

ORACLE®

PEOPLESOFT

PeopleTools 8.53: System and Server Administration

October 2014

ORACLE®

PeopleTools 8.53: System and Server Administration
CDSKU pt853pbr1_r03
Copyright © 1988, 2014, Oracle and/or its affiliates. All rights reserved.

License Restrictions Warranty/Consequential Damages Disclaimer

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

Warranty Disclaimer

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

Restricted Rights Notice

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

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

Hazardous Applications Notice

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

Trademark Notice

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

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

Third Party Content, Products, and Services Disclaimer

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

Documentation Accessibility

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

Access to Oracle Support

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

Contents

- Preface.....xix**
 - Understanding the PeopleSoft Online Help and PeopleBooks.....xix
 - PeopleSoft Hosted Documentation.....xix
 - Locally Installed Help.....xix
 - Downloadable PeopleBook PDF Files.....xix
 - Common Help Documentation.....xix
 - Field and Control Definitions.....xx
 - Typographical Conventions.....xx
 - ISO Country and Currency Codes.....xxi
 - Region and Industry Identifiers.....xxi
 - Access to Oracle Support.....xxi
 - Documentation Accessibility.....xxii
 - Using and Managing the PeopleSoft Online Help.....xxii
 - System and Server Administration.....xxii
 - Contact Us.....xxii
 - Follow Us.....xxii
- Chapter 1: Getting Started with System and Server Administration..... 23**
 - System and Server Administration Overview..... 23
 - PSADMIN..... 23
 - Analytic Servers..... 24
 - Web Servers..... 24
 - Verity Search Indexes..... 25
 - PeopleSoft Configuration Manager..... 26
 - PeopleTools Utilities..... 26
 - Tracing and Debugging..... 27
 - Jolt Configuration Options..... 27
 - Timeout Settings..... 28
 - System and Server Administration Implementation..... 28
- Chapter 2: Understanding PeopleSoft Internet Architecture..... 29**
 - PeopleSoft Architecture Fundamentals..... 29
 - Database Server..... 30
 - Application Servers..... 30
 - Application Servers..... 31
 - Oracle Tuxedo and Oracle Jolt..... 31
 - Domains..... 31
 - PeopleSoft Server Processes..... 32
 - Services..... 33
 - Listeners, Handlers, and Queues..... 34
 - PeopleSoft Process Scheduler Server..... 36
 - Web Server..... 36
 - Web Server Software..... 37
 - PeopleSoft Servlets..... 38
 - Oracle Jolt..... 39
 - Web Browser..... 40
 - Server Configuration Options..... 41
 - Logically Separate Server Configuration..... 41
 - Physically Separate Server Configuration..... 43

Implementation Options.....	44
Chapter 3: Working with Server Domain Configurations.....	47
Understanding PS_HOME and PS_CFG_HOME.....	47
Implementing Flexible Server Installations.....	48
Applying Security Restrictions.....	48
Working with the Default PS_CFG_HOME.....	49
Locating the Default PS_CFG_HOME.....	49
Using PSADMIN with the Default PS_CFG_HOME.....	50
Working with Alternate PS_CFG_HOME Locations.....	50
Specifying Alternate PS_CFG_HOME Locations.....	50
Using the %V and %K Meta Variable.....	51
Configuring Domains in Alternate Locations of PS_CFG_HOME.....	52
Managing Domains.....	53
Working with PS_APP_HOME.....	53
Understanding PS_APP_HOME.....	53
Managing Environments with PS_APP_HOME.....	54
Working with PS_CUST_HOME.....	56
Understanding PS_CUST_HOME.....	56
Managing Environments with PS_CUST_HOME.....	57
Chapter 4: Using the PSADMIN Utility.....	61
Understanding PSADMIN.....	61
Starting PSADMIN.....	62
Using PSADMIN.....	63
Using Configuration Templates.....	63
Using the Quick-Configure Menu.....	64
Setting Domain-Level Environment Variables.....	65
Understanding Domain-Level Environment Settings.....	65
Working with Domain-Level Environment Settings.....	66
Setting Domain-Level Environment Variables.....	67
Using the PSADMIN Command-Line Interface.....	70
Understanding the PSADMIN Command-Line Interface.....	70
Using the Miscellaneous Commands.....	70
Using the Pure Internet Architecture Commands.....	70
Using the Application Server Commands.....	72
Using the Process Scheduler Commands.....	78
Using the Search Server Commands.....	82
Working with PSADMIN Exit Codes.....	82
Using PSADMIN Configuration Files.....	84
Understanding PSADMIN Configuration Files.....	84
Configuring a Domain.....	85
Loading a Configuration.....	86
Archiving Application Server Configuration Files.....	86
Booting a Domain.....	87
Stopping a Domain.....	87
Monitoring a Domain.....	87
Configuring the Application Server to Handle Cache Files and Replay Files.....	87
Chapter 5: Using PSADMIN Menus.....	91
Understanding PSADMIN Menus.....	91
Using the Application Server Administration Menu.....	92
Accessing the Application Server Options.....	92
Administering a Domain.....	93

Importing Domain Configurations.....	93
Booting a Domain.....	94
Shutting Down a Domain.....	95
Checking the Domain Status.....	95
Purging the Domain Cache.....	97
Configuring a Domain.....	99
Editing Configuration and Log Files.....	99
Creating a Domain.....	100
Deleting a Domain.....	101
Configuring an Application Server Domain to Preload Cache.....	101
Cleaning Domain IPC Resources.....	104
Using the Process Scheduler Menu.....	104
Using the Search Server Menu.....	105
Using the Web (PIA) Server Menu.....	105
Understanding PeopleSoft Web Server Administration Using PSADMIN.....	106
Accessing the Web Server Options.....	106
Creating a PIA Domain.....	107
Deleting a PIA Domain.....	110
Booting a PIA Domain.....	110
Shutting Down a PIA Domain.....	112
Checking PIA Domain Status.....	112
Configuring a PIA Domain.....	113
Editing PIA Configuration Files.....	114
Viewing Log Files.....	115
Administering a PIA Site.....	115
Deleting a PIA Site.....	117
Managing Windows Services for PIA Domains.....	118
Switching Configuration Homes (PS_CFG_HOME).....	119
Setting Up the PeopleSoft Windows Service to Start an Application Server Domain.....	120
Understanding Microsoft Windows Services.....	120
Configuring the PeopleSoft Service.....	121
Testing the Service.....	122
Editing the Service Configuration File.....	123
Chapter 6: Setting Application Server Domain Parameters.....	125
Startup Options.....	125
DBName.....	125
DBType.....	125
UserID.....	125
UserPswd.....	125
Connect ID.....	126
Connect Password.....	126
ServerName.....	126
StandbyDBName.....	126
StandbyDBType.....	126
StandbyUserId.....	126
StandbyUserPswd.....	126
Database Options.....	127
SybasePacketSize.....	127
UseLocalOracleDB.....	127
EnableDBMonitoring.....	127
PSDB Maximum Cursors.....	127

Security Options.....	128
Validate Signon With Database.....	128
DomainConnectionPwd.....	129
Workstation Listener Options.....	129
Address.....	129
Port.....	130
Encryption.....	130
Min Handlers.....	130
Max Handlers.....	130
Max Clients per Handler.....	130
Client Cleanup Timeout.....	130
Init Timeout.....	131
Tuxedo Compression.....	131
Jolt Listener Options.....	131
Address.....	131
Port.....	131
Encryption.....	131
Min Handlers.....	132
Max Handlers.....	132
Max Clients per Handler.....	132
Client Cleanup Timeout.....	132
Init Timeout.....	132
Client Connection Mode.....	132
Jolt Compression Threshold.....	132
Additional Prompt.....	133
Jolt Relay Adapter Options.....	133
Listener Address.....	133
Listener Port.....	133
Domain Settings.....	134
Domain ID.....	134
Add to PATH.....	134
Spawn Threshold.....	134
Log Directory.....	135
Restartable.....	135
Allow Dynamic Changes.....	135
LogFence.....	136
LogFieldSeparator.....	136
AppLogFence.....	137
Trace-Log File Character Set.....	137
PeopleCode Debugger Options.....	137
Trace Options.....	137
TraceSQL.....	137
TraceSQLMask.....	138
TracePC.....	138
TracePCMask.....	138
TracePPR and TracePPRMask.....	138
TracePIA and TracePIAMask.....	140
TraceAE.....	140
TraceAnalytic and Trace AnalyticMask.....	140
TracePPM.....	140
DumpMemoryImageAtCrash.....	140

DumpManagerObjectsAtCrash.....	140
Log Error Report, Mail Error Report.....	141
Cache Settings.....	141
EnableServerCaching.....	141
CacheBaseDir.....	141
ServerCacheMode.....	141
MaxCacheMemory.....	142
EnableDBCACHE.....	143
PreLoadCache and PreLoadMemoryCache.....	143
Remote Call Options.....	144
RCCBL Redirect.....	144
RCCBL PRDBIN.....	146
PSAPPSRV Options.....	146
Min Instances.....	146
Max Instances.....	146
Service Timeout.....	147
Recycle Count.....	147
Percentage of Memory Growth.....	147
Allowed Consec Service Failures.....	148
Max Fetch Size.....	149
Auto Select Prompts.....	149
AutoLoad JVM.....	149
Serial Recycle.....	149
PSANALYTICSRV Options.....	150
Min Instances.....	150
Max Instances.....	150
Analytic Instance Idle Timeout.....	150
PSSAMSRV Options.....	150
Min Instances.....	150
Max Instances.....	150
Service Timeout.....	151
Recycle Count.....	151
Allowed Consec Service Failures.....	151
Max Fetch Size.....	151
PSQCKSRV Options.....	151
Min Instances.....	151
Max Instances.....	152
Service Timeout.....	152
Recycle Count.....	152
Allowed Consec Service Failures.....	152
Max Fetch Size.....	152
Serial Recycle.....	152
PSQRYSRV Options.....	152
Min Instances.....	153
Max Instances.....	153
Service Timeout.....	153
Recycle Count.....	153
Allowed Consec Service Failures.....	153
Max Fetch Size.....	153
Use Dirty-Read.....	153
Serial Recycle.....	154

Integration Broker Server Processes.....	154
SMTP Settings.....	154
SMTPServer.....	155
SMTPPort.....	155
SMTPServer1.....	155
SMTPPort1.....	155
SMTPSender.....	155
SMTP BlackBerryReplyTo.....	155
SMTPSourceMachine.....	155
SMTPCharacterSet.....	155
SMTPEncodingDLL.....	155
SMTPGuaranteed.....	156
SMTPTrace.....	156
SMTPSendTime.....	156
SMTPUserName.....	156
SMTPUserPassword.....	156
SMTPUserName1.....	156
SMTPUserPassword1.....	156
SMTPTimeToWaitForResult.....	157
SMTPUseSSL.....	157
SMTPSSLPort.....	157
SMTPClientCertAlias.....	157
SMTPSSLPort1.....	157
SMTPUseSSL1.....	157
SMTPClientCertAlias1.....	157
SMTPDNSTimeoutRetries.....	157
Implementing NT LAN Manager.....	158
Implementing Simple Authentication and Security Layer.....	158
SMTP Further Considerations.....	158
Interface Driver Options.....	159
SCP_LOCALE.....	159
PSTOOLS Options.....	159
EnablePPM Agent.....	159
Add to CLASSPATH.....	159
JavaVM Options.....	160
Proxy Host.....	160
Proxy Port.....	161
Non Proxy Hosts.....	161
Usage Monitoring State.....	161
Character Set.....	162
Suppress App Error Box (Microsoft Windows Only).....	162
DbFlags.....	162
Suppress SQL Error.....	163
Integration Broker Options.....	164
Min Message Size for Compression.....	164
Thread Pool Size.....	164
Search.....	164
Search Indexes.....	165
PSRENSRV Options.....	165
log-severity_level.....	165
io_buffer_size.....	165

default_http_port.....	165
default_https_port.....	165
default_auth_token.....	165
PSPPMRSV Options.....	165
Min Instances.....	166
Max Instances.....	166
Select Server Process Options.....	166
Do you want the Publish/Subscribe servers configured?.....	166
Move quick PSAPPSRV services into a second server (PSQCKSRV)?.....	166
Move long-running queries into a second server (PSQRYSRV)?.....	166
Do you want JOLT configured?.....	167
Do you want JRAD configured?.....	167
Do you want WSL Configured?.....	167
Do you want to enable PeopleCode Debugging (PSDBGSRV)?.....	167
Do you want Event Notification configured (PSRENSRV)?.....	167
Do you want MCF Servers configured?.....	167
Do you want Performance Collators configured?.....	167
Do you want Analytic Servers configured?.....	167
Do you want Domains Gateway configured?.....	167
Chapter 7: Working with Oracle WebLogic.....	169
Understanding WebLogic.....	169
Working With the WebLogic Server Administration Console.....	169
Starting WebLogic.....	170
Starting WebLogic on Microsoft Windows.....	171
Starting WebLogic on UNIX.....	172
Stopping WebLogic.....	172
Stopping WebLogic Using the Administration Console.....	172
Stopping WebLogic Using the Command Line.....	173
Using WebLogic Server Administration Console to Monitor PeopleSoft Sessions.....	173
Setting Up Reverse Proxy Servers.....	174
Understanding Reverse Proxy Servers For PeopleSoft Implementations.....	175
Configuring Microsoft IIS 7.0 as an RPS.....	175
Configuring WebLogic as an RPS.....	177
Configuring Oracle Sun Java System Web Server as an RPS.....	179
Configuring Apache HTTP as an RPS.....	183
Setting The HTTP Session Timeout.....	184
Setting Authentication Failure Timeout.....	184
Enabling or Disabling HTTP Keep Alive.....	184
Changing WebLogic User Passwords.....	185
Implementing WebLogic SSL Keys and Certificates.....	187
Understanding SSL Encryption with WebLogic.....	187
Obtaining Encryption Keys.....	187
Preparing Keys and Certificates for the Keystore.....	189
Importing Keys and Certificates Into the Keystore.....	192
Configuring WebLogic SSL Encryption Keys.....	193
Enabling TLS-Only on WebLogic.....	195
Configuring TLS-Only on WebLogic Server.....	195
Configuring Browsers for TLS.....	196
Configuring Reverse Proxy Servers for TLS.....	196
Working With WebLogic Session Cookies.....	196
Securing Servlets on WebLogic.....	197

Adjusting the JVM Heap Size.....	199
Determining the Service Pack Level.....	200
Enabling HTTP Access Log.....	201
Chapter 8: Working with IBM WebSphere.....	203
Understanding WebSphere Application Server within Your PeopleSoft Implementation.....	203
Deploying PeopleSoft Applications With WebSphere.....	203
Using The Integrated Solutions Console.....	203
WebSphere Application Server Profiles.....	205
IBM HTTP Server.....	206
Starting and Stopping WebSphere Application Servers.....	206
Starting the WebSphere Server.....	206
Stopping the WebSphere Server.....	206
Managing WebSphere Server with PSADMIN.....	207
Configuring Reverse Proxy Servers For WebSphere.....	207
Understanding Reverse Proxy Servers With IBM WebSphere.....	207
Configuring IBM HTTP Server as a Reverse Proxy Server.....	208
Configuring Microsoft IIS as a Reverse Proxy Server.....	209
Configuring Oracle Sun Java System Web Server as a Reverse Proxy Server.....	211
Setting Up SSL For WebSphere.....	213
Understanding WebSphere Key Stores.....	213
Generating a Certificate Using pskeymanager.....	213
Configuring the WebSphere Container to Support SSL.....	215
Enabling TLS-Only on WebSphere.....	216
Configuring WebSphere for TLS.....	217
Configuring Browsers for TLS.....	217
Testing TLS.....	217
Configuring Reverse Proxy Servers for TLS.....	218
Securing The Administrative Console and Applications For WebSphere.....	218
Understanding WebSphere Security.....	218
Securing the Administrative Console.....	219
Configuring Application Security.....	221
Setting HTTP Session Timeout.....	224
Setting Authentication Failure Timeout.....	224
Working With JVM Heap Size.....	224
Working with Logging and Tracing Options.....	225
Enabling HTTP Access and HTTP Error Logging.....	225
Enabling General Logging and Tracing.....	225
Chapter 9: Configuring Verity Search and Building Verity Search Indexes.....	227
Understanding Verity Search Indexes.....	227
Overview of Verity Search Indexes.....	227
Types of Verity Indexes.....	228
Components of the Verity and PeopleSoft Search Architecture.....	228
Building Verity Indexes.....	230
Verity Search Index Limitations.....	230
User Search Strategies for Verity.....	231
Configuring Verity Search Options.....	232
Understanding PeopleSoft Search Configurations.....	232
Configuring Verity Search to run within the Application Server (Type-1).....	235
Configuring Verity Search to Run as a Separate Process (Type-2).....	235
Configuring a Separate Verity Search Server (Type-3).....	235
Verity Search Server Administration.....	239

Working with Verity Indexes.....	242
Understanding Common Verity Controls.....	242
Understanding Supported MIME Types.....	242
Opening Existing Verity Collections.....	244
Creating New Verity Collections.....	244
Building Record-Based Verity Indexes.....	244
Modifying Record-Based Verity Index Properties.....	245
Adding Subrecords to Verity Search Indexes.....	247
Building File System (Spider) Verity Indexes.....	248
Setting File System Options.....	249
Defining What to Index.....	250
Building HTTP Spider Verity Indexes.....	251
Defining HTTP Gateway Settings.....	252
Defining What to Index.....	253
Administering Verity Search Indexes.....	253
Specifying the Verity Index Location.....	253
Administering the Verity Search Index.....	255
Editing Properties.....	255
Scheduling Administration.....	256
Sharing Indexes Between Application Servers and PeopleSoft Process Scheduler.....	256
Modifying the VdkVgwKey Key.....	257
Building Verity-Based Registry Search Indexes.....	257
Associating Keywords with Content References.....	258
Building Search Indexes for the Portal Registry.....	258
Chapter 10: Using PeopleSoft Configuration Manager.....	261
Understanding PeopleSoft Configuration Manager.....	261
Common Elements in PeopleSoft Configuration Manager.....	261
Starting PeopleSoft Configuration Manager.....	262
Launching Configuration Manager.....	262
Managing PeopleTools 8.4x and PeopleTools 8.5x Configuration Manager Settings.....	262
Specifying Startup Settings.....	263
Specifying Display Settings.....	265
Specifying Crystal Report, Business Interlink, and JDeveloper Settings.....	267
Specifying Trace Settings.....	268
Specifying Workflow Settings.....	269
Specifying Remote Call/AE Settings.....	269
Configuring Developer Workstations.....	270
Importing and Exporting Environment Settings.....	271
Configuring User Profiles.....	271
Defining a Profile.....	272
Specifying Databases and Application Servers.....	272
Configuring Process Scheduler.....	274
Configuring nVision.....	276
Specifying Common Settings.....	278
Specifying Command Line Options.....	279
Setting Up the PeopleTools Development Environment.....	281
Understanding the PeopleTools Development Environment.....	281
Understanding the Client Setup Process.....	281
Verifying PS_HOME Access.....	281
Verifying Connectivity.....	281
Verify Supporting Applications.....	282

Using the Configuration Manager.....	282
Running the Client Setup Process.....	282
Chapter 11: Using PeopleTools Utilities.....	285
Understanding the PeopleTools Utilities.....	285
Using the System Information Page.....	285
Understanding the System Information Page.....	285
Viewing the System Information Page.....	286
Using Administration Utilities.....	288
Signon Event Message.....	289
PeopleTools Options.....	292
Message Catalog.....	303
Spell Check System Dictionary.....	305
Translate Values.....	306
Load Application Server Cache.....	307
Tablespace Utilities.....	314
Tablespace Management.....	315
DDL Model Defaults.....	315
BOE Integration Administration.....	317
Strings Table.....	318
Lookup Exclusion.....	319
XML Link Function Registry.....	319
Merchant Integration Utilities.....	319
TableSet IDs.....	319
Record Group.....	320
TableSet Control.....	320
Convert Panels to Pages.....	321
Update Utilities.....	323
Remote Database Connection.....	324
URL Maintenance.....	325
File Extension List.....	339
Copy File Attachments.....	343
Query Administration.....	345
Sync ID Utilities.....	345
nVision Report Request Admin.....	345
Analytic Server Administration.....	345
Upgrade Conversion.....	345
Analytic Model Viewer.....	345
Analytic Instance Load/Unload.....	345
Analytic Instance Create/Del/Copy/.....	345
Pre-Load Cache Utilities.....	345
Gather Utility.....	346
QAS Administration.....	347
Oracle Resource Management.....	347
Using Audit Utilities.....	348
Using the Record Cross Reference Component.....	348
Performing a System Audit.....	349
Performing Database Level Auditing.....	350
Using Debug Utilities.....	350
Using the PeopleTools Test Utilities Page.....	351
Replay Appserver Crash.....	354
Using the Trace PeopleCode Utility.....	354

Using the Trace SQL Utility.....	354
Using International Utilities.....	354
Setting International Preferences.....	355
Setting Process Field Size.....	355
Administering Time Zones.....	355
Managing Languages.....	356
Using Optimization Utilities.....	357
Using PeopleSoft Ping.....	357
PeopleSoft Ping Chart.....	358
PeopleSoft Ping Delete.....	359
PeopleSoft Ping Options.....	359
Chapter 12: Tracing, Logging, and Debugging.....	361
Working with PeopleSoft Server Process Logs.....	361
Setting Up the PeopleCode Debugger.....	362
Debugging for a Two-Tier Connection.....	362
Debugging for a Three-Tier Connection.....	363
Using the PeopleCode Debugger.....	365
Configuring PeopleCode Trace.....	365
Configuring SQL Trace.....	366
Enabling IDDA Logging.....	367
Understanding IDDA Logging.....	367
Enabling IDDA Logging.....	368
Working with IDDA Functional Categories.....	369
Configuring Logging Options.....	369
Viewing IDDA Logging Output.....	370
Chapter 13: Working with Jolt Configuration Options.....	373
Configuring Jolt Failover and Load Balancing.....	373
Configuring Weighted Load Balancing.....	373
Configuring Jolt Failover.....	373
Configuring Dynamic Reloading of the Application Server Connection String.....	374
Configuring Jolt Session Pooling.....	374
Configuring Parallel Pagelet Loading.....	374
Configuring Domain Connection Password.....	375
Using Jolt Internet Relay.....	375
Jolt Internet Relay Architecture.....	376
Implementation Considerations.....	377
Configuring JRLY.....	377
Configuring JRAD.....	379
Running Jolt Internet Relay.....	380
Appendix A: Securing PS_HOME and PS_CFG_HOME.....	383
Understanding PS_HOME and PS_CFG_HOME Security.....	383
Understanding PS_HOME Security.....	383
Understanding Minimum Access Required by the User Starting Domains.....	384
Understanding PS_CFG_HOME Security.....	385
Securing PS_HOME on UNIX.....	385
Managing a Secure PS_HOME on UNIX.....	386
Working with User Accounts.....	386
Configuring Partial PS_HOME Access.....	386
Securing PS_HOME on Windows.....	386
Multiple Administrator User Accounts.....	387
Local User Accounts.....	389

Managing a Secure PS_HOME on Windows.....	391
Working With Mapped Drives, UNC Paths, and TM_TUXIPC_MAPDRIVER.....	391
Working With Oracle ProcMGR Windows Service.....	392
Managing TM_TUXIPC_MAPDRIVER.....	392
Resolving Initialization Timeout Issues.....	392
Working With the Tuxedo TM_CPAU Setting.....	393
Implementing PS_CFG_HOME Security.....	394
Securing PS_CFG_HOME on UNIX.....	394
Securing PS_CFG_HOME on Windows.....	395
Appendix B: WebLogic Managed Server Architecture.....	397
PeopleSoft Internet Architecture Servlets and Applications.....	397
WebLogic Domain Types.....	397
Understanding WebLogic Domain Types.....	398
Single-Server Domains.....	398
Multi-Server Domains.....	399
Distributed Managed Servers.....	403
Common Default Settings.....	404
WebLogic Domain Directory Structure and Files.....	409
WebLogic Domain Directory Structure.....	409
WebLogic Domain File Listing by Type.....	409
J2EE Application Files.....	414
PIA Install and Reinstall Options.....	415
Administering a WebLogic Server Life Cycle.....	415
Understanding the WebLogic Server Life Cycle.....	415
Starting and Stopping Single-Server Processes.....	416
Starting and Stopping Multi-Server Processes.....	417
Starting and Stopping a Distributed Managed Server.....	421
Tuning Performance and Monitoring Resources.....	421
Managing JVM Heap Size.....	421
Monitoring HTTP Session Count for PeopleSoft Interaction Hub.....	423
Changing Configuration Settings.....	423
Understanding the WebLogic Server Configuration Files.....	423
Changing the WebLogicAdmin Server's Listen Ports.....	424
Changing Application and Server Deployment Targets.....	424
Appendix C: PeopleSoft Timeout Settings.....	427
Web Server Timeouts.....	427
Session-Timeout.....	430
Web Server Default System Timeout.....	430
Application Server Timeouts.....	431
Process Scheduler Timeouts.....	433
Search Server Timeouts.....	434
PIA Timeouts.....	435
Appendix D: Ensuring Session Stickiness.....	437
Working With Session Cookie Names.....	437
Using PeopleSoft Interaction Hub and Content Providers.....	437
Working With Absolute URLs.....	438
Setting The Cookie Domain.....	439
Working with Load Balancer Timeouts.....	439
Configuring The Session Cookie Name for WebSphere.....	439
Configuring The Cookie Domain for WebSphere.....	440
Configuring the Session Cookie Name for WebLogic.....	440

Configuring Session Cookie Domain for WebLogic.....	440
Appendix E: Troubleshooting Server Issues.....	441
Uploading Files Using Non-Latin Characters.....	441
Solution For UNIX.....	441
Solution For Windows.....	441
WebSphere: Port Set In the Host Header of a Request Returned Incorrectly.....	442
Scenario.....	442
Solution.....	442
Appendix F: Replicating an Installed Environment.....	445
Understanding Environment Replication.....	445
Replicating a Web Server Environment.....	446
Replicating PS_CFG_HOME.....	446
Reconfiguring Replicated Environment Management Components.....	448
Reconfiguring an Environment Management Agent.....	448
Reconfiguring the Environment Management Hub.....	449
Reconfiguring the Environment Management Viewer.....	449
Updating the GUID Value of Cloned Databases.....	450
Appendix G: Working with PSADMUTIL.....	451
Understanding PSADMUTIL.....	451
Using PSADMUTIL from the Command Line.....	451

Preface

Understanding the PeopleSoft Online Help and PeopleBooks

The PeopleSoft Online Help is a website that enables you to view all help content for PeopleSoft Applications and PeopleTools. The help provides standard navigation and full-text searching, as well as context-sensitive online help for PeopleSoft users.

PeopleSoft Hosted Documentation

You access the PeopleSoft Online Help on Oracle's PeopleSoft Hosted Documentation website, which enables you to access the full help website and context-sensitive help directly from an Oracle hosted server. The hosted documentation is updated on a regular schedule, ensuring that you have access to the most current documentation. This reduces the need to view separate documentation posts for application maintenance on My Oracle Support, because that documentation is now incorporated into the hosted website content. The Hosted Documentation website is available in English only.

Locally Installed Help

If your organization has firewall restrictions that prevent you from using the Hosted Documentation website, you can install the PeopleSoft Online Help locally. If you install the help locally, you have more control over which documents users can access and you can include links to your organization's custom documentation on help pages.

In addition, if you locally install the PeopleSoft Online Help, you can use any search engine for full-text searching. Your installation documentation includes instructions about how to set up Oracle Secure Enterprise Search for full-text searching.

See *PeopleTools 8.53 Installation* for your database platform, "Installing PeopleSoft Online Help." If you do not use Secure Enterprise Search, see the documentation for your chosen search engine.

Note: Before users can access the search engine on a locally installed help website, you must enable the Search portlet and link. Click the Help link on any page in the PeopleSoft Online Help for instructions.

Downloadable PeopleBook PDF Files

You can access downloadable PDF versions of the help content in the traditional PeopleBook format. The content in the PeopleBook PDFs is the same as the content in the PeopleSoft Online Help, but it has a different structure and it does not include the interactive navigation features that are available in the online help.

Common Help Documentation

Common help documentation contains information that applies to multiple applications. The two main types of common help are:

- Application Fundamentals

- Using PeopleSoft Applications

Most product lines provide a set of application fundamentals help topics that discuss essential information about the setup and design of your system. This information applies to many or all applications in the PeopleSoft product line. 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 help. They provide the starting points for fundamental implementation tasks.

In addition, the *PeopleTools: PeopleSoft Applications User's Guide* introduces you to the various elements of the PeopleSoft Pure Internet Architecture. It also explains how to use the navigational hierarchy, components, and pages to perform basic functions as you navigate through the system. While your application or implementation may differ, the topics in this user's guide provide general information about using PeopleSoft Applications.

Field and Control Definitions

PeopleSoft documentation includes definitions for most fields and controls that appear on application pages. These definitions describe how to use a field or control, where populated values come from, the effects of selecting certain values, and so on. If a field or control is not defined, then it either requires no additional explanation or is documented in a common elements section earlier in the documentation. For example, the Date field rarely requires additional explanation and may not be defined in the documentation for some pages.

Typographical Conventions

The following table describes the typographical conventions that are used in the online help.

<i>Typographical Convention</i>	<i>Description</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.
... (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.
⇒	This continuation character has been inserted at the end of a line of code that has been wrapped at the page margin. The code should be viewed or entered as a single, continuous line of code without the continuation character.

ISO Country and Currency Codes

PeopleSoft Online Help topics use International Organization for Standardization (ISO) country and currency codes to identify country-specific information and monetary amounts.

ISO country codes may appear as country identifiers, and ISO currency codes may appear as currency identifiers in your PeopleSoft documentation. Reference to an ISO country code in your documentation does not imply that your application includes every ISO country code. The following example is a country-specific heading: "(FRA) Hiring an Employee."

The PeopleSoft Currency Code table (CURRENCY_CD_TBL) contains sample currency code data. The Currency Code table is based on ISO Standard 4217, "Codes for the representation of currencies," and also relies on ISO country codes in the Country table (COUNTRY_TBL). The navigation to the pages where you maintain currency code and country information depends on which PeopleSoft applications you are using. To access the pages for maintaining the Currency Code and Country tables, consult the online help for your applications for more information.

Region and Industry Identifiers

Information that applies only to a specific 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 region-specific heading: "(Latin America) Setting Up Depreciation"

Region Identifiers

Regions are identified by the region name. The following region identifiers may appear in the PeopleSoft Online Help:

- 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 the PeopleSoft Online Help:

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

Access to Oracle Support

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

Documentation Accessibility

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

Using and Managing the PeopleSoft Online Help

Click the Help link in the universal navigation header of any page in the PeopleSoft Online Help to see information on the following topics:

- What's new in the PeopleSoft Online Help.
 - PeopleSoft Online Help accessibility.
 - Accessing, navigating, and searching the PeopleSoft Online Help.
 - Managing a locally installed PeopleSoft Online Help website.
-

System and Server Administration

This book includes several chapters relating to administration tools for the PeopleSoft Pure Internet Architecture, including application servers and web servers. This content assists you in configuring and maintaining your server environment.

Note: PeopleSoft supports a number of versions of UNIX and Linux in addition to Microsoft Windows. Throughout this book, there are references to operating system configuration requirements. Where necessary, this book refers to specific operating systems by name (for example, Solaris, HP/UX, Linux, and so on). However, for simplicity the word UNIX is used to refer to all UNIX-like operating systems, including Linux.

Contact Us

[Send us your suggestions](#) Please include release numbers for the PeopleTools and applications that you are using.

Follow Us



Get the latest PeopleSoft updates on [Facebook](#).



Follow PeopleSoft on [Twitter@PeopleSoft_Info](#).

Chapter 1

Getting Started with System and Server Administration

System and Server Administration Overview

This section discusses:

- PSADMIN.
- Analytic servers.
- Web servers.
- Verity Search indexes.
- PeopleSoft Configuration Manager.
- PeopleTools utilities.
- Tracing and debugging.
- Jolt Internet Relay.
- Timeout settings.

PSADMIN

You use PSADMIN for managing application server domains, PeopleSoft Process Scheduler domains, integration server processes, search domains, and so on. PSADMIN also enables you to configure and manage the behavior of servers with respect to a wide range of PeopleTools infrastructure elements, including:

- Tuxedo and Jolt.
- PeopleCode debugging.
- Caching.
- Analytic server framework.
- Transactional SQL requests.
- Performance enhancement.
- PeopleSoft Query.
- Integration Broker.

- Email.
- Real time event notification.
- Performance Monitor.
- MultiChannel Framework.

You launch and run PSADMIN using a command line interface.

Related Links

[Understanding PSADMIN](#)

[Understanding PSADMIN Menus](#)

Analytic Servers

The *analytic server framework* provided by PeopleSoft is a general server infrastructure designed to meet the needs of PeopleSoft products that process large amounts of data in memory. It provides a stateful model of client/server connectivity that these products require to be part of the PeopleTools system, by keeping track of configuration settings, transaction information, and other data for a session. For example, client software could request that an analytic model or optimization model be recalculated in one transaction, then retrieve the results of the calculation on that model at a later time. A server process handles these requests, and maintains the model state and calculated data in memory between the requests. Additional transactions can then modify the model and perform recalculations on it without shuffling all of the data between the client and the server or dumping all the data to a database, thus preserving in-memory performance.

When a program doesn't "maintain state" or when the infrastructure of a system prevents a program from maintaining state, it's known as a *stateless* program or system. It can't take information about the last session into the next session, such as settings the user makes or conditions that arise during processing. All session state is maintained by the client and is transferred to the server with each request. As long as an application server is up and running, a user's session remains active and functional, and any application server can perform requested transactions.

However, with some products, such as Analytic Calculation Engine or PeopleSoft Optimization Framework, running a calculation on a multi-dimensional model is likely to produce far more data than is reasonable to shuttle between a client and server to maintain a stateless connection. For performance reasons, the calculations are performed completely in memory. If these calculations were to be synchronized and stored in the database so that a stateless connection could be maintained, performance would suffer significantly.

