes (performed in a defined order) that users can initiate in real time, directly from a transaction entry page. |
| process definition | Process definitions define each run request. |
| process instance | A unique number that identifies each process request. This value is automatically incremented and assigned to each requested process when the process is submitted to run. |
| process job | You can link process definitions into a job request and process each request serially or in parallel. You can also initiate subsequent processes based on the return code from each prior request. |
| process request | A single run request, such as a Structured Query Report (SQR), a COBOL or Application Engine program, or a Crystal report that you run through PeopleSoft Process Scheduler. |
| process run control | A PeopleTools variable used to retain PeopleSoft Process Scheduler values needed at runtime for all requests that reference a run control ID. Do not confuse these with application run controls, which may be defined with the same run control ID, but only contain information specific to a given application process request. |
| product category | In PeopleSoft Enterprise Incentive Management, indicates an application in the Enterprise Incentive Management suite of products. Each transaction in the PeopleSoft Enterprise Incentive Management system is associated with a product category. |
| programs | In PeopleSoft Enterprise Learning Management, a high-level grouping that guides the learner along a specific learning path through sections of catalog items. PeopleSoft Enterprise Learning Systems provides two types of programs—curricula and certifications. |
| progress log | In PeopleSoft Services Procurement, tracks deliverable-based projects. This is similar to the time sheet in function and process. The service provider contact uses the progress log to record and submit progress on deliverables. The progress can be logged by the activity that is performed, by the percentage of work that is completed, or by the completion of milestone activities that are defined for the project. |
| project transaction | In PeopleSoft Project Costing, an individual transaction line that represents a cost, time, budget, or other transaction row. |
| promotion | In PeopleSoft Promotions Management, a trade promotion, which is typically funded from trade dollars and used by consumer products manufacturers to increase sales volume. |
| prospects | In PeopleSoft Enterprise Campus Solutions, students who are interested in applying to the institution. In PeopleSoft Enterprise Contributor Relations, individuals and organizations that are most likely to make substantial financial commitments or other types of commitments to the institution. |
| publishing | In PeopleSoft Enterprise Incentive Management, a stage in processing that makes incentive-related results available to participants. |

| | |
|--------------------------------|---|
| rating components | In PeopleSoft Enterprise Campus Solutions, variables used with the Equation Editor to retrieve specified populations. |
| record group | A set of logically and functionally related control tables and views. Record groups help enable TableSet sharing, which eliminates redundant data entry. Record groups ensure that TableSet sharing is applied consistently across all related tables and views. |
| record input VAT flag | Abbreviation for <i>record input value-added tax flag</i> . Within PeopleSoft Purchasing, Payables, and General Ledger, this flag indicates that you are recording input VAT on the transaction. This flag, in conjunction with the record output VAT flag, is used to determine the accounting entries created for a transaction and to determine how a transaction is reported on the VAT return. For all cases within Purchasing and Payables where VAT information is tracked on a transaction, this flag is set to Yes. This flag is not used in PeopleSoft Order Management, Billing, or Receivables, where it is assumed that you are always recording only output VAT, or in PeopleSoft Expenses, where it is assumed that you are always recording only input VAT. |
| record output VAT flag | Abbreviation for <i>record output value-added tax flag</i> . See <i>record input VAT flag</i> . |
| recname | The name of a record that is used to determine the associated field to match a value or set of values. |
| recognition | In PeopleSoft Enterprise Campus Solutions, the recognition type indicates whether the PeopleSoft Enterprise Contributor Relations donor is the primary donor of a commitment or shares the credit for a donation. Primary donors receive hard credit that must total 100 percent. Donors that share the credit are given soft credit. Institutions can also define other share recognition-type values such as memo credit or vehicle credit. |
| reference data | In PeopleSoft Sales Incentive Management, system objects that represent the sales organization, such as territories, participants, products, customers, channels, and so on. |
| reference object | In PeopleSoft Enterprise Incentive Management, this dimension-type object further defines the business. Reference objects can have their own hierarchy (for example, product tree, customer tree, industry tree, and geography tree). |
| reference transaction | In commitment control, a reference transaction is a source transaction that is referenced by a higher-level (and usually later) source transaction, in order to automatically reverse all or part of the referenced transaction's budget-checked amount. This avoids duplicate postings during the sequential entry of the transaction at different commitment levels. For example, the amount of an encumbrance transaction (such as a purchase order) will, when checked and recorded against a budget, cause the system to concurrently reference and relieve all or part of the amount of a corresponding pre-encumbrance transaction, such as a purchase requisition. |
| regional sourcing | In PeopleSoft Purchasing, provides the infrastructure to maintain, display, and select an appropriate vendor and vendor pricing structure that is based on a regional sourcing model where the multiple ship to locations are grouped. Sourcing may occur at a level higher than the ship to location. |
| relationship object | In PeopleSoft Enterprise Incentive Management, these objects further define a compensation structure to resolve transactions by establishing associations between compensation objects and business objects. |
| remote data source data | Data that is extracted from a separate database and migrated into the local database. |
| REN server | Abbreviation for <i>real-time event notification server</i> in PeopleSoft MultiChannel Framework. |
| requester | In PeopleSoft eSettlements, an individual who requests goods or services and whose ID appears on the various procurement pages that reference purchase orders. |

| | |
|-----------------------------|--|
| reversal indicator | In PeopleSoft Enterprise Campus Solutions, an indicator that denotes when a particular payment has been reversed, usually because of insufficient funds. |
| role | Describes how people fit into PeopleSoft Workflow. A role is a class of users who perform the same type of work, such as clerks or managers. Your business rules typically specify what user role needs to do an activity. |
| role user | A PeopleSoft Workflow user. A person's role user ID serves much the same purpose as a user ID does in other parts of the system. PeopleSoft Workflow uses role user IDs to determine how to route worklist items to users (through an email address, for example) and to track the roles that users play in the workflow. Role users do not need PeopleSoft user IDs. |
| roll up | In a tree, to roll up is to total sums based on the information hierarchy. |
| run control | A run control is a type of online page that is used to begin a process, such as the batch processing of a payroll run. Run control pages generally start a program that manipulates data. |
| run control ID | A unique ID to associate each user with his or her own run control table entries. |
| run-level context | In PeopleSoft Enterprise Incentive Management, associates a particular run (and batch ID) with a period context and plan context. Every plan context that participates in a run has a separate run-level context. Because a run cannot span periods, only one run-level context is associated with each plan context. |
| search query | You use this set of objects to pass a query string and operators to the search engine. The search index returns a set of matching results with keys to the source documents. |
| search/match | In PeopleSoft Enterprise Campus Solutions and PeopleSoft Enterprise Human Resources Management Solutions, a feature that enables you to search for and identify duplicate records in the database. |
| seasonal address | In PeopleSoft Enterprise Campus Solutions, an address that recurs for the same length of time at the same time of year each year until adjusted or deleted. |
| section | In PeopleSoft Enterprise Incentive Management, a collection of incentive rules that operate on transactions of a specific type. Sections enable plans to be segmented to process logical events in different sections. |
| security event | In commitment control, security events trigger security authorization checking, such as budget entries, transfers, and adjustments; exception overrides and notifications; and inquiries. |
| serial genealogy | In PeopleSoft Manufacturing, the ability to track the composition of a specific, serial-controlled item. |
| serial in production | In PeopleSoft Manufacturing, enables the tracing of serial information for manufactured items. This is maintained in the Item Master record. |
| service impact | In PeopleSoft Enterprise Campus Solutions, the resulting action triggered by a service indicator. For example, a service indicator that reflects nonpayment of account balances by a student might result in a service impact that prohibits registration for classes. |
| service indicator | In PeopleSoft Enterprise Campus Solutions, indicates services that may be either withheld or provided to an individual. Negative service indicators indicate holds that prevent the individual from receiving specified services, such as check-cashing privileges or registration for classes. Positive service indicators designate special services that are provided to the individual, such as front-of-line service or special services for disabled students. |

| | |
|--------------------------------|---|
| session | <p>In PeopleSoft Enterprise Campus Solutions, time elements that subdivide a term into multiple time periods during which classes are offered. In PeopleSoft Contributor Relations, a session is the means of validating gift, pledge, membership, or adjustment data entry. It controls access to the data entered by a specific user ID. Sessions are balanced, queued, and then posted to the institution's financial system. Sessions must be posted to enter a matching gift or pledge payment, to make an adjustment, or to process giving clubs or acknowledgements.</p> <p>In PeopleSoft Enterprise Learning Management, a single meeting day of an activity (that is, the period of time between start and finish times within a day). The session stores the specific date, location, meeting time, and instructor. Sessions are used for scheduled training.</p> |
| session template | In PeopleSoft Enterprise Learning Management, enables you to set up common activity characteristics that may be reused while scheduling a PeopleSoft Enterprise Learning Management activity—characteristics such as days of the week, start and end times, facility and room assignments, instructors, and equipment. A session pattern template can be attached to an activity that is being scheduled. Attaching a template to an activity causes all of the default template information to populate the activity session pattern. |
| setup relationship | In PeopleSoft Enterprise Incentive Management, a relationship object type that associates a configuration plan with any structure node. |
| share driver expression | In PeopleSoft Business Planning, a named planning method similar to a driver expression, but which you can set up globally for shared use within a single planning application or to be shared between multiple planning applications through PeopleSoft Enterprise Warehouse. |
| single signon | With single signon, users can, after being authenticated by a PeopleSoft application server, access a second PeopleSoft application server without entering a user ID or password. |
| source key process | In PeopleSoft Enterprise Campus Solutions, a process that relates a particular transaction to the source of the charge or financial aid. On selected pages, you can drill down into particular charges. |
| source transaction | In commitment control, any transaction generated in a PeopleSoft or third-party application that is integrated with commitment control and which can be checked against commitment control budgets. For example, a pre-encumbrance, encumbrance, expenditure, recognized revenue, or collected revenue transaction. |
| speed key | See <i>communication key</i> . |
| SpeedChart | A user-defined shorthand key that designates several ChartKeys to be used for voucher entry. Percentages can optionally be related to each ChartKey in a SpeedChart definition. |
| SpeedType | A code representing a combination of ChartField values. SpeedTypes simplify the entry of ChartFields commonly used together. |
| staging | A method of consolidating selected partner offerings with the offerings from the enterprise's other partners. |
| standard letter code | In PeopleSoft Enterprise Campus Solutions, a standard letter code used to identify each letter template available for use in mail merge functions. Every letter generated in the system must have a standard letter code identification. |
| statutory account | Account required by a regulatory authority for recording and reporting financial results. In PeopleSoft, this is equivalent to the Alternate Account (ALTACCT) ChartField. |

| | |
|------------------------------|--|
| step | In PeopleSoft Sales Incentive Management, a collection of sections in a plan. Each step corresponds to a step in the job run. |
| storage level | In PeopleSoft Inventory, identifies the level of a material storage location. Material storage locations are made up of a business unit, a storage area, and a storage level. You can set up to four storage levels. |
| subcustomer qualifier | A value that groups customers into a division for which you can generate detailed history, aging, events, and profiles. |
| Summary ChartField | You use summary ChartFields to create summary ledgers that roll up detail amounts based on specific detail values or on selected tree nodes. When detail values are summarized using tree nodes, summary ChartFields must be used in the summary ledger data record to accommodate the maximum length of a node name (20 characters). |
| summary ledger | An accounting feature used primarily in allocations, inquiries, and PS/nVision reporting to store combined account balances from detail ledgers. Summary ledgers increase speed and efficiency of reporting by eliminating the need to summarize detail ledger balances each time a report is requested. Instead, detail balances are summarized in a background process according to user-specified criteria and stored on summary ledgers. The summary ledgers are then accessed directly for reporting. |
| summary time period | In PeopleSoft Business Planning, any time period (other than a base time period) that is an aggregate of other time periods, including other summary time periods and base time periods, such as quarter and year total. |
| summary tree | A tree used to roll up accounts for each type of report in summary ledgers. Summary trees enable you to define trees on trees. In a summary tree, the detail values are really nodes on a detail tree or another summary tree (known as the <i>basis</i> tree). A summary tree structure specifies the details on which the summary trees are to be built. |
| syndicate | To distribute a production version of the enterprise catalog to partners. |
| system function | In PeopleSoft Receivables, an activity that defines how the system generates accounting entries for the general ledger. |
| TableSet | A means of sharing similar sets of values in control tables, where the actual data values are different but the structure of the tables is the same. |
| TableSet sharing | Shared data that is stored in many tables that are based on the same TableSets. Tables that use TableSet sharing contain the SETID field as an additional key or unique identifier. |
| target currency | The value of the entry currency or currencies converted to a single currency for budget viewing and inquiry purposes. |
| tax authority | In PeopleSoft Enterprise Campus Solutions, a user-defined element that combines a description and percentage of a tax with an account type, an item type, and a service impact. |
| template | A template is HTML code associated with a web page. It defines the layout of the page and also where to get HTML for each part of the page. In PeopleSoft, you use templates to build a page by combining HTML from a number of sources. For a PeopleSoft portal, all templates must be registered in the portal registry, and each content reference must be assigned a template. |
| territory | In PeopleSoft Sales Incentive Management, hierarchical relationships of business objects, including regions, products, customers, industries, and participants. |
| 3C engine | Abbreviation for <i>Communications, Checklists, and Comments engine</i> . In PeopleSoft Enterprise Campus Solutions, the 3C engine enables you to automate business processes that involve additions, deletions, and updates to communications, checklists, |