Web Servers

PeopleSoft supports Oracle WebLogic and IBM WebSphere web servers, which provide the same basic functionality to support PeopleSoft applications, including a console interface, secure sockets layer (SSL), and reverse proxy servers (RPS).

Each web server has its own way of accomplishing its functionality, and each adds its own extra features that you might find useful to your PeopleSoft system. This PeopleBook provides supplemental information about configuring and administering the web servers where it has particular relevance to PeopleSoft.

Note: The information in this PeopleBook is not intended to replace any Oracle WebLogic or IBM WebSphere documentation. Always refer to the manufacturer's documentation for detailed information about your web server.

Related Links

[Understanding WebLogic](#)

[Understanding WebSphere Application Server within Your PeopleSoft Implementation](#)

Verity Search Indexes

A search index is a collection of files that is used during a search to quickly find documents of interest. You build a search index to enable searching on a given set of documents. The set of files that make up the index is a *collection*. This collection contains a list of words in the indexed documents, an internal documents table containing document field information, and logical pointers to the actual document files.

Important! The Verity search engine is not supported with use of the PeopleSoft Search Framework. If you intend to use the PeopleSoft Search Framework, or any of the features based on the Search Framework, such as Global Search or Keyword Search, you must install Oracle Secure Enterprise Search and configure PeopleSoft Search Framework. If you are implementing PeopleSoft applications at the 9.2 version level or greater, Verity is no longer supported. Verity can be used only by PeopleSoft applications 9.1x and below.

Fields contain metadata about a document. For example, Author and Title might be fields in an index. *VdkVgwKey* is a special field that identifies each document and is unique to all of the documents in the collection.

Every search index can be modified by changing the configuration files that are associated with the index. These configuration files are known as *style* files and reside in the style directory under the database directory. A typical configuration of style files define fields for a particular index.

PeopleSoft software supports these types of search indexes:

- Record-based indexes.

Record-based indexes are used to create indexes of data in PeopleSoft tables. For example, if the PeopleSoft application has a catalog record that has two fields (Description and PartID), you can create a record-based index to index the contents of the Description and PartID fields.

- HTTP spider indexes.

HTTP spider indexes index a web repository by accessing the documents from a web server. You typically specify the starting uniform resource locator (URL). The indexer walks through all documents by following the document links and indexes the documents in that repository. You can control to what depth the indexer should traverse.

- File system indexes.

File system indexes are similar to HTTP spider indexes, except that the repository that is indexed is a file system. You typically specify the path to a file directory, then the indexer indexes all documents within that folder. HTTP spider indexes and file system indexes are sometimes collectively referred to

as *spider* indexes. The indexer recognizes a wide variety of document formats, such as Word or Excel documents. Any document in an unknown format is skipped by the indexer.

Related Links

[Understanding Verity Search Indexes](#)

PeopleSoft Configuration Manager

PeopleSoft Configuration Manager is a Microsoft Windows application that simplifies development workstation administration by enabling you to adjust PeopleSoft registry settings from a central location. You can set up one development workstation to reflect the environment at your site, then export the configuration file, which can be shared among other workstations. You can also define separate profiles for connecting to different PeopleSoft databases.

Note: The PeopleSoft Configuration Manager applies only to development environment workstations, such as workstations used to launch Application Designer and Data Mover on Windows.

PeopleSoft configuration parameters are grouped on the Configuration Manager pages according to the function, feature, or tool that they control, including:

- Startup settings.
- Display settings.
- Crystal report and Business Interlink settings.
- Trace settings.
- Workflow settings.
- Remote call settings.
- Developer workstations.
- Importing and exporting environment settings.
- Defining configuration profiles.

Related Links

[Understanding PeopleSoft Configuration Manager](#)

PeopleTools Utilities

The PeopleTools utilities are a set of various configuration and administration interfaces that serve as a browser-based method of setting numerous system settings. These utilities, most of which are available through the PeopleTools Utilities menu, provide the ability to configure, maintain, or launch a wide range of features, including:

- The System Information page.
- The message catalog.

- The spell check dictionary.
- Translate values.
- Application server caching.
- SQR customization.
- Table management and sharing.
- Backward compatibility.
- Remote database connection.
- File attachments.
- Stored URLs.
- Mobile data synchronization (deprecated).
- Update tracking.
- Platform-specific database features.
- Database auditing.
- International settings.
- Optimization utilities.
- PeopleSoft Ping.

Related Links

[Understanding the PeopleTools Utilities](#)

Tracing and Debugging

You can use the PeopleCode Debugger to interactively debug a PeopleCode program's configurations of a two-tier connection to the database or a three-tier connection to the database. You can temporarily override the PeopleSoft Configuration Manager trace settings for PeopleCode and SQL programs.

Jolt Configuration Options

With Jolt, PeopleSoft provides the options of configuring load balancing, session pooling, domain connection passwords, and (for some special configurations) Jolt Internet Relay. Load balancing enables you to route requests to servers according to the ability of a server to handle a given request load. Powerful, dedicated servers can take a higher load while less powerful servers can take a lighter load. Session pooling enables user sessions to share web server connections, which is a more efficient use of system resources. The domain connection password enables you to add a connection password to the session between the PeopleSoft Internet Architecture and the application server. Jolt Internet Relay enables you to route connections from one web server to another, perhaps through a fire wall, for specific configuration or security needs.

Timeout Settings

This topic lists the delivered default timeout settings for the web server, application server, PeopleSoft Process Scheduler, search servers, and PeopleSoft Internet Architecture (PIA).

System and Server Administration Implementation

The functionality of system and server administration for your PeopleSoft applications is delivered as part of the standard installation of PeopleTools, which is provided with all PeopleSoft products.

Several activities must be completed before you administer the system and servers for your implementation:

- Install PeopleTools according to the installation guide for your database platform and operating system.
- Install your PeopleSoft application according to the installation guide for your database platform and application.
- Establish a user profile that gives you access to the tools, pages, and processes that you'll use.

Other Sources of Information

In addition to implementation considerations presented in this section, take advantage of all PeopleSoft sources of information, including the installation guides, release notes, PeopleBooks, and Red Papers.

Chapter 2

Understanding PeopleSoft Internet Architecture

PeopleSoft Architecture Fundamentals

Your PeopleSoft application runs within the PeopleSoft Internet Architecture (PIA), which requires a variety of software and hardware elements, including:

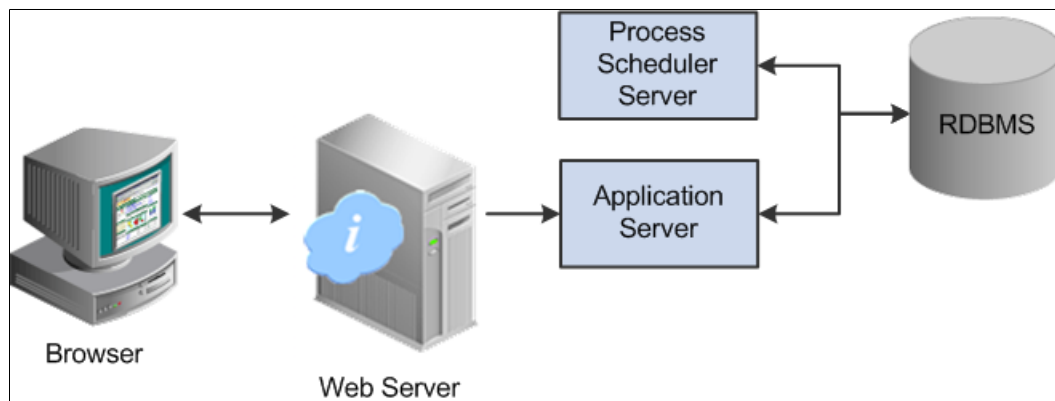
- relational database management system (RDBMS).
- application server(s).
- Process Scheduler server(s).
- web server(s).
- web browsers.

It's important to understand the role of each element before you can decide which configuration options will work best for your implementation. The following diagram illustrates, at a high level, the physical relationship between the basic elements of the PeopleSoft architecture:

Note:

Image: Distinct tiers of the architecture starting with the database, moving to the application server and Process Scheduler server, to the web server, then out to the browser

The browser sends requests to the web server, which transmits the request to the application server, which generates the SQL to run on the database.



Configuring the PeopleSoft infrastructure is not just about enabling internet application deployment through a browser. PeopleSoft enables you to take advantage of numerous PeopleSoft intranet, internet, and back-end solutions, including:

- Integration Broker, our service oriented architecture (SOA).
- PeopleSoft Interaction Hub.
- Feeds Framework.
- Search Framework.
- Performance Monitor.

All of these additional technologies stem from the basic architecture depicted in the previous diagram.

Note: Network connectivity between the PeopleSoft server elements support Internet Protocol versions 4 and 6 (IPv4 and IPv6).

Database Server

The database server houses a database engine and your PeopleSoft application database, which includes all the PeopleTools metadata, application definitions, system tables, application tables, and application data. The database server simultaneously handles the application server connections, development environment connections, and batch programs running against it.

The PeopleSoft database is the repository for all information managed by your PeopleSoft application. Both application data and PeopleSoft metadata are stored and maintained in the database. Application Designer, the main tool of the development environment, enables you to define, modify, and maintain this metadata, which the system uses to drive the runtime architecture. This collection of metadata defines a PeopleSoft application.

With Application Designer you can create dozens of different types of application definitions, such as fields, records, and pages. When an application developer saves an application definition, the system saves this definition to the metadata repository in the PeopleSoft database. At runtime, the application server retrieves the most recent application object definitions from the metadata repository, compiles and caches the application definition into memory, and runs the business rules based on the most current definitions.

Application Servers

This section discusses:

- Application servers.
- Tuxedo and Jolt.
- Domains.
- PeopleSoft server processes.
- Services.
- Listeners, handlers, and queues

- Database connectivity.

Application Servers

The application server is the core of PeopleSoft Internet Architecture. It runs business logic and submits SQL to the database server. An application server consists of numerous PeopleSoft server processes, grouped in domains. Each server process within a domain provides unique processing abilities, enabling the application server to respond effectively to a multitude of transaction requests generated throughout the PeopleSoft architecture.

Application servers require database connectivity software installed locally to maintain the SQL connection with the RDBMS. You must install the required connectivity software and associated utilities for your RDBMS on any server on which you intend to run the PeopleSoft Application Server.

After the application server establishes a connection to the database, any device that initiates a transaction request through the application server takes advantage of the application server's direct connection to the database.

Oracle Tuxedo and Oracle Jolt

PeopleSoft uses Oracle Tuxedo, a middleware framework and transaction monitor, to manage transactions between the application server and the database. PeopleSoft also uses Oracle Jolt, a Java API and class library, to facilitate communication between Tuxedo running on the application server and the PeopleSoft software running on the web server. Tuxedo and Jolt are required elements of the PeopleSoft application server.

Note: Tuxedo doesn't actually perform processing on the application server; it schedules PeopleSoft server processes to perform the transactions.

Domains

A domain is the collection of server processes, supporting processes, and resource managers that enable the database connections required to fulfill application requests. You manage each domain with a separate configuration file, and you configure each application server domain to connect to a single database. A single application server machine can support multiple application server domains running on it. You configure an application server domain using the PSADMIN utility.

There can be a one-to-one or a many-to-one relationship between application server domains and a database. In the simplest case, you configure a single application server domain to connect to a single PeopleSoft database. In a more sophisticated environment, you may configure multiple application server domains, with each domain connecting to the same PeopleSoft database. The opposite is not valid; a single application server domain cannot be used to connect to multiple PeopleSoft databases.

For example, suppose you have installed three application databases. In this case, you must configure at least three application server domains, one for each database. As demand increases, you may need to configure multiple application server domains per database, for redundancy, fail-over, and performance reasons.

You can configure multiple application server domains under a single PeopleSoft configuration home directory, or *PS_CFG_HOME*. In this context, *PS_CFG_HOME* refers to the PeopleSoft high-level directory on the application server. *PS_CFG_HOME* is the directory to which you installed the PeopleSoft application server configuration files creating a domain.

PSADMIN creates a directory beneath *PS_CFG_HOME*\appserv for each application server domain that you configure. For example, suppose you create three domains: HCMDMO1, HCMDMO2, and HCMDMO3. In this case, PSADMIN creates subdirectories \HCMDMO1, \HCMDMO2, and \HCMDMO3 beneath the *PS_CFG_HOME*\appserv directory.

PeopleSoft Server Processes

When you boot an application server domain, it starts the set of server processes associated with that domain. Numerous server processes run in a domain. Each server process establishes a persistent connection to a PeopleSoft database, and this connection acts as a generic SQL pipeline that the server process uses to send and receive SQL. Each server process uses its own, exclusive, SQL connection to facilitate requests from multiple sources. From the RDBMS perspective, each server process within a domain represents a connected user.

A server process is executable code that receives incoming transaction requests. The server process carries out a request by making calls to a service, such as MgrGetObject. The server process waits for the service to complete, then returns information to the device that initiated the request, such as a browser. While a server process waits for a service to complete, other transaction requests wait in a queue until the current service completes. A service may take a fraction of a second to complete or several seconds, depending on the type and complexity of the service. When the service completes, the server process is then available to process the next request in the corresponding queue.

The number of server processes of each type is configurable and typically varies within a domain, depending on the requirements of your application or the main purpose of the domain. For example, if a domain's primary function is to handle application requests, you might see more of the server processes devoted to that task within that domain, such as PSAPPSRV, PSQCKSRV, and PSQRYSRV. Likewise, if a domain's primary function is to handle Integration Broker SOA requests, you might see more of the server processes devoted to that task in the domain, such as the messaging server processes.

You need to configure only those server processes that your implementation requires per domain. The minimum server processes that a domain requires are PSAPPSRV, PSSAMSRV, and PSWATCHSRV.

Note: PSWATCHSRV is a process that always starts in a domain. It is not an optional server process.

You can configure multiple instances of the same server processes to start when you boot the application server domain. This helps you to handle predicted workloads.

The following tables describes the PeopleSoft server processes. Not all of the server processes will necessarily be a part of every domain as that depends on the configuration options you select.

The required server processes are:

Server Process	Description
PSAPPSRV	This process performs functional requests, such as building and loading components. It also provides the memory and disk-caching feature for PeopleTools objects on the application server. PSAPPSRV is required to be running in <i>any</i> domain.
PSSAMSRV	This SQL application manager process handles the conversational SQL that is mainly associated with Application Designer. This process is required to be running on any domain.

Server Process	Description
PSWATCHSRV	Monitors the domain and detects any orphaned application server processes.

Other server processes that you can elect to configure based on your system requirements, include:

Server Process	Description
PSQCKSRV	Handles quick, read-only SQL requests. This is an optional process designed to improve performance by reducing the workload of PSAPPSRV.
PSQRYSRV	Handles any query run by PeopleSoft Query. This is an optional process designed to improve performance by reducing the workload of PSAPPSRV.
PMSGDSP, PMSGHND, PSPUBDSP, PSPUBHND, PSSUBDSP, PSSUBHND	Used for "Publish-Subscribe" processing for Integration Broker and the PeopleSoft service oriented architecture (SOA).
PSANALYTICSRV	Performs processing and requests required for Analytic Calculation Engine.
PSRENSRV	Enables real-time event notification (REN) used in various PeopleTools technology, such as MultiChannel Framework and report distribution.
PSUQSRV, PSMCFLOG	Used within the MultiChannel Framework to manage queues and transaction logs.
PSPPMSRV	Handles the processing of all the data recorded by the Performance Monitor.
PSDBGSRV	When debugging PeopleCode, this server process maintains an independent connection to the database to avoid conflict.

Services

When a PeopleSoft application sends a request to the application server, it sends a service name and a set of parameters, such as MgrGetObject and its parameters. Tuxedo then queues the transaction request to a specific server process that is designed to handle certain services.

When a server process boots, it advertises to the system the predefined services it handles.

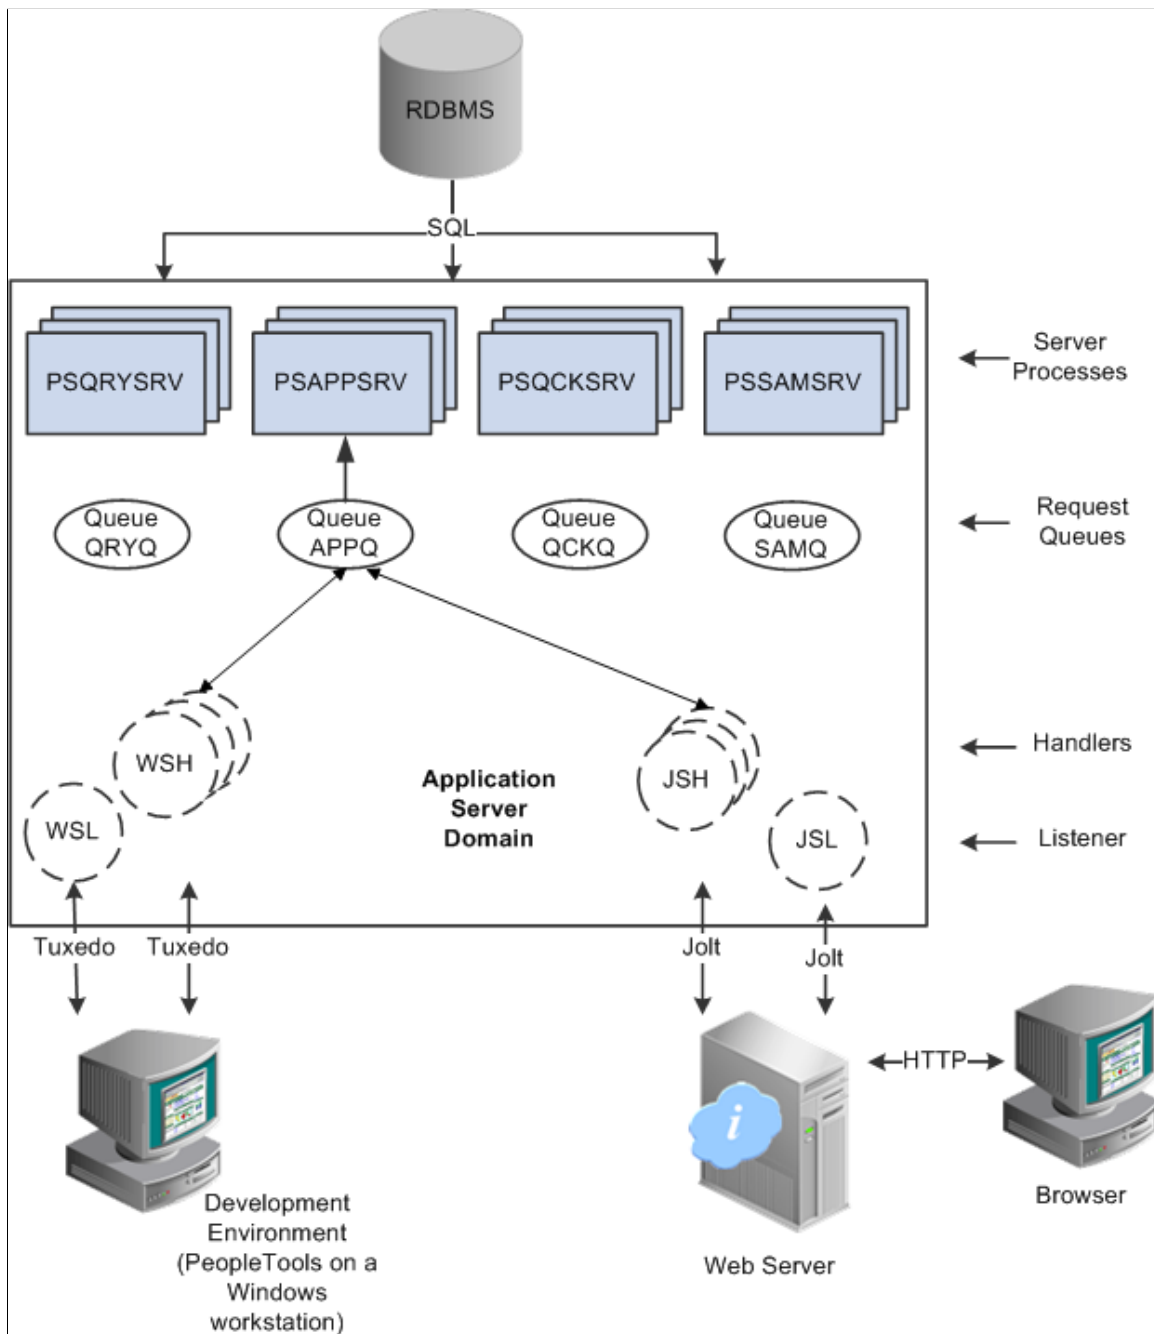
Note: When discussing PeopleSoft architecture mechanics and system features, the term *service* is used in various contexts. The following statement may help to clarify this term within the context of an application server domain: An application server domain calls server processes, such as PSAPPSRV, which in turn invoke services, such as MgrGetObject, which run against the database. The services called by domain server processes are not to be confused with the services and service operations associated with the service oriented architecture (SOA) provided by Integration Broker.

Listeners, Handlers, and Queues

Listeners, handlers, and queues provide the basis of a domain's functionality. They receive requests, route requests, store requests, monitor requests, and return request responses.

Image: Listeners route application requests to handlers where they are placed in queues monitored by server processes that submit requests to the database

The following diagram illustrates how listeners, handlers, and queues interact with incoming requests and server processes:



Note: For simplicity, the diagram does not depict every server process that runs on the application server.

The following table describes each component depicted in the previous diagram:

<i>Item</i>	<i>Description</i>
Browser	Enables internet and intranet access to all PeopleSoft applications, including many PeopleTools system administration interfaces.
Development Environment	<p>A Microsoft Windows workstation used for specific application development and system administration tasks. For example, an application developer accesses Application Designer in the development environment to modify a page or record definition. Likewise, a system administrator would access PeopleSoft Data Mover in the development environment to import a data file. Using the development environment you can connect directly to the database (a two-tier connection) or through the application server (a three-tier connection).</p> <hr/> <p>Note: Using the PeopleTools development environment in Microsoft Windows, you can connect directly to the database, or through an application server.</p> <hr/> <p>Note: No end user applications are deployed to the C++/Windows platform. All end user applications are deployed through the browser.</p> <hr/>
Workstation listener (WSL)	The workstation listener monitors Tuxedo ports for initial connection requests sent from the PeopleTools development environment (as in Application Designer). After the workstation listener accepts a connection from a workstation, it directs the request to a workstation handler. From that point, the workstation interacts only with the workstation handler to which it is assigned.
Workstation handler (WSH)	The workstation handler processes the requests it receives from the workstation listener. A unique port number identifies a workstation handler. The port numbers for the workstation handler are selected (internally by Tuxedo) from a specified range of numbers. You can configure multiple workstation handlers to start automatically if demand increases and other handlers become overloaded.
Jolt server listener (JSL)	The Jolt server listener applies only to browser requests. The Jolt server listener monitors the Jolt port for connection requests sent from the browser through the web server. After the Jolt server listener accepts a connection, it directs the request to a Jolt server handler. From that point, the browser interacts with the Jolt server handler. This is analogous to the relationship between the workstation server listener and workstation server handler.
Jolt server handler (JSH)	The Jolt server handler applies only to browser requests. The Jolt server handler processes the requests it receives from the Java server listener. The port numbers for the Jolt server handler are selected internally by Tuxedo in sequential order.

<i>Item</i>	<i>Description</i>
Request queues	Each type of server process has a service request queue that it shares with other server processes of the same type. For example, PSAPPSRV processes use the queue, APPQ, and PSQCKSRV processes use the queue, QCKQ. The workstation handler and Jolt server handler insert requests into the appropriate queue, and then the individual server processes complete each request in the order that it arrives.
Server processes	The server processes act as the core of the application server domain. They maintain the SQL connection and make sure that each transaction request gets processed on the database and that the results are returned to the appropriate origin.

PeopleSoft Process Scheduler Server

The PeopleSoft Process Scheduler environment can also be thought of as the "batch" environment. It is the location where many of your batch programs, such as Application Engine programs, run, and in most situations, this is also where you have COBOL and SQR executables installed.

In a multiserver environment, you can decide where your Process Scheduler environment resides based on available servers and performance requirements. In the PeopleSoft topology, you can install Process Scheduler on a separate server, or it can run on either the application server or the database server.

While you can install PeopleSoft Process Scheduler on any supported application server, database server, or batch server, it's important that you choose a location that's supported in your PeopleSoft environment, which varies per RDBMS platform. In most database environments, you have at least two options.

You use PSADMIN to configure and administer both the application server and PeopleSoft Process Scheduler server. The PeopleSoft Process Scheduler setup procedure in PSADMIN provides a menu-driven interface to configure PeopleSoft Process Scheduler parameters and administer the Process Scheduler server agent, and it is very similar to the PSADMIN menus for configuring an application server domain.

Even though the application server and PeopleSoft Process Scheduler have PSADMIN as a common interface, take advantage of Tuxedo middleware, and share directories under *PS_HOME* and *PS_CFG_HOME*, they are separate entities. For instance, you boot, configure, and shut down the application server and the PeopleSoft Process Scheduler server separately, and the two environments can exist on separate server machines.

Related Links

[PeopleSoft Process Scheduler](#)

Web Server

A Java-enabled web server is required to extend the PeopleSoft architecture to the internet and intranet. When you install the PeopleSoft Internet Architecture on the web server, you install a collection of

PeopleSoft Java servlets designed to handle a wide range of PeopleSoft transactions originating from the internet or an intranet.

This section discusses:

- Web server software
- PeopleSoft servlets
- Jolt

Web Server Software

PeopleTools provides and supports the following industry-standard web servers for use within your PeopleSoft implementation:

- Oracle WebLogic
- IBM WebSphere


The following software runs on the PeopleSoft web server:

Software Component	Description
Web server services	Web server services software manages the web server and provides the HTTP/S 'listener' for browser and remote system requests.
J2EE Web container	The web container (or servlet engine) is the J2EE environment in which the PeopleSoft servlets run. This component is embedded within the web server software.

Software Component	Description
Java servlets	<p>Java is the platform-independent programming language used widely for web-based programs, web applications, and servlets. Servlets are Java programs that run on the web server, unlike 'applets' which are downloaded to a client browser to run. During the PeopleSoft installation, a variety of PeopleSoft Java servlets are installed on the web server.</p> <hr/> <p>Note: By default, PeopleTools installs PIA domains into the current PS_CFG_HOME location.</p> <hr/>

Image: Web server running web server services and the web container where PeopleSoft servlets are installed

The following diagram explains about Web server running the web server services and the web container that in turn contain PeopleSoft Servlet installations.



PeopleSoft Servlets

The following PeopleSoft servlets reside on the web server:

Servlet	Description
PORTAL	<p>The portal servlet handles all of the requests and formatting for the users accessing PeopleSoft through PeopleSoft portal technologies. It manages various functionality, such as browser requests, search, content management, and homepage personalization.</p>

Servlet	Description
PSIGW	The gateway servlet transmits service requests and responses between defined nodes as part of the PeopleSoft Service Oriented Architecture. The gateway handles PeopleSoft-to-PeopleSoft, PeopleSoft-to-third party, and third party-to-PeopleSoft services.
PSEMHUB	Supports the Environment Management Framework, used by various lifecycle management tools, such as Change Assistant, to collect and monitor system information.
Report Repository	This report repository servlet enables users to easily access and distribute the output of batch reports, such as Crystal and SQR, run through PeopleSoft Process Scheduler. This servlet retrieves the report output in the report repository and serves it to the browser.
PSINTERLINKS	Used with the deprecated product PeopleSoft Business Interlinks.

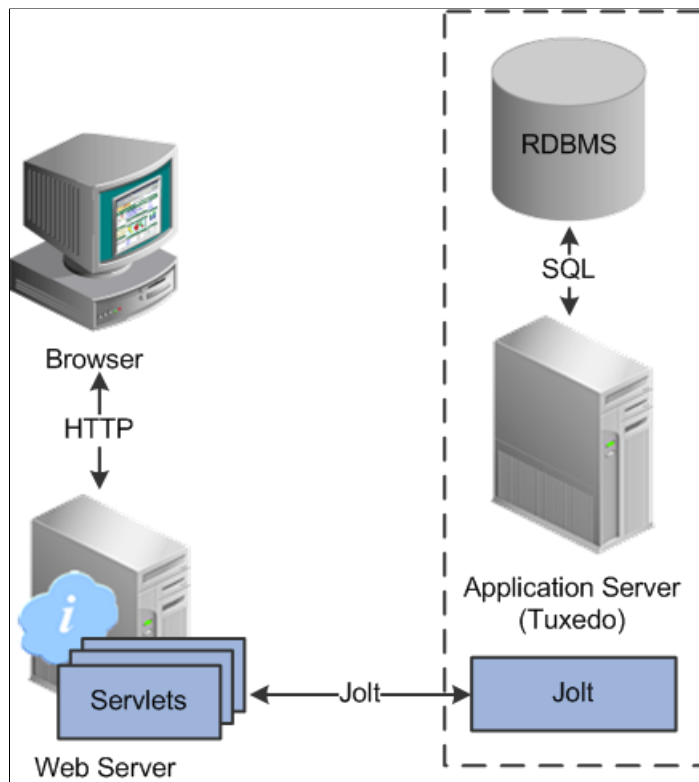
Oracle Jolt

When you install Tuxedo, Jolt gets installed by default. The PeopleSoft servlets on the web server transmit requests and data through a connection to Jolt, which runs on the application server. Jolt extends Tuxedo capabilities by acting as the communication layer between the Java-based environment of the

servlets and the C++ environment of the application server. You configure the servlets to direct requests from the web server to a predefined Jolt port on the application server.

Image: PeopleSoft servlets on the web server sending messages to the application server through Jolt

The following diagram shows the relationship between Jolt and servlets on the web server:



Web browsers and integrated systems don't send requests directly to the application server. Instead, they send HTTP/S requests to the PeopleSoft servlets running on the web server. The web server translates the HTTP/S request into a Jolt request that is sent to a specified Jolt port. Then the application server, running on Tuxedo, submits the appropriate SQL to the database.

Web Browser

A supported web browser is the primary means by which end users access PeopleSoft applications and system administrators access administrative tools. You do not need to install other software on the workstation running the browser, such as connectivity software or downloaded applets.

The PeopleSoft system sends only the following to the browser:

- HTML
- JavaScript
- Cookies

Because the browser processes only this basic internet content, the client workstation is not burdened with unnecessary processing responsibility. All processing occurs at the server level.

After the system authenticates a user during signon, PeopleTools security deploys web browser cookies to store a unique access token for each user session. As the user navigates around the system, perhaps to a separate PeopleSoft application, the token in the browser cookie is used to reauthenticate the user and bypass the sign-in process. The browser authentication cookie is:

- an in-memory cookie, never written to disk.
- encrypted to prevent snooping.
- check-summed to prevent tampering.

Note: To access PeopleSoft applications, the browser option to allow session cookies *must* be enabled.

Server Configuration Options

Depending on a variety of variables, such as your server operating system, your hardware resources, and your site's performance requirements, you can configure your environment to support physically separate or logically separate servers. In some cases, the PeopleSoft standard installation procedure recommends one or the other depending on, for example, the combination of your database type and operating system. Any platform-dependent configuration requirements are discussed in your PeopleTools installation documentation.

See *PeopleTools 8.53 Installation* for your platform.


Logically Separate Server Configuration

A logically separate server environment means that multiple servers share the same physical machine. The servers are *logically*, but not *physically*, separate.

The following diagram depicts a logical configuration with two server machines. One server contains the web server, and the other contains the application server and database server:

Image: Logical separation between the RDBMS and the application server with the two elements running on the same physical server machine

The following diagram depicts a logical configuration with two server machines. One server contains the web server, and the other contains the application server and database server:



The solid line surrounding the application server and the database server represents one physical machine.

Although this diagram depicts a separate web server, the web server software could also reside on the same machine with both the application server and the database server. When installing multiple PeopleSoft architecture elements onto the same machine, the only requirement is that each element be supported by the underlying operating system.


If all servers are located on the same machine, however, you should consider security and performance issues. If you're deploying PeopleSoft applications to the internet, you will most likely want your web server outside of your network firewall and not on the same machine as the database server. Generally, having your application server on the same physical machine as the database server provides the best performance as this configuration has no network layer between the application server and the database.

Note: For development, testing, or training purposes, you might want to have all PeopleSoft architecture elements (or as many as possible) on the same server machine.

Physically Separate Server Configuration

Image: Physical separation between web server, application server, and RDBMS server with each server residing on different physical server machines

A physically separate server configuration means that the web server, application server, and database server each reside on separate machines. The following diagram depicts a physically separate server configuration:




If the application server and database server don't reside on the same machine, then the application server and the database server should be connected to the same high-performance, backbone network. This ensures optimum performance.

Regarding performance, one advantage to keeping the architecture elements separate is that as demand on your system increases, you can add more server machines at each 'tier' to house more application servers or web servers in a modular fashion. If the server elements are installed onto a single machine, your options of increasing system performance are limited by the resources of that machine.

Image: Physical separation between web server, application server, and RDBMS promotes scalability, enabling the addition of multiple web servers and application servers to meet increased system demand

The following diagram illustrates how you can add multiple web servers and application servers when increased demand has pushed previous hardware to its limits.