and comments. You define events and triggers to engage the engine, which runs the mass change and processes the 3C records (for individuals or organizations) immediately and automatically from within business processes.

| | |
|------------------------------------|--|
| 3C group | Abbreviation for <i>Communications, Checklists, and Comments group</i> . In PeopleSoft Enterprise Campus Solutions, a method of assigning or restricting access privileges. A 3C group enables you to group specific communication categories, checklist codes, and comment categories. You can then assign the group inquiry-only access or update access, as appropriate. |
| TimeSpan | A relative period, such as year-to-date or current period, that can be used in various PeopleSoft General Ledger functions and reports when a rolling time frame, rather than a specific date, is required. TimeSpans can also be used with flexible formulas in PeopleSoft Projects. |
| trace usage | In PeopleSoft Manufacturing, enables the control of which components will be traced during the manufacturing process. Serial- and lot-controlled components can be traced. This is maintained in the Item Master record. |
| transaction allocation | In PeopleSoft Enterprise Incentive Management, the process of identifying the owner of a transaction. When a raw transaction from a batch is allocated to a plan context, the transaction is duplicated in the PeopleSoft Enterprise Incentive Management transaction tables. |
| transaction state | In PeopleSoft Enterprise Incentive Management, a value assigned by an incentive rule to a transaction. Transaction states enable sections to process only transactions that are at a specific stage in system processing. After being successfully processed, transactions may be promoted to the next transaction state and “picked up” by a different section for further processing. |
| Translate table | A system edit table that stores codes and translate values for the miscellaneous fields in the database that do not warrant individual edit tables of their own. |
| tree | The graphical hierarchy in PeopleSoft systems that displays the relationship between all accounting units (for example, corporate divisions, projects, reporting groups, account numbers) and determines roll-up hierarchies. |
| tuition lock | In PeopleSoft Enterprise Campus Solutions, a feature in the Tuition Calculation process that enables you to specify a point in a term after which students are charged a minimum (or <i>locked</i>) fee amount. Students are charged the locked fee amount even if they later drop classes and take less than the normal load level for that tuition charge. |
| unclaimed transaction | In PeopleSoft Enterprise Incentive Management, a transaction that is not claimed by a node or participant after the allocation process has completed, usually due to missing or incomplete data. Unclaimed transactions may be manually assigned to the appropriate node or participant by a compensation administrator. |
| universal navigation header | Every PeopleSoft portal includes the universal navigation header, intended to appear at the top of every page as long as the user is signed on to the portal. In addition to providing access to the standard navigation buttons (like Home, Favorites, and signoff) the universal navigation header can also display a welcome message for each user. |
| update access | In PeopleSoft Enterprise Campus Solutions, a type of security access that permits the user to edit and update data. See also <i>inquiry access</i> . |
| user interaction object | In PeopleSoft Sales Incentive Management, used to define the reporting components and reports that a participant can access in his or her context. All Sales Incentive Management user interface objects and reports are registered as user interaction objects. User interaction objects can be linked to a compensation structure node through a compensation relationship object (individually or as groups). |

| | |
|---------------------------|---|
| variable | In PeopleSoft Sales Incentive Management, the intermediate results of calculations. Variables hold the calculation results and are then inputs to other calculations. Variables can be plan variables that persist beyond the run of an engine or local variables that exist only during the processing of a section. |
| VAT exception | Abbreviation for <i>value-added tax exception</i> . A temporary or permanent exemption from paying VAT that is granted to an organization. This terms refers to both VAT exoneration and VAT suspension. |
| VAT exempt | Abbreviation for <i>value-added tax exempt</i> . Describes goods and services that are not subject to VAT. Organizations that supply exempt goods or services are unable to recover the related input VAT. This is also referred to as exempt without recovery. |
| VAT exoneration | Abbreviation for <i>value-added tax exoneration</i> . An organization that has been granted a permanent exemption from paying VAT due to the nature of that organization. |
| VAT suspension | Abbreviation for <i>value-added tax suspension</i> . An organization that has been granted a temporary exemption from paying VAT. |
| warehouse | A PeopleSoft data warehouse that consists of predefined ETL maps, data warehouse tools, and DataMart definitions. |
| work order | In PeopleSoft Services Procurement, enables an enterprise to create resource-based and deliverable-based transactions that specify the basic terms and conditions for hiring a specific service provider. When a service provider is hired, the service provider logs time or progress against the work order. |
| worker | A person who is part of the workforce; an employee or a contingent worker. |
| workset | A group of people and organizations that are linked together as a set. You can use worksets to simultaneously retrieve the data for a group of people and organizations and work with the information on a single page. |
| worksheet | A way of presenting data through a PeopleSoft Business Analysis Modeler interface that enables users to do in-depth analysis using pivoting tables, charts, notes, and history information. |
| worklist | The automated to-do list that PeopleSoft Workflow creates. From the worklist, you can directly access the pages you need to perform the next action, and then return to the worklist for another item. |
| XML schema | An XML definition that standardizes the representation of application messages, component interfaces, or business interlinks. |
| yield by operation | In PeopleSoft Manufacturing, the ability to plan the loss of a manufactured item on an operation-by-operation basis. |
| zero-rated VAT | Abbreviation for <i>zero-rated value-added tax</i> . A VAT transaction with a VAT code that has a tax percent of zero. Used to track taxable VAT activity where no actual VAT amount is charged. Organizations that supply zero-rated goods and services can still recover the related input VAT. This is also referred to as exempt with recovery. |