Within a PeopleSoft system, you can configure multiple PeopleSoft Internet Architecture installations on a web server as well as configure multiple domains on a single application server machine. If you reach the resource limits of a single server machine, you can incorporate more server machines to house

additional PeopleSoft Internet Architecture installations or application server domains. Incorporating multiple, physically separate server machines provides increased:

- scalability
- fail-over
- availability

Implementation Options

Once you have the basic PeopleSoft Internet Architecture elements configured, you can elect to incorporate additional PeopleTools technology to run within that architectural foundation, including:

Technology	Description
Portal	<p>Oracle PeopleTools provides a variety of portal options ranging from the base portal technology provided by PeopleTools, to PeopleSoft Interaction Hub and Portal Pack solutions. These portal solutions enable you to provide simple end user navigation, aggregate web-based content, share content with other portals, and more.</p> <p>See the product documentation for <i>PeopleTools: Portal Technology</i>.</p> <p>See the product documentation for <i>PeopleSoft Interaction Hub PeopleBooks</i>.</p> <p>See Your PeopleSoft application portal pack documentation.</p>
Integration Broker	<p>Integration Broker provides a service oriented architecture (SOA) for exposing PeopleSoft business logic as services and consuming external web services for PeopleSoft applications to invoke. Integration Broker supports synchronous and asynchronous messaging with other PeopleSoft applications and with third-party systems.</p> <p>See the product documentation for <i>PeopleTools: PeopleSoft Integration Broker</i>.</p>
Performance Monitor	<p>Performance Monitor enables you to capture and store detailed performance information for any transaction occurring within your PeopleSoft implementation. You can monitor performance in real-time as well as analyze stored, historical data to help identify trends and problem areas.</p> <p>See the product documentation for <i>PeopleTools: Performance Monitor</i>.</p>

Technology	Description
MultiChannel Framework	<p>MultiChannel Framework delivers an integrated infrastructure to support multiple interaction channels for call center agents or other PeopleSoft users who must respond to incoming requests and notifications on these channels. MultiChannel Framework supports voice, email, web-based chat, instant messaging, and generic event channels.</p> <p>See the product documentation for <i>PeopleTools: PeopleSoft MultiChannel Framework</i>.</p>
Analytic Calculation Engine	<p>Analytic Calculation Engine enables application developers to define both the calculation rules and the display of calculated data within PeopleSoft applications for the purposes of multidimensional reporting, data editing, and analysis.</p> <p>See the product documentation for <i>PeopleTools: Analytic Calculation Engine</i>.</p>
Search server	<p>Configuring a search server enables you to off-load Verity search processing onto a remote server, rather than configuring search capabilities per domain. In this configuration, Tuxedo routes search requests from application server domains to the search domain running on the remote search server. Multiple application server domains may use the same search server to execute search requests.</p> <p>See Configuring Verity Search Options.</p>

Chapter 3

Working with Server Domain Configurations

Understanding PS_HOME and PS_CFG_HOME

On any server that you install the PeopleTools software, the installation program installs the required files for that server into one high-level directory structure, referred to as PS_HOME. After creating a domain, the configuration files associated with that domain reside in a directory structure referred to as "PS_CFG_HOME".

By default, the system separates the binary files (executables and libraries) stored in PS_HOME from the ASCII files (configuration and log files) associated with a domain stored in PS_CFG_HOME. This separation of the binary and ASCII files applies only to these servers:

- PeopleSoft Application Server.
- PeopleSoft Process Scheduler Server.
- PeopleSoft Search Server.
- PeopleSoft Internet Architecture.

Note: Decoupling binary files from ASCII files does not apply to any other PeopleSoft servers, such as file servers or database servers.

The following table describes the two file types and provides examples of these types within the PeopleSoft server.

<i>Location</i>	<i>File Type</i>	<i>Description</i>	<i>PeopleSoft Examples</i>
PS_HOME	Binary	Compiled, non-modifiable executables and libraries.	PSADMIN.EXE PSAPPSRV.EXE PSAESRV.EXE PSAPPENG.DLL PSODBC.DLL
PS_CFG_HOME	ASCII	Text files associated with the configuration and administration of a domain that can be viewed, modified, or generated by the system.	PSAPPSRV.CFG PSAPPSRV.CFX PSTUXCFG APPSRV_<DATE>.LOG TUXLUG.<DATE>

Note: While some ASCII files are present in PS_HOME, these should be considered non-modifiable. For example, all of the files that are in PS_CFG_HOME when it is initially created, come from the associated PS_HOME.

The decoupling of these file types enables system administrators to:

- Streamline and provide more flexible PeopleSoft server installations.
- Apply unique security restrictions to the binary file and configuration file locations.


Note: Although domains always contain their base template, the CFX and UBX templates (such as small, medium, and large) remain in PS_HOME. This means that when you create a new domain the template that you choose comes from PS_HOME\appserv, not PS_CFG_HOME\appserv.

Implementing Flexible Server Installations

Image: Multiple PS_CFG_HOME locations, containing multiple domains, all referencing the same PS_HOME installation on a remote server

With the binary files separate from the domain configuration, you have the option of installing multiple domains on multiple separate servers all leveraging the binary files of a single PS_HOME.

This diagram illustrates an example to implement multiple PS_CFG_HOME locations, containing multiple domains, all referencing the same PS_HOME installation on a remote server



Creating domains on separate servers enables you to:

- Install the PS_HOME binaries a single time.
- Incorporate additional server machines for additional domains per demand and performance requirements.
- Apply updates and upgrades to a single PS_HOME, reducing the time required for upgrading your system.

Applying Security Restrictions

With the server binary files and configuration files in separate locations, you can now apply uniform restrictions per file type. For example, some sites might prefer to have the binary files under read-only

security, while providing write access to the configuration files to specific users for administrative tasks and certain server processes, such as logging.

See [Understanding PS_HOME and PS_CFG_HOME](#).

Working with the Default PS_CFG_HOME

This section discusses how to:

- Locate the default PS_CFG_HOME.
- Use PSADMIN with the default PS_CFG_HOME.

Locating the Default PS_CFG_HOME

When you launch PSADMIN, if a PS_CFG_HOME does not exist, the system creates the PS_CFG_HOME directory in the “user” directory of the current user (the owner of the domain). The system assumes the presence of the following environment variables:

<i>Operating System</i>	<i>Required Environment Variable</i>
UNIX/Linux	HOME
Windows	USERPROFILE

For example, depending on the operating system of the server, the system creates PS_CFG_HOME in the following location on the same drive as PS_HOME.

<i>Operating System</i>	<i>PS_CFG_HOME Location</i>
UNIX/Linux	\$HOME/psft/pt/<version>
Windows	%USERPROFILE%\psft\pt/<version>

After you create a domain, the domain exists under \$PS_CFG_HOME\appserv\<>domain>.

With a user of *tsawyer*, on UNIX/Linux this would appear as:

```
/home/tsawyer/peopletools/8.53
    peopletools.properties
    /appserv
        /CRM
        /HR
        /PRCS
            /CRM_PRCs
            /HR_PRCs
        /Search
    /webserv
    /ver_dom
```

With a user of *tsawyer*, on Windows this would appear as:

```
C:\Documents and Settings\tsawyer\psft\pt\8.53
    peopletools.properties
    \appserv
```

```
\CRM
\HR
\PRCS
  \CRM_PRCs
  \HR_PCRS
\Search
\Temp
\webserv
pswinsrv.cfg
```

Note: The previous examples show a situation in which CRM, HR, CRM_PRCs, HR_PRCs and ver_dom are all domain directories. They are not in PS_CFG_HOME by default, and appear only after the domains are created.

To display the default PS_CFG_HOME location, you can submit the following command to PSADMIN:

```
psadmin -defaultPS_CFG_HOME
```

Note: These commands are not case sensitive.

Using PSADMIN with the Default PS_CFG_HOME

Launching PSADMIN requires no extra steps or variables defined when domains exist on the same machine as PS_HOME in the default PS_CFG_HOME location.

When you launch PSADMIN, it will either create (if one doesn't exist) PS_CFG_HOME or search for PS_CFG_HOME, based on the current environment settings.

To start PSADMIN, the following conditions need to be fulfilled:

- PS_CFG_HOME must be a valid location. That is, the base directory must exist (UNIX), or the drive must exist (Windows).
- The PS_CFG_HOME must be writeable, and the user running PSADMIN must have write access to the PS_CFG_HOME location.

Working with Alternate PS_CFG_HOME Locations

Domains can exist on the same physical machine or a different physical machine than where the PS_HOME directory resides. That is, multiple domains on different machines can leverage a single installation of the PeopleSoft binary files (PS_HOME) installed on a location accessible through your network.

This section discusses how to:

- Specify alternate PS_CFG_HOME locations.
- Configure domains in alternate locations of PS_CFG_HOME.

Specifying Alternate PS_CFG_HOME Locations

The value of the PS_CFG_HOME environment variable determines where PSADMIN installs a domain. If you accept the default location of PS_CFG_HOME, this environment variable does not need to be

specified. However, if you intend to install your domains in a different location, you need to override the default by setting the `PS_CFG_HOME` environment variable prior to launching PSADMIN.

For UNIX and Linux, you can:

- Set the `PS_CFG_HOME` environment variable in the `psconfig.sh` file. However, if you hard code the `PS_CFG_HOME` environment variable in `psconfig.sh` other users will not be able to use the same `PS_HOME` for different `PS_CFG_HOMES`. All users of the same `PS_HOME` would invoke the same `psconfig.sh`, and therefore see their domains in the same `PS_CFG_HOME`.
- Set `PS_CFG_HOME` in the startup script for the platform/shell, resulting in the automatic inheritance of the environment variable when a user signs on.

For Windows, you can set `PS_CFG_HOME` through the Control Panel in the user variables interface, or issue a `SET` command prior to starting PSADMIN. For example,

```
C:\>SET PS_CFG_HOME=n:\psftdomains
C:\>cd pt850\appserv
C:\pt850\appserv>psadmin
```

In this case, any domains created during this PSADMIN session would be created in `n:\psftdomains\appserv`.

Note: If you elect to operate in the same fashion as previous PeopleTools releases (where the configuration files and the binary files exist within the same directory structure) set `PS_CFG_HOME = PS_HOME`.

Using the %V and %K Meta Variable

Use the `%V` meta variable if you wish to keep all of your `PS_CFG_HOMES` together without needing to set `PS_CFG_HOME` each time you install a new PeopleTools version.

For example, you could set the `PS_CFG_HOME` environment variable as follows:

```
PS_CFG_HOME=C:\PeopleTools\installs\%V
```

If using the `%V` meta variable, assume you have installed two versions of PeopleTools: PT8.50 and PT8.51. In this case, PSADMIN automatically maps the `%V` to the PeopleTools version, creating the following `PS_CFG_HOME` locations:

```
C:\PT\installs
    \8.50-00
    \8.51-00
```

Use the `%K` variable if you wish to store `PS_CFG_HOMES` by patch level. That is, while `%V` stores the files by version (8.50-00, 8.51-00), `%K` stores files down to the patch level. For example, using `%K`, a `PS_CFG_HOME` may look similar to:

```
C:\PT\installs
    \8.53
        \8.53-01
        \8.53-03
```

Configuring Domains in Alternate Locations of PS_CFG_HOME

When using PSADMIN with alternate locations of PS_CFG_HOME, make note of the requirements discussed in this section.

Working with Remote PS_HOME Locations

If you intend to install your domains on a separate machine from where PS_HOME resides, keep these items in mind:

- The network location to where PS_HOME resides must be mapped to from the machine where PS_CFG_HOME resides.
- If a domain references a PS_HOME on a mapped drive, problems may arise if the drive letter or path is changed after initially creating the domain. If a drive mapping changes, the domain definitions may reference invalid path information. In the event of mapped drive changes, shut down the existing domains and reconfigure (or recreate) the domains.
- If PS_HOME resides on a remote server, the operating system of the PS_HOME server and the PS_CFG_HOME server must match. PeopleTools binaries are native to a platform.

Installing Necessary Components

While you can leverage a single, remote installation of PeopleTools, the server on which PS_CFG_HOME resides must have any additional components installed locally, such as Tuxedo, database drivers, ODBC connectivity information, and so on, depending on your implementation.

Ensuring that PS_CFG_HOME is Set Appropriately

Only domains installed within the current PS_CFG_HOME directory can be administered by PSADMIN. That is, the list of domains to administer that PSADMIN displays depends on the value of the PS_CFG_HOME variable. PSADMIN does not aggregate a domain list across multiple locations.

Assume that you have domains installed under these two PS_CFG_HOME directories:

- N:\psftdomains\appserv
- M:\psftfomains\appserv

Assume also that on server N domain 1 and domain 2 are installed, and on server M domain 3 and domain 4 are installed.

If you launch PSADMIN from server N, where PS_CFG_HOME, by default, is set to N:\psftdomains\appserv, you will only be able to view and administer domain 1 and domain 2.

Domains created on machine N, can only be configured on machine N. There are settings in the domain PSTUXCFG file that bind a domain to its host. You can not boot or configure a domain from a different machine.

Note: The user creating and configuring domains in PSADMIN must have write access to the PS_CFG_HOME location.

Note: You can have multiple PS_CFG_HOME locations on the same host server, but you will need to make sure the PS_CFG_HOME is set appropriately prior to starting PSADMIN, or by using the Switch Config Home PSADMIN option.

Managing Domains

When working with decoupled PS_HOME and PS_CFG_HOME directory structures, keep these recommendations in mind:

- Use distinct PS_CFG_HOME locations for each PS_HOME. When a PS_CFG_HOME is created, PSADMIN brings content from the current PS_HOME into the PS_CFG_HOME, which effectively binds that PS_CFG_HOME to that PS_HOME. That is, all domains within a given PS_CFG_HOME should reference the same PS_HOME. Domains in the same PS_CFG_HOME that reference different PS_HOMES can encounter unpredictable behavior, especially when applying updates to the system.
- Install as few PS_HOMES as possible. Ideally, each PS_CFG_HOME requiring the same PeopleTools release level points to the same PS_HOME. This approach helps reduce the number of locations that require patches to be applied.
- After applying updates (patches) to the PS_HOME, recreate domains associated with that PS_HOME.
- A domain configured on one machine cannot be copied manually and run on another machine. Use the Replicate Config Home PSADMIN option to perform this task.
- The Windows service configuration file (PSWINSRV.CFG) will be invalid if it is shared between multiple PS_CFG_HOMES on different network drives.
- As you maintain your system and upgrade domains, keep in mind that dormant domains will consume disk space. Make sure to keep track of where your retired domains reside and take action to remove them when no longer required. It is generally recommended to remove old log files and trace files for both efficiency and security purposes.

Working with PS_APP_HOME

This topic contains an overview and discusses how to manage environments where PS_APP_HOME has been implemented.

Understanding PS_APP_HOME

The PS_APP_HOME environment variable refers to the location where you have installed the contents of your PeopleSoft application. In all previous PeopleSoft releases, you installed PeopleTools into a high-level directory referred to as PS_HOME, and then you installed your PeopleSoft application into the existing PS_HOME location such that PeopleTools and your PeopleSoft application co-exist within the same directory structure.

In the current release, PeopleTools provides the option of installing your PeopleSoft application into a separate directory, outside of the PS_HOME location. The location where you install your PeopleSoft application files is represented by the PS_APP_HOME environment variable. For example, an environment where PS_APP_HOME is implemented may have the following environment variables set.

Environment Variable	Description
PS_HOME	Location where PeopleTools is installed. For example: C:\PT_853.
PS_CFG_HOME	Location where the server domain configurations are installed. For example: PS_CFG_HOME=D:\pt_server
PS_APP_HOME	Location where your PeopleSoft application files are installed. For example: PS_APP_HOME=N:\HCM_920

Important! Deciding whether to install your PeopleSoft application into a separate PS_APP_HOME should be considered while planning your installation.

See *PeopleTools 8.53 Installation* and your PeopleSoft application installation guide.

Installing your applications into PS_HOME is still assumed to be the default approach, and, if you elect to stay with this model, there is no change to your existing PeopleSoft system management experience. However, sites that take advantage of the PS_APP_HOME option can expect a variety of benefits, including:

- PeopleTools patches and updates can be applied knowing that no application content has been disturbed. Likewise, PS_APP_HOME application updates can be applied independently.
- Increased flexibility and modularity in that a single PS_HOME can be associated with multiple PS_APP_HOME locations.
- Disk space savings, due to shared PS_HOME amongst multiple PS_APP_HOME locations. Using the traditional method, each application installation requires its own PeopleTools installation.
- When troubleshooting it will be easier to isolate application and PeopleTools issues.

Managing Environments with PS_APP_HOME

If you elect to install your PeopleSoft application outside of PS_HOME, then you need to make sure to set the PS_APP_HOME environment variable prior to running PSADMIN, which can be done manually or by using a.BAT or shell script. Also, as you configure your system, you need to be mindful of the PS_APP_HOME location.

Note: If you do not set the PS_APP_HOME environment variable explicitly, PeopleSoft runtime retrieves all runtime content from PS_HOME. That is, unless you set PS_APP_HOME to a value different than PS_HOME, the system assumes PS_APP_HOME=PS_HOME.

In many of the configuration settings where PS_APP_HOME might be involved, PeopleTools includes both PS_APP_HOME and PS_HOME references. For example,

```
RCCBL PRDBIN=%PS_APP_HOME%/cblbin:%PS_HOME%/cblbin
```

In environments where PS_APP_HOME has been implemented, system elements, such as server domains have access to content supplied by both PS_HOME and PS_APP_HOME, but files added by the

PeopleSoft application installation (those residing in PS_APP_HOME) that have the same name as those in the PeopleTools PS_HOME take preference.

The following table outlines key areas in the PeopleSoft system, where you need to consider the PS_APP_HOME variable if you have installed your PeopleSoft applications into a location outside of PS_APP_HOME. When you are installing, configuring, and maintaining your system, make sure to keep these items in mind and review the documentation associated with these areas.

PeopleSoft Infrastructure Element	PS_APP_HOME Consideration
Application Server	<p>PS_APP_HOME must be set explicitly before running PSADMIN, similar to setting PS_CFG_HOME manually.</p> <p>Remote Call COBOL can look for COBOL programs in both PS_APP_HOME or PS_HOME, so be sure to review those settings.</p> <p>See RCCBL PRDBIN.</p>
Process Scheduler	<p>PS_APP_HOME must be set explicitly before running PSADMIN, similar to setting PS_CFG_HOME manually.</p> <p>The Process Scheduler domains reference PS_APP_HOME and PS_HOME locations for the following processes:</p> <ul style="list-style-type: none"> • COBOL • Crystal Reports • SQR • nVision <p>See "Understanding the PeopleSoft Process Scheduler Configuration File" (PeopleTools 8.53: PeopleSoft Process Scheduler).</p>
Search Server (Verity)	<p>PS_APP_HOME must be set explicitly before running PSADMIN, similar to setting PS_CFG_HOME manually.</p> <p>As you would normally, if there are any style files or indexes in custom locations in your PS_APP_HOME, those need to be referenced appropriately in configuration files or domain-level settings.</p>
COBOL	<p>During your installation, you will need to perform different tasks when compiling COBOL.</p> <p>See the product documentation for <i>PeopleTools 8.53 Installation</i> "Installing and Compiling COBOL on <your operating system>"</p>
SQR	<p>When specifying locations where the system should look for SQR reports, for example, in Process Scheduler or in Configuration Manager, make sure to consider reports that may exist in PS_APP_HOME.</p>

PeopleSoft Infrastructure Element	PS_APP_HOME Consideration
Lifecycle Management	<p>For applying updates and performing upgrades, PS_APP_HOME also needs to be considered for various tasks, including:</p> <ul style="list-style-type: none"> • Crawling the environment with Environment Management Framework. • Creating change packages • Defining environments for upgrades. <p>See Change Assistant and Update Manager.</p>
Windows Development Environment	<p>When setting up Windows Development workstation profiles using Configuration Manager, you set the PS_APP_HOME variable on the Process Scheduler tab, and you'll also consider your PS_APP_HOME location when specifying:</p> <ul style="list-style-type: none"> • Crystal Reports locations. • SQR report locations. • COBOL location. • nVision directory paths. • Data Mover directories. <p>See Understanding PeopleSoft Configuration Manager.</p>

Working with PS_CUST_HOME

This topic contains an overview and discusses how to manage environments where PS_CUST_HOME has been implemented.

Understanding PS_CUST_HOME

PS_CUST_HOME (PeopleSoft Customization Home) enables you to further clarify the ownership of files within your PeopleSoft system and separate them for more efficient lifecycle maintenance and system administration. While PS_APP_HOME enables you to separate PeopleSoft application files from PeopleTools files, PS_CUST_HOME enables you to identify your site's customized files and store them in a location separate from your PeopleTools and PeopleSoft application files. This creates a clear distinction between code and files delivered by Oracle (PeopleTools and PeopleSoft applications) and code produced or customized by individual customers.

Important! You can continue to store your custom and modified files within PS_HOME or PS_APP_HOME if desired. Using PS_CUST_HOME is an option that you can take advantage of if the additional flexibility provided by PS_CUST_HOME is required.

If your site does not create custom scripts, reports, or programs, then PS_CUST_HOME would not apply to your system. However, if your site does create custom scripts, reports, or programs, or if your site modifies delivered application files, then implementing PS_CUST_HOME should be considered.

Examples of customized files that might be stored in PS_CUST_HOME include:

- Data Mover scripts.
- DAT files.
- COBOL programs.
- SQR programs.
- Crystal Reports.
- Java class files.

With your custom files stored in the PS_CUST_HOME location, you can freely take PeopleTools or PeopleSoft application upgrades and maintenance updates, knowing that your customized files will not be affected. This approach also keeps PS_HOME and PS_APP_HOME consistent with the content provided by Oracle, being an exact replication of the files as they were installed.

In addition to lifecycle management benefits, PS_CUST_HOME also offers a significant advantage with sharing PS_HOME locations among multiple PeopleSoft application systems. If you are not using PS_CUST_HOME, customizations stored in PS_HOME or PS_APP_HOME that apply to only one application make it impossible to share that PS_HOME or PS_APP_HOME location with other PeopleSoft applications. With PS_CUST_HOME, all customizations can be consolidated into a single, isolated PS_CUST_HOME, and multiple runtime environments can utilize the resources within that PS_CUST_HOME. If separate environments require different customizations, they can share the same PS_HOME and PS_APP_HOME but point to separate PS_CUST_HOME locations.

An environment where PS_CUST_HOME is implemented may have the following environment variables set.

Environment Variable	Description
PS_HOME	Location where PeopleTools is installed. For example: C:\PT_853
PS_CFG_HOME	Location where the server domain configurations are installed. For example: PS_CFG_HOME=D:\PT_SERVER
PS_APP_HOME	Location where your PeopleSoft application files are installed. For example: PS_APP_HOME=N:\HCM_920
PS_CUST_HOME	Location where your customized PeopleSoft application files are installed. For example: PS_CUST_HOME=N:\PSFT_CUSTOM

Managing Environments with PS_CUST_HOME

Managing an environment where PS_CUST_HOME exists is similar to adding PS_APP_HOME to the list of environment variables, except with PS_APP_HOME you install the application files

using the delivered installation program. As such, the application files are installed automatically into PS_APP_HOME using the standard directory structure based on PS_HOME. With PS_CUST_HOME, you need to create the directory structure manually, adding any required folder within the PS_CUST_HOME using the same structure and naming found in PS_HOME or PS_APP_HOME.

For example, if you have a custom SQR report that you want to store in PS_CUST_HOME, you need to create this directory structure in your PS_CUST_HOME and place your customized file within the \sqr directory:

```
N:\PS_CUST_HOME\sqr
```

Likewise, if you have custom SQL or Data Mover scripts, you need to create this directory structure in your PS_CUST_HOME and place your customized files within the \scripts directory:

```
N:\PS_CUST_HOME\scripts
```

If using PS_CUST_HOME, you need to make sure to set the PS_CUST_HOME environment variable prior to running PSADMIN. This can be done permanently in the environment variables interface on a Windows server, manually on the command line prior to running PSADMIN, or by using a.BAT or shell script.

Note: If you do not set the PS_CUST_HOME environment variable explicitly, PeopleSoft runtime retrieves all runtime content from PS_HOME or PS_APP_HOME, if set.

When you configure your system, you need to be mindful of the PS_CUST_HOME location. In the PeopleTools configuration settings where the system would need to find files within PS_CUST_HOME, you need to include PS_CUST_HOME as a potential location for such files. For example,

```
RCCBL PRDBIN=%PS_CUST_HOME%/cblbin;%PS_APP_HOME%/cblbin;%PS_HOME%/cblbin
```

In environments where PS_CUST_HOME has been implemented, system elements, such as server domains have access to content supplied by PS_HOME, PS_APP_HOME, and PS_CUST_HOME, but files residing in PS_CUST_HOME take preference over those with the same name in either PS_HOME or PS_APP_HOME. If PS_CUST_HOME, PS_APP_HOME, and PS_HOME all exist within a PeopleSoft system, at runtime, the system checks for files in this order:

1. PS_CUST_HOME
2. PS_APP_HOME
3. PS_HOME

The following table outlines key areas in the PeopleSoft system, where you need to consider the PS_CUST_HOME variable if you added custom files to that location. When you are installing, configuring, and maintaining your system, make sure to keep these items in mind and review the documentation associated with these areas.

PeopleSoft Infrastructure Element	PS_CUST_HOME Consideration
Application Server	<p>PS_CUST_HOME must be set explicitly before running PSADMIN, similar to setting PS_CFG_HOME manually.</p> <p>Remote Call COBOL can look for COBOL programs in PS_CUST_HOME, PS_APP_HOME, or PS_HOME, so be sure to review those settings.</p> <p>See RCCBL PRDBIN.</p>
Process Scheduler	<p>PS_CUST_HOME must be set explicitly before running PSADMIN, similar to setting PS_CFG_HOME manually.</p> <p>The Process Scheduler domains reference PS_CUST_HOME, PS_APP_HOME, and PS_HOME locations for the following processes:</p> <ul style="list-style-type: none"> • COBOL • Crystal Reports • SQR • nVision <p>See "Understanding the PeopleSoft Process Scheduler Configuration File" (PeopleTools 8.53: PeopleSoft Process Scheduler).</p>
Search Server (Verity)	<p>PS_CUST_HOME must be set explicitly before running PSADMIN, similar to setting PS_CFG_HOME manually.</p> <p>As you would normally, if there are any style files or indexes in custom locations in your PS_CUST_HOME, those need to be referenced appropriately in configuration files or domain-level settings.</p>
COBOL	<p>During your installation, you will need to perform different tasks when compiling COBOL, which are described in detail in your PeopleTools installation guide.</p> <p>If you do have custom COBOL programs or files, you must set the PS_CUST_HOME environment variable prior to compiling the source code. If PS_CUST_HOME is set, the system compiles the source code in PS_HOME, PS_APP_HOME, and PS_CUST_HOME simultaneously, despite the source code residing in separate locations.</p> <p>See <i>PeopleTools 8.53 Installation for <your platform></i>: "Installing and Compiling COBOL on <your operating system>"</p>
SQR	<p>When specifying locations where the system should look for SQR reports, for example, in Process Scheduler or in Configuration Manager, make sure to consider reports that may exist in PS_CUST_HOME.</p>

PeopleSoft Infrastructure Element	PS_CUST_HOME Consideration
Lifecycle Management	<p>For applying updates and performing upgrades, PS_CUST_HOME also needs to be considered for various tasks, including:</p> <ul style="list-style-type: none"> • Crawling the environment with Environment Management Framework. • Creating change packages • Defining environments for upgrades. <p>See Change Assistant and Update Manager.</p> <hr/> <p>Note: By placing files within PS_CUST_HOME, you are taking complete responsibility for maintaining the source of that file and ensuring any changes are made to it for upgrade or maintenance purposes.</p>
Windows Development Environment	<p>When setting up Windows Development workstation profiles using Configuration Manager, you set the PS_CUST_HOME variable on the Process Scheduler tab, and you'll also consider your PS_CUST_HOME location when specifying:</p> <ul style="list-style-type: none"> • Crystal Reports locations. • SQR report locations. • COBOL location. • nVision directory paths. • Data Mover directories. <p>See Understanding PeopleSoft Configuration Manager.</p>

Chapter 4

Using the PSADMIN Utility

Understanding PSADMIN

PSADMIN simplifies the process of configuring and administering all of the server processes and features that are available for PeopleSoft servers. For example, you use PSADMIN to configure application server domains, Process Scheduler servers, PIA domains, and search servers.

Note: *PS_HOME* is the directory where you install PeopleTools.

Accessing Network Drives in Microsoft Windows Server

This section applies only if all of the following are true:

- You've upgraded to the current PeopleTools release, including the required Tuxedo version and rolling patch level, from PeopleTools 8.45 or older.
- You plan to administer your application server domains on Microsoft Windows Server.
- One or more PeopleSoft processes need to access a mapped network drive directly. Activities requiring this can include:
 - Using an instance of PSADMIN that was launched from the network drive.
 - Accessing a database on the network drive.
 - Outputting reports to a location on the network drive.

Note: This section applies only to Tuxedo-managed servers (application server, Process Scheduler server, and search servers). It does not apply to PIA domains.

Any PeopleSoft processes that reference mapped network drives by their drive letters in this environment must be able to find the drives, and must have appropriate permission to access them. In Windows Server, the operating system does not provide this access directly.

PeopleSoft uses Tuxedo's Oracle ProcMGR service and an associated environment variable, *TM_TUXIPC_MAPDRIVER*, to gain access to the network drives. You must configure these elements to provide the appropriate access before you start any PeopleSoft servers or other processes.

To configure access to mapped network drives:

1. Determine which shared network directories your PeopleSoft system will need to access with a drive and directory path.
2. In your Microsoft Windows system, make sure that the shared network directories are available, and grant the domain administrator privileges to access them.

3. Access the Services utility in Windows Control Panel.
4. Stop the Oracle ProcMGR service.
5. In the Oracle ProcMGR service properties, ensure that the “log on as” account is the account with domain administrator privileges for the required shared directories.
6. Define the TM_TUXIPC_MAPDRIVER environment variable to specify the drive mappings and paths of the required shared directories, with the following format:

```
drive1:=\\machine_name1\dirpath1[;drive2:=\\machine_name2\dirpath2[...]]
```

For example:

```
U:=\\myMachine\e$;V:=\\myMachine\PSFT
```

7. Start the Oracle ProcMGR service.

The service uses the value of TM_TUXIPC_MAPDRIVER to create the necessary drive mappings, and uses its own log on settings to provide your PeopleSoft system with access to those locations.

Note: Every time your PeopleSoft system needs to access a new mapped network location, you must repeat this procedure, including the new drive mapping along with the others.

Starting PSADMIN

This section assumes that you have already installed and configured the PeopleSoft application server.

See *PeopleTools Installation for your database platform*.

To start the PSADMIN utility:

1. Launch the psadmin executable.

PSADMIN is in PS_HOME\appserv.

2. Select the server that you want to configure, administer, or monitor from the PSADMIN menu.

```
-----
PeopleSoft Server Administration
-----

Config Home:  D:\PT_SERVER\8.53

1) Application Server
2) Process Scheduler
3) Search Server
4) Web (PIA) Server
5) Switch Config Home
6) Service Setup
7) Replicate Config Home
q) Quit

Command to execute (1-7, q):
```

Using PSADMIN

Using PSADMIN involves selecting the number of the menu item that reflects the action that you want to take, entering the correct number at the command line, and pressing `Enter`. However, in some cases, you may want to use the command-line options that PSADMIN offers.

See [Using the PSADMIN Command-Line Interface](#).

Each PSADMIN menu has the same look and feel. To select a menu item, enter the corresponding number at the prompt and press `Enter`. To return to the previous menu enter `q` (quit) at the prompt.

Note: Because the numbers corresponding to the PSADMIN menu commands can change as minor releases and patches provide different features and capabilities to PSADMIN, this documentation generally refers to the names of the commands rather than their menu item numbers.

Using Configuration Templates

The initial values that you see in PSADMIN are derived from the configuration template that you select when you create your domain. The delivered templates provide a range of possible implementations. Each configuration template includes a number of server processes, such as PSAPPSRV, that is sufficient for its intended load.

These are the delivered templates:

Configuration Template	Default PSAPPSRV Server Processes	Example Usage
Small	2	Use for 1–50 users.
Medium	8	Use for 50–500 users.
Large	25	Use for 500–1000 users.
Developer	2	Use for development and demonstration environments only.

Note: The usage examples above are *sample* ranges. Each site must determine the optimal template and number of server processes to suit the types of transactions of a particular user base. The delivered templates provide reasonable starting points for your implementations. You can add additional server processes through PSADMIN or by editing the PSAPPSRV.CFG file. For optimum performance, the number of users per configuration template can vary significantly from application to application.

Note: Oracle does not support creating custom CFX or UBX templates nor modifying delivered CFX or UBX templates.

Using the Quick-Configure Menu

When you create a domain for the first time, PSADMIN presents you with the most commonly changed parameters on the Quick-Configure menu, so that you can get up and running quickly. After the initial setup, you may at any time select *Configure this domain* on the PeopleSoft Domain Administration menu to access the Quick-Configure menu.

```

-----
Quick-configure menu -- domain: ORCLDMO1
-----
      Features                               Settings
      =====                               =====
1) Pub/Sub Servers      : Yes   16) DBNAME       : [ORCLDMO1]
2) Quick Server        : No    17) DBTYPE       : [ORACLE]
3) Query Servers       : No    18) UserId       : [QEDMO]
4) Jolt                : Yes   19) UserPswd    : [QEDMO]
5) Jolt Relay          : No    20) DomainID    : [TESTSERV]
6) WSL                 : No    21) AddToPATH   : [c:\oracle\product\11.1.0\dbhome_
1\bin]
7) PC Debugger         : Yes   22) ConnectID   : [people]
8) Event Notification : Yes   23) ConnectPswd: [people]
9) MCF Servers         : No    24) ServerName  : []
10) Perf Collator      : No    25) WSL Port    : [7000]
11) Analytic Servers   : Yes   26) JSL Port    : [9000]
12) Domains Gateway    : No    27) JRAD Port   : [9100]

      Actions
      =====
13) Load config as shown
14) Custom configuration
15) Edit environment settings
   h) Help for this menu
   q) Return to previous menu

```

Enter selection (1-27, h, or q):

The Quick-Configure menu shows which features are currently set for the newly created domain. The menu contains the values that are most commonly changed when setting up a demonstration or test domain.

To change the value of a parameter under Features, just enter the number corresponding to the feature to toggle the feature on or off.

To change the value of a parameter under Settings, enter the number corresponding to the setting and enter the new value at the prompt.