Index

A

- accessing log files
 - examples 13
- accessing Server Administration Workbench
 - iSeries 169
 - Windows 103
- accessing the Web Server Monitor 154
- accounts
 - Windows 86, 328
- adding a host to the PATROL console 180
- adding a user to the administrators group
 - printer setup
 - Windows 86
- adding administrators
 - database security
 - iSeries 40
- adding an EnterpriseOne server for querying using SAW for Windows 103
- adding OWKM files to the PATROL environment 180
- additional documentation xx
- administering batch processes
 - iSeries 16, 17
 - encoding passwords 18
 - monitoring processes 17
 - reviewing batch output files 16
 - scheduling reports 17
 - UNIX 55
 - monitoring processes 57
 - reviewing batch output files 56
 - running reports 58
 - UNIX & Linux
 - scheduling reports 58
 - Windows 90, 91, 92
- administrators 3
 - hardware 3
 - third-party software 3
 - Windows
 - printer setup 86
- AIX
 - clusters 201, 204
 - EnterpriseOne setup 202
 - HACMP software 200
 - kernel parameter settings 70

- maxuproc 69
 - system parameters 70
 - tune parameters 70
 - viewing system parameters 70
- application fundamentals xix
- applications
 - consultants 3
 - developers 3
- applying security on the iSeries 20
- architecture
 - enterprise servers 76
 - UNIX 49
- authorization lists
 - iSeries
 - database security 41
 - samples 37
- automatic start
 - enterprise servers 88
 - iSeries 12
 - UNIX (HP 9000) 52
 - UNIX (RS/6000) 53
 - UNIX (Sun Solaris) 53

B

- backing up EnterpriseOne tables 235
 - creating a backup 243
 - enterprise server requirements 236
 - Oracle on UNIX or Windows 245
 - restoring a backup file
 - iSeries 246
 - Oracle on UNIX or Windows 246
 - SQL Server 247
 - SQL Server on Windows 247
 - SQL Server 245
- backing up servers
 - EnterpriseOne tables and object owner IDs 238
 - requirements 235, 236
- batch output files
 - iSeries 16
 - UNIX
 - reviewing 56
 - Windows 91
- batch processes
 - iSeries 16

- administration 16
- encoding passwords 18
- monitoring processes 17
- running reports 17
- scheduling reports 17
- log files 292
 - enterprise server 300, 305, 309
 - workstation event flow 292
 - workstation setup 292
- troubleshooting 283
- UNIX
 - administration 55
 - monitoring processes 57
 - running reports 58
 - scheduling reports 58
- Windows 90, 91, 92, 93
- batch queues
 - settings for starting batch queues 80
- BSFNLIB security parameter 25
- business functions
 - security files 94
 - UNIX 67
 - troubleshooting development workstations 289

C

- checking EnterpriseOne server configurations using SAW for Windows 104
- checking EnterpriseOne server connectivity using SAW for Windows 105
- Citrix Metaframe
 - incorporating with Terminal Server Edition 186
- cleaning up the enterprise server
 - iSeries 14
 - cleanup 14
 - Windows 89
- clusters 195
 - AIX 201, 204
 - EnterpriseOne setup 202
 - HACMP software 200
 - HP-UX 195, 198
 - MC/ServiceGuard 197
 - multiple instances of
 - EnterpriseOne 196
 - Oracle package 197
 - troubleshooting 208
 - registering EnterpriseOne with SUNClustering 207
 - Solaris 206
 - Windows
 - failover and failback parameters 216
 - setting up EnterpriseOne 212
 - troubleshooting 217
- Clusters
 - Windows
 - advanced options for EnterpriseOne resources 216
- comments, submitting xxiv
- common elements xxiv
- compile error log
 - workstation 291
- Configurable Network Computing consultant 3
- Configurable Network Computing Implementation guide 1
- configuring EnterpriseOne server monitoring settings using SAW for Windows
 - selecting a server, port, and processes to monitor 107
 - setting up event notification 107
 - setting up run-time parameters 108
- configuring OWKM 177
- configuring SAW 126
- configuring server monitoring settings using SAW for Windows 106
- configuring the Web Server Monitor 143
- consultants
 - application 3
 - CNC 3
 - custom solution 3
 - hardware, network, and third-party software 3
- contact information xxiv
- CPU binding for AIX 70
- creating a backup
 - iSeries 243
 - Oracle on UNIX or Windows 245
 - SQL Server 245
- creating an application server
 - AIX 204
 - cluster resources 205
 - creating an EnterpriseOne subsystem
 - iSeries 28
- creating the OUTQ
 - iSeries printer 15
- cross-references xxiii
- Customer Connection website xx

D

- data sources
 - troubleshooting
 - Server Map 312
- databases
 - iSeries
 - security 29
 - setting up security 39
 - UNIX
 - increasing the maxprocess limit 52
 - starting the enterprise server 52
- debugging
 - enterprise server 310
- deleting a SnapShot file 250
- deleting server log files
 - iSeries 170
 - UNIX 172
- deployment servers
 - backup requirements 235
- developers
 - application 3
- development workstation
 - troubleshooting
 - business functions 289
 - event rules 289
- development workstations
 - troubleshooting 288
- directory structures
 - UNIX 47
- disk space
 - displaying resources 170
- disk striping
 - AIX 70
- displaying disk space resources
 - iSeries 170
- displaying log files for server processes
 - UNIX 173
- displaying server log files
 - iSeries 170
 - UNIX 172
- displaying user profile information
 - database security
 - iSeries 41
- DLL files
 - troubleshooting 287
- documentation
 - printed xx
 - related xx
 - updates xx