Note: All of the values that you change remain in effect until you modify them again.

Note: If you select a Settings parameter, then press Enter without entering a new value at the prompt, the existing value of that parameter is deleted. If you then enter *q* to quit the Quick-Configure menu, your changes are discarded, and the original values remain. However, if you load the configuration as shown, the parameter is saved without a value.

If the parameter is required, you'll see an error message indicating that the configuration could not be completed. The next time you access the Quick-Configure menu, PSADMIN reloads any empty required parameter with its original default value, just as it would appear if you were creating a new domain.

To configure the rest of the parameters that are not presented on the Quick-Configure menu, select *Custom configuration* to view the full list.

The Quick-Configure menu is not intended to replace the series of configuration sections in the custom configuration interface. In most cases, your site requires the parameters and tuning options that are available only through the full custom configuration menu. For this reason, the Quick-Configure menu is provided primarily for situations where you're setting up a demonstration domain for testing or for development purposes.

Note: When you use custom configuration, pressing Enter instead of entering a new value for a parameter does not delete the parameter's value. PSADMIN interprets Enter to mean that you want to retain the parameter's existing value. If you want to remove the value, you can enter one or more "white space" characters by pressing the space bar, or you can enter five underscore characters, and press ENTER.

To apply settings to the configuration files, you must select *Load config as shown*.

To add and specify environment variables at the domain level, select *Edit environment settings*.

For more information on editing domain environment variables, see [Setting Domain-Level Environment Variables](#).

Setting Domain-Level Environment Variables

This section provides an overview and discusses:

- Working with domain-level environment settings
- Setting domain-level environment variables.

Understanding Domain-Level Environment Settings

Environment settings can be set at the following levels:

- **Operating System:** At this level, the environment settings apply to all processes running at that time on the host, such as PS_CFG_HOME.
- **Domain:** At this level, you can specify settings that apply to only that domain. That is you can override any operating system environment settings, as well as any that come by default from the higher-level configuration file (.UBX).

Note: Setting a variable at the domain-level overrides the variable set in the parent environment. Variables like PS_HOME and PS_CFG_HOME cannot be set at the domain level.

The settings you modify or add for that domain will be saved in the *PS_ENVFILE section one of the following files, depending on the domain type:

- psappsrv.ubx
- psprcsrv.ubx
- pssrchrsv.ubx

When the configuration settings are loaded for that domain, then these settings will also be reflected in the .env file for that domain.

Note: It is not recommended to modify the .env file manually because it will be overwritten by .ubx values if the domain is ever re-configured (selecting the Load config as shown option on the Quick-configure menu). Environment settings should be set using the PeopleSoft Domain Environment Settings menu.

Working with Domain-Level Environment Settings

To set domain level environment settings, select *Edit Environment Settings* from the Quick-configure menu for an application server, Process Scheduler server.

Note: There is no Quick-configure menu for the search server (Verity), however when you configure the domain, PSADMIN prompts to determine if you want to manage your domain environment variables.

The following example displays application server domain settings.

```

-----
PeopleSoft Domain Environment Settings
-----
      Domain Name: ORCLDMO1

      TM_BOOTTIMEOUT           : [120]
      TM_RESTARTSRVTIMEOUT     : [120]
      TM_BOOTPRESUMEDFAIL      : [Y]
      FLDTBLDIR32              : [{TUXDIR}{FS}udataobj]
      FIELDTBLS32              : [jrep.f32,tpadm]
      ALOGPFX                  : [{LOGDIR}{FS}TUXACCESSLOG]
      COBPATH                  : [{PS_APP_HOME}\CBLBIN%PS_COBOLTYPE%;
{$PS_HOME}\CBLBIN%PS_COBOLTYPE%]
      INFORMIXSERVER           : [{Startup}\ServerName]
      # IPC_EXIT_PROCESS        : [1]
      IPC_TERMINATE_PROCESS     : [1]
      PATH                     : [{PS_HOME}\verity\{VERITY_OS}
\{VERITY_PLATFORM}\bin;{SPATH};{$Domain Settings\Add to PATH}]

```

- 1) Edit environment variable
- 2) Add environment variable
- 3) Remove environment variable
- 4) Comment / uncomment environment variable
- 5) Show resolved environment variables
- 6) Save
- h) Help for this menu
- q) Return to previous menu

Edit environment variable

Enables you to edit the values of the listed environment settings. After selecting this option, enter the row number of the environment setting you want to edit, and modify the value as needed.

```

Environment variable: TM_BOOTTIMEOUT
Enter value [120]:150

```

For any settings that you modify, the system displays an asterisk next to it (*).

Add environment variable

Use this option to add new environment variables for the domain. After selecting this option, enter the name of the environment variable and its value when prompted.

```
Enter name of environment variable: MY_VARIABLE
Enter value: D:\PSFT_HCM
```

For any settings that you add, the system displays an asterisk next to it (*).

Remove environment variable

Removes any of the currently listed environment settings. After selecting this option, enter the row number of the environment setting you want to remove.

Comment / uncomment environment variable

Use to comment or uncomment the environment setting. When you comment a setting, the system ignores that domain-level setting. Any default, or operating system-level settings would take effect. When you uncomment the setting, the system recognizes the domain-level setting value.

After selecting this option, enter the row number of the environment setting you want to comment or uncomment.

Environment variables that are commented (ignored by the system), have a # symbol next to it.

Show resolved environment variables

Expands the meta variables for system settings, such as PS_HOME, TUXDIR, and so on, so that you can view the actual file path.

For example, the following is unresolved:

```
: [{LOGDIR} {FS} TUXACCESSLOG]
```

And the following is resolved:

```
: [D:\PT_SERVER\8.53\appserv\FSDMO853\LOGS
\TUXACCESSLOG]
```

Save

Select to save the settings you've added or modified.

Help for this menu

Select to display online help information for this interface.

Setting Domain-Level Environment Variables

In most cases, you may not need to modify or add domain-level environment variables.

Some reasons to set a domain-level environment variable are:

- custom PeopleCode may require a domain-level environment variable.
- additional tracing may require additional Tuxedo environment variables, if specified in your Tuxedo documentation.
- a GCS analyst may request a specific environment setting for troubleshooting issues.

Note: PS_HOME, PS_CFG_HOME, PS_APP_HOME and PS_CUST_HOME environment variables *should not* be set at the domain-level and can cause undesirable results if set at the domain level.

The following table describes the default environment settings exposed at the domain level:

Server Type	Environment Setting	Description
Application Server	TM_BOOTTIMEOUT	Time allowed before Tuxedo performs a timeout for a domain boot.
	TM_RESTARTSRVTIMEOUT	Time allowed before Tuxedo performs a timeout for a server process restart.
	TM_BOOTPRESUMEDFAIL	<p>TM_BOOTPRESUMEDFAIL is an indicator to the Tuxedo system that dictates how the domain should behave if a process fails to start before TM_BOOTTIMEOUT (above) has been reached.</p> <p>When set to <i>Y</i>, no further domain processes will be started and any existing domain processes will be terminated.</p> <p>If set to <i>N</i>, the boot process will proceed. By default, this is set to <i>Y</i> because domains cannot be relied upon to function correctly if any of the domain processes fail to start in a timely fashion.</p>
	FLDTBLDIR32	<p>The directory in which field buffer descriptor files are located. This is an internal Tuxedo setting.</p> <p>Modification of this value could result in unpredictable behavior as Tuxedo will be unable to locate its buffer descriptor files.</p>
	FIELDTBLS32	<p>The list of the buffer descriptor files that are required by the domain runtime. These files are located in the directory defined by the FLDTBLDIR32 variable. This is an internal Tuxedo setting.</p> <p>Modification of this value could result in unpredictable behavior as Tuxedo may allocate buffers incorrectly.</p>
	ALOGPFX	Path to the Tuxedo access log for that domain.
	COBPATH	COBOL path for that domain.
	INFORMIXSERVER	(Applies only to Informix database connections) Specifies the default database server to which an explicit or implicit connection is made.

Server Type	Environment Setting	Description
	IPC_EXIT_PROCESS	<p>Indicates that the process should be shutdown in a graceful fashion, allowing cleanup functions to execute.</p> <p>This is analogous to normal shutdown and is achieved by using the MS C++ ExitProcess function.</p> <p>Typically this value or IPC_TERMINATE_PROCESS is set. These values should be considered to be mutually exclusive.</p>
	IPC_TERMINATE_PROCESS	<p>IPC_TERMINATE_PROCESS is an indication that the process should exit immediately.</p> <p>This is analogous to a process kill or forced shutdown. The exit handlers do not have an opportunity to run. This would typically be used in a situation where the process has failed to respond to a normal shutdown instruction.</p> <p>Typically this value or IPC_EXIT_PROCESS is set. These values should be considered to be mutually exclusive.</p>
	PATH	Specifies a set of directories where executable programs are located for that domain.
Process Scheduler Server	COBPATH	See COBPATH explanation for application server.
	INFORMIXSERVER	See INFORMIXSERVER explanation for application server.
	PATH	See PATH explanation for application server.
Search Server (Verity)	TM_RESTARTSRVTIMEOUT	Time allowed before Tuxedo performs a timeout on a server process restart.
	TM_RESTARTSRVTIMEOUTKILL	TM_RESTARTSRVTIMEOUTKILL is an indication to the Tuxedo system that a process that hangs on restart should be killed.
	IPC_EXIT_PROCESS	See IPC_EXIT_PROCESS explanation for application server.
	IPC_TERMINATE_PROCESS	See IPC_TERMINATE_PROCESS explanation for application server.
	PATH	See PATH explanation for application server.

Using the PSADMIN Command-Line Interface

This section provides an overview of the PSADMIN command-line interface and discusses how to:

- Use the miscellaneous commands.
- Use the application server commands.
- Use the Process Scheduler commands.
- Use the Search Server commands.
- Work with PSADMIN exit codes.

Understanding the PSADMIN Command-Line Interface

In some cases, you might want to use the PSADMIN command-line interface rather than starting the PSADMIN interface and navigating to a particular menu. The command line offers a direct method of carrying out certain server administration tasks. It also enables you to include PSADMIN actions in scripts, and simplifies the task of creating numerous domains that use default server settings.

PSADMIN has several variations of its basic command-line syntax for miscellaneous activities and server administration, which are described in the following sections.

Note: Before you begin using the PSADMIN commands, you should become familiar with PSADMIN and the components that it controls.

Note: When using the PSADMIN command line and specifying paths containing spaces (such as directory paths on Windows), the paths must be enclosed within quotes.

Using the Miscellaneous Commands

Following are the available miscellaneous PSADMIN commands:

Command	Description
<code>psadmin -h</code>	Displays command help and syntax.
<code>psadmin -v</code>	Displays the PSADMIN version number, as in <i>Version 8.53</i> .
<code>psadmin -env</code>	Displays your current environment variables.
<code>psadmin -defaultps_cfg_home</code>	Displays the current PS_CFG_HOME environment variable.
<code>psadmin -replicate -ch <source PS_CFG_HOME></code>	Replicates a PS_CFG_HOME and its configured domains to a new location, where they can be started without manual reconfiguration.

Using the Pure Internet Architecture Commands

The Pure Internet Architecture commands follow this basic syntax:

```
psadmin -w command - d PIA_domain
```

For example:

```
psadmin -w start - d peoplesoft
```

Command	Example	Result of the Example
start	psadmin -w start -d peoplesoft	Starts the peoplesoft domain.
shutdown	psadmin -w shutdown -d peoplesoft	Stops the peoplesoft domain.
shutdown!	psadmin -w shutdown! -d peoplesoft	Performs a forced shutdown of the peoplesoft domain.
status	psadmin -w status -d peoplesoft	Displays the status of the domain (started, stopped, and so on).
remove	psadmin -w remove -d peoplesoft	Deletes the domain.

Specifying PIA Domain Configuration Settings

You can use the command line to set selected domain settings using the following syntax:

```
psadmin -w configure - d PIA_domain -c c_set -p p_set
```

Where *c_set* (configuration settings) must be provided in the following order and format:

minHeapSize/maxHeapSize/maxThreads/authTokenDomain

And *p_set* (port settings) must be provided in the following order and format:

httpPort/httpsPort

For example:

```
psadmin -w configure - d peoplesoft -c 512/512/50/urdomain.com -p  
80/443
```

Specifying PIA Site Configuration Settings

You can use the command line to set selected site settings using the following syntax:

```
psadmin -w configure - d PIA_domain -s site -c c_set
```

The *c_set* (configuration settings) must be provided in the following order and format:

appSrvConnString/webProfile/pooling/reportPath/webProfUser/webProfPword/appSrvConnPword

c_set option	Description
appSrvConnString	Application server connect string, such as host:jolt port.

c_set option	Description
webProfile	The web profile the site will use (PROD, DEV, and so on).
pooling	Indicate whether Jolt pooling should be enabled. (Enabled/Disabled)
reportPath	The reports path, such as D:\psreports.
webProfUser	The user required for accessing the web profile.
webProfPword	The password for the web profile user.
appSrvConnPword	If using a domain connection password, enter that value.

For example:

```
psadmin -w configure -d peoplesoft -s ps -c URSERVER:9000/PROD/
Enabled/D:\psreports/PS/PS/Liverpool05
```

Using the Application Server Commands

For application server administration, PSADMIN has two syntax formats — one for creating new application server domains, and the other for administering existing domains.

Using the Application Server Create Command

Use the following syntax to create a new application server domain:

```
psadmin -c create -d domain -t template -s s_set [-p p_set] [-env env_set]
```

The **create** command creates an application server configuration file with the specified domain name, using the specified configuration template.

The *domain* parameter must be the name of an application server domain that you want to create, for example, HRDMO.

The *template* parameter must have one of the following values:

- *small*
- *medium*
- *large*
- *developer*

The *s_set* parameter is an optional string of startup values which provide initial configuration settings that you would otherwise specify on the PSADMIN application server Quick-Configure menu. You must enter the startup string as follows:

- In Windows, the values must be separated by slashes:

```
DBNAME/DBTYPE/OPR_ID/OPR_PSWD/DOMAIN_ID/ADD_TO_PATH/DB_CNCT_ID/DB_CNCT_PSWD/SE⇒
RVER_NAME/DOM_CONN_PWD/{ENCRYPT|NOENCRYPT}
```


- In UNIX, the values must be separated by percent signs:

```
DBNAME%DBTYPE%OPR_ID%OPR_PSWD%DOMAIN_ID
%ADD_TO_PATH%DB_CNCT_ID%DB_CNCT_PSWD%SERVER_NAME%DOM_CONN_PWD%{ENCRYPT|NOENCRYP
PT}
```

Important! You must enter these values in the order shown. You can omit required values only by truncating the string from right to left. For example, you can specify *DBNAME/DBTYPE*, but you can't specify *DBNAME/DOMAIN_ID*.

These startup settings all have default values if you omit any of them. The default values are generally the values you provided when setting up your PeopleSoft environment, and are the same as they would initially appear on the PSADMIN application server Quick-Configure menu.

The following table describes the startup settings:

Startup (s_set) Setting	Description
<i>DBNAME</i>	<p>Enter the name of the database to which the application server will connect. This is the same as the <i>DBName</i> parameter in the Startup section of the <i>psappsrv.cfg</i> file.</p> <hr/> <p>Note: If you don't include the <i>s_set</i> parameter, the value of this setting is the same as the domain name that you specify in the command.</p>
<i>DBTYPE</i>	<p>Enter the database type. Valid values are ORACLE, INFORMIX, SYBASE, MICROSOFT, DB2ODBC, and DB2UNIX. This is the same as the <i>DBType</i> parameter in the Startup section of the <i>psappsrv.cfg</i> file.</p> <hr/> <p>Note: Notice the spelling of MICROSOFT. DB2ODBC is the database type for DB2 z/OS.</p>
<i>OPR_ID</i>	<p>Enter the user ID, such as QEDMO, for the domain to use to connect to the database. This is the same as the <i>UserId</i> parameter in the Startup section of the <i>psappsrv.cfg</i> file.</p>
<i>OPR_PSWD</i>	<p>Enter the user password that is associated with the specified user ID. This is the same as the <i>UserPswd</i> parameter in the Startup section of the <i>psappsrv.cfg</i> file.</p>
<i>DOMAIN_ID</i>	<p>Enter a domain ID, such as TESTSRV1, TESTSRV2, and so on. This does not need to match the domain name. This name is important only because the Tuxedo Web Monitor uses it to identify application server domains on each machine. This is the same as the <i>Domain ID</i> parameter in the Domain Settings section of the <i>psappsrv.cfg</i> file.</p>

Startup (s_set) Setting	Description
<i>ADD_TO_PATH</i>	<p>(Optional) Enter the directory path that contains your connectivity software or database drivers. This is the same as the <i>Add to PATH</i> parameter in the Domain Settings section of the psappsrv.cfg file.</p> <hr/> <p>Note: If this value contains spaces, it must be in double quotes (" "). For example: "c:\Program Files".</p> <hr/> <p>Important! If you want this setting to be blank, but you can't truncate the string to this point (you still need to specify a value for <i>CNCT_ID</i>), you can specify a value of "_____" (five underscores without the quotes) in this position. PSADMIN interprets this as a blank value.</p> <hr/>
<i>CNCT_ID</i>	<p>Enter the connect ID, which is required for all platforms. This is the same as the <i>ConnectId</i> parameter in the Startup section of the psappsrv.cfg file.</p> <p>See "Connect ID" (PeopleTools 8.53: Security Administration).</p>
<i>CNCT_PSWD</i>	<p>Enter the password that is associated with the connect ID. This is the same as the <i>ConnectPswd</i> parameter in the Startup section of the psappsrv.cfg file.</p>
<i>SERV_NAME</i>	<p>(Optional) If your RDBMS requires that you specify the server name on which the database resides, enter the appropriate server name. This is the same as the <i>ServerName</i> parameter in the Startup section of the psappsrv.cfg file.</p> <hr/> <p>Important! If you want this setting to be blank, but you can't truncate the string to this point (you still need to specify the ENCRYPT setting), you can specify a value of "_____" (five underscores without the quotes) in this position. PSADMIN interprets this as a blank value.</p> <hr/>
ENCRYPT NOENCRYPT	<p>Specify ENCRYPT to encrypt the values of the <i>UserPswd</i> and <i>ConnectPswd</i> parameters in the psappsrv.cfg file. If you specify NOENCRYPT (the default value), these values appear in clear text in the file.</p>

The *p_set* parameter is an optional string of port numbers that you would otherwise specify on the PSADMIN application server Quick-Configure menu. Typically, you include this parameter only if you have more than one domain on the same application server machine or if you need to provide a specific value due to your environment or testing needs. Otherwise, you should accept the defaults for easy configuration.

You must specify the port numbers as follows:

- In Windows, the values must be separated by slashes.

WSL_PORT/JSL_PORT/JRAD_PORT

- In UNIX, the values must be separated by percent signs.

WSL_PORT%JSL_PORT%JRAD_PORT

Important! You must enter these values in the order shown. You can omit values only by truncating the string from right to left. For example, you can specify *WSL_PORT/JSL_PORT*, but you can't specify *WSL_PORT/JRAD_PORT*. These port numbers all have default values if you omit any of them. The default values are the values you provided when setting up your PeopleSoft environment, and are the same as they would initially appear on the PSADMIN application server Quick-Configure menu.

The following table describes the port settings:

Port (p_set) Setting	Description
<i>WSL_PORT</i>	<p>Workstation listener port number. This is the same as the <i>Port</i> parameter in the Workstation Listener section of the PSAPPSRV.CFG file.</p> <hr/> <p>Note: Specify this value only if you intend for the domain to support Windows workstations connecting in the development environment.</p> <hr/>
<i>JSL_PORT</i>	<p>Jolt port number. This is the same as the <i>Port</i> parameter in the JOLT Listener section of the PSAPPSRV.CFG file.</p>
<i>JRAD_PORT</i>	<p>Jolt internet relay port number. This is the same as the <i>Listener Port</i> parameter in the JOLT Relay Adapter section of the PSAPPSRV.CFG file.</p> <hr/> <p>Note: Specify this value only if you intend for the domain to support browser deployment, <i>and</i> your web server resides on a separate machine from the application server.</p> <hr/>

The *env_set* specifies domain environment settings, having the following format:

```
ENVAR_NAME1=ENVVAR_VALUE1#ENVAR_NAME2=ENVVAR_VALUE2
```

Following is an example of the application server **create** command:

```
D:\PT\appserv>psadmin -c create -d HRDOM01 -t small -s PT853803/ORACLE/QEDMO/QEDMO
/TESTSRV2/"c:\oracle\product\11.1.0\client_1\bin"/people/people/_____/ENCRYPT
-p 7200/9020/9200
```

When you launch the command, you'll see progress messages similar to the following:

```
Copying application server configuration files...
Copying [D:\PT\appserv\small.cfx] to [D:\PT_SERVER\8.53\appserv\HRDOM01\
psapprv.cfg]
Stripping Annotations...
Copying [D:\PT\appserv\small.cfx] to [D:\PT_SERVER\8.53\appserv\HRDOM01\
psapprv.cfx]
Copying Jolt repository file...
```

```
Domain created.
Performing load prechecks ...
Loading validation table...
  setting DBName=PT853803
  setting DBType=ORACLE
  setting UserId=QEDMO
  setting UserPswd=sgDcfYICragaN1Bz+MTRxf9CAk21Bqlkn/DFpUQAaTs=
  setting ConnectId=people
  setting ConnectPswd=kyD3QPxnrag=
  setting ServerName=
  setting Port=7200
  setting Port=9020
  setting Listener Port=9200
  setting Domain ID=TESTSRV2
  setting Add to PATH=c:\oracle\product\11.1.0\client_1\bin
New CFG file written with modified Startup parameters

Log Directory entry not found in configuration file.
Setting Log Directory to the default... [PS_SERVDIR\LOGS]
Spawning disabled for server PSAPPSRV.
WARNING: PSSAMSRV is configured with Min instance set to 1. To avoid loss of
service, configure Min instance to atleast 2.
Configuration file successfully created.
CFG setting changes completed, loading configuration...
Domain configuration complete.
```

Using the Application Server Administration Commands

Use the following syntax to administer an existing application server domain:

```
psadmin -ccommand -ddomain
```

The *domain* parameter must be the name of an application server domain that you want to administer, for example, HR846DMO.

The valid values of the *command* parameter are as follows:

Command	Example	Result of the Example
boot	<code>psadmin -c boot -d PSDMO</code>	Boots an application server domain named PSDMO.
parallelboot	<code>psadmin -c parallelboot -d PSDMO</code>	Boots an application server domain named PSDMO, using the parallel boot option.
configure	<code>psadmin -c configure -d PSDMO</code>	Reloads the domain configuration for the PSDMO domain.
pslist	<code>psadmin -c pslist -d PSDMO</code>	Displays the processes that have been booted for the PSDMO domain. This includes the system process ID for each process.
shutdown	<code>psadmin -c shutdown -d PSDMO</code>	Shuts down the PSDMO application server domain, by using a normal shutdown method. In a normal shutdown, the domain waits for users to complete their tasks and turns away new requests before terminating all of the processes in the domain.
shutdown!	<code>psadmin -c shutdown! -d PSDMO</code>	Shuts down the PSDMO application server domain by using a forced shutdown method. In a forced shutdown, the domain <i>immediately</i> terminates all of the processes in the domain.
sstatus	<code>psadmin -c sstatus -d PSDMO</code>	Displays the Tuxedo processes and PeopleSoft server processes that are currently running in the PSDMO application server domain.
cstatus	<code>psadmin -c cstatus -d PSDMO</code>	Displays the currently connected users in the PSDMO application server domain.
qstatus	<code>psadmin -c qstatus -d PSDMO</code>	Displays status information about the individual queues for each server

Command	Example	Result of the Example
		process in the PSDMO application server domain.
preload	<code>psadmin -c preload -d PSDMO</code>	Preloads the server cache for the PSDMO domain.
cleanipc	<code>psadmin -c cleanipc -d PSDMO</code>	Cleans the IPC resources for the PSDMO domain.
purge	<code>psadmin -c purge -d PSDMO</code>	Purges the cache for the PSDMO domain.
import	<code>psadmin -c import c:\ptinstalls\pt85x\...\psappsrv.cfg` -n NEWSRVR</code>	Imports a domain configuration. See PSADMIN command line help for all possible options.

Using the Process Scheduler Commands

For Process Scheduler administration, PSADMIN has two syntax formats — one for creating new Process Scheduler configurations, and the other for administering existing configurations.

Using the Process Scheduler Create Command

Use the following syntax to create a new Process Scheduler configuration:

```
psadmin -p create -d domain -t template -ps ps_set [-env env_set] [-s silent]
```

The PSADMIN **create** command creates a Process Scheduler configuration file for the specified database, using the specified configuration template.

The *template* parameter must be the name of a .cfx file located in *PS_HOME*\appserv\prcs, without the extension. This represents the operating system platform on which you're running PeopleSoft Process Scheduler. For example, to use the template file called windows.cfx on a Windows machine, specify the value *windows*.

The *ps_set* parameter is an optional string of startup values which provide initial configuration settings that you would otherwise specify on the PSADMIN Process Scheduler Quick-Configure menu. You must enter the startup string as follows:

- In Windows, the values must be separated by slashes.

```
DBNAME/DBTYPE/PRCSSERVER/OPR_ID/OPR_PSWD/DB_CNCT_ID/DB_CNCT_PSWD/SERVER_NAME/LO⇒  
GOUT_DIR/  
SQRBIN/ADD_TO_PATH/DBBIN/CRYSTALBIN/DOM_CONN_PWD/{ENCRYPT|NOENCRYPT}
```

- In UNIX, the values must be separated by commas.

```
DBNAME,DBTYPE,PRCSSERVER,OPR_ID,OPR_PSWD,DB_CNCT_ID,DB_CNCT_PSWD,SERVER_NAME,LO⇒  
GOUT_DIR,  
SQRBIN,ADD_TO_PATH,{ENCRYPT|NOENCRYPT}
```

Note: The UNIX syntax does not include the `DBBIN` setting.

Important! You must enter these values in the order shown. You can omit required values only by truncating the string from right to left. For example, you can specify `DBNAME/DBTYPE`, but you can't specify `DBNAME/LOGOUT_DIR`.

These startup settings all have default values if you omit any of them. The default values are generally the values you provided when setting up your PeopleSoft environment, and are the same as they would initially appear on the PSADMIN Process Scheduler Quick-Configure menu.

Note: Because these PeopleSoft Process Scheduler settings are already documented in the *PeopleSoft Process Scheduler PeopleBook*, this section provides only a basic overview of the relationship between the settings on the command line and the equivalent settings on the PSADMIN Process Scheduler Quick-Configure menu.

The following table describes the startup settings:

Startup (<i>ps_set</i>) Setting	Description
<i>DBNAME</i>	This is the equivalent of the <i>DBName</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu. Note: If you don't include the <i>ps_set</i> parameter, the value of this setting is the same as the database name that you specify in the command.
<i>DBTYPE</i>	This is the equivalent of the <i>DBType</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.
<i>PRCSSERVER</i>	This is the equivalent of the <i>PrCsServer</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.
<i>OPR_ID</i>	This is the equivalent of the <i>UserId</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.
<i>OPR_PSWD</i>	Enter the user password that is associated with the specified user ID. This is the equivalent of the <i>UserPswd</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.
<i>CNCT_ID</i>	This is the equivalent of the <i>ConnectId</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.
<i>CNCT_PSWD</i>	This is the equivalent of the <i>ConnectPswd</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.

Startup (ps_set) Setting	Description
<i>SERV_NAME</i>	<p>(Optional) This is the equivalent of the <i>ServerName</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.</p> <hr/> <p>Important! If you want this setting to be blank, but you can't truncate the string to this point (you still need to specify a value for <i>LOGOUT_DIR</i>), you can specify a value of "_____" (five underscores without the quotes) in this position. PSADMIN interprets this as a blank value.</p>
<i>LOGOUT_DIR</i>	<p>This is the equivalent of the <i>Log/Output Dir</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.</p> <hr/> <p>Note: If this value contains spaces, it must be in double quotes (" "). For example: "c:\psft app\log_output".</p>
<i>SQRBIN</i>	<p>This is the equivalent of the <i>SQRBIN</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.</p> <hr/> <p>Note: If this value contains spaces, it must be in double quotes (" "). For example: "C:\PeopleTools\bin\sqr\MSS\binw".</p>
<i>ADD_TO_PATH</i>	<p>(Optional) This is the equivalent of the <i>AddToPATH</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu.</p> <hr/> <p>Note: If this value contains spaces, it must be in double quotes (" "). For example: "%WINDIR%\SYSTEM32;c:\Program Files".</p>
<i>CRYSTALBIN</i>	<p>(Windows only) Enter the Crystal installation path. This is equivalent to the <i>Crystal Path</i> parameter in the PSADMIN Process Scheduler Quick-Configure menu.</p> <hr/> <p>Note: If a valid path is not entered, error messages related to PS ODBC and Crystal Runtime will appear when PSADMIN loads the configuration.</p> <hr/> <p>Note: If this value contains spaces, it must be in double quotes (" ").</p> <hr/> <p>Important! If you want this setting to be blank, but you can't truncate the string to this point, you can specify a value of "_____" (five underscores without the quotes) in this position. PSADMIN interprets this as a blank value.</p>

Startup (ps_set) Setting	Description
<i>DBBIN</i>	(Windows only) This is the equivalent of the <i>DBBIN</i> parameter on the PSADMIN Process Scheduler Quick-Configure menu. Note: If this value contains spaces, it must be in double quotes (" "). For example: "C:\my apps\db\MSSQL\Binn".
ENCRYPT NOENCRYPT	Specify ENCRYPT to encrypt the values of the <i>UserPswd</i> and <i>ConnectPswd</i> parameters in the <i>psprcs.cfg</i> file. If you specify NOENCRYPT (the default value), these values appear in clear text in the file.

The *env_set* specifies domain environment settings, having the following format:

```
ENVAR_NAME1=ENVVAR_VALUE1#ENVAR_NAME2=ENVVAR_VALUE2
```

Following is an example of the Process Scheduler **create** command:

```
psadmin -p create -d PSHRDB1 -t nt -ps PSHRDB1/MICROSFT/PSNT/
PS/PS/people/people/____/"c:\psft app\log_output"/c:\psfthr\bin\sqr\MSS\binw/
c:\WINNT\SYSTEM32/C:\Program Files\BusinessObjects\BusinessObjects Enterprise 12.0\
win32_x86
/c:\apps\db\mssql\bin/ENCRYPT
```

Using the Process Scheduler Administration Commands

Use the following syntax to administer an existing Process Scheduler configuration:

```
psadmin -pcommand -ddatabase
```

The *database* parameter must be the name of a database that's associated with a PeopleSoft Process Scheduler Server Agent, for example, PSHRDMO.

The valid values of the *command* parameter are as follows:

Command	Example	Result of the Example
start	psadmin -p start -d psdmo	Starts a Process Scheduler.
stop	psadmin -p stop -d psdmo	Stops a Process Scheduler.
configure	psadmin -p configure -d psdmo	Configures a Process Scheduler.
status	psadmin -p status -d psdmo	Displays the status of a Process Scheduler.
cleanipc	psadmin -p cleanipc -d psdmo	Cleans the IPC resources for specified domain
kill	psadmin -p kill -d psdmo	Kills the domain (similar to forced shutdown).

Related Links

[Using the Process Scheduler Menu](#)

Using the Search Server Commands

Use the following syntax to administer an existing search server domain:

```
psadmin -s command -d domain
```

The domain parameter must be the name of the search server domain that you want to administer, for example, PSSRCH. The valid values of the command parameter are as follows:

Command	Example	Result of the Example
create	<pre>psadmin -s create -d PSSRCH -t <template></pre>	Creates a domain with the specified name and template.
boot	<pre>psadmin -s boot -d PSSRCH</pre>	Boots a search server.
configure	<pre>psadmin -s configure -d PSSRCH</pre>	Configures a search server.
shutdown	<pre>psadmin -s shutdown -d PSSRCH</pre>	Shuts down the domain, by using a normal shutdown method. In a normal shutdown, the domain waits for current transactions to complete and turns away new requests before terminating all of the processes in the domain.
shutdown!	<pre>psadmin -s shutdown! -d PSSRCH</pre>	Shuts down the domain by using a forced shutdown method. In a forced shutdown, the domain immediately terminates all of the processes in the domain.
sstatus	<pre>psadmin -s sstatus -d PSSRCH</pre>	Displays the Tuxedo processes and PeopleSoft server processes that are currently running in the domain.
cstatus	<pre>psadmin -s cstatus -d PSSRCH</pre>	Displays the currently connected users/clients.
qstatus	<pre>psadmin -s qstatus -d PSSRCH</pre>	Displays status information about the individual queues for each server process in the application server domain.
cleanipc	<pre>psadmin -p cleanipc -d PSSRCH</pre>	Cleans the IPC resources for the domain.

Working with PSADMIN Exit Codes

The following table lists most PSADMIN command line options for the application server and Process Scheduler server and their return codes. This can provide insight into the return codes displayed when using PSADMIN from the command line or in scripts. In general, PSADMIN returns 0 for success and non-zero when an error occurs.

Note: *N/A* indicates the command is not applicable to that server.