domains

Windows 84

- DTAPATH datapath security parameter 22

E

editing the owenv script

AIX 202

editing the server jde.ini file

UNIX 175

email

troubleshooting 335

enterprise server components

monitored by OWKM 179

enterprise server initialization

iSeries 7

enterprise servers

backup requirements 236

creating a separate instance 71

debugging 310

iSeries 10, 12

architecture 5

backup requirements 236

cleaning up the server 14

initialization 7

installation problems 313

interprocess communication

(IPC) 321

jde.ini file 322

manual start 10

PPAT 319

setting up a printer 15

shutting down the server 12

Linux 54

automatic start 54

log files

batch processes 305

batch processing 300

jde.log 300

jdedebug.log 302

logic processing 299

viewing files 307

monitoring using KM commands 181

running multiple instances 71

Server Map data source

troubleshooting 312

troubleshooting 309

types of problems 299

UNIX 51, 54

architecture 49

- automatic start (HP 9000) 52
- automatic start (RS/6000) 53
- automatic start (Sun Solaris) 53
- backup requirements 237
- directory structure 47
- initialization 50
- jde.ini file 324
- jde.ini file security 68
- manual start 52
- setting up a printer 55
- Windows 75, 83, 85, 87, 88, 89, 94
 - backup requirements 237
 - testing by submitting report 332
- EnterpriseOne
 - backing up tables 235
- UNIX
 - directory structure 47
- EnterpriseOne DB admin profile security
 - parameter 24
- error messages
 - memory violations 280
 - workstations 279
 - communication failures 281
 - form and grid add failures 280
- event rules
 - troubleshooting 289
- executable file security
 - UNIX 67
 - Windows 94

F

- failover and failback parameters
 - Windows cluster 216
- failures
 - communication 281
 - form and grid add 280
- file descriptors
 - kernel parameter settings
 - HP-UX 65
 - Sun Solaris 65
- file security
 - UNIX
 - business function files 67
 - EnterpriseOne executables 67
 - jde.ini file (enterprise server) 68
 - specification files 67
 - Windows 94
- files
 - restoring a backup file
 - SQL Server 247

- viewing batch output files 18

G

- generating serialized web objects 258
- global tables
 - workstations 290
- glossary 339

H

- HACMP software
 - AIX 200
 - application server 204
 - cluster resources 205
 - EnterpriseOne setup 202
 - groups 201
 - installation 201
 - Oracle database 201
 - owenv script 202
 - start resource control script 203
 - stop resource control script 204
- hardware
 - administrators 3
 - consultants 3
- HP 9000
 - enterprise server
 - automatic start 52
- HP-UX
 - clusters 195
 - EnterpriseOne package 198
 - multiple instances of
 - EnterpriseOne 196
 - Oracle package 197
 - troubleshooting 208
 - kernel parameter settings 61
 - file descriptors 65
 - message queues 62
 - processes 65
 - semaphores 63
 - shared memory 64

I

- identifying an EnterpriseOne instance 20
 - on the iSeries 19
- implementation methodology
 - roles 3
- INILIB security parameter 22
- initialization
 - troubleshooting 79
 - UNIX 50

- installation
 - EnterpriseOne 54
 - HACMP software 201
 - Oracle database 201
 - verifying for Windows 90
- IPC resources
 - Linux 66
- IPCS
 - setup for database security
 - iSeries 40
- iSeries
 - administering servers 5
 - authorization lists
 - database security 41
 - samples 37
 - backing up tables 243
 - batch output files 16
 - batch processes 16, 17
 - administration 16
 - encoding passwords 18
 - monitoring processes 17
 - scheduling reports 17
 - database security
 - administrators 40
 - removing administrative authority 41
 - user profile information 41
 - enterprise server 10, 14
 - architecture 5
 - automatic start 12
 - backup requirements 236
 - email 319
 - initialization 7
 - installation problems 313
 - interprocess communication (IPC) 321
 - JDBNET 320
 - jde.ini file 322
 - manual start 10
 - process flow 5
 - setting up a printer 15
 - shutting down the server 12
 - EnterpriseOne database security 29
 - integrated file system logging support 13
 - jde.log 14
 - jddebug.log 14
 - library structure 8
 - restoring a backup file 246
 - Server Administration Workbench 168

- accessing SAW 169
- disk space resources 170
- server log files 170
- setting up EnterpriseOne database security 39
- iSeries database security parameters 21
 - BSFNLIB 25
 - DTAPATH datapath 22
 - EnterpriseOne DB admin profile 24
 - INILIB 22
 - modifying system profile 23
 - modifying JDE profile 23
 - modifying security profile 24
 - secure log path 25

J

- JDBNET
 - troubleshooting
 - iSeries 320
- jde.ini file
 - [NETWORK QUEUE SETTINGS] 91
 - troubleshooting
 - iSeries 322
 - UNIX 324
 - UNIX
 - changing the server port number 175
 - [NETWORK QUEUE SETTINGS] 56
 - security 68
 - turning on or off debugging 175
 - Windows 80, 94
- jde.log
 - disabling 294
 - enterprise server 299, 300
 - setup 307
 - iSeries 14
 - locating the file 296
 - workstation 290
 - setup 293
- jdecpy.log
 - workstation 291, 297
- jddebug.log
 - clearing the file 14
 - disabling 294
 - enterprise server 299, 302, 304
 - disabling 303
 - enabling 303
 - naming conventions 305
 - server locations 304
 - setting up 308

- location on a workstation 294
- workstation 291
- jdeinst.log
 - definition 291
- JDENET
 - iSeries
 - troubleshooting 317
- JITI
 - reducing the frequency 191

K

- kernels
 - AIX 70
 - maxuproc 69
 - parameter settings 69
 - tune parameters 70
 - viewing system parameters 70
 - HP-UX
 - file descriptors 65
 - message queues 62
 - parameter settings 61
 - processes 65
 - semaphores 63
 - shared memory 64
 - Linux
 - parameter settings 66
 - Sun Solaris
 - file descriptors 65
 - message queues 62
 - parameter settings 61
 - processes 65
 - semaphores 63
 - shared memory 64

L

- libraries
 - iSeries 8
- Linux
 - administering servers 47
 - batch processes
 - scheduling reports 58
 - enterprise server 54
 - architecture 49
 - starting automatically 54
 - file limits 66
 - file security 60
 - IPC resources 66
 - kernel parameter settings 66
 - listing batch output files