Command	Application Server	Process Scheduler
-v	0	N/A
-h	0	N/A
-env	0	N/A
-defaultPS_CFG_Home	0	N/A
-invalid	-10	N/A
-create -d <domain_name> -t <template>	0	0
-create -d <domain_name> -t <invalid_template>	275	275
create -d <already_existing_domain> -t <template>	265	1
create -d <domain_name> -t <template> -s <s_set>	0	N/A
create -d <domain_name> -t <invalid_template> -s <s_set>	1	N/A
create -d <domain_name> -t <template> -s <s_set> -p <p_set>	0	N/A
create -d database -t template -ps ps_set	N/A	0
create -d database -t template -silent	N/A	0
delete -d <domain_name>	0	0
delete -d <nonexistent_domain>	235	235
configure -d <domain_name>	0	0
configure -d <nonexistent_domain>	235	235
boot -d <domain_name>	0	N/A
boot -d <nonexistent_domain>	235	N/A
parallelboot -d <domain_name>	-1	N/A
parallelboot -d <nonexistent_domain>	235	N/A
shutdown -d <domain_name>	0	N/A
shutdown -d <nonexistent_domain>	235	N/A
shutdown! -d <domain_name>	0	N/A
shutdown! -d <nonexistent_domain>	235	N/A
cleanipc -d <domain_name>	0	0
cleanipc -d <nonexistent_domain>	235	235

Command	Application Server	Process Scheduler
PreloadFileCache not set in PSAPPSRV.CFG.		
preload -d <domain_name> PreloadFileCache set in PSAPPSRV.CFG but not in database.	1	N/A
preload -d <domain_name> PreloadFileCache set in PSAPPSRV.CFG and in database.	0	N/A
preload -d <nonexistent_domain>	235	N/A
stop -d <domain_name> Domain is not started.	1	0
pslist -d <domain_name>	0	0
pslist -d <domain_name> Domain is not started.	0	0
pslist -d <nonexistent_domain>	235	235
purge -d <domain_name> -noarch -log <log_comments>	0	235
purge -d <domain_name> -arch <archive_dir> -log <log_comments>	0	235

Using PSADMIN Configuration Files

This section provides an overview of PSADMIN executables and configuration files and discusses how to:

- Configure a domain.
- Load a configuration.
- Archive application server configuration files.
- Boot a domain.
- Stop a domain.
- Monitor a domain.

Understanding PSADMIN Configuration Files

You can create, configure, and boot an application server domain from the PSADMIN interface or through its command-line options.

The configuration and data files on which the executables rely all reside in *PS_CFG_HOME*\appserv*domain_name*. Each domain has its own set of these files:

- PSAPPSRV.CFG
This is the catch-all configuration file that contains the entire collection of configuration values for a given application server domain.
- PSAPPSRV.UBX
This is the template or model file for the PSAPPSRV.UBB file. This is where the system stores environment settings after being edited through PSADMIN.
- PSAPPSRV.UBB
This file stores and passes all of the domain values to the Tuxedo load configuration program (tmloadcf.exe).
- PSAPPSRV.PSX
This is the template or model file specifically for the messaging server configuration sections.
- PSAPPSRV.ENV
This contains environment information, such as the PS_HOME referenced by a domain.
- PSAPPSRV.VAL
This contains the format specification for the configuration parameters and, for some parameters, a set of valid values that can be assigned. This helps to prevent administrators from entering invalid values.
- PSTUXCFG
This contains PeopleSoft and Tuxedo information regarding the location of executables, files, and command lines for server processes. This file is required to boot a domain.
- JREPOSITORY
This file contains a list of the services handled by the application server on behalf of Jolt.

Note: Oracle does not support creating custom configuration files, such as CFX or UBX templates, nor modifying delivered templates. However, modifying PSAPPSRV.CFG and PSPRCS.CFG can be done through PSADMIN or manually through a text editor.

Configuring a Domain

Regardless of how you specify domain values, ultimately you must run PSADMIN to generate some necessary files that include your specific values. In the following example, PSADMIN invokes another PeopleSoft executable, UBBGEN, which reads the values and format in the psappsrv.cfg, psappsrv.val, and psappsrv.ubx files, and generates the psappsrv.ubb and psappsrv.env files.

Where you see *Do you want to change any config values? (y/n)*, regardless of what you enter, PSADMIN calls UBBGEN.

If you have already entered values manually in the `psappsrv.cfg` file and enter `n`, UBBGEN reads those values and writes to the necessary files.

If you enter `y`, you see the PSADMIN prompt interface, which is actually a wrapper to UBBGEN. UBBGEN reads the previous values in the `psappsrv.cfg`, presents those values, and allows you to change them. It presents the values in the format that is derived from reading the `PSAPPSRV.UBX` file, and it validates selected values based on criteria in the `PSAPPSRV.VAL` file.

Note: In the previous example, UBBGEN both reads from and writes to the `psappsrv.cfg` file. It reads the previous values or defaults and, if any values are modified, it writes the new values to the new `psappsrv.cfg` file.

Here are the scenarios by which you can configure a domain:

- Start PSADMIN, and enter values at all of the prompts.

This generates all of the necessary files automatically.

- Edit the `psappsrv.cfg` file.

If you decide not to use PSADMIN you must complete the following tasks in order:

- From the command line, create a domain based on a particular template.
- Edit the `psappsrv.cfg` file in a text editor.
- Issue the `configure` command from the PSADMIN command line. This is the command that calls UBBGEN. You see the following after issuing this command:

```
n:\<ps_home>\appserv>psadmin -c configure -d HRDOM01

Performing load prechecks ...
Loading validation table...

Loading new configuration...
Domain configuration complete.
```

Loading a Configuration

After you configure a domain and PSADMIN creates the new configuration file, PSADMIN loads the new configuration settings into `PSTUXCFG` so that the domain can properly boot. This occurs automatically after you have completed all of the prompts for values in PSADMIN.

To load the new configuration, PSADMIN calls the Tuxedo executable, `TMLOADCF.EXE`, which populates the `PSTUXCFG` file. `TMLOADCF.EXE` reads the newly entered values that appear in the `PSAPPSRV.UBB` file and writes them to the `PSTUXCFG` file.

Archiving Application Server Configuration Files

To track changes made to the `psappsrv.cfg` file and the history of the changes, a subdirectory, named "archive," stores various versions of the CFG file. You can find this subdirectory in the domain name directory, as in `C:\PS_CFG_HOME\appserv\domain name\archive`, where the current version of `psappsrv.cfg` resides.

When you boot the application server domain for the first time, PSADMIN places a copy of `psappsrv.cfg` in the archive directory. In subsequent boots, if PSADMIN detects a change in `psappsrv.cfg` based on the time stamp values, it replaces the current `psappsrv.cfg` with the latest version. The file name of the new version is then `psappsrv_mmddyy_hhmm_ss.cfg`, as displayed on the time stamp.

Booting a Domain

When you select Boot this domain, PSADMIN calls the Tuxedo executable, `TMBOOT.EXE`, which uses the information that resides in the `PSAPPSRV.ENV` and `PSTUXCFG` files to boot the appropriate domain.

Note: You can track the server processes that start by using PSADMIN, `ps -ef` command in UNIX, or Task Manager in Microsoft Windows.

Stopping a Domain

When you select Domain shutdown menu and select one of the shutdown options, PSADMIN calls the Tuxedo executable, `TMSHUTDOWN.EXE`, which also uses the information that resides in the `PSAPPSRV.ENV` and `PSTUXCFG` files to shut down the appropriate domain.

Following a successful domain shutdown, PSADMIN checks and stops orphaned processes in the domain. If PSADMIN identifies and stops any orphaned server processes, it displays a screen message at the end of the shutdown operation.

Monitoring a Domain

To detect any orphaned application server processes, a server process, `PSWATCHSRV`, monitors the application server domain. Every two minutes, `PSWATCHSRV` identifies and stops any hung or orphaned server processes. If any hung or orphaned processes are found, it writes a message to the application server log file. The `PSWATCHSRV` process is the first process to start when you boot up the domain and the last one to stop when you shut down the domain.

Domain ID Name

To identify orphaned application server processes, all server processes within a server's domain must be uniquely identified. Therefore, the system appends a unique number to the domain ID in the `psappsrv.cfg` file. If you refer to domain IDs in scripts or processes, you may need to change those to reflect the new naming convention.

The command line varies slightly depending on the application server process, but it looks like this:

```
PSAPPSRV -C dom=pt84_52692 ...
```

Configuring the Application Server to Handle Cache Files and Replay Files

When an application server instance crashes, cache files and replay files are generated automatically. Over time, the size of these files can consume a large amount of disk space if there are recurring crashes in a domain. To minimize the buildup of cache files and replay files, you can modify the `psappsrv.cfg` file based on the following rules:

- When a crash occurs, the system creates a directory in the domain's LOGS directory.
- The dump file is saved in a directory within the domain's LOGS directory.
- The DumpMemoryImageAtCrash setting in the Trace section of the psappsrv.cfg file saves the memory image of the failed process in Microsoft Windows.

This functionality is only available for Windows. If the value of DumpMemoryImageAtCrash is MINI, a miniature memory image (with a size less than or equal to 64K) is generated. If the value is FULL, then a full memory image is created. Depending on how much memory is consumed by the application, this full memory image can be quite large. The location of the memory image is the same as the replay file.

- If DumpManagerObjectsAtCrash is set to Y, then the application server instance:
 1. Generates the replay file.
 2. Dumps the customized objects being used by the current service request into the special cache directory.

The cache directory resides in the same location as the replay file.

3. If the value of DumpMemoryImageAtCrash is NONE and the platform is set to MS Windows, a miniature memory image is created.
- The settings for DumpManagerObjectsAtCrash and DumpMemoryImageAtCrash are dynamic. That is, the application server doesn't need to be restarted for these settings to be effective.
 - There is no separate setting for generating the replay file.

This file is generated as mentioned previously.

- Regardless of the setting in DumpManagerObjectsAtCrash, a summary report of objects in each managed type for which at least one object is loaded in memory is written to the dump file or application log file.

The summary report resembles the following example:

```
PDM Definitions: Total=36   Customized=0   In-Use=10
RDM Definitions: Total=53   Customized=52  In-Use=50
MDM Definitions: Total=1    Customized=0   In-Use=0
PCM Definitions: Total=199  Customized=0   In-Use=3
PGM Definitions: Total=1    Customized=0   In-Use=1
CRM Definitions: Total=67   Customized=0   In-Use=0
SSM Definitions: Total=1    Customized=0   In-Use=1
CLM Definitions: Total=1    Customized=0   In-Use=0
UPM Definitions: Total=1    Customized=0   In-Use=0
```

Total indicates the total number of in-memory definitions being used by the current service.

Customized indicates how many of those objects are customized, and In-Use indicates how many of those objects were being used at the time of the crash.

If DumpManagerObjectsAtCrash is set to Y, the summary for each managed object type follows the list of configured objects that are being dumped as part of the crash information gathering. If a configured object is in use, its name is prefixed with an asterisk.

A sample report for a managed object type follows:

```
RDM (PSOPTIONS/ENG)
*RDM (PSTREEDEFNLABELS/ENG)
RDM Definitions: Total=10   Customized=2   In-Use=1
```

Note: The asterisk that precedes the object name indicates that this object is being used by the current service request.

Chapter 5

Using PSADMIN Menus

Understanding PSADMIN Menus

You use PSADMIN to manage all of the main PeopleSoft server types. After launching PS_HOME \appserv\psadmin.exe from the command line, the PeopleSoft Server Administration (PSADMIN) main menu appears.

```
-----  
PeopleSoft Server Administration  
-----  
  
Config Home: D:\PT_SERVER\8.53  
  
1) Application Server  
2) Process Scheduler  
3) Search Server  
4) Web (PIA) Server  
5) Switch Config Home  
6) Service Setup  
7) Replicate Config Home  
q) Quit
```

Command to execute (1-7, q):

Config Home

Displays the PS_CFG_HOME the current PSADMIN session. If it is not the PS_CFG_HOME desired, use the *Switch Config Home* option to select the appropriate location.

Application Server

Select to manager your application server domains.

Process Scheduler

Select to manager your Process Scheduler domains.

Search Server

Select to manager your Verity—based search server domains.

Note: This menu applies only to the Verity-based search servers. It does not apply to *any* PeopleSoft Search Framework components.

Switch Config Home

Select to switch to an alternative PS_CFG_HOME location. Determine the current PS_CFG_HOME from the Config Home value displayed.

Service Setup

(Applies to Microsoft Windows only). Select to configure domains to start using Windows Services.

Replicate Config Home

Select to replicate configured PS_CFG_HOME locations so that you can share these pre-created, tuned, and configured server configurations on the local server or on remote servers.

Note: As a configuration option, you can configure a domain to *spawn* server processes according to the volume of transaction requests. There is no explicit parameter that you must set to enable spawning. In the following configuration section descriptions, some servers enable you to specify a minimum and maximum number of server processes. To enable spawning for these server processes, the maximum value must exceed the minimum value by an increment of at least one. As needed, the domain spawns server processes up to the maximum value. As the volume of transactions decreases, the number of spawned server processes decreases, or decays, until the minimum value is reached. By default, spawning is *disabled*.

Note: Generally, the documentation reflects the order in which the configuration sections appear in the PSADMIN interface or the PSAPPSRV.CFG file.

Using the Application Server Administration Menu

This section discusses how to:

- Access the application server options.
- Administer a domain.
- Boot a domain.
- Shut down a domain.
- Perform a normal shutdown.
- Perform a forced shutdown.
- Check the domain status.
- Purge the domain cache.
- Configure a domain.
- Edit configuration and log files.
- Create a domain.
- Delete a domain.
- Configure an application server domain to preload cache.
- Clean domain IPC resources.

Accessing the Application Server Options

To access the menu options for configuring and administering an application server, select *Application Server* from the PeopleSoft Server Administration (PSADMIN) menu.

```
-----  
PeopleSoft Application Server Administration  
-----
```

- ```

1) Administer a domain
2) Create a domain
3) Delete a domain
4) Import domain configuration
q) Quit

```

Command to execute (1-4, q) : 4

The Administer a domain menu offers numerous configuration, administration, and logging parameters that you may access frequently.

The menu options and parameters within the Create a domain and Delete a domain menus are one-time tasks (per domain).

The Import domain configuration menu enables you to import existing configurations.

## Administering a Domain

To administer a domain, you must have already created a domain. After you have created a domain, specify environment-specific settings for the application server to function correctly with your system. The following sections describe all of the menus and menu options that you use to administer and configure an application server domain.

To administer a domain:

1. Select *Administer a domain* from the PeopleSoft Application Server Administration menu.
2. In the Select domain number to administer command line, enter the number that corresponds to the previously created domain that you want to administer that appears in the Tuxedo domain list.
3. Select the option that you want to perform from the PeopleSoft Domain Administration menu.

PSADMIN transparently sets several environment variables before invoking any Tuxedo administrative commands. You don't need to set these variables manually. These environment variables are:

- TUXCONFIG = *PS\_CFG\_HOME/appserv/domain\_name/PSTUXCFG*
- APPDIR = *PS\_CFG\_HOME/appserv/domain\_name*
- PATH = TUXDIR/bin; *PS\_HOME/bin/server/winx86*; PATH
- APP\_PW = Application Password (initialize)

## Importing Domain Configurations

You import existing domain configurations by selecting the Import domain configuration option from the PeopleSoft Application Server Administration menu. From the PeopleSoft Import Application Server Configuration menu, you have the option to import a regular domain or an Integration Broker Master Configuration.

```

PeopleSoft Import Application Server Configuration

```

- ```

1) Import regular domain
2) Import IB Master Configuration
q) Quit

```

Command to execute (1-2, q) : 1

Importing a Regular Application Server Domain

When importing a regular application server domain, you can import from a file or from a domain.

```
-----
PeopleSoft Import Application Server Configuration
-----
```

```
1) Import from file
2) Import from application domain
q) Quit
```

Command to execute (1-2, q) :

To import configuration settings from a file:

1. From the PeopleSoft Import Application Server Configuration menu, select *Import from file*.
2. Enter the complete file path to the PSAPPSRV.CFG file containing the desired configuration settings.

For example,

```
c:\documents and settings\TSAWYER.DOMAINX\PSFT\PT\8.53\appserv\appserv\PT850TST
```

3. Enter a name for the new domain.

To import configuration settings from an existing domain:

1. From the PeopleSoft Import Application Server Configuration menu, select *Import from application domain*.
2. Enter the path to the appropriate PS_CFG_HOME.

For example,

```
c:\documents and settings\TSAWYER.DOMAINX\PSFT\PT\8.53\appserv
```

3. From the Tuxedo domain list, select the domain containing the desired configuration settings.
4. Enter a name for the new domain.

Importing an Integration Broker Master Configuration

Use this option to import a master configuration, if you are implementing a "master-slave" integration server configuration. This implementation option is documented in detail within the Integration Broker Administration PeopleBook.

See "Implementing Master-Slave Dispatchers" (PeopleTools 8.53: PeopleSoft Integration Broker Administration).

Booting a Domain

This boots the Tuxedo domain (the application server) by using the tmbboot command. This command will start all of the server processes that have been configured for your domain.

```
-----
PeopleSoft Domain Boot Menu
```

```

-----
Domain Name: DOC

1) Boot (Serial Boot)
2) Parallel Boot
q) Quit

Command to execute (1-2, q) [q]:

```

You have two booting options: a serial boot and a parallel boot.

Running a Serial Boot

A serial boot starts server processes in a sequential order, with one process beginning to start after the previous process has completely started.

Running a Parallel Boot

A parallel boot starts server processes at the same time, rather than having each process start sequentially. This option typically provides shorter boot durations.

Shutting Down a Domain

The PeopleSoft Domain Shutdown menu offers two options: a normal shutdown and a forced shutdown.

```

-----
PeopleSoft Domain Shutdown Menu
-----
Domain Name: psdmo

1) Normal shutdown
2) Forced shutdown
q) Quit

Command to execute (1-2, q) [q]:

```

Performing a Normal Shutdown

A normal shutdown is a quiescent shutdown that waits for users to complete their tasks and turns away new requests before terminating all of the processes in the domain.

Performing a Forced Shutdown

A forced shutdown is a nonquiescent shutdown that *immediately* terminates all of the processes in the domain. Normally, you use the forced shutdown only when a Bulletin Board Liaison (BBL) process encounters errors and cannot be shut down by using a normal shutdown.

Note: The BBL is a primary Tuxedo process that controls the domain.

Checking the Domain Status

Use the PeopleSoft Domain Status menu to view the status of the server, queues, or clients connected to the domain.

```

-----
PeopleSoft Domain Status Menu
-----
Domain Name: psdmo

```

- 1) Server status
- 2) Client status
- 3) Queue status
- q) Quit

Command to execute (1-3, q) [q]:

Server Status

Select *Server status* to invoke the Tuxedo `tadmin psr` subcommand (print server processes), which displays the Tuxedo processes and PeopleSoft server processes that are currently running. For example:

```
> Prog Name      Queue Name  Grp Name      ID RqDone Load Done Current Service
-----
BBL.exe          62645      RTDC796+      0  42801  2140050 ( IDLE )
PSAPPSRV.exe    APPQ       APPSRV        1   422   21100 ( IDLE )
PSMONITORSRV.e  MONITOR    MONITOR       1     0     0 ( IDLE )
PSWATCHSRV.exe  WATCH     WATCH         1     0     0 ( IDLE )
PSAPPSRV.exe    APPQ       APPSRV        2     4    200 ( IDLE )
JSL.exe          00095.00200 JSLGRP       200     0     0 ( IDLE )
PSMSGDSP.exe    00098.00100 PUBSUB       100     4    200 ( IDLE )
PSSAMSRV.exe    SAMQ       APPSRV        100     0     0 ( IDLE )
PSMSGHND.exe    MBHQ       PUBSUB        101     0     0 ( IDLE )
JREPSVR.exe     00094.00250 JREPGRP      250    29   1450 ( IDLE )
```

The number of items appearing depends on the number of server processes that you have configured.

Client Status

Select *Client status* to invoke the Tuxedo `tadmin pctl` subcommand (printclient), which displays connected users. For example:

```
> LMID          User Name      Client Name      Time      Status  Bgn/Cmmt/Abrt
-----
RTDC79623VMC   NT             JSH              **:**:**  IDLE    0/0/0
RTDC79623VMC   NT             JSH              **:**:**  IDLE    0/0/0
RTDC79623VMC   JPOOL_15      dhcp-adc-twvpn+  0:02:03  IDLE/W  0/0/0
RTDC79623VMC   NT             tadmin           0:00:00  IDLE    0/0/0
```

Queue Status

Examining the status of the individual queues for each server process provides valuable tuning information. Check the queues by using the Queue status option. In the following example, the results of the Queue status option show the individual server processes, the associated queue, the number of server processes currently running, and the number of requests waiting to be processed:

```
> Prog Name      Queue Name  # Serve Wk Queued # Queued Ave. Len Machine
-----
PSAPPSRV.exe    APPQ       2         -         0         - RTDC79623+
JSL.exe          00095.00200 1         -         0         - RTDC79623+
JREPSVR.exe     00094.00250 1         -         0         - RTDC79623+
BBL.exe          62645      1         -         0         - RTDC79623+
PSMONITORSRV.e  MONITOR    1         -         0         - RTDC79623+
PSWATCHSRV.exe  WATCH     1         -         0         - RTDC79623+
PSSAMSRV.exe    SAMQ       1         -         0         - RTDC79623+
PSMSGHND.exe    MBHQ       1         -         0         - RTDC79623+
PSMSGDSP.exe    00098.00100 1         -         0         - RTDC79623+
```

The results alert you to any bottlenecks that may be occurring on your application server. With this information, you can make more informed performance decisions. For instance, if the bottlenecks appear to be persistent, it may indicate that you need to add more instances of a particular server process, such as PSAPPSRV for example. Or the results may indicate that you need to start either a PSQCKSRV or a PSQRYSRV.

Purging the Domain Cache

A proven technique for resolving problem application server environments is to purge the application server domain cache located in `PS_CFG_HOME\appserv\domain_name\CACHE`.

Important! You should purge the cache only after due consideration, and in after consultation with PeopleSoft support or consultants.

Please keep the following in mind:

- You cannot purge shared file cache.
- You can purge the cache regardless of whether the application server domain is running; there's no need to shut it down and reboot. However, the procedure is less disruptive and runs more quickly if the domain is shut down or its activity level is low.
- Purging the cache can take five minutes or more on a large or busy domain, depending on the domain configuration.
- If database cache is enabled, the database cache will be purged. If the database cache is purged, it affects all domains using that cache. The database cache is shared across all domains that use the same byte-order (little-endian or big-endian).

To purge the domain cache:

1. On the PeopleSoft Domain Administration menu, select Purge Cache.

If the cache is currently empty, the purge operation is cancelled, and the PeopleSoft Domain Administration menu reappears.

If the cache is not empty, the following prompt appears:

```
Enter log comments about this purge, if any (maximum 256 characters):
```

2. Enter any information (up to 256 characters) that you want recorded explaining the circumstances of this cache purge operation, and press `Enter`. Your comments will be saved to a purge log file.

The following prompt appears:

```
Do you wish to archive the contents of the current cache? (y/n) [n] :
```

3. Enter `y` to archive the cache contents, or `n` to delete them permanently. The default response is `n`.

If you enter `y`, the following prompt appears:

```
Cache contents will be archived to
PS_CFG_HOME\appserv\domain_name\Archive\CACHE_mmddyy_hhmm_ss.
Hit Enter to continue or provide a new location:
```

Note: At runtime, `PS_CFG_HOME` and `domain_name` are replaced with values appropriate to your system, and `mmddyy_hhmm_ss` represents the date and time of the cache purge operation.

4. (If you chose to archive the cache contents) Enter a different archive location if desired, and press `Enter`.

If the location you enter is rejected, the following message appears, and you're prompted to continue:

```
Failed to archive cache to location.
```

Note: Continuing this procedure with an invalid location will purge the cache without archiving.

The default location is a unique directory name. Keep in mind that the location you enter might have been rejected for the following reasons:

- The directory can't be created due to an invalid drive or network mount.
 - The directory can't be created due to insufficient user privileges.
 - The directory has insufficient space for the cache files.
5. When prompted to continue, enter *y* to continue the purge operation, *orn* to cancel the operation and return to the PeopleSoft Domain Administration menu.

Note: Archiving the cache increases the time required to complete the purge, because the cache files must be copied to the archive location.

If the application server domain is running, you might see INFO: messages related to version, patch level, serial number and so on. These are normal and don't require any action. When the cache is successfully purged, the following message appears:

```
Purge Cache operation completed successfully.
```

```
You may notice that the cache directory is non-empty. Cache files have been
invalidated and will be refreshed from the database.
```

If the cache was archived, you'll also see the following:

```
You may also have noticed a number of Sharing Violation messages during the
Cache Purge.
These messages are no cause for alarm and are expected as part of the
cache archival.
```

If the application server domain is running, an entry is written to the application server log file to indicate that the cache has been purged.

The purge log file is saved (including any comments you entered in step 2) as *PS_CFG_HOME*\appserv*domain_name*\LOGS\PurgeCache_*mmddy_hhmm_ss*.log.

Note: At runtime, *PS_CFG_HOME* and *domain_name* are replaced with values appropriate to your system, and *mmddy_hhmm_ss* represents the date and time of the cache purge operation.

Following is an example of the purge log file contents:

```
Date:02/17/10 11:47
User Explanation: Processes appeared to take a long time to recycle.
Cache Contents archived to C:\ptservers\appserv\QEDMO\Archive\CACHE_021710_1147_01
```

Configuring a Domain

This option prompts you with a model configuration file to gather such parameters as port numbers, the number of various server processes that are needed, encryption enabling, and so forth. PSADMIN then invokes a subprogram, UBBGEN, which takes the configuration parameters, builds the file `PS_CFG_HOME/appserv/domain-name/psappsrv.ubb`, and carries out the `tmloadcf - y psappsrv.ubb` command to generate the following file: `PS_CFG_HOME/appserv/domain-name/PSTUXCFG`.

The following topics describe all of the parameters that you encounter while configuring an application server. Either read this section before you fine tune the configuration of your application server or have it available while you are doing it.

To configure a domain:

1. Select *Configure this domain* from the PeopleSoft Domain Administration menu.

Enter *n* (No), if you do not want to continue. This returns you to the previous menu. Otherwise, enter *y* (Yes).

2. When prompted to change configuration values, enter *y*.

If you don't need to change any of the values, enter *n*. By doing so, you create a new configuration file with the same values that were previously specified. Enter *n*, or elect not to modify the PSADMIN parameters, if:

- You have changed only the location of TUXDIR.
- You would rather edit the PSAPPSRV.CFG file manually.
- You installed a new Tuxedo patch.

Note: If you edit the `psappsrv.cfg` file directly, it is recommended to reload your domain configuration. This is necessary because some settings in `psappsrv.cfg` are transferred to the `PSTUXCFG` file for the domain. This transfer of settings can only be achieved by running UBBGEN and `tmloadcf`, which the "Configure this domain" option performs.

Editing Configuration and Log Files

Use the Edit Configuration/Log Files menu to view the application server and Tuxedo log files. You can also manually edit the `PSAPPSRV.CFG` file if you do not want to use the PSADMIN interface.

To have PSADMIN start your text editor (such as Notepad or KEDIT) so that you can manually edit or view application server configuration and log files, you must specify the text editor in the environment settings. For example, to use KEDIT, the editor environment setting should look like this:

```
set EDITOR=c:\apps\kedit\keditw32.exe
```

To use Notepad, it should look like this:

```
set EDITOR=c:\Windows\Notepad.exe
```

Note: You can view and edit a domain's `PSAPPSRV.CFG` file while the domain is running, but the changes that you specify do not take effect until the next time you reconfigure the domain.

For the following options, you must enter your operator ID to view and edit the files:

```
Edit PSAPPSRV.tracesql (PSAPPSRV SQL trace file)
Edit PSSAMSRV.tracesql (PSSAMSRV SQL trace file)
```

For example:

```
Command to execute (1-7, q) [q]: 5
Enter the operator ID : PTXYZ
```

Note: PeopleSoft secures the Structured Query Language (SQL) traces because, in some instances, the SQL that is traced may involve sensitive information.

Note: Server configuration files (.CFG) support the use of environment variables, such as %PS_HOME%, %TEMP%, and so on.

Edit PSAPPSRV.CFG

The PSAPPSRV.CFG file contains all of the configuration settings for an application server domain. The PSADMIN interface provides prompts so that you can edit and modify this file within a structured format. In many cases, and perhaps due to personal preference, you may opt to edit the PSAPPSRV.CFG file manually in a text editor. When editing this configuration file manually, note that it is similar to editing an INI file, because all of the parameters are grouped in sections.

Edit APPSRV.LOG

This log file contains PeopleTools specific logging information.

Edit TUXLOG

The TUXLOG file enables you to trace the Tuxedo component for troubleshooting information.

Edit PSAPPSRV.tracesql

You can specifically trace the activity of the PSAPPSRV server process by setting the PSAPPSRV.tracesql option.

Edit PSSAMSRV.tracesql

You can specifically trace the activity of the PSSAMSRV server process by setting the PSSAMSRV.tracesql option.

Creating a Domain

The Create a domain option creates a subdirectory under `PS_CFG_HOME/appserv` using the domain name you specify as the directory name and copies the required domain files to that directory.

To create an application server domain:

1. Select *Create a domain* from the PeopleSoft Application Server Administration menu.
2. Enter the name of the domain that you want to create; the name must not exceed eight characters.
3. Select a configuration template from the Configuration template list.

The configuration templates are preconfigured sets of application server processes.

Deleting a Domain

Use the Delete a domain option to shut down the domain, if running, and delete the domain's subdirectory.

Note: Before you delete a domain, make sure that it is not running.

To delete a domain:

1. Select *Delete a domain* from the PeopleSoft Application Server Administration menu.
2. From the Tuxedo domain list, select the number that corresponds to the domain that you want to delete.
3. When prompted to continue, enter *y* and press Enter.
4. If the application server domain is currently running, PSADMIN informs you of that, and you will need to determine if you want to continue.

```
Domain processes are currently running. This option will shutdown
and delete the domain.
```

```
Do you want to continue? (y/n) [n] :
```

If you enter *y*, PSADMIN performs a forced shutdown on the domain prior to deleting the domain.

If you enter *n*, PSADMIN returns you to the PeopleSoft Application Server Administration menu, allowing you to shutdown the domain manually, if needed.

Configuring an Application Server Domain to Preload Cache

This section provides an overview and discusses how to:

- Create cache projects.
- Delete cache projects.
- Preloading File Cache.

See [ServerCacheMode](#).

See [Load Application Server Cache](#).

Understanding Preload Cache Projects

To improve performance, the application server uses a caching mechanism that keeps commonly used objects in memory and either file form on the application server or in a database cache. Caching reduces the need for a complete database query each time a definition is accessed. As more definitions are accessed, more data becomes stored in the cache. However, if a page, for example, has not already been accessed, it does not exist in the current cache, and the user may experience a slower response time as the system requests the page from the database for the first time. To prevent this initial performance degradation, you can preload file, database, and memory cache with commonly used definitions inserted into a cache project.

Preloading cache involves creating a project containing commonly used definitions and then referring to these projects in the PSADMIN settings PreloadMemoryCache and PreloadCache. By default, PreloadMemoryCache and PreloadCache are commented out because the two parameters need to be set to a specific name of a project that you create. You can set the parameters to reference separate projects. You use the Select Preload Objects page to select frequently used definitions and build the preload project containing the objects selected.

The cache project is intended to be used for a new domain, where the cache is not yet built. Prior to providing the domain for production use, use the Preload Cache PSADMIN option to build the cache containing the objects specified in the project. The domain starts the PSAPPSRV process, builds the cache, and shuts down.

The memory cache project is intended to be used during a server process recycle. When a the system starts a new process to replace an old one, the new process loads memory cache based on the project specified by PreloadMemoryCache so that the new process will not have delays when processing the first few service requests. Because it is desirable to have new processes start as quickly as possible, there is a timeout (or limit) of 60 seconds for PreloadMemoryCache. That is, PreloadMemoryCache preloads as many definitions as possible before the timeout of 60 seconds.

Note: In general, it is not recommended to create large projects for the PreloadMemoryCache project. For all cache projects (used for file, database, and memory cache) the optimum selections for the projects will require tuning and testing at your site.

Creating Cache Projects

To create cache projects:

1. In a browser, select PeopleTools, Utilities, Administration, PreLoad Cache, Select Preload Objects.
2. Select the Add a New Value tab, and in the Project Name edit box, enter the name of the project that will contain the definitions you select, and clickAdd.

Note: All project names used to contain definitions for preloaded cache *must* contain the "PLC_" prefix.

3. On the Select Preload Objects page, enter a Description, and select the Object Type to add.

Object Type refers to the definitions you create with PeopleTools, such as components, menus, pages, Application Engine programs, and so on.

4. Use the Key fields, to refine the selection of the definition as needed.

The keys will differ depending on the Object Type selected.

Note: Project definitions support up to four keys for identification, so you can use up to four keys when selecting objects for the preload cache project.

5. Click Save.
6. Click Build Project Definition.
7. On the Process Scheduler Request page, click OK.

This invokes an Application Engine program (PTCHPLC_PRJ) that creates the project definition in the database and populates it with the definitions you selected. The PTCHPLC_PRJ also automatically includes all related definitions in the project.

Note: Cache projects can be created manually in Application Designer, without running the PTCHPLC_PRJ program. The project name must include "PLC_".

Deleting Cache Projects

To delete a preload file cache project:

1. Select PeopleTools, Utilities, Administration, Preload Cache, Delete Preload Project.
2. On the Find an Existing Value page, click the appropriate project name.
3. On the Delete Preload Project page confirm that you have selected the appropriate project and click Delete the preload project.

Preloading Cache Projects

To preload a cache project:

1. Shut down the application server domain.
2. Edit the PSAPPSRV.CFG configuration file for the appropriate domain.

In the [Cache Settings] section, uncomment the `PreloadCache=` parameter, and enter the name of the pre-load project for this application server domain.

For example:

```
PreloadCache=PLC_PROJECTA
```

3. On the PeopleSoft Domain Administration menu in PSADMIN, select 9) Preload Cache.

When preloading cache, keep the following in mind:

- Preload Cache does not work when shared caching is enabled.
- If database caching is enabled, Preload Cache loads the cache project into the database cache.

Preloading Memory Cache

To preload the memory cache:

Edit the PSAPPSRV.CFG configuration file for the appropriate domain.

In the [Cache Settings] section, uncomment the `PreloadMemoryCache=` parameter, and enter the name of the preload project that should be preloaded on this application server for memory cache.