- UNIX 56
- loading OWKM files 180
- locating the jde.log 296
- log files
 - clearing jde.log 14
 - clearing jdedebug.log 14
 - enterprise server
 - batch processes 305
 - batch processing logs 300
 - jde.log 300
 - jdedebug.log 302
 - viewing logs 307
 - enterprise server batch process log 300
 - enterprise servers 299
 - EnterpriseOne workstations 290
 - jde.log
 - locating 296
 - jdedebug.log 291
 - locating 294
 - viewing 293
 - workstation
 - application development 291
 - batch processes 292
 - compile errors 296
 - global tables 290
 - jde.log 293
 - jdecpy.log 297
 - jdeinst.log 298
 - logic processing logs 290, 299
 - sql.log 297, 298
 - troubleshooting 290
 - troubleshooting strategies 292
 - viewing server logs 307
 - workstation batch process log 309
- logging
 - integrated file system logging support
 - iSeries 13
- logic processing logs
 - enterprise server 299

M

- maintaining file security
 - executable files 94
 - Linux
 - overview 60
 - UNIX
 - business function files 67
 - jde.ini (enterprise server) 68
 - overview 60
 - specification files 67

- Windows 94
- maintaining multiple instances of EnterpriseOne
 - HP-UX 196
- manual start
 - enterprise servers
 - iSeries 10
 - UNIX 52
 - Windows enterprise servers 89
- MC/ServiceGuard
 - setting up an Oracle package 197
- memory
 - kernel parameter settings
 - HP-UX 64
 - Sun Solaris 64
- memory violations
 - workstations 280
- message queues
 - kernel parameter settings
 - HP-UX 62
 - Sun Solaris 62
- MMA Partners xx
- modify JDE profile security parameter 23
- modify security profile security parameter 24
- modify system profile security parameter 23
- modifying the server jde.ini file
 - iSeries 170
 - UNIX 174
- monitoring batch processes
 - iSeries 17
 - UNIX 57
 - Windows 91
- monitoring enterprise server statistics using SAW for Windows 110
 - printing log and debug log files 113
 - turning logging on and off 113
 - viewing log and debug log files 112
 - viewing process user statistics 113
 - viewing statistics for batch processes 114
- monitoring enterprise servers
 - using KM commands 181
- monitoring EnterpriseOne
 - Knowledge Module 176, 178
 - adding a host to the PATROL console 180
 - adding OWKM files to the PATROL environment 180
 - loading OWKM files 180
 - monitoring business functions running on connected enterprise servers 145
 - monitoring database connection pools 146
 - monitoring JDENET connection pools 145
 - monitoring the Web server system 148
 - monitoring Web server log files 151
 - monitoring Web server users 144
 - troubleshooting the web server 153
 - Web Server Monitor 143
- monitoring EnterpriseOne enterprise server statistics using SAW for Windows
 - changing process settings 112
 - viewing statistics for server disks 114
- monitoring EnterpriseOne server processes 108
- monitoring EnterpriseOne servers
 - from the web 126
- monitoring EnterpriseOne web server statistics using SAW for Windows 115
- monitoring EnterpriseOne with Server Administration Workbench
 - iSeries 168
 - accessing SAW 169
 - IPC and disk space resources 168
 - modifying the jde.ini file 170
 - server log files 170
 - UNIX 171
 - accessing SAW 172
 - server log files 172
 - Windows 103
- monitoring EnterpriseOne with the Web Server Monitor 143
 - configuring 143
 - monitoring JDENET connection pools 145
 - monitoring Web server users 144
- monitoring server process statistics using SAW for Windows 110
- monitoring virtual clients 150
- monitoring web server CallObjects 116
- monitoring web server statistics using SAW for Windows 115
- monitoring Web server statistics using SAW for Windows

- monitoring Web server user connection statistics 119
- Monitoring web server statistics using SAW for Windows
 - monitoring web server connection statistics 116
 - monitoring web server user logs 121
- monitoring Web server user statistics 117
- multi-user mode
 - restrictions with Terminal Server Edition 187

N

- network
 - administrators 3
 - consultants 3
 - services for Windows 87, 88
 - Terminal Server Edition 187
- notes xxiii

O

- object owner IDs
 - EnterpriseOne tables 238
- ODBC
 - troubleshooting
 - sql.log 298
- Oracle
 - backing up tables on UNIX or Windows 245
 - HACMP software 201
 - restoring a backup file 246
- output files
 - viewing batch output files 18
- OUTQ
 - creating for iSeries printer 15
 - starting for iSeries printer 15
- ownev 206

P

- Package Management guide 2
- packages
 - EnterpriseOne for a cluster 198
- parameter settings
 - AIX 70
 - maxuproc 69
 - tune parameters 70
 - viewing system parameters 70
 - HP-UX memory 64
 - Sun Solaris memory 64

- passwords
 - encoding passwords 18
- PDF
 - iSeries output location 16
- PeopleBooks
 - ordering xx
- PeopleCode, typographical
 - conventions xxii
- PeopleSoft application fundamentals xix
- performance
 - AIX 70
 - kernel parameter settings 70
 - Terminal Server Edition 188
- PORTTEST
 - iSeries 315
 - Windows
 - verifying the EnterpriseOne installation 215
- prerequisites xix
- printed documentation xx
- printers
 - iSeries 15
 - creating the OUTQ 15
 - printing multiple copies to a remote printer 15
 - setup 15
 - troubleshooting 287
- UNIX
 - setup 55
- Windows 83, 84, 85, 86
 - administrators 86
 - setup problems 329
- printing log and debug log files using SAW for Windows 113
- printing multiple copies to a remote printer 15
- printing server log files
 - iSeries 170
 - UNIX 172
- processes
 - iSeries
 - monitoring batch processes 17
 - kernel parameter settings
 - HP-UX 65
 - Sun Solaris 65
 - UNIX
 - monitoring batch processes 57
 - Windows 91
- production workstation
 - preliminary troubleshooting 283

troubleshooting 282

Q

- querying enterprise servers using SAW for Windows
 - querying enterprise server disk components 124
 - querying enterprise server IPC resources 123
 - querying enterprise server log files 125
 - querying enterprise server process components 122
 - querying enterprise server users 123
- Querying enterprise servers using SAW for Windows
 - querying the enterprise server job queue 125
- querying EnterpriseOne enterprise servers using SAW for Windows 122
- queue services
 - Windows 87, 88