For example:

```
PreloadMemoryCache=PLC_PROJECTB
```

Related Links

[Load Application Server Cache](#)

Cleaning Domain IPC Resources

Use the Clean IPC Resources of this domain option to clear the interprocess communication (IPC) resources utilized by a domain. When a domain shuts down normally, the IPC resources it was using get released as part of the shut down process. However, if a domain terminates abnormally, in many cases the IPC resources are still assigned to the previous domain instance. This option enables you to clean any orphaned IPC resources assigned to a domain.

Using the Process Scheduler Menu

PeopleSoft Process Scheduler is used to schedule and run batch processes against the database. Process Scheduler domains are managed by Tuxedo, just as application server domains are. You only need to configure PeopleSoft Process Scheduler on a server where you intend to run batch processes. The menu options are identical to those provided for administering an application server domain, except where the options are not applicable to a Process Scheduler domain.

```
-----
PeopleSoft Process Scheduler Administration
-----
```

```
Domain Name: QEDMO
```

- 1) Boot this domain
- 2) Domain shutdown menu
- 3) Domain status menu
- 4) Configure this domain
- 5) TUXEDO command line (tmadmin)
- 6) Edit configuration/log files menu
- 7) Clean IPC resources of this Domain
- q) Quit

Command to execute (1-7, q) :

The Process Scheduler PSADMIN options differ from the application server PSADMIN options for these items.

Function	Difference
Booting	The Process Scheduler domain boot does not provide the option to select a serial or parallel boot. All Process Scheduler domains start in serial mode.
Viewing//Editing log files and configuration files	The Process Scheduler Edit/View Configuration/Log Files Menu does not display options for viewing the PSAPPSRV and PSSAMSRV SQL trace files.
Configuring Integration Broker (messaging) servers	For Process Scheduler domains there is no option to import a Master IB Domain nor is there the option to configure messaging servers.
Caching	For Process Scheduler domains there are not options to purge or preload file cache.

All of the Process Scheduler server configuration information for a specific domain is contained in the PSPRCS.CFG file, and PSADMIN provides an interface for and prompts you to edit the PSPRCS.CFG file.

Although you typically edit the PSPRCS.CFG file through PSADMIN, you can find the PSPRCS.CFG file in the following directory:

- Windows: *PS_CFG_HOME*\APPSERV\PRCS\domain
- UNIX: *PS_CFG_HOME*/appserv/prcs/domain

Related Links

[Using the Application Server Administration Menu](#)

Using the Search Server Menu

If you are setting up a remote search domain for Verity, you use the Search Server menu options to configure your search domain. The configuration and administration options used in implementing a search domain are identical to those used in setting up an application server domain.

Important! You use the Search Server menu to configure remote search domains for the Verity search engine only. These options do not apply to any Oracle Secure Enterprise Search (SES) components.

Related Links

[Configuring Verity Search Options](#)

[Using the Application Server Administration Menu](#)

Using the Web (PIA) Server Menu

This section provides an overview and discusses how to:

- Access the PeopleSoft web server options.
- Create a PIA domain.
- Delete a PIA domain.
- Boot a PIA domain.
- Shut down a PIA domain.
- Check PIA domain status.
- Configure a PIA domain.
- Edit PIA configuration files.
- View PIA log files.

- Administer a PeopleSoft site.
- Delete a PeopleSoft site.
- Manage Windows Services to start PIA domains.

Understanding PeopleSoft Web Server Administration Using PSADMIN

Using PSADMIN you can complete the majority of the web server administrative tasks related to your PeopleSoft Internet Architecture installation. While you can perform these tasks outside PSADMIN, using PSADMIN provides a centralized interface from which you can install, delete, configure, tune, and manage your PeopleSoft web domains and PeopleSoft sites.

For example, you can install PeopleSoft Internet Architecture domains and sites using the PeopleSoft install program (setup.exe), but you would need to run a separate BAT or SH file to start it, and navigate to log files manually to confirm it has started. Using PSADMIN you can install the domain and site, but also configure that domain, boot that domain, confirm it has started, and view log files all from the same interface. The centralized PSADMIN options provide a convenient alternative for system administrators.

Important! By default, the system installs your PeopleSoft Internet Architecture installation within the current PS_CFG_HOME location on the web server host. If you have the PIA_HOME environment variable set, this overrides the use of PS_CFG_HOME for PIA installations.

Note: PSADMIN only manages single-server PIA installations. If you intend to implement multi-server installations, those need to be created and managed using traditional methods, such as the PsMpPIAInstall program and the delivered scripts for starting and stopping domains.

Related Links

[Understanding WebLogic](#)

[Understanding WebSphere Application Server within Your PeopleSoft Implementation](#)

Accessing the Web Server Options

To access the PSADMIN web server options, select *4) Web (PIA) Server* from the main PeopleSoft Server Administration menu.

```

-----
PeopleSoft PIA Administration
-----

PIA Home:   D:\PT_SERVER\8.53

1) Administer a domain
2) Create a domain
3) Delete a domain

q) Quit

Command to execute:
```

PIA Home

Displays the location of your PeopleSoft Internet Architecture domains. The default location is PS_CFG_HOME. If your application server and web server run on the same machine, they will share the PS_CFG_HOME location, by default.

Create a domain

Select to create a PeopleSoft Internet Architecture domain.

Delete a domain

Select to delete a PeopleSoft Internet Architecture domain

Creating a PIA Domain

You create a PIA domain by selecting *Create a domain* from the PeopleSoft PIA Administration menu.

```
-----
PeopleSoft PIA Administration - Create Domain
-----
```

```
PIA Home: D:\PT_SERVER\8.53
```

- 1) Create a domain
- 2) Replicate a domain

- q) Quit

Command to execute:

Creating a Domain

Creating a PIA domain using PSADMIN provides an option to using the PeopleSoft Multiplatform PIA install program.

Note: You cannot create a multi-server PIA installation using PSADMIN.

```
-----
PeopleSoft PIA Administration - Create Domain
-----
```

```
PIA Home: D:\PT_SERVER\8.53
```

- 1) Domain name : [peoplesoft]
- 2) Web server : [weblogic]
- 3) Web server root directory : [c:\oracle]
- 4) Web server login id : [system]
- 5) Web server password : [Passw0rd]
- 6) Website name : [ps]
- 7) Application server : [RTDC79623VMC]
- 8) JSL Port : [9000]
- 9) HTTP Port : [80]
- 10) HTTPS Port : [443]
- 11) Authentication Token Domain : []
- 12) Web profile : [PROD]
- 13) Web profile user id : [PTWEBSEVER]
- 14) Web profile password : [PTWEBSEVER]
- 15) PeopleSoft reports directory : [c:\psreports]

- c) Create the domain as shown
- q) Quit

Command to execute:

Enter the number of the menu option to edit the values.

PIA Home

Ensure that the appropriate location is displayed. By default, PSADMIN installs the PIA domain into the PS_CFG_HOME environment setting.

Domain name

Enter the name for your domain, such as *PeopleSoft_HR*.

Web server	Select the appropriate web server type. Enter 2 to toggle between <i>weblogic</i> and <i>websphere</i> .
Web server root directory	Enter the location where you have installed your web server, such as <i>D:\oracle\WebLogic103464bit</i> .
Web server login id	Enter the administrative user ID for logging into the web server. This would be the same user ID you use to log in to the administrative console.
Web server password	Enter the password associated with the web server login ID.
Website name	Enter the name of the initial website, such as <i>PSHR</i> .
<hr/>	
	Note: Using PSADMIN you can only add one site. To add additional sites, you need to use the PsMpPIAInstall program.
Application server	Enter the application server host machine, such as <i>SERVER01</i> .
JSL port	Enter the JSL port on which the appropriate application server domain listens for Jolt requests.
HTTP port	Enter the HTTP port for your web server installation
HTTPS port	Enter the HTTPS port for your web server installation.
Authentication Token Domain	Enter the domain used for your authentication tokens, if implemented, such as <i>.yourcompany.com</i> .
	See See the product documentation for <i>PeopleTools 8.53 Installation</i> for your database platform: "Setting Up the PeopleSoft Pure Internet Architecture in GUI Mode," Using Authentication Domains in the PeopleSoft Pure Internet Architecture Installation
	See "Understanding the Authentication Domain" (PeopleTools 8.53: Portal Technology).
Web profile	Choose the name of the web profile you intend to use. A web profile is a named group of configuration property settings that the portal applies throughout your PeopleSoft system to control all portal-related behavior. The web profile name will be used to configure this web site. You can specify one of the predelivered web profiles, PROD, DEV, TEST, or KIOSK.
	See "Configuring Web Profiles" (PeopleTools 8.53: Portal Technology).
Web profile user ID	Enter the user ID associated with the web profile you've entered.
Web profile password	Enter the password associated with the web profile user ID.
PeopleSoft reports directory	Specify the root directory for the Report Repository, where Process Scheduler writes report output. Make sure that the

report repository directory is shared. The user that starts PIA must have write access to the Report Repository directory.

Replicating a Domain

Replicating a PIA domain enables you to copy a PIA domain configuration from another location into the current PS_CFG_HOME. This enables you to tune and configure a single PIA domain, and replicate that configuration in multiple locations. Or, you may simply need to move a configuration to a newer or different server.

Note: You can copy domains from one operating system to another, as in from Windows to Linux, but you cannot copy domains between web server types, as in from Oracle WebLogic to IBM WebSphere.

```
-----
PeopleSoft PIA Administration - Create Domain
-----
```

```
PIA Home:      D:\PT_SERVER\8.53

1) Source Domain home           : [D:\PT_SERVER\8.53]
2) Source Domain name          : [PSHR1]
3) Target Domain name          : [PSHR2]
4) Target Web server           : [weblogic]
5) Target Web server home      : [c:\oracle]
6) Target Web server user id   : [system]
7) Target Web server password  : [Passw0rd]
8) PeopleSoft reports directory : [c:\psreports]

r) Replicate the domain as shown
q) Quit
```

Command to execute:

Source Domain home

Enter the PS_CFG_HOME where the domain you want to copy is installed.

Note: The source domain home can be on the local host or a remote host.

The source domain can come from a mapped drive or mounted file system. The target domain will be created within the current PS_CFG_HOME.

Source Domain name

Enter the name of the domain to replicate.

Target Domain name

Enter the name for the new, replicated domain.

Note: The name must be unique if you are replicating within the same PS_CFG_HOME.

Target Web server

Specify the web server type installed on both the target and source host.

Enter 4 to toggle between *weblogic* and *websphere*.

Target Web server home

Enter the location where the web server software is installed.

Target Web server user ID	Enter the administrative user ID for the target web server (the user ID used to login to the administrative console).
Target Web server password	Enter the password associated with the administrative user ID on the target web server.
PeopleSoft reports directory	Enter the location of the Report Repository on the target web server.

Deleting a PIA Domain

You delete PIA domains by selecting *Delete a domain* from the PeopleSoft PIA Administration menu.

```
-----
PeopleSoft PIA Domain Administration - Choose a Domain
-----
```

```
1) PT853802
2) ptools853
```

```
q) Quit
```

Command to execute:

To delete a PIA domain:

1. From the PeopleSoft PIA Administration menu, select 3) *Delete a domain*.
2. From the PeopleSoft PIA Domain Administration - Choose a Domain menu, select the PIA domain to delete.
3. When prompted to continue, enter *y*.
4. View the screen messages to verify the system has deleted the domain successfully.

```
-----
PeopleSoft PIA Domain Administration - Choose a Domain
-----
```

```
1) PT853802
2) ptools853
```

```
q) Quit
```

Command to execute: 1

```
This operation will delete the configuration files for the domain.
If the domain is currently running it will be stopped.
Do you want to continue? [Y/N]: y
```

```
Verifying domain status.....
Removing the domain.....
The domain has been removed.
```

Booting a PIA Domain

To boot a PIA domain:

1. From the PeopleSoft PIA Administration menu, select 1) *Administer a domain*.

2. Choose the PIA domain from the domain list.
3. From the PeopleSoft PIA Domain Administration menu, select *1) Boot this domain*.

This will boot the domain and verify its status.

4. View the screen messages to verify the system has booted the domain successfully.

```
-----
PeopleSoft PIA Domain Administration
-----

PIA Home:      D:\PT_SERVER\8.53
PIA Domain:    ptools853

1) Boot this domain
2) Shutdown this domain
3) Get the status of this domain
4) Configure this domain
5) Edit configuration files
6) View log files
7) Administer a site
8) Delete a site
9) Windows Service Setup

q) Quit

Command to execute: 1

Starting the domain.....
Server state changed to STARTING.....
Server state changed to STANDBY.
Server state changed to STARTING.....
.
Server state changed to ADMIN.
Server state changed to RESUMING..
Server state changed to RUNNING.
Verifying domain status.....
The domain has started.
```

Note that the subsequent display of the PeopleSoft PIA Domain Administration menu indicates Domain Status: started.

```
-----
PeopleSoft PIA Domain Administration
-----

PIA Home:      D:\PT_SERVER\8.53
PIA Domain:    ptools853
Domain Status: started

1) Boot this domain
2) Shutdown this domain
3) Get the status of this domain
4) Configure this domain
5) Edit configuration files
6) View log files
7) Administer a site
8) Delete a site
9) Windows Service Setup

q) Quit

Command to execute:
```

Shutting Down a PIA Domain

To shut down a PIA domain:

1. From the PeopleSoft PIA Administration menu, select *1) Administer a domain*.
2. Choose the PIA domain from the domain list.
3. From the PeopleSoft PIA Domain Administration menu, select *2) Shutdown this domain*.

This will stop the domain and verify its status.

4. View the screen messages to verify the system has shut down the domain successfully.

```

-----
PeopleSoft PIA Domain Administration
-----

PIA Home:      D:\PT_SERVER\8.53
PIA Domain:    ptools853

1) Boot this domain
2) Shutdown this domain
3) Get the status of this domain
4) Configure this domain
5) Edit configuration files
6) View log files
7) Administer a site
8) Delete a site
9) Windows Service Setup

q) Quit

Command to execute: 2

Stopping the domain.....
Verifying domain status.....

Command to execute: Server "PIA" was force shutdown successfully ...

Done

```

Checking PIA Domain Status

To verify status of a PIA domain:

1. From the PeopleSoft PIA Administration menu, select *1) Administer a domain*.
2. Choose the PIA domain from the domain list.
3. From the PeopleSoft PIA Domain Administration menu, select *3) Check the status of this domain*.
4. View the `Domain Status:` line at the top of the menu to verify the domain status.

For example,

```
Domain Status: started
```

or

```
Domain Status: stopped
```


Configuring a PIA Domain

To configure a PIA domain, select *4) Configure this domain* from the PeopleSoft PIA Domain Administration menu. You use the PeopleSoft PIA Administration - Quick Configure Domain menu to modify PIA domain settings.

```
-----
PeopleSoft PIA Administration - Quick Configure Domain
-----
```

```
PIA Home:   D:\PT_SERVER\8.53
PIA Domain: ptools853
Domain Status: stopped
```

```
1) HTTP Port           : [80]
2) HTTP SSL Port      : [443]
3) Min Heap Size      : [512m]
4) Max Heap Size      : [512m]
5) Max Threads        : [50]
6) Auth Token Domain  : []
```

```
s) Save
q) Quit
```

Command to execute:

HTTP Port

The HTTP port for your web server.

HTTPS Port

The HTTPS port for your web server.

Min Heap Size

The minimum JVM Heap size for the web server instance occupied by your PIA domain.

Note: For WebLogic include the 'm' as in *512m*, but for WebSphere exclude the 'm' as in *512*.

Max Heap Size

The maximum JVM Heap size for the web server instance occupied by your PIA domain.

Note: For WebLogic include the 'm' as in *512m*, but for WebSphere exclude the 'm' as in *512*.

Max Threads

Specify the maximum number of simultaneous transactions that the PIA domain web server instance can handle.

Auth Token Domain

Enter the Authentication Token Domain.

See "Understanding the Authentication Domain" (PeopleTools 8.53: Portal Technology).

After making configuration changes and saving, PSADMIN responds differently, displaying different prompts, depending on the state of your domain, and the configuration options you've modified.

Configuration Option	Description
Standard	Can be modified while the server is running. A reboot is required for the change to take effect.

Configuration Option	Description
Dynamic	Can be modified while the server is running. Takes effect immediately with no reboot required.
Volatile	Cannot be modified while the server is running. The server must be stopped before the change can be saved.

Editing PIA Configuration Files

To view and edit PIA configuration files, select 5) *Edit configuration files* from the PeopleSoft PIA Domain Administration menu.

Note: These files are provided for informational purposes. In most cases, you will not be required to modify these files unless asked to do so by a Global Customer Support representative.

Note: To open the file editor and view or edit the PIA configuration files on a non-Windows environment, such as UNIX or Linux, you need to have an X Window System configured and running.

```
-----
PeopleSoft PIA Domain Administration - Choose a Config File
-----
```

- ```

1) setEnv
2) config.xml
3) portal/weblogic.xml
4) pspc/weblogic.xml
5) logging.properties

```

```
q) Quit
```

Command to execute:

|                                           |                                                                                                                                                 |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>setEnv</b>                             | Opens the setEnv file in your text editor. This file contains environment properties that the system initializes when you start the web server. |
| <b>config.xml</b>                         | Opens the config.xml file for the PIA domain web server instance. The config.xml file contains numerous web server settings.                    |
| <b>portal/&lt;web server type&gt;.xml</b> | Opens the <web server type>.xml file associated with the PeopleSoft Internet Architecture portal deployment.                                    |
| <b>pspc/&lt;web server type&gt;.xml</b>   | Opens the <web server type>.xml file associated with the PeopleSoft Portlet Container.                                                          |
| <b>logging.properties</b>                 | Opens the logging.properties file for adjusting global and handler-specific logging levels.                                                     |

## Viewing Log Files

To view PIA domain log files, select *6) View log files* from the PeopleSoft PIA Domain Administration menu.

```

PeopleSoft PIA Domain Administration - Choose a Log File

```

```
1) PIA_access.log 0
2) PIA_stderr.log 0
3) PIA_stdout.log 7951
4) PIA_weblogic.log 384999
5) ptools853.log 57793
```

```
q) Quit
```

Command to execute:

|                                |                                                           |
|--------------------------------|-----------------------------------------------------------|
| <b>PIA_access.log</b>          | Opens the PIA_access.log in your text editor.             |
| <b>PIA_stderr.log</b>          | Opens the PIA_stderr.log in your text editor.             |
| <b>PIA_stdout.log</b>          | Opens the PIA_stdout.log in your text editor.             |
| <b>PIA_weblogic.log</b>        | Opens the PIA_weblogic.log in your text editor.           |
| <b>&lt;domain name&gt;.log</b> | Opens the log for the current domain in your text editor. |

## Administering a PIA Site

To administer a PIA site:

1. From the PeopleSoft PIA Administration menu, select *7) Administer a site*.
2. Choose the PIA site from the site list.

```

PeopleSoft PIA Site Administration

```

```
PIA Home: D:\PT_SERVER\8.53
PIA Domain: ptools853: stopped
PIA Site: PT853803
```

```
1) Configure this site
2) Edit configuration file
```

```
q) Quit
```

Command to execute:

### Configuring a PIA Site

To configure a PIA site, select *1) Configure this site* from the PeopleSoft PIA Site Administration menu.

```

PeopleSoft PIA Administration - Quick Configure Site

```

```
PIA Home: D:\PT_SERVER\8.53
PIA Domain: ptools853
Domain Status: stopped
```

PIA Site: PT853803

```

1) Application Server Connect String : [RTDC79623VMC:9010]
2) Jolt Pooling : [Enabled]
3) AppServer Password : [{V1.1}605vpwGxd5o=]
4) Dynamic Config Reload Setting : [0]
5) Web Profile : [PROD]
6) Web User Id : [PTWEBSERVER]
7) Web Password : [PTWEBSERVER]
8) Report Repository File Path : [d:/psreports]
9) OBIEE Analytics Server URL : []
s) Save
q) Quit

```

Command to execute:

### Application Server Connect String

Specifies the application server(s) from which this web site will accept connections if Jolt failover and load balancing is implemented.

This modifies the `psserver` setting in the configuration.properties file.

See [Configuring Jolt Failover and Load Balancing](#).

### Jolt Pooling

Select to toggle Jolt Pooling [Enabled | Disabled].

This modifies the `joltpooling` setting in the configuration.properties file.

See [Configuring Jolt Session Pooling](#).

### AppServer Password

Set the domain connection password, which is required for connecting to the application server domain.

This modifies the `DomainConnectionPwd` setting in the configuration.properties file.

See [Configuring Domain Connection Password](#).

### Dynamic Config Reload Setting

This setting enables the `psserver` string (Application Server Connect String) to be reloaded dynamically. That is, the domain does not need to be restarted for the changed values to take effect.

Set Dynamic Config Reload Setting to 1 enables the feature, while setting it to 0 disable the feature. Default value is 0.

This modifies the `DynamicConfigReload` setting in the configuration.properties file.

---

**Note:** This setting applies only to the Application Server Connect String.

---

See [Configuring Dynamic Reloading of the Application Server Connection String](#).

### Web Profile

Specify the web profile used to store the configuration settings for this web site.

|                                    |                                                                                                                                                                                                                                                                                           |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                    | This modifies the <code>WebProfile</code> setting in the configuration.properties file.                                                                                                                                                                                                   |
| <b>Web User ID</b>                 | The user ID that the Portal uses to access the web profile.<br><br>This modifies the <code>WebUserId</code> setting in the configuration.properties file.                                                                                                                                 |
| <b>Web Password</b>                | The user password that the Portal uses to access the web profile.<br><br>This modifies the <code>WebPassword</code> setting in the configuration.properties file.                                                                                                                         |
| <b>Report Repository File Path</b> | The location of the Report Repository.<br><hr/> <b>Note:</b> The system uses this setting value only if a Report Repository location is not specified in the current web profile. <hr/> This modifies the <code>ReportRepositoryPath</code> setting in the configuration.properties file. |
| <b>OBIEE Analytics Server URL</b>  | Enter the location of the OBIEE analytics server, if implemented.                                                                                                                                                                                                                         |
| <b>Save</b>                        | After you have made your changes, select <i>S) Save</i> to save your changes to the PIA site.                                                                                                                                                                                             |

## Editing PIA Configuration Files

To edit the configuration.properties file directly in your text editor, select *1) Edit configuration file* from the PeopleSoft PIA Site Administration menu.

## Deleting a PIA Site

To delete a PIA site:

1. From the PeopleSoft PIA Domain Administration menu, select *8) Delete a site*.
2. Choose the PIA site from the site list.

---

**Note:** The domain must have at least two sites for the delete operation to succeed. If only one site exists in the domain, the system will not allow you to delete the only remaining site. A domain must contain at least one site.

---

3. When prompted to make sure you want to delete the site, enter *y*.

```

PeopleSoft PIA Domain Administration - Choose a Site

```

- ```

1) PSHR
2) PT853803

q) Quit
```

```

Command to execute: 1

Are you sure you want to remove the site [Y/N]: y

Removing the site.....
The site has been removed.

```

Managing Windows Services for PIA Domains

This section discusses:

- Installing a Windows Service for a PIA Domain.
- Uninstalling a Windows Service for a PIA Domain.

Installing a Windows Service for a PIA Domain

To install a Windows Service for starting a PIA domain, select 9) *Windows Service Setup* from the PeopleSoft PIA Domain Administration.

To install the Windows Service, select 1) *Install Service* from the Windows Service Setup menu.

```

-----
Windows Service Setup
-----

PIA Home:    D:\PT_SERVER\8.53
PIA Domain:  ptools853: stopped

1) Install Service
2) Uninstall Service

q) Quit

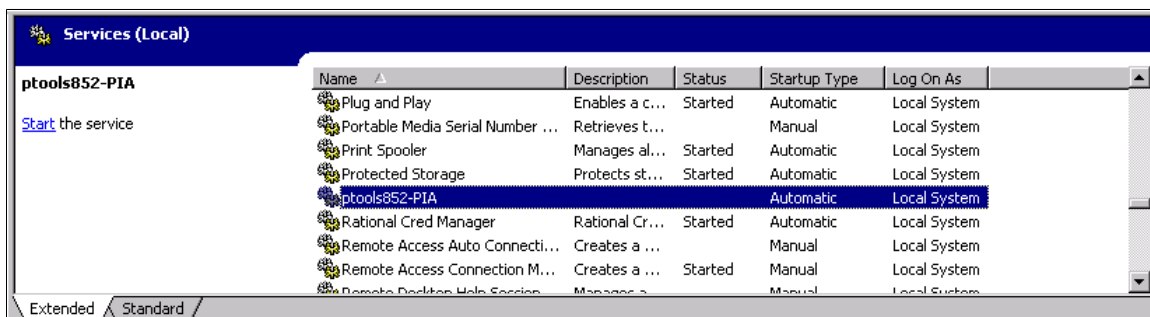
```

Command to execute: 1

Defining NT service for WebLogic admin server PIA
ptools853-PIA installed.

Image: Services console showing the newly installed Windows Service

This example illustrates the fields and controls on the Services console showing the newly installed Windows Service. You can find definitions for the fields and controls later on this page.



Uninstalling a Windows Service for a PIA Domain

To uninstall a Windows Service for starting a PIA domain, select 9) *Windows Service Setup* from the PeopleSoft PIA Domain Administration.

To uninstall the Windows Service, select 2) *Uninstall Service* from the Windows Service Setup menu.

```

-----
Windows Service Setup
-----

PIA Home:    D:\PT_SERVER\8.53
PIA Domain:  ptools853: stopped

1) Install Service
2) Uninstall Service

q) Quit

Command to execute: 2

ptools853-PIA removed.

```

Switching Configuration Homes (PS_CFG_HOME)

By default, when you install application servers, Process Scheduler servers, search servers (Verity), and PIA domains, the system installs them in the current PS_CFG_HOME location. If you do not modify or override the PS_CFG_HOME location, and you install all of these server types on the same host, all of these server types will exist and be managed by PSADMIN in the same PS_CFG_HOME location. With all of the server types installed in the same PS_CFG_HOME, you would not need to use the Switch Config Home feature. This is the recommended approach.

PSADMIN supports multiple PS_CFG_HOME locations, but you need to make sure to set the PS_CFG_HOME variable correctly before installing domains, and use the Switch Config Home option to specify the appropriate PS_CFG_HOME when performing administrative and configuration tasks.

Note: You can set the appropriate PS_CFG_HOME outside of PSADMIN manually, or from a BAT or shell script, or you can set it using the Switch Config Home feature.

To switch configuration homes:

1. Navigate to the main PSADMIN PeopleSoft Server Administration menu.
2. Select option 5) *Switch Config Home*.

```

-----
PeopleSoft Server Administration
-----

Config Home:  D:\HR_SERVER\8.53

1) Application Server
2) Process Scheduler
3) Search Server
4) Web (PIA) Server
5) Switch Config Home
6) Service Setup
7) Replicate Config Home
q) Quit

```

Command to execute (1-7, q): 5

3. At the prompt, enter the complete path to the desired PS_CFG_HOME location, and press ENTER.

Note: If the directory structure does not currently exist, PSADMIN creates it.

Please provide new Config Home location or q to quit: D:\PT_SERVER\8.53

4. Verify that the correct Config Home value appears on the PeopleSoft Server Administration menu.

```
-----
PeopleSoft Server Administration
-----
```

```
Config Home:  D:\PT_SERVER\8.53
```

```
1) Application Server
2) Process Scheduler
3) Search Server
4) Web (PIA) Server
5) Switch Config Home
6) Service Setup
7) Replicate Config Home
q) Quit
```

Command to execute (1-7, q):

Note: For PIA domains, you can set the PIA_HOME environment variable prior to starting PSADMIN. If this variable is set, PSADMIN recognizes it, and will install your PIA domains in the PIA_HOME location. You can switch to this location using the Switch Config Home feature.

Setting Up the PeopleSoft Windows Service to Start an Application Server Domain

This section provides an overview and discusses how to:

- Configure the PeopleSoft service.
- Test the service.
- Edit the service configuration file.

Note: This section applies only to Microsoft Windows servers. It involves setting application server, Process Scheduler server, and search server domains to start as PeopleSoft Windows services.

Understanding Microsoft Windows Services

A Microsoft Windows service is a Microsoft-standard package that automatically starts and stops a process when you boot or shut down a Windows system. You can also start and stop Microsoft Windows services manually through the Administrative Tools - Services utility, which you can access through the Control Panel. A service uses a standard application programming interface (API) so that it can interact with the Control Panel and log messages to the standard event log.

For PeopleSoft domains, the service starts in an environment that is separate from any users who are signed in to the system (or to the server machine). If using the Windows service, administrators do not need to log on to a machine, start PSADMIN, and enter the proper commands to start a domain. In

addition, if you use the PeopleSoft service, an administrator's logon session does not need to remain open while the domain runs.

If you have multiple application server domains, Process Scheduler servers, or search servers on the same machine, you can start them all by using the same service setup.

Note: The PeopleSoft service supersedes the method that is provided in the Microsoft Windows resource kit. Do not use SRVANY.EXE or AT commands to start PeopleSoft domains.

Configuring the PeopleSoft Service

The following procedure assumes that you have already installed and configured a domain on the Microsoft Windows server.

After completing this procedure, the specified PeopleSoft domains start and shut down automatically when the operating system recycles.

To set up the Microsoft Windows service:

1. Open the System utility within the Control Panel, and make sure the following variables are set on the Environment tab in the System Variables section:

Variable	Value
TEMP	Specify the location of the TEMP directory on the Microsoft Windows server, as in C:\TEMP.

2. In PSADMIN utility, select *Service Setup* from the main menu.
3. Select *Configure a Service* from the PeopleSoft Services Administration menu.
4. Enter *y* to indicate that you want to change configuration values.
5. Enter the names of the application server domains, the Process Scheduler databases, or search server domains that you want to include as part of the Microsoft Windows service.

To add multiple values, separate each value with a comma and a space.

For example,

```
Application Server Domains=HRDOM1, HRDOM2, HRDOM3
```

6. Select *Install a Service* from the PeopleSoft Services Administration menu.

Note: All of the domains and databases that you specified are part of a single Windows service called PeopleSoft *PS_CFG_HOME*, where *PS_CFG_HOME* is the location of the domain installations.

7. Return to the Windows Control Panel, and start the Administrative Tools, Services utility.
8. In the Services utility, scroll to find the entry that adheres to the PeopleSoft *PS_CFG_HOME* naming convention and access its properties.
9. On the General tab of the service properties, select a startup type of *Automatic*.

Note: The default startup mode is *Manual*.

10. On the Log On tab, the Log On As setting must match the Log On As setting that's defined for the Oracle ProcMGR service, which was created when you installed Tuxedo.

Both services should either be configured to *Log On As Local System Account*, or to *Log On As This Account* (referring to the same account).

11. On the General tab of the service properties, click Start.

Note: If using a secure, remote PS_HOME, the user ID that starts the Oracle ProcMGR service and the PeopleSoft service must be same domain user ID and have sufficient privileges to at least read-execute the remote resources. Also, the TM_IPC_MAPDRIVER environment variable needs to be set appropriately. The PeopleSoft Windows Service must run under the same user id as the ProcMGR service because the PeopleSoft Windows service executes the PSADMIN command line options to start and stop domains.

See [Understanding PS_HOME and PS_CFG_HOME](#).

Service Start Failure

It's possible that one or more of the domains or databases that are configured as part of the PeopleSoft service will fail to start, for reasons unrelated to the service.

The service is marked as started even if only one of its assigned domains starts. A message is written to the Windows event log for each domain, indicating whether it has started or not. If you experience problems with any domain or database, check the event log to see if it started successfully.

If all of the assigned domains and databases fail to start, the service is marked as stopped, and the following message is written to the event log:

```
Unable to start any of the domains configured for service service_name.
```

Testing the Service

To test the Microsoft Windows service, reboot the server, and make sure that the appropriate server executables are running.

For example, for the application server, use the Microsoft Windows Task Manager or the Server status option from the Domain status menu to see that the following executables are running:

- PSAPPSRV.EXE
- PSSAMSRV.EXE
- BBL.EXE
- WSL.EXE

Also make sure that any additional server processes that you have configured, such as PSQCKSRV.EXE, are running.

Editing the Service Configuration File

The Windows Services section of PSADMIN modifies the PSWINSRV.CFG file in the *PS_CFG_HOME*\appserv directory. You can edit the file directly by selecting *Edit a Service Configuration File* from the PeopleSoft Services Administration menu. This opens the PSWINSRV.CFG file in a text editor, where you can enter and save your changes.

The following sections describe each parameter.

Service Start Delay

When a domain resides on the same machine as the database server, consider using the Service Start Delay setting. By using this feature, you can avoid the situation where the database server is booting and is not ready to process requests at the time that the service attempts to boot the domain. In this scenario, without a delay set, the connection fails.

You can configure a Service Start Delay parameter that specifies a delay, in seconds, that elapses before a service attempts to start any domains. This allows the RDBMS enough time to boot and become available to accept requests.

The default is 60 seconds.

Application Server Domains

Specify the names of the domains that you want to start automatically when you boot the application server machine.

If you specify multiple domains, separate each domain with a comma and a space.

Process Scheduler Databases

Enter the databases with which a Process Scheduler server is associated. For each database that you specify, the associated Process Scheduler server starts when you boot the Microsoft Windows server.

If you specify multiple databases, separate each database with a comma and a space.

Search Server Domains

Specify the names of the domains that you want to start automatically when you boot the search server machine.

If you specify multiple domains, separate each domain with a comma and a space.

Chapter 6

Setting Application Server Domain Parameters

Startup Options

Set database sign-in values in the Startup section, so that the domain can connect to the database.

Note: When specifying passwords, such as UserPswd, Connect Password, and StandbyUserPswd you will be prompted with: `Do you want to encrypt this password?` Selecting `y` encrypts the password so that it does not appear in clear text in the configuration file (PSAPPSRV.CFG) or PSADMIN.

DBName

Enter the PeopleSoft database name, such as FSDMO or HRDMO. This parameter is case sensitive.

DBType

Enter the PeopleSoft database type, such as DB2ODBC, DB2UNIX, INFORMIX, MICROSOFT, ORACLE, or SYBASE. If you enter an invalid database type, PSADMIN prompts you with a valid list.