R

- reducing JITI frequency 191
- related documentation xx
- remote printers
 - printing multiple copies on the iSeries 15
- removing administrative authority from user profiles
 - iSeries 41
- reports
 - batch processes 92, 93
 - iSeries 17
 - UNIX 58
 - UNIX 58
 - Windows
 - file location problems 332
- restoring a backup file
 - iSeries 246
 - Oracle on UNIX or Windows 246
 - SQL Server 247
 - SQL Server on Windows 247
- restoring a SnapShot file 250
- reviewing batch output files
 - iSeries 16
 - UNIX 56
 - Windows 91
- RS/6000

- enterprise server
 - starting automatically 53
- running batch process reports from the server command line
 - UNIX 58
- running multiple instances of EnterpriseOne on the iSeries 19
- running SnapShot from the command line 250

S

- saving EnterpriseOne in a SnapShot file 249
- SAW for Windows
 - configuring EnterpriseOne server monitoring settings 107
 - monitoring data on enterprise servers connected to the web server 115
 - monitoring enterprise server statistics
 - printing log and debug log files 113
 - turning logging on and off 113
 - viewing process user statistics 113
 - viewing statistics for batch processes 114
 - monitoring EnterpriseOne enterprise server statistics 112
 - monitoring server process statistics 110
 - viewing log and debug log files 112
 - viewing statistics for server disks 114
 - monitoring web server CallObjects 116
 - monitoring web server connection statistics 116
 - monitoring Web server user connection statistics 119
 - monitoring web server user logs 121
 - monitoring Web server user statistics 117
 - querying enterprise server disk components 124
 - querying enterprise server IPC resources 123
 - querying enterprise server job queue 125
 - querying enterprise server log files 125
 - querying enterprise server process components 122
 - querying enterprise server users 123
 - setting up event notification 107

- setting up run-time parameters 108
- scheduling reports
 - iSeries
 - command line 17
 - UNIX 58
 - Windows 93
- scripts
 - Solaris clustering 206
- secure log path security parameter 25
- security
 - iSeries
 - EnterpriseOne databases 29
 - setup for EnterpriseOne databases 39
 - maintaining file security for
 - Windows 94
 - specification file 94
 - UNIX
 - business function files 67
 - file security 60
 - jde.ini file (enterprise server) 68
 - specification files 67
 - Windows 94
- semaphores
 - kernel parameter settings
 - HP-UX 63
 - Sun Solaris 63
- Server Administration Workbench
 - checking EnterpriseOne server configurations 104
 - checking EnterpriseOne server connectivity 105
 - configuring server monitoring settings 106
 - jde.ini file 104
 - monitoring enterprise server statistics 110
 - querying EnterpriseOne enterprise servers 122
 - removing a server from the query list 105
 - Windows 103
 - monitoring EnterpriseOne web server statistics 115
 - Windows monitoring 108
- Server Administration Workbench (SAW) 101
 - accessing SAW 103
 - iSeries 168
 - accessing SAW 169
 - disk space resources 170
 - modifying the jde.ini file 170
 - server log files 170
 - UNIX
 - accessing SAW 172
 - editing the jde.ini file 175
 - log files 173
 - modifying the jde.ini file 174
 - monitoring EnterpriseOne 171
 - server log files 172
 - server port number 175
 - stopping server processes 174
 - turning on or off debugging 175
 - Windows 103
- Server and Workstation Administration guide 2
- Server Map data source
 - troubleshooting 312
- server monitoring interface
 - accessing 130
- servers
 - AIX
 - application server for cluster 204
 - setting value of maxuproc 69
 - backing up tables 243
 - backup requirements 236
 - deployment server 235
 - command line 17, 58
 - batch process reports for iSeries 17
 - batch process reports for UNIX 58
 - iSeries 10
 - administration 5
 - cleaning up the enterprise server 14
 - IPC and disk space resources 168
 - log files 170
 - logging support 13
 - modifying the jde.ini file 170
 - server processes 168
 - UNIX
 - administration 47
 - jde.ini file security 68
 - log files 172
 - setting up a printer 55
 - shutting down the enterprise server 54
 - starting the enterprise server 51
 - Windows 75, 83, 89
- services
 - Windows 83, 87
- setting business function file security
 - UNIX 67

- Windows 94
- setting failover and failback parameters 216
- Windows cluster 216
- setting specification file security
 - UNIX 67
- setting up a EnterpriseOne package 198
- setting up a printer
 - iSeries 15
 - printing multiple copies to a remote printer 15
 - Windows 84, 85, 86
 - administrators 86
 - troubleshooting 329
- setting up EnterpriseOne on a Windows cluster
 - failover and failback parameters 216
- setting up EnterpriseOne on the terminal server 188
- setting up iSeries EnterpriseOne database security 29, 39
 - administrators 40
 - displaying user profile information 41
 - removing administrative authority
 - user profiles 41
- shared memory
 - kernel parameter settings
 - HP-UX 64
 - Sun Solaris 64
- shutting down the enterprise server
 - Linux 54
 - UNIX 54
- single occurrence report
 - scheduling for UNIX 58
- SMC views
 - working with 128
- SMIT application
 - setting the value of maxuproc 69
- SnapShot 249
 - deleting a file 250
 - restoring a file 250
 - running SnapShot from the command line 250
 - saving EnterpriseOne in a file 249
- Solaris
 - clusters 206
 - enterprise server
 - automatic start 53
 - kernel parameter settings 61
 - file descriptors 65

- message queues 62
 - processes 65
 - semaphores 63
 - shared memory 64
- specification file security
 - UNIX 67
- SQL Server
 - backing up tables 245
 - restoring a backup file 247
 - Windows
 - restoring a backup file 247
- sql.log
 - reading 292
 - workstation 291, 297, 298
 - troubleshooting ODBC 298
 - turning on the log 297
- standalone executable files 272
- starting the enterprise server
 - iSeries 10
 - automatic 12
 - manual 10
 - Linux
 - automatic 54
 - UNIX 51, 53
 - automatic (HP 9000) 52
 - automatic (RS/6000) 53
 - database settings 52
 - manual 52
 - Windows 89
- stopping server processes
 - UNIX 174
- suggestions, submitting xxiv
- SUNClustering
 - registering EnterpriseOne 207
- SunOracleMgr.sh 206
- SunStartResource.sh 206
- SunStopResource.sh 206
- System Administration guide 1