UserID

Enter the PeopleSoft user ID that is authorized to start the application server. For the application server to boot, the appropriate user ID with the correct authorizations must be assigned to this parameter. This is the ID that the application server passes to the database for authentication and connection. The user ID that you enter here is not related to the actual user (administrator) that carries out the boot command.

The Can Start Application Server permission must be set in the permission list. The authorization to start an application server does not (directly or indirectly) grant any authorizations or privileges beyond the ability to start the application server. Each user who attempts to sign in enters a unique user ID and password, which the application server uses to authenticate each user.

Related Links

"Defining Permissions" (PeopleTools 8.53: Security Administration)

UserPswd

Enter the password to be used by the specified user ID that will gain access to the database.

The password:

- must be specified in uppercase to simplify administration of the system.

- should not exceed 32 characters.
- (Windows) should not contain any forward-slash characters (/).
- (UNIX) should not contain any percent characters (%).

Connect ID

Required for all database platforms. Enter the database-level ID that the PeopleSoft system uses to make the initial connection to the database. This user name must have authority to select from PSACCESPRFL, PSLOCK, PSOPRDEFN, and PSSTATUS.

Connect Password

Enter the password for the connect ID. For instance, this might be the UNIX user's password (either uppercase or lowercase).

The password:

- should not exceed eight characters.
- (Windows) should not contain any forward-slash characters (/).
- (UNIX) should not contain any percent characters (%).

ServerName

Required for Sybase and Informix. Enter the name of the server on which the PeopleSoft database is installed. This value is case sensitive.

StandbyDBName

(Used only for by Oracle databases with Oracle Active Data Guard implemented). Specify the standby database name.

StandbyDBType

(Used only for by Oracle databases with Oracle Active Data Guard implemented). Specify the standby database type, such as *ORACLE*.

StandbyUserId

(Used only for by Oracle databases with Oracle Active Data Guard implemented). Specify the user ID required to connect to the standby database.

StandbyUserPswd

(Used only for by Oracle databases with Oracle Active Data Guard implemented). Specify the user password required to connect to the standby database.

Database Options

Use the Database Options section to specify environment variables that may improve the performance of the system. These options do not apply to every database.

SybasePacketSize

Enter a Transmission Control Protocol (TCP) packet size. The minimum value is 512 and the maximum value is 65538. The default packet size is 512. If you change the packet size, make the corresponding changes to the Sybase database server.

See Your Sybase documentation.

UseLocalOracleDB

(Applies only to Oracle databases.) Use this option to enable processes to initiate a local connection to a PeopleSoft database that is running on the same machine. You should use this option for all PeopleSoft Process Scheduler (batch) and application server configurations that are local (on the same server) to the PeopleSoft Oracle instance. This type of connection enables processing to complete significantly quicker.

Enter 1 to enable this option, and enter 0 to disable it.

Note: Using the local Oracle connection disables the Query Kill function.

Note: On Microsoft Windows, UseLocalOracleDB is not supported when using a 32-bit client connecting to a 64-bit database.

Related Links

"PeopleSoft Servers and the Oracle Connection String" (PeopleTools 8.53: Data Management)

EnableDBMonitoring

Required for database-level auditing. How this works varies slightly, depending on the database platform. Use this option to view more information regarding the clients that are connected to a database server through the application server. For instance, with this enabled, you can view the client machine name or user ID that is associated with a particular connection. Without this option enabled, all connections appear somewhat anonymously, as in PSFT or APPSERV.

The default value is 1 (enabled). Enter 0 to disable it.

Note: This parameter isn't supported for Informix or DB2 LUW.

PSDB Maximum Cursors

Enables you to configure the maximum number of cursors opened at one time. In general, the maximum number of cursors that can be open at one time within PeopleTools is 64,000. However, some database platforms place greater restrictions on the number of cursors that can be open at one time. By default, if no value is provided, the system allows a maximum of 1024 cursors.

To increase or decrease the default value, assign your custom value to this parameter. The value you set should be determined on multiple factors including:

- Any predefined limits imposed by your database platform.
- Internal testing of your PeopleSoft application(s).

Always refer to your database documentation for the latest recommendations related to maximum cursors settings. The following information can be used as a guideline.

Database	Maximum Cursors Limit
Oracle	No predefined limit.
IBM DB2 z/OS	Limited only by available storage.
IBM DB2 LUW	Limited only by available storage.
Microsoft SQL Server	No predefined limit.
Sybase	No predefined limit.
Informix	No predefined limit, but SQL statement length limited to 64 KB.

Note: When running Application Engine programs or Data Mover scripts you may see the following message in your log files:

```
Initializing PSDB maximum cursor = 1024
This is an informational message, not an error.
```

Security Options

Use the Security section to set an additional layer to the sign-in process.

Validate Signon With Database

Use this option to set an additional level of authorization-checking to be performed at the database level. Enter 1 to enable this option, and enter 0 to disable it.

With this option disabled, if a PeopleSoft user attempts to connect to an application server, the application server ensures that the user's PeopleSoft user ID and password exist on PSOPRDEFN. If it does not exist, the request to connect fails. This is PeopleTools-level authentication.

With this option enabled, the application server first attempts to connect to the database by using the user ID and password as part of the database connection string. If the authorization is successful, it disconnects, and then the normal PeopleSoft sign-in procedure occurs.

With this option enabled, to connect successfully to the database, the user must be defined on either the operating system or the database and within PeopleSoft.

Note: For DB2 z/OS, the user ID and password must be defined as z/OS user logon IDs.

DomainConnectionPwd

The domain connection password adds an extra layer of security between the application server domain and any connections made to it. This password enables you to further prevent unauthorized clients from establishing connections to an application server domain. It is recommended to use PSADMIN to update this value. The value can be up to 30 characters.

All domains, PeopleSoft Internet Architecture, and three-tier workstations used for a particular database, must use the same domain connection password.

Note: It is not required to add a domain connection password. It is an additional security layer for use if desired. If you add or change this value in the domain configuration, you *must* also update any PeopleSoft Internet Architecture sites and three-tier Windows workstations to reflect the new password expected by the application server domain.

For the PeopleSoft Internet Architecture configuration, you enter the password in the configuration.properties file using the same setting, DomainConnectionPwd.

For a three-tier Windows workstation connection, you enter the password in the Configuration Manager profile using the Domain Connection Password field on the Database/Application Server tab of the Edit Profile dialog box.

Related Links

[Configuring Domain Connection Password](#)

Workstation Listener Options

The workstation listener is the component to which PeopleSoft development environments running on Windows send Tuxedo messages. Development environments run Application Designer, for example.

By default, the workstation listener is disabled.

Address

%PS_MACH% resolves automatically to the machine name that PSADMIN obtains by using a system application programming interface (API) call. You can also specify the machine's Internet Protocol (IP) address (dotted notation) or its resolvable name (domain name server [DNS] name).

You should not change this value except in the following rare cases. If you are configuring files to run an application server on another machine (that is, you plan to copy PSAPPSRV.CFG and PSAPPSRV.UBB to a domain on another machine), you must overlay %PS_MACH% with the other machine's name.

Note: If you enter a literal IP address or machine name in place of the %PS_MACH% system variable, PSADMIN automatically prepends '/' to the value during the configure process. For example, if you enter *10.831.248.117* in place of %PS_MACH%, after configuring the domain, the value appears as *//10.831.248.117* in both PSADMIN and the psappsrv.cfg file.

Port

Enter the 4-digit port number to assign to the WSL. Port numbers are arbitrary numbers between 1000 and 64,000 and must not be in use by another service. The default value is 7000.

Encryption

Use this option to enable the encryption of data messages between client workstations and the application server. Specify one of the following values:

- *0* — No encryption.

Important! This is the default value.

- *40* — 40-bit encryption.
- *128* — 128-bit encryption.

Note: Because this is a dynamic parameter, you must modify it by selecting Custom Configuration on the Quick-Configure menu, and reboot the application server domain for it to take effect.

Min Handlers

Enter the number of workstation handlers (WSHs) to be started at boot time. The minimum number of handlers required depends on the domain. A smaller domain, handling less connections, may only need one WSH to start at boot time. A larger domain, handling more connections, may require ten WSHs to start at boot time.

Max Handlers

Enter the maximum number of WSHs that can be started for a domain. As connections increase, the system spawns more WSHs (up to the Max Handler value) to keep up with system demand. If the Min Handlers value equals the Max Handlers value, Tuxedo does not spawn incremental WSHs.

Max Clients per Handler

Enter the maximum number of client workstation connections that each WSH can manage. Each WSH can handle approximately 60 client connections. Numbers vary depending upon the resources of the server. In most cases, you should decrease the default as opposed to increasing it. The default is *40*.

Client Cleanup Timeout

Enter the amount of time, in minutes, that a client connection can remain idle (no work requested) before Tuxedo terminates the client connection. Client disconnects are transparent to a client, and a user just clicks the mouse to cause a reconnection. The default value for this setting is *60* minutes.

Init Timeout

This value, when multiplied by SCANUNIT (a UBB parameter value that is defined in the PSAPPSRV.UBB file) specifies the amount of time, in seconds, that Tuxedo allows for a client connection request to bind to a WSH before terminating the connection attempt.

Tuxedo Compression

Enter the minimum length of a data message for which the application server initiates data compression. While compression results in favorable performance gains for transactions over a wide area network (WAN), testing reveals that compression can degrade performance slightly over a local area network (LAN) due to the compression and decompression overhead.

You should use the default threshold of 5000, which sets a balance between WAN and LAN environments. This means that only network request and response messages over 5000 bytes are compressed, and those 5000 and under are uncompressed. If you support both WAN and LAN users, you can configure a hybrid environment by configuring two application servers: one to support WAN users (with compression set to 100) and another to support LAN users (with compression set to 100000, effectively turning compression off).

Jolt Listener Options

Use this section to configure PeopleSoft Internet Architecture connections. The Jolt listener enables Tuxedo to exchange messages with the web server.

Address

See the equivalent parameter for the workstation listener.

Port

Enter the port number that is used for the Jolt server listener (JSL). This value can be any port number that is not already in use by another service on the machine that runs the application server domain. By default, the JSL port is enabled.

Encryption

Use this option to enable the encryption of data messages between client workstations and the web server. Specify one of the following values:

- *0* — No encryption.

Important! This is the default value. Incoming Jolt requests from the web server (portal, PIA, and Integration Broker) are not encrypted.

- *40* — 40-bit encryption.
- *128* — 128-bit encryption.

Min Handlers

Enter the number of Jolt server handlers (JSH) to be started at boot time. Each JSH multiplexes up to 50 connections.

Max Handlers

Enter the maximum number of JSHs.

Note: JSHs spawn by using successive port numbers starting at the port number for the JSL in the PSAPPSRV.CFG file. Make sure that the additional ports are free before configuring spawning.

Max Clients per Handler

Enter the maximum number of client connections that each JSH can manage.

Client Cleanup Timeout

Enter the amount of time, in minutes, that a client connection can remain idle (no work requested) before Tuxedo terminates the client connection. Client disconnects are transparent to a client, and a user just clicks the mouse to cause a reconnection. The default value for this setting is 10 minutes.

Init Timeout

See the equivalent parameter for the workstation listener.

Client Connection Mode

Enter one of these options to control the allowed connection modes from clients:

- **RETAINED:** The network connection is retained for the full duration of a session.
- **RECONNECT:** The client establishes and brings down a connection when an idle timeout is reached and reconnects for multiple requests within a session. The reconnection is transparent to the user.
- **ANY:** (Default) The server allows client code to request either a **RETAINED** or **RECONNECT** type of connection for a session. Whereas, with the other two options, the server dictates from which type of client it accepts connections. This option translates to the **-c Connection Mode** parameter for the JSL section in the PSAPPSRV.UBB file.

Jolt Compression Threshold

Jolt compression can significantly improve performance. Jolt compression enables messages that are transmitted through a Jolt connection to be compressed as they flow over the network. You are likely to see the most significant performance improvements over a WAN.

For compression, the configuration files contain a default compression threshold. This default value should provide the best results for most situations. However, your application server administrator can adjust this value to suit your implementation.

The compression threshold indicates to the server how large a packet must be to require compressing. In other words, the value that you set is the minimum number of bytes that a single packet must be before the server compresses it.

Many of the XML messages being sent around the system are greater than 100,000 bytes. These messages contain HTML in compressed states, so it's generally not required that these messages be compressed. Because of this, the PeopleSoft default is set to 1,000,000 bytes.

Be careful when adjusting compression settings. If you set the threshold too high, then no packets will be large enough to be compressed. If you set the threshold too low, you may greatly reduce network traffic, but be aware that the server will have an increased workload from compressing numerous packets. Typically, you should decrease the threshold according to the bandwidth of the workstation hardware as described in the following paragraphs.

If you are handling only LAN connections, you may want to disable compression by setting the threshold to 99999999 so that only packets larger than 99,999,999 bytes are compressed. Of course, such a large value effectively disables compression so that no packets are compressed. This means no extra work for the server compressing packets.

Alternatively, if you have mostly low bandwidth, as in 56-kilobyte (KB) modem connections over a WAN, then you most likely want to compress the packets as much as possible. When decreasing the compression threshold, keep in mind that the law of diminishing returns applies. Setting the threshold much below 1000 puts an increasing load on the server, and this can nullify any performance increases that you may have gained from reduced network traffic.

Additional Prompt

After you finish all of the configuration sections, PSADMIN prompts you to configure Jolt (which is enabled by default).

If you are using the PeopleSoft Internet Architecture, you must have Jolt enabled for browser access.

Jolt Relay Adapter Options

The Jolt relay adapter is disabled by default. Unless you have a specific need for JRAD, you should skip this section.

Related Links

[Using Jolt Internet Relay](#)

Listener Address

The default is %PS_MACH%. Enter the machine on which the application server is running. See the equivalent parameter for the workstation listener.

Listener Port

This option is for advanced configurations requiring the Jolt internet relay (JRLY). The listener port listens for JRLY requests and must match the JRLY “OUT” port setting in the JRLY configuration file of

the sending machine. The port number, as in 9100, is not used unless you enter y at the prompt that asks if you want to configure JRAD.

Domain Settings

Use this section to specify general settings for the entire domain—not just for a specific component of the domain.

Domain ID

Enter the name of the application server domain. It does not need to match the name that you specified when you created the domain. This name is important only because the Tuxedo Web Monitor and PeopleSoft Watch Server (PSWATCHSRV) use it to identify application server domains and the processes associated with each machine. It should not exceed 8 characters. Generally, you should use the database name in lowercase.

Add to PATH

Enter the directory that contains your 64-bit database connectivity software, as in /apps/db/oracle/bin. If the database connectivity directory is not already specified in the path, you can set it by specifying this parameter. The value is added to the path.

On UNIX, if you don't enter a value, it uses the current directory—not the current path. To have it set by default to the current path, enter a period (.).

Note: On Windows, entries that contain a space must be surrounded by quotes. If you don't enter a value, it uses the current path.

Spawn Threshold

Enter a parameter that's supplied to Tuxedo for control of process spawning by using the -p command-line option for all server processes. The default setting (1,600:1,1) rarely needs to be changed.

This setting enables the dynamic decay of spawned server processes as the transaction volume decreases. The value can be loosely translated to mean that if, in 600 seconds, there is less than or equal to one job in the queue, the decay process begins. This is described in more detail in the timeout settings topic of this product documentation.

New server processes will be spawned according to the rule defined here. By default, if there is one outstanding request in the queue for one second or more, an additional process is spawned. Additional processes will be spawned all the way up to the Max Instances defined for that server type. If Max Instances and Min Instances are identical, this setting has no effect.

Note: This parameter applies only if, for PSAPPSRV, the value of *Max Instances* is greater than that of *Min Instances*. By default, spawning is disabled.

For more information, also see *servopts(s)* in the Oracle Tuxedo reference manual

Related Links

[Application Server Timeouts](#)

Log Directory

The log directory contains log files the system generates for a domain, such as Tuxedo logs (TUXLOG) and APPSRV logs.

The default log directory for a domain is %PS_SERVDIR\logs.

To specify a custom log directory:

1. Uncomment the Log Directory setting in the domain's PSAPPSRV.CFG file.
2. Enter the desired log directory location either directly into the PSAPPSRV.CFG file or through PSADMIN.
3. Reconfigure the domain using PSADMIN for the new setting to take effect.

When entering custom log directory locations, keep the following length restrictions in mind.

Domain Type	Log Directory Character Length Limit
Application Server	60
Process Scheduler Server	70
Search Server	80

Restartable

Enter *y* to have Tuxedo restart server processes (except the BBL process) if the server dies abnormally, as in a 'kill' on UNIX or through the Task Manager on Microsoft Windows. Otherwise, enter *n*.

Allow Dynamic Changes

Often, administrators must set a trace or performance parameter while the domain is up and running. If you enable this option, then you don't need to reboot the domain for the modified parameter value to take effect.

Enter *y* or *n* to enable or disable dynamic changes. When disabled, you must reboot (or cycle the processes) for changes to take effect.

When enabled, the server checks an internal time stamp for a particular service request to see if any values have changed for the parameters for which dynamic changes are valid. If values have changed, the system uses the modified parameter value.

You should enable this option in your test and development domains. For production environments, you should enable dynamic changes selectively.

These parameters allow dynamic changes:

- Recycle Count.

- Consecutive service failures.
- Trace SQL and Trace SQL Mask.
- Trace PC and Trace PC Mask.
- Trace PPR and Trace PPR Mask.
- Log Fence.
- Enable DB Monitoring.
- Enable Debugging.
- Dump Memory Image at Crash.
- Dump Managed Objects at Crash.
- Log Error Report.
- Mail Error Report.
- SMTP Settings (all except SMTPGuaranteed, SMTPTrace, and SMTPSendTime).
- Analytic Instance Idle Timeout.
- Analytic Per Server Log.

Note: The parameters that allow dynamic changes are also identified through comments in the PSAPPSRV.CFG file. Look for the phrase “Dynamic changes allowed for X,” where *X* is the parameter name. This option does not apply to configuration parameters that Tuxedo relies on, such as the number of processes, whether restart is enabled, the port numbers, the amount of handlers, and so on.

LogFence

Enter a level of network tracing, ranging from –100 (suppressing) to 5 (all). The default is 3.

The trace file is generated in *PS_CFG_HOME*\appserv*domain*\LOGS\psappsrv.log.

LogFieldSeparator

Separate individual fields of data within the domain log file(s). If this setting is not explicitly defined, the default value is a single SPACE character. This setting must be uncommented in the configuration file to be enabled.

These meta values can be used as the field separator:

- *TAB*
- *SPACE*

While the actual key strokes can be used to specify a separator, the escape sequence equivalent of them is not supported. For example \t is not recognized as a tab character. For purpose of clarity, avoid key strokes that are not visible to the naked eye, instead use their meta value equivalents.

AppLogFence

This setting is not available through the PSADMIN interface, but can be entered directly into the PSAPPSRV.CFG file.

You can use this parameter conditionally to determine whether you want to do specific logging from your application. You can implement this parameter from PeopleCode using the %AppLogFence system variable. Complete documentation for this option is in the PeopleCode Developer's Guide.

See "Using Application Logging" (PeopleTools 8.53: PeopleCode Developer's Guide).

Trace-Log File Character Set

Enter the character set (ANSI or UNICODE) of the machine to which you typically write and read the traces and log files. If the character sets are not matched between the file and the machine, the file is unreadable.

PeopleCode Debugger Options

Use this section to enable and configure the PeopleCode debugging environment. Configuring PeopleCode debugging is discussed in detail in another section of this PeopleBook.

Related Links

[Setting Up the PeopleCode Debugger](#)

Trace Options

This section enables you to specify the tracing options that you can enable on the application server to track the Structured Query Language (SQL) and PeopleCode of the domains. You can also set all of the trace parameters from the PeopleSoft sign-in page. Just beneath the Sign In button, click the link that opens the trace flags page. This enables you to set the trace options and then sign in to the system.

Note: With many of the following trace options, you need to view the comments in the PSAPPSRV.CFG to understand what to enter to return the trace information you require.

TraceSQL

Enter the logging level for SQL tracing for all clients. Traces are written to *PS_CFG_HOME/appserv/domain/LOGS/domain_user_IDservername.tracesql*.

Enter 0 to disable tracing; enter 7 to enable a modest tracing level for debugging. For other levels of tracing, set this option to a value that equals the sum of the needed options. For example, to trace only SQL, enter 1; to trace SQL statements and connect statements enter 7 (1 + 2 + 4 = 7). A setting of 7 is recommended for troubleshooting connection and other basic problems. Tracing can consume large amounts of disk space over time, so be sure to reset this option to 0 when you finish troubleshooting.

Important! The trace file stores elapsed times for SQL events to a precision of one microsecond (six decimal places). However, due to limitations of the operating system, Windows precision is actually in milliseconds (three decimal places), so the last three digits in a Windows trace will always be zero. Elapsed times in UNIX are accurate to one microsecond.

TraceSQLMask

Enter the logging level ceiling for SQL tracing for individual clients. Traces are written to *PS_CFG_HOME/appserv/domain/LOGS/client_user_IDservername.tracesql*.

For Windows clients, you specify the necessary SQL tracing level by using the PeopleSoft Configuration Manager on the Trace tab. To prevent clients from turning on the application server trace and consuming resources, the application server uses TraceSQLMask as an administrative control facility.

If a client transmits a request to trace SQL, the application server compares the value that is transmitted to the TraceSQLMask value. If the client value is less than or equal to the TraceSQLMask value, the application server enables the trace. However, if the client value is greater, the application server enables the trace up to the TraceSQLMask value. Trace files are written on the application server; no trace shows up on the client workstation.

TracePC

Enter a level for PeopleCode tracing for activity that is generated by all clients on a domain. Eligible values are defined in the configuration file. TracePC values are displayed in the PeopleSoft Configuration Manager on the Trace tab. You can find the results in *PS_CFG_HOME/appserv/domain/LOGS/domain.log*.

Important! The trace file stores elapsed times for PeopleCode events to a precision of one microsecond (six decimal places). However, due to limitations of the operating system, Windows precision is actually in milliseconds (three decimal places), so the last three digits in a Windows trace will always be zero. Elapsed times in UNIX are accurate to one microsecond.

TracePCMask

Enter which PeopleCode trace options that are requested by client machines will be written to the trace file. You can find the results in *PS_CFG_HOME/appserv/domain/LOGS/client_machine.domain.log*.

TracePPR and TracePPRMask

Use these options to trace the activity in the page processor. Typically, these options are used internally only by PeopleSoft developers; however, you may need to view the results of this trace when troubleshooting.

Tracing-related display processing is useful for seeing when and if related displays are being updated and if they are updating successfully. Some sample output in the log file from setting this flag includes:

```
Starting Related Display processing
Related Display processing - PPR_RELDSPVALID not set
Related Display processing - All Rows
  Starting Related Display processing for - PSACLMENU_VW2.MENUNAME
  Related Display processing for - PSACLMENU_VW2.MENUNAME - completed successfully
Finished Related Display processing
```

By using the keylist generation tracing in addition to the related display tracing, you can determine why the related displays have the wrong value. It shows where the keys are coming from. The following is a sample of keylist generation tracing:

```
Starting Keylist generation
  Keylist generation - FIELDVALUE is a key
  FIELDVALUE is low key
  Low key value was supplied =
  Key FIELDVALUE =
  Keylist generation - FIELDNAME is a key
  Keylist generation - Finding value for USRXLATTABLE_VW.FIELDNAME
  Not Found in key buffer
  Searching for field FIELDNAME in component buffers
  Scanning level 1
    Scanning record DERIVED_USROPTN for field FIELDNAME
    Field FIELDNAME found in record DERIVED_USROPTN
  Found in component buffers, value = PT_TIME_FORMAT
  Key FIELDNAME = PT_TIME_FORMAT
  Keylist generation - USEROPTN is a key
  Keylist generation - Finding value for USRXLATTABLE_VW.USEROPTN
  Not Found in key buffer
  Searching for field USEROPTN in component buffers
  Scanning level 1
    Scanning record DERIVED_USROPTN for field USEROPTN
    Scanning record PSUSROPTLIST_VW for field USEROPTN
    Field USEROPTN found in record PSUSROPTLIST_VW
  Found in component buffers, value = TFRMT
  Key USEROPTN = TFRMT
  Keylist Generation complete
  FIELDNAME = PT_TIME_FORMAT
  FIELDVALUE =
  USEROPTN = TFRMT
```

In this example, you can see how the system builds the keylist by first searching in the current record (key buffer), then searching the buffers in the current level, and then searching up a level, and so on. It also indicates exactly what record the key value is taken from. This is useful on complex components where there are often several instances of a particular field; a common problem is that the value is derived from an unexpected location.

Combining the keylist tracing and the related display tracing provides a good view of the system behavior. For example:

```
Starting Related Display processing

Related Display processing - All Rows
  Starting Related Display processing for - PSACLMENU_VW2.MENUNAME
  Starting Keylist generation
  Keylist generation - MENUNAME is a key
  MENUNAME is low key
  Low key value was supplied = APPLICATION_ENGINE
  Key MENUNAME = APPLICATION_ENGINE
  Keylist Generation complete
  MENUNAME = APPLICATION_ENGINE
  Related Display processing for - PSACLMENU_VW2.MENUNAME - completed successfully
```

Each related display goes through the keylist generation process, and you can see exactly what key values are used to populate the related displays and where those key values came from.

The work record flag is a performance feature. If every field in a level-0 record has a value from the keylist and is display-only, then it is marked as a work record because the values can't be changed. After it is marked as a work record, that affects how the record behaves. For example, PeopleCode for fields in

the record but not in the component don't run, data isn't saved, and so on. By enabling this tracing option, you can see which records are flagged as work records. The output looks like this:

```
Work flag cleared for record PSCLASSDEFN_SRC
Work flag cleared for record PSCLASSDEFN_SRC
Work flag cleared for record PSCLASSDEFN
Work flag cleared for record PSPRCSPRFL
Work flag cleared for record SCRTY_QUERY
Work flag set for record PSCLASSDEFN
Work flag set for record PSPRCSPRFL
Work flag set for record SCRTY_QUERY
```

Because the flag is turned on and off at various points, the last setting (set or cleared) is the most important. In the previous trace, PSCLASSDEFN, which is marked as a work record, is cleared and then set again.

TracePIA and TracePIAMask

Use these options for tracing PeopleSoft page (PIA page) generation.

TraceAE

Use this parameter to activate specific Application Engine traces for tracing Application Engine programs.

TraceAnalytic and Trace AnalyticMask

The bits enable logging for analytic servers beyond the standard LogFence setting.

TracePPM

The Performance Monitor agent is a thread that reports performance metrics for each instrumented server if monitoring is enabled for the database. Select *1* to enable and *0* to disable.

See Performance Monitor.

DumpMemoryImageAtCrash

This parameter determines whether or not a memory image of the failing process is created when a crash occurs. By default, the value is 'NONE' which means that a memory image will not be written during a crash. This value can be set to 'MINI' or 'FULL'. MINI means that a shorter memory image is written. This may be a better option if you are leaving this option turned on permanently. FULL may be appropriate when you are debugging a known issue. Typically, this option is used internally only by PeopleTools.

DumpManagerObjectsAtCrash

This parameter assists PeopleSoft in debugging any crash problems at your site by providing insight into the customized object definitions to reproduce the crash on another database.

Log Error Report, Mail Error Report

If you enter `y` (enabled) and runtime errors are detected (nonfatal error conditions), the system writes a message and information regarding the runtime error to the current log file. If you assign the `MailErrorReport` parameter an email address, an individual, such as a system administrator, can be alerted whenever the system writes an error to the log. If `MailErrorReport` is enabled but `LogErrorReport` is set to `n`, the system still sends a message for application server crashes, which always cause data to be written to the log. The following is an example of setting this parameter to send notifications to an email address: `MailErrorReport=tom.sawyer@bigcompany.com`.

Cache Settings

Use this section to specify how to handle caching at your site. Enabling caching on the application server improves performance, but you need to review the available cache options and determine which option best suits your site.

Note: The majority of the cache settings need to be uncommented in the `PSAPPSRV.CFG` file.

EnableServerCaching

With `EnableServerCaching`, you specify what objects the system stores in cache on the application server. To enable application server disk caching the value must be set to 1 or 2.

If you enter 1 the system caches only the most used classes of objects, and if you enter 2, the system caches all object types regardless of the frequency of use. Which option you select depends on internal testing at your site.

To disable application server caching, set this value to 0. In most cases there is no reason to disable server caching. Doing so significantly degrades performance, because it requires the application server to retrieve an object from the database each time the system needs it.

CacheBaseDir

This setting is the location where cache files will be written and stored for this domain.

Note: You must preload your shared cache before you enable shared caching on the application server. Application Engine processes are independent from application server domains, directories, and configuration files. Therefore, Application Engine processes do not share cache with application server domain processes.

ServerCacheMode

Caching enables the system to store definitions locally on the application server, eliminating unnecessary trips to the database server. If caching were disabled, the system would need to retrieve definitions from the database with each request, every time. This adds significant overhead to each transaction and affects system response times.

If server caching is enabled on the application server, which is usually the recommended approach, there are two modes of caching from which to choose.

Cache Mode	Description
Non-shared cache	<p>By default, non-shared cache mode is enabled (<code>ServerCacheMode</code> set to <code>0</code>).</p> <p>With the non-shared cache mode, each server process that starts within a domain maintains its own separate cache file.</p> <p>For example, assume you had two PSAPPSRV processes configured in a domain. In non-shared cache mode, there is one cache directory per PSAPPSRV server process, which each individual PSAPPSRV process uses separately.</p> <p>In this case, you can find the cache files in <code>PS_CFG_HOME\appserv\domain\cache\PSAPPSRV_1and\PSAPPSRV_2</code>.</p> <p>While the cache directories will grow over time to include the most used definitions, you have the option to <i>preload</i> the non-shared cache directories with the most used system definitions.</p> <p>See Configuring an Application Server Domain to Preload Cache.</p>
Shared cache	<p>To set shared caching for the domain, enter 1 at the Set <code>ServerCacheMode</code> prompt. With this option enabled, you can find the cache files in <code>PS_CFG_HOME\appserv\domain\cache\share</code>.</p> <p>The system assumes that a preloaded cache exists in the share directory. The preloaded cache contains most instances of the managed object types that are cached to file. When you boot the application server, if shared cache files are enabled but no cache files exist in the expected location, the system reverts to unshared caching.</p> <p>To preload shared cache, you run delivered Application Engine programs that build your shared cache.</p> <p>See Load Application Server Cache.</p>

MaxCacheMemory

PeopleTools stores metadata in a memory cache on the server machine to increase system performance. However, in rare situations if the memory cache becomes too large, it can reduce the amount of memory available to other processes running on the server, which can degrade general performance. Use the `MaxCacheMemory` setting in certain situations where it is required to specify the maximum allowable size of the memory cache.

Note: Typically, this setting only needs to be implemented in rare circumstances, or when instructed to do so by GCS support staff. In most situations, it is recommended that sufficient memory be installed on the server machine to ensure optimal performance.

The `MaxCacheMemory` setting enables you to establish a size limit in megabytes, that the memory cache should not exceed. Every time the system retrieves an object from the database, such as a page, it updates the object's `LastUsedDate` value. When your system reaches the memory cache threshold controlled by

the MaxCacheMemory setting, the system "prunes" the oldest objects in the memory cache until the desired memory cache size is reached.

Note: If a memory cache limit has been set, the system prunes the memory cache to keep it roughly 10% below the specified limit.

The default value of this setting is 0 (disabled), which disables memory cache pruning altogether, allowing for an unlimited memory cache. That is, when MaxCacheMemory is set to 0, the system continues to store a memory cache, but the system will not attempt to prune or reduce the size of the memory cache.

The default setting (disabled) might not be optimal for your application. You can adjust this setting to achieve the best trade-off between speed and available memory based on your system requirements. If you are in a situation where you need to limit the size of the memory cache, enter that limit in megabytes. Make sure to test your memory cache limit thoroughly in realistic scenarios, reflecting your production environment.

EnableDBCACHE

If enabled, the metadata (cache) is accessed from database, rather than the file system.

To implement database caching, uncomment the parameter and enter *Y* to enable database caching, or *N* to disable database caching. After changing this parameter, the domain does not need to be reconfigured.

The database cache is shared by all domains that enable this option.

When database caching is enabled, these settings are ignored:

- EnableServerCaching
- ServerCacheMode
- CacheBaseDir

You can load the database cache using the Load Application Server Cache utility or the preload cache utility.

Note: Database caching is also available for Process Scheduler domains.

Related Links

[Load Application Server Cache](#)

[Configuring an Application Server Domain to Preload Cache](#)

PreLoadCache and PreLoadMemoryCache

If you have created a cache project, specify the project name. Before booting a domain, you have the option to preload the cache, and this option creates the cache based on the load project you specify.

If database caching is enabled, the cache is preloaded into the database.

Related Links

[Configuring an Application Server Domain to Preload Cache](#)

Remote Call Options

RCCBL Redirect

Enter *0* to disable redirection and *1* to enable redirection. Redirection causes the server process to retain intermediate work files that are used to pass parameter values between the server process and a Remote Call/COBOL program for debugging purposes, and it also generates Remote Call trace information in the COBOL output files.

If set to *0*, the system generates no COBOL output files. If set to *1* (enabled), the system generates:

- rmtcall_in.<pid>
- rmtcall_out.<pid>
- rmtcall_cblmsg.<pid>
- Remote Call Tracing entries for COBOL output files.

Note: Where <pid> is the Process ID.

Note: RCCBL Redirect should be enabled only when debugging because writing Remote Call trace entries to the COBOL output files can add to system overhead and overall performance.

Remote Call Trace Information

The information captured for Remote Call tracing follows this general structure:

<i>Trace Field</i>	<i>Description</i>
Remote Call trace identifier (RCCT)	Remote Call COBOL Trace (RCCT) helps you to identify the Remote Call tracing in the COBOL output files.
Remote Call Component	Enables you to identify the type of Remote Call being traced: <ul style="list-style-type: none"> • PSPCM • PTPECOBL • PTPNETRT • PTPSQLRT • Application COBOL

Trace Field	Description
Trace information	<p>The the trace record can be different depending on what is being traced:</p> <ul style="list-style-type: none"> In some cases there are parameter/value pairs written in the record, such as: <pre>RCCT: PSPCM: pSamParms->szDBName QE850GA RCCT: PTPNETRT: DATA-FIELD-REG BEGIN ->a<- END</pre> In other cases there will be trace information that reports the program flow: <pre>RCCT: PTPNETRT: RemoteCall PSCOBNET function CBL_ GET_OS_INFO</pre>

Remote Call Output Files

The location and naming of the COBOL output files differ depending on whether the Remote Call is invoked on the application server or the Process Scheduler server (by the PSAESRV server process).

Location	Description
Application Server	<p>If the Remote Call is invoked on the application server, the COBOL output file(s) are written to:</p> <pre>PS_CFG_HOME/appserv/<APP_SERVER_DOMAIN>/LOGS</pre> <p>The output file with the Remote Call trace entries follows this naming convention:</p> <pre><COBOL_PGM>_<OPERID>_<TIMESTAMP>.out</pre>
Process Scheduler Server	<p>If Remote Call is invoked on the Process Scheduler server (PSAESRV), the COBOL output file will be written to:</p> <pre>PS_CFG_HOME/appserv/prcs/<PROCESS_SCHEDULER_DOMAIN>/log_output/CLB_ <Program Name>_<Process Instance></pre> <p>The output file with the Remote Call trace entries follows this naming convention:</p> <pre><COBOL_PGM>_<TIMESTAMP>.out</pre> <p>For example, if a program named PTPSTAE runs it generates the following:</p> <pre>PS_CFG_HOME/appserv/prcs/HCMPCRS/log_output/CBL_PTPSTAE_94/PTPECOBL_ _1121155222.out</pre> <hr/> <p>Note: PTPECOBL is the initial target COBOL program invoked by the Remote Call function because PTPECOBL handles the parameter passing from the PeopleCode program to the target application Remote Call COBOL program.</p> <hr/>

RCCBL PRDBIN

Use this parameter to specify where RemoteCall can find the COBOL executables. By default, RemoteCall looks in a predefined location. This might not be the location where you've installed them on your system:

- In UNIX, RemoteCall looks in `$PS_HOME/cblbin`.
- In Windows, RemoteCall looks in `%PS_HOME%\cblbin%PS_COBOLTYPE%`. The `%PS_COBOLTYPE%` variable contains a single letter that indicates the character encoding for the database platform. It's automatically set to one of the following values when the application server starts:
 - *U* — Unicode.
 - *A* — Non-Unicode.
 - *E* — EBCDIC.

To override this default behavior, set `RCCBL PRDBIN` to the absolute path of your COBOL executables, for example:

- In Windows: `RCCBL PRDBIN=c:\pscobol\MYDOMAIN\cblbin`
- In UNIX: `RCCBL PRDBIN=/app/pssoft/MYDOMAIN/cblbin`

Note: This parameter doesn't appear in the PSADMIN custom configuration interface if it's not already set. You must define it by editing the application server configuration file directly. On the PeopleSoft Domain Administration menu, select *Edit configuration/log files menu*, then select *Edit psappsrv.cfg (current configuration file)* to open `psappsrv.cfg` in a text editor. Define the `RCCBL PRDBIN` parameter in the RemoteCall section of the file.

PSAPPSRV Options

The PSAPPSRV server process performs the functional requests, such as building and loading panel groups. It also provides the in-memory-caching feature for PeopleTools objects on the application server. Each server process maintains its own cache.

Min Instances

Enter the minimum number of application server instances that start when you boot the domain. There's always at least this number of instances running. This translates to the PSAPPSRV server's `-m (min)` parameter in the UBB file.

Max Instances

Enter the maximum number of server instances that can be started. This translates to the PSAPPSRV server's `-M (Max)` parameter in the UBB file.

Service Timeout

Enter the number of seconds that a PSAPPSRV waits for a service request, such as MgrGetObj or PprLoad, to complete before timing out. Service timeouts are recorded in the TUXLOG and APPSRV.LOG. In the event of a timeout, PSAPPSRV is terminated and a replacement process is started by Tuxedo.

Recycle Count

Enter the number of service requests that each server has carried out before being terminated (intentionally) and then immediately restarting. Servers must be intermittently recycled to clear buffer areas. The time that is required to recycle a server is negligible, occurring in milliseconds. The recycle count does not translate into a native Tuxedo parameter in the PSAPPSRV.UBB file. Instead, the value is stored in memory and is managed by a PeopleSoft server.

Note: The default value for the small, medium, and large templates is 10,000.

Note: Serial Recycle or Dynamic Recycle parameters can delay the recycle even if the number of the requests is greater than or equal to the Recycle Count value.

Percentage of Memory Growth

The Percentage of Memory Growth option, enables you to perform dynamic recycling, in a test environment, so that you can arrive at a static Recycle Count value suited to your production system. The Percentage of Memory Growth parameter indicates the percentage of memory growth to reach before the PSAPPSRV process will automatically restart (dynamically recycle). The default is 20, meaning an additional 20% of memory growth will be incurred after the process has established a baseline memory cache.

A dynamic recycling configuration is intended to be used in testing environment where a usage load, representative of your production usage load, can be run against a system. During this load test, administrators can monitor log files to determine when the dynamic recycle occurs. The intervals and trends related to the dynamic recycling can help to identify an appropriate Recycle Count value for a production environment. When the load test is complete, the Recycle Count value should be adjusted appropriately, and the dynamic recycle feature should be disabled.

When dynamic recycling is configured, the system evaluates the memory size after every 100 service requests to determine if the process needs to be recycled. If any new PeopleTools objects have been loaded into the memory cache, or if a JVM has been initialized since the last memory evaluation, the memory baseline is reset to the current value and no recycling will occur.

If no PeopleTools objects have been loaded, and if the JVM has not been initialized since the last memory evaluation, and if the memory footprint has grown by more than the specified percentage, the system recycles the PSAPPSRV process and adds messages to the APPSRV_xxyy.log file.

Note: Due to the overhead involved in measuring the memory usage, dynamic recycling is *not* recommended for use in a production environment.

Note: Unless you can emulate, in your test environment, a usage load *representative* of a typical production usage load, the results of your test will be of little value for determining the optimal Recycle Count value.

Configuring Dynamic Recycling

To configure dynamic recycling:

1. Open PSAPPSRV.CFG.
2. In the PSAPPSRV section, set Recycle Count to 0 to disable the fixed recycle interval.
3. Uncomment the Percentage of Memory Growth parameter, and adjust the value as needed to a value between 1–100, where 100 = 100%.

```
Percentage of Memory Growth=20
```

4. Restart the domain.

Viewing Logs for Dynamic Recycle Count Information

Recycle messages are logged with the service count indicated at that point. The system logs messages based on the following scenarios:

Memory Evaluation Scenario	Message Logged
Memory could not be determined.	Use Recycle Count for recycling due to unable to obtain virtual memory size In this case, the message is logged at the "Status Level". This message appears at LogFence=3.
The current service count is less than the Recycle Count, although the memory usage exceeds the Percentage of Memory Growth value.	This indicates that the Recycle Count was set to a value other than 0. Should this condition occur, the system logs this message: delay dynamic recycle(services=100, Recycle Count=500)
At intervals of 100 and if the LogFence is set to at least 4.	dynamic recycle results [recycle=true, orig_mem=1000000, current_mem=2000000, max_mem=3000000, check_count=100, obj_loaded=100, jvm_loaded=true]

Allowed Consec Service Failures

Enter a number greater than 0 to enable dynamic server processes to restart for service failures. To disable this option, enter 0. The default is 2. The value that you enter is the number of consecutive service failures that will cause a recycle of the server process. This is a catchall error handling routine that enables PSAPPSRV, PSQCKSRV, and PSAMSRV to terminate themselves if they receive multiple, consecutive, fatal error messages from service routines. Such errors should not occur consecutively, but if they do, the server process must be recycled or cleansed. A retry message appears on the client browser when this occurs.

Max Fetch Size

The default is 5000 (K). Enter the maximum memory that is used by the server to store fetched rows for a transaction before sending the result set back to a client. If the memory limit is exceeded, the client receives the rows retrieved with a memory buffer exceeded warning. You should use the default value. PSAPPSRV supports nonconversational transactions, so this parameter provides a way to balance high-volume throughput with the needs of users working with large volumes of data. A value of 0 means unlimited memory is used. The memory is not preallocated—it is acquired as needed for each transaction.

Auto Select Prompts

Enter 1 (the default) to enable automatic prompting on lookup pages. When the user selects the prompt lookup button, the application server automatically returns all values for that field, up to 300 rows. If necessary, the user can refine the search further by entering partial data in the Search By field.

Enter 0 to require the user to enter a partial value before the automatic prompt list appears.

AutoLoad JVM

You can add this parameter, if needed, manually to the [PSAPPSRV] section of the PSAPPSRV.CFG file. AutoLoad JVM controls whether the JVM gets loaded automatically when the domain boots.

By default, domain behavior reflects a setting of AutoLoad JVM=0 (not enabled).

If you have configured EnablePPM Agent=1, then JVM will be loaded at the domain boot time. If you have configured EnablePPM Agent=0, then to load JVM when the domain boots, you need add AutoLoad JVM=1 to the [PSAPPSRV] section of the PSAPPSRV.CFG file.

Serial Recycle

When serial recycling is enabled for a server process, the system recycles server processes of that type within a domain on a serial basis—one after another—to allow processing to continue uninterrupted. By default, the domain behavior reflects Serial Recycle=Y (enabled).

When Serial Recycle=Y for PSAPPSRV, for example, then only one PSAPPSRV process will recycle during the recycle time of 60 seconds when the Recycle Count limit is reached. Then the next PSAPPSRV will recycle.

When serial recycling is not enabled, all the server processes of that type recycle simultaneously when the Recycle Count limit is reached, which can cause throughput to pause.

Serial recycling applies to these server process types: PSAPPSRV, PSQCKSRV, PSQRYSRV, and the Integration Broker server processes.

To disable serial recycling, manually add the Serial Recycle parameter, and assign the value *N*. For example:

```
Serial Recycle=N
```

The recycle time is 60 seconds.

PSANALYTICSRV Options

PSANALYTICSRV relates to the server processes that are associated with the *analytic server framework*. Most of the parameters for PSANALYTICSRV need to be reviewed in the PSAPPSRV.CFG file and uncommented as needed.

Related Links

PeopleSoft MultiChannel Framework

Min Instances

Enter the minimum number of analytic server instances that start when you boot the application server domain. There's always at least this number of instances running. The default value of this parameter is 3.

Max Instances

Enter the maximum number of analytic server instances that can result from dynamically spawning new processes. The default value of this parameter is 3.

Analytic Instance Idle Timeout

Enter the number of minutes of inactivity before the analytic instance times out and is unloaded.

This value takes effect only if the PeopleCode AnalyticInstance class Load method specifies a value of *-1* for its IdleTimeOut parameter when loading an analytic instance. This includes Load PeopleCode that's launched from an analytic grid, which enables you to avoid having to explicitly specify a timeout.

The default value of this parameter is *0* (no timeout limit) for domains that are configured with a developer template, and *30* minutes for other domains.

PSSAMSRV Options

The PSSAMSRV server process communicates through the Tuxedo conversational mode. It performs transactional SQL requests (updates).

Min Instances

Enter how many servers are started at boot time. This translates to the PSSAMSRV server's *-m* (min) parameter in the UBB file.

Max Instances

Enter the maximum number of servers that can be started. This translates to the PSSAMSRV server's *-M* (Max) parameter in the UBB file.

Service Timeout

Enter the number of seconds that the server processes waits for a request before timing out. This stops runaway processes, like an rcbl timeout.

Recycle Count

Enter the number of service requests that each server carries out before being terminated (intentionally). Tuxedo immediately restarts the server. Servers must be intermittently recycled to clear buffer areas. The time that is required to recycle a server is negligible, occurring in milliseconds. The recycle count does not translate into a native Tuxedo parameter in the PSAPPSRV.UBB file. Instead, the value is stored in memory and is managed by a PeopleSoft server.

Allowed Consec Service Failures

Enter a number greater than zero to enable dynamic server process restarts for service failures. To disable this option, enter 0. The default is 2. The value that you enter is the number of consecutive service failures that cause a recycle of the server process. This is a catchall error handling routine that enables PSAPPSRV, PSQCKSRV, and PSSAMSRV to terminate themselves if they receive multiple, consecutive, fatal error messages from service routines. Such errors should not occur consecutively, but if they do, the server process must be recycled or cleansed. A retry message appears on the client browser when this occurs.

Max Fetch Size

The default is 32 (K). Enter the maximum memory that is used by the server to store fetched rows for a transaction before sending results to the client and refilling the memory buffer. When the memory limit is reached, the server sends rows to the client, but then resumes refilling the buffer and sending results to the client until the query is complete. You should leave the default value unchanged.

PSSAMSRV supports conversational transactions, so this parameter enables users to tune performance by adjusting the number of network round-trips that are required for the average transaction. A value of 0 causes unlimited memory to be used, which means one round-trip no matter how large the result set. The memory is not preallocated, but is acquired as needed.

PSQCKSRV Options

The PSQCKSRV is an optional server process to improve performance. Essentially, the PSQCKSRV, or quick server, is a copy of the PSAPPSRV. It performs quick requests, such as nontransactional (read-only) SQL requests. The PSQCKSRV improves overall performance by enabling the PSAPPSRV process to direct a portion of its workload to PSQCKSRV.

Min Instances

Enter how many servers are started at boot time. This translates to the PSQCKSRV server's `-m` (min) parameter in the UBB file.

Max Instances

Enter the maximum number of servers that can be started. This translates to the PSQCKSRV server's `-M (Max)` parameter in the UBB file.

Service Timeout

Enter the number of seconds that a PSQCKSRV waits for a request before timing out. This stops runaway processes, like an `rcdbl` timeout. This applies to incremental PSQCKSRV servers that are dynamically started by the Max Instances parameter.

Recycle Count

Use the PSAPPSRV specifications.

Allowed Consec Service Failures

Enter a number greater than zero to enable dynamic server process restarts for service failures. To disable this option, enter 0. The default is 2. The value that you enter is the number of consecutive service failures that will cause a recycle of the server process. This is a catchall error handling routine that enables PSAPPSRV, PSQCKSRV, and PSAMSRV to terminate themselves if they receive multiple, consecutive, fatal error messages from service routines. Such errors should not occur consecutively, but if they do, the server process must be recycled or cleansed. A retry message appears on the client browser when this occurs.

Max Fetch Size

Use the PSAPPSRV specifications.

Serial Recycle

Use the PSAPPSRV specifications.

PSQRYSRV Options

PSQRYSRV handles the SQL that is generated by PeopleSoft Query (PSQED.EXE). With PSQRYSRV configured, SQL-intensive, complicated, user-defined queries are offloaded to a dedicated server process, thus freeing PSAPPSRV and PSQCKSRV to handle the SQL requests for which they are more suited.

Note: When running PS/nVision reports from a three-tier, Windows client connection, the system also routes the SQL generated by both matrix (ledger-based) and tabular (PS Query-based) reports through PSQRYSRV if it is enabled.

PSQCKSRV also processes SQLRequest services; however, if PSQRYSRV is configured, it processes all SQLRequests that are initiated specifically by PSQuery (SQLQuery:SQLRequest) or PS/nVision.

Like the PSQCKSRV server process, PSQRYSRV is an optional server process. However, if you allow users to initiate queries from PeopleSoft Query, you should take advantage of this server process.

Min Instances

Enter how many servers are started at boot time. This translates to the PSQRYSRV server's `-m (min)` parameter in the UBB file.

Max Instances

Enter the maximum number of servers that can be started. This translates to the PSQRYSRV server's `-M (Max)` parameter in the UBB file.

Service Timeout

Enter the number of seconds that PSQRYSRV waits for a request before timing out. This stops runaway processes.

Recycle Count

Enter the number of service requests that each server carries out before being terminated (intentionally) by Tuxedo and then immediately restarted. Servers must be intermittently recycled to clear buffer areas. The time that is required to recycle a server is negligible, occurring in milliseconds.

If the recycle count is set to 0, PSQRYSRV is never recycled.

Allowed Consec Service Failures

Enter a number greater than 0 to enable dynamic server process restarts for service failures. To disable this option, enter 0. The default is 2. The value that you enter is the number of consecutive service failures that will cause a recycle of the server process. This is a catchall error handling routine that enables PSAPPSRV, PSQCKSRV, PSQRYSRV, and PSSAMSRV to terminate themselves if they receive multiple, consecutive, fatal error messages from service routines. Such errors should not occur consecutively, but if they do, the server process must be recycled or cleansed. A retry message appears on the client browser when this occurs.

Max Fetch Size

Enter the maximum size (in KB) of a result set that is returned from a SELECT query. Use 0 for no limit.

Use Dirty-Read

(Applies only to DB2 systems.) Enter 1 to enable the application server to read uncommitted data from a table. Enter 0 to disable dirty reads.

This parameter can be used for general reporting and PeopleSoft Query. You can run dirty read queries through the application server, the Process Scheduler, and in a two-tier connection. The Use Dirty-Read setting in the application server configuration controls the behavior of PSAPPSRV, PSQCKSRV, and PSQRYSRV.

Note: Dirty reads are not recommended if you are reading data and doing subsequent processing based on the disposition of the data at the time that it is read. Between the time the data is read by a subsequent process and the time the unit of work is completed by the first process, any activity affecting the table data at the time a subsequent process read could be rolled back, invalidating the accuracy of the data that a subsequent process read.

Related Links

PeopleSoft Query

"Setting Parameters for the Application Engine Server" (PeopleTools 8.53: PeopleSoft Process Scheduler)

Serial Recycle

Use the PSAPPSRV specifications.

Integration Broker Server Processes

A variety of server processes are devoted to Integration Broker processing. These servers are configured in these sections:

- Publish&Subscribe
- PSBRKDSP
- PSBRKHND
- PSSUBDSP
- PSSUBHND
- PSPUBDSP
- PSPUBHND

These server processes act as dispatchers and handlers of the messages being distributed by the Integration Broker.

Related Links

PeopleSoft Integration Broker Administration

SMTP Settings

You can send electronic mail requests, issued with workflow or PeopleCode, to the application server, which passes the requests to the specified mail server (SMTPServer). To specify the appropriate SMTP server and port to receive the email requests, you must edit the SMTP Settings section.

When set in the PSAPPSRV.CFG file, these SMTP settings are not dynamic: SMTPGuaranteed, SMTPTrace, SMTPSendTime. They require a domain reboot to take effect.

Note: You can also control most of these settings using the PeopleCode SMTPSession class, which temporarily overrides them without changing the values in PSAPPSRV.CFG. You use this class primarily when you want to send multiple emails in a single session of the SMTP server, instead of having to change the permanent SMTP settings for every email. In the event that your application PeopleCode does not specify its own SMTP settings, the system uses the settings in the PSAPPSRV.CFG file.

See "SMTPSession Class" (PeopleTools 8.53: PeopleCode API Reference).

SMTPServer

Enter the host name and IP address of the mail server machine.

SMTPPort

Enter the port number on the mail server machine.

SMTPServer1

Enter the host name and IP address of the failover mail server machine in case the other specified server is down.

SMTPPort1

Enter the port number on the failover mail server machine.

SMTPSender

Enter the sender's internet address. This must be a valid address, such as user1@xyzcorp.com.

SMTP BlackBerryReplyTo

Enter the internet address that you want to be the reply to address for Blackberry email responses. This must be a valid address such as user1@xyzcorp.com.

SMTPSourceMachine

Enter the sender's source machine name and internet address in the form of MACHINE.XYZCORP.COM. This value is required in some, but not all environments.

SMTPCharacterSet

Enter the character set that is used on the sender's machine.

SMTPEncodingDLL

Enter the name of a dynamic-link library (DLL) that is used to translate the mail message from the sender's character set (such as latin1, sjis, big5, gb, ks-c-5601-1987, or ks-c-5601-1992) to a 7-bit safe character set for transmission.

SMTPGuaranteed

Set this parameter to *1* if you want TriggerBusinessEvent email PeopleCode to be delivered through the Integration Broker system, which provides some additional administration capabilities for ensuring delivery of the message. By enabling this feature you implement a mechanism to ensure that emails get routed to the appropriate place in case SMTP mail fails for reasons such as network timeouts, downed mail servers, invalid parameters, and so on.

This setting enables the system to track email messages sent using Integration Broker queues. When sending an email with this option enabled, the system performs an asynchronous local-to-local publish, and for the subscription the system calls MCFOutboundMail.send to email the message.

SMTPTrace

Enter *1* to enable the tracing of email details to the log file. Enter *0* to disable SMTP tracing.

LogFence should be set to *5*.

The system writes the log information to SMTP<DDMM>.log in %PS_SERVDIR%/LOGS, by default, or the custom value set for Log Directory. For example:

```
PS_CFG_HOME\APPSERV\domain\LOGS\SMTP6_27.log
```

SMTPSendTime

Enter *1* to have messages contain a send time that is populated by the application server. Enter *0* to leave the send time blank and have it populated by the receiving gateway (depending on the gateway).

SMTPUserName

Enter the user name to log in to the SMTP server. This applies only when authentication is enabled on the SMTP server.

SMTPUserPassword

Enter the password for the user specified by *SMTPUserName* to access the SMTP server. This applies only when authentication is enabled on the SMTP server.

Use PSCipher to encrypt the SMTPUserPassword value.

See "Using the PSCipher Utility" (PeopleTools 8.53: Security Administration).

SMTPUserName1

Enter the user name to log in to the failover SMTP server. This applies only when authentication is enabled on the failover SMTP server.

SMTPUserPassword1

Enter the password for the user specified by *SMTPUserName1* to access the failover SMTP server. This applies only when authentication is enabled on the failover SMTP server.

Use PSCipher to encrypt the SMTPUserPassword1 value.

See "Using the PSCipher Utility" (PeopleTools 8.53: Security Administration).

SMTPTimeToWaitForResult

Enter the time in milliseconds for the mail system to wait for the result of sending each email. If the time is set to 0, the system doesn't wait, and the returned result will be always be *%ObEmail_SentButResultUnknown* (= -1). If you set this parameter to -1, the system will wait for the completion of the send process. The default value of this setting is 10000 (ten seconds).

SMTPUseSSL

Indicates that SSL connections are enabled. Enter *Y* for yes or *N* for no. The default value is *N*.

SMTPSSLPort

If using SSL, specify the SSL port on the SMTP server. The default is 465.

SMTPClientCertAlias

If the SMTP server is configured for client authentication, enter the alias name of the client certificate.

SMTPSSLPort1

Applies to the failover SMTP server.

If using SSL, specify the SSL port on the SMTP server. The default is 465.

SMTPUseSSL1

Applies to the failover SMTP server.

Indicates that SSL connections are enabled. Enter *Y* for yes or *N* for no. The default value is *N*.

SMTPClientCertAlias1

Applies to the failover SMTP server.

If the SMTP server is configured for client authentication, enter the alias name of the client certificate.

SMTPDNSTimeoutRetries

Enables you to configure the number of times the `IsDomainNameValid` method of the `MCFMailUtil` class retries to verify that the domain of an email address submitted to the method is valid. If you need to override the system default, manually add this parameter to the [SMTP Settings] section of the `PSAPPSRV.CFG` file.

Related Links

"`IsDomainNameValid`" (PeopleTools 8.53: PeopleCode API Reference)

Implementing NT LAN Manager

PeopleTools supports NT LAN Manager (NTLM) as a primary authentication mechanism. NTLM is a suite of authentication and session security protocols used in various Microsoft network protocol implementations. To use NTML with PeopleTools, you need to add these parameters manually to the PSAPPSRV.CFG file and PSPRCS.CFG file in the SMTP settings section.

Parameter	Description
SMPNTLMEnable	Specify whether NTLM is enabled. Add <i>true</i> or <i>false</i> .
SMPNTLMDomain	Specify the NTLM domain.
SMPNTLMEnable1	Specify whether NTLM is enabled for the backup server.
SMPNTLMDomain1	Specify the NTLM domain for the backup server.

Implementing Simple Authentication and Security Layer

PeopleTools supports Simple Authentication and Security Layer (SASL) as a primary authentication mechanism. SASL is a framework for authentication and data security in many internet protocols. While there are several SASL mechanisms, PeopleTools supports PLAIN SASL only. To use SASL with PeopleTools, you need to add these parameters manually to the PSAPPSRV.CFG file and PSPRCS.CFG file in the SMTP settings section.

Parameter	Description
SMTPSASLEnable	Specify whether SASL is enabled.
SMTPSASLMechanism	Specify the SASL mechanism.
SMTPSASLAuthorizationid	Specify the authorization ID
SMTPSASLRealm	Specify the realm.
SMTPSASLEnable1	Specify whether SASL is enabled for the backup server.
SMTPSASLMechanism1	Specify the SASL mechanism for the backup server.
SMTPSASLAuthorizationid1	Specify the authorization ID for the backup server.
SMTPSASLRealm1	Specify the realm for the backup server.

SMTP Further Considerations

Keep in mind the following considerations:

- PeopleSoft mail integration is on the application server only.

PeopleSoft software does not support VIM/MAPI, because this option is client-side-only integration, and PeopleSoft Internet Architecture applications run on the server-side.

- The application server communicates directly with the SMTP server through telnet by using standard SMTP commands with Multipurpose Internet Mail Extensions (MIME) 1.0 messages.
- PeopleSoft software currently supports UTF-8 encoding of the e-mail messages by default, and you can encode e-mail messages in other ways.

With server-side integration, you do not have to certify any specific e-mail client application. You can use any application to read e-mail.

- PeopleSoft recommends using the Multichannel Framework mail classes for all e-mail sent from a PeopleSoft application.

Related Links

"Understanding PeopleSoft MultiChannel Framework Mail Classes" (PeopleTools 8.53: PeopleCode API Reference)

Interface Driver Options

Set the following parameter for configuring the interface driver for business interlinks.

Note: PeopleSoft Business Interlinks is a deprecated product and these options currently exist only for upgrade compatibility and transition.

SCP_LOCALE

Enter the RPS_LOCALE string, which the driver sends to the Supply Chain Planning (SCP) server.

PSTOOLS Options

You may need to set the following parameters in advanced configurations.

EnablePPM Agent

Controls whether the Performance Monitor agent runs and collects performance monitor data. Enter *1* to enable the Performance Monitor agent, and *0* to disable the Performance Monitor agent. This setting overrides the value for this parameter that is set in the database.

Related Links

Performance Monitor

Add to CLASSPATH

The Add To CLASSPATH environment variable parameter enables the Java Virtual Machine (JVM) and other Java applications where to find the Java class libraries, including any user-defined class libraries. Because PeopleTools automatically generates CLASSPATH entries for core, delivered class libraries, use

this variable to specify additional class libraries that the PeopleSoft software needs to access. To use this parameter, you need to uncomment it in the PSAPPSRV.CFG file.

The PeopleCode API Reference provides details on where you can place custom and third-party Java classes.

See "System Setup for Java Classes" (PeopleTools 8.53: PeopleCode API Reference).

JavaVM Options

Specify additional options to be passed to the JVM loaded by the domain's server processes. Separate the options with spaces, for example:

```
-Xrs -Xmx256m -Xms256m
```

If the domain will run as a Windows service, you must specify at least the default option, `-Xrs`.

Note: If you are using the AIX operating system, these options may be required: `JavaVM Options=-Xrs -Djava.awt.headless=true -Xcheck:jni`.

The JavaVM Options parameter specified in the [PSTOOLS] section specifies global JavaVM options used by every server process in a domain. To override this global value for a particular server process, you can apply custom JavaVM options to individual server processes by adding the JavaVM Options parameter manually to the configuration section for that server process.

JavaVM Options can appear multiple times in a single PSAPPSRV.CFG or PSPRCS.CFG file. While the JavaVM Options value in the [PSTOOLS] section applies to all server processes governed by a particular configuration file, the system only uses the JavaVM Options value in the [PSTOOLS] section for server processes that do not have the JavaVM Options parameter added to its configuration settings section.

For example, if the JavaVM Options parameter has been added to the [PSAPPSRV] section of the PSAPPSRV.CFG file and has been assigned a value, then that value will be used when loading the JVM as a thread of that PSAPPSRV process. If the JavaVM Options parameter does not appear, or has no value, in the [PSAPPSRV] section, then the system uses the value specified in the [PSTOOLS] section when loading the JVM as a thread of the PSAPPSRV process. This applies to any server process: PSAPPSRV, PSQRYSRV, PSAESRV, and so on.

See your JRE documentation for valid JVM options.

Note: If the server runs behind a proxy server, and Java code will be invoked, you must set the `-D` Java options to reflect your proxy connection settings. For example, Java code invoked by any Tuxedo-managed server, such as the application server or Process Scheduler server, requires the following Java options: `-Dhttp.proxyHost=<value>`, `-Dhttp.nonProxyHosts=<value>`, and `-Dhttp.proxyPort=<value>`.

Proxy Host

If the HTTP destination, such as the gateway host, is behind a proxy server for security reasons, enter the distinguished name of the proxy server, as in `proxy.oracle.com`.

Note: The proxy settings configured in the PSAPPSRV.CFG file do not apply to the embedded JVM. For Java code invoked on a PeopleSoft server (application server, Process Scheduler server, and so on) that requires a connection to the proxy, set the Java proxy settings using the Java VM Options setting in the PSAPPSRV.CFG file or PSPRCS.CFG file.

Proxy Port

Enter the port number on which the proxy server is listening for transmissions. For instance, 80 is a typical default port number.

Non Proxy Hosts

Enter a list of the hosts that should be connected to directly, not through a proxy server. Separate the hostnames with a pipe symbol (|). You can use an asterisk (*) as a wildcard character to specify a pattern of similar hostnames.

For example, `localhost|*.oracle.com` indicates that the local host and all hosts with names ending in `.oracle.com` will be accessed directly.

Note: The length of this string cannot exceed 1024 characters.

Usage Monitoring State

Required only if you are enabling the Usage Monitor, which generates system usage metrics using Performance Monitor technology. The usage metrics can then be incorporated with the PeopleSoft Testing Framework to enable you to design more efficient test plans that focus efforts on the elements of the system most used or most affected by and update or upgrade. Usage Monitor must be enabled in the domain before usage metrics can be captured.

Value	Description
0	Usage Monitor is disabled. No usage information will be captured.
1	Usage Monitor is enabled, all users are monitored, and all users are anonymous, as in usage activity is not associated with a particular user ID. Anonymous data is collected unless the user configures the Test Name and Test Case Name fields on the Usage Monitoring page.
2	Usage Monitor is enabled, all users are monitored, and user <i>are not</i> anonymous. That is, usage activity <i>is</i> associated with a particular user ID.

Note: There are other options that need to be enabled within Performance Monitoring in addition to the Usage Monitor State parameter.

Related Links

"Usage Monitor" (PeopleTools 8.53: Performance Monitor)

Character Set

Enter the character set to be used as the system locale. This setting is used to convert text to and from UNICODE when using interfaces that do not support UNICODE. Such interfaces include file names, text file contents, and other operating system calls that require non-UNICODE text. This setting controls how files used by the PeopleTools file attachment feature are named. The default values are:

Operating System	Value
UNIX	<i>latin1</i>
z/OS	<i>ccsid1047</i>
Windows	<i>CP_ACP</i>

Note: By default, the parameter is commented out. With no value specified, invalid values or "" (blank value), the system assumes the above default values as the character set.

To override the default, uncomment the Character Set parameter, and select a character set from the following list corresponding to the languages that this application server will process. For example, to set the character set for Simplified Chinese:

```
Character Set=gb2312
```

Note: The character set of the application server and the character set of any Microsoft Windows workstations connecting to that application server must match. The *utf8* option is valid only when the locale character set is UTF-8.

Related Links

"Character Sets Across the Tiers of the PeopleSoft Architecture" (PeopleTools 8.53: Global Technology)

Suppress App Error Box (Microsoft Windows Only)

Enter *y* to suppress an application error box or message from appearing after an application error occurs. Enter *n* to view error dialogs and message boxes.

Note: If the system generates an error box for an application server process and this parameter is set to *n*, Tuxedo can't restart the down process until you close the error box.

DbFlags

The following values are valid for the DbFlags parameter:

Value	Description
<i>0</i>	Enable the %UpdateStats meta-SQL construct.
<i>1</i>	Disable the %UpdateStats meta-SQL construct.

Value	Description
2	Ignore the Truncate command for DB2 LUW. Use Delete instead.
4	<p>Disable a secondary database connection (used with the GetNextNumberWithGapsCommit PeopleCode function).</p> <p>This prevents the creation of a secondary database connection, bundling all SQL into a single unit of work. Without the additional database connection, the database row lock is held for a longer time, reducing concurrency in a multiple-user environment.</p> <hr/> <p>Important! Analytic instance processing requires a secondary database connection, so if you're using analytic servers, ensure that this value is not set. In addition, the CallAppEngine PeopleCode function requires a secondary connection when used with the allowcommit parameter set to <i>True</i>. Therefore, ensure that this value is not set when using CallAppEngine with allowcommit set to <i>True</i>.</p>
8	<p>Disables persistent secondary database connections (used by either the GetNextNumberWithGapsCommit PeopleCode function or the CallAppEngine PeopleCode function when the allowcommit parameter is set to <i>True</i>).</p> <p>When this value is specified, any secondary database connection created by GetNextNumberWithGapsCommit or by CallAppEngine is immediately closed after the commit completes execution. This keeps the number of database connections to a minimum, but requires each call to create a new secondary database connection on demand.</p> <hr/> <p>Important! The performance impact of creating each new secondary database connection is significant, especially in high volume user production environments. Therefore, don't use this setting without carefully considering its effect.</p>

DbFlags uses a bit mask so that you can specify one or more of these values. You set this parameter to the total of the values that you want to apply