T

- tables
 - backing up EnterpriseOne tables 235
 - backing up servers 243
 - object owner IDs 238
- Terminal Server Edition 185
 - incorporating Citrix Metaframe 186
 - network considerations 187
 - performance considerations 188
 - restrictions in multi-user mode 187
 - setting up EnterpriseOne 188

- troubleshooting 189
 - EnterpriseOne UBE output security 190
- terms 339
- testing
 - iSeries
 - PORTTEST 315
 - submitting a report 318
- third-party software
 - administrators 3
 - consultants 3
- troubleshooting
 - accessing EnterpriseOne 192
 - batch processes 283
 - data selection and sequencing criteria lost 192
 - development workstation 288
 - enterprise servers 299, 309
 - batch process log 300, 305
 - jde.log 300
 - jddebug.log 302
 - Server Map data source 312
 - UNIX 324
 - Windows 327
 - EnterpriseOne development tools 192
 - HP-UX on a cluster 208
 - import/export with Microsoft Excel 191
 - iSeries enterprise server 313, 318
 - architecture 5
 - initialization 7
 - interprocess communication (IPC) 321
 - JDBNET 320
 - jde.ini file 322
 - PPAT 319
 - running JDENET 317
 - log path is incorrect 192
 - logging in to EnterpriseOne 193
 - logging off versus disconnecting 191
 - ODBC
 - sql.log 298
 - production workstations 282
 - restarting EnterpriseOne 191
 - run-time error occurs during server connection test 192
 - setting up the jddebug.log 291
 - shortcut path is incorrect 193
 - shortcuts do not work in email messages 192
 - specification files are locked 191
 - submit UBE locally to TSE 190
 - Terminal Server Edition 189
 - EnterpriseOne UBE output security 190
 - UBE output security on TSE 189
 - UNIX enterprise server
 - architecture 49
 - directory structure 47
 - initialization 50
 - viewing log files 293
 - Windows enterprise server 75, 76, 79, 327
 - clusters 217
 - email 335
 - finding data 333
 - log file location 329
 - report file location 332
 - running EnterpriseOne manually 331
 - setting up a printer 329
 - setting up EnterpriseOne accounts 328
 - stopping EnterpriseOne as run manually 334
 - testing by submitting a report 332
 - using Visual C++ to stop processes 335
 - workstations 279, 281
 - batch process log 292
 - business functions 289
 - compile error log 296
 - DLL files (production) 287
 - event rules 289
 - interactive applications 282
 - jdecpy.log 297
 - jdeinst.log 298
 - log files 290
 - preliminary troubleshooting 283
 - printing 287
 - setting up the jde.log 293
 - sql.log 297, 298
 - standalone installation 286
 - strategies using log files 292
 - viewing server logs 307
- troubleshooting EnterpriseOne servers
 - web servers 336
- troubleshooting using the Web Server Monitor 153
- tune parameters

- AIX 70
 - kernel parameter settings 70
- turning logging on and off using SAW for Windows 113
- type security parameter 21
- typographical conventions xxii

U

- uninstalling the network and queue services
 - Windows 87
- UNIX
 - administering servers 47
 - AIX
 - kernel parameter settings 69
 - backing up Oracle tables 245
 - batch output files
 - reviewing 56
 - batch processes
 - administration 55
 - monitoring processes 57
 - running reports 58
 - scheduling reports 58
 - business function files
 - security 67
 - creating a separate instance of the enterprise server 71
 - enterprise server 51, 54
 - architecture 49
 - backup requirements 237
 - directory structure 47
 - initialization 50
 - jde.ini file 324
 - manual start 52
 - setting up a printer 55
 - EnterpriseOne executable files
 - security 67
 - file security 60
 - jde.ini file
 - security 68
 - kernel parameter settings
 - HP-UX 61
 - Linux 66
 - Sun Solaris 61
 - restoring an Oracle backup file 246
 - running multiple EnterpriseOne enterprise servers 71
 - Server Administration Workbench
 - accessing SAW 172
 - editing the jde.ini file 175
 - log files 173

- modifying the jde.ini file 174
 - monitoring EnterpriseOne 171
 - server port number 175
 - stopping a server process 174
 - turning on or off debugging 175
- specification files
 - security 67
- verifying the EnterpriseOne installation 54
- user accounts
 - Windows 85
- user profiles
 - iSeries database security 41
 - displaying information 41

V

- viewing batch output files 18
- viewing system parameters
 - AIX 70
- virtual clients
 - monitoring 150
- visual cues xxiii

W

- warnings xxiii
- Windows
 - administering the server 75
 - backing up Oracle tables 245
 - batch 90
 - batch output files 91
 - batch processes 91, 92, 93
 - business function files 94
 - clusters
 - advanced options for EnterpriseOne resources 216
 - failover and failback parameters 216
 - setting up EnterpriseOne 212
 - troubleshooting 217
 - enterprise server 75, 76, 79, 83, 87, 88, 89
 - backup requirements 237
 - EnterpriseOne on Terminal Server Edition 185
 - file security 94
 - jde.ini file 80, 94
 - printer setup 83
 - restoring an Oracle backup file 246
 - Server Administration Workbench 103, 104

- specification files 94
- SQL Server
 - restoring a backup file 247
 - verifying the EnterpriseOne installation 90
- working with EnterpriseOne network and queue services
 - Windows 89
- working with SMC views 128
- workstations
 - development (troubleshooting)
 - business functions 289
 - event rules 289
 - error messages 279
 - global tables 290
 - log files
 - application development 291
 - batch processes 292
 - compile errors 296
 - jde.log 293
 - jdecpy.log 297
 - jdeinst.log 298
 - sql.log 297, 298
 - troubleshooting strategies 292
 - viewing 293
 - logic processing logs 290
 - preliminary troubleshooting 283
 - production (troubleshooting) 282
 - batch processes 283
 - DLL files 287
 - printing 287
 - standalone installation 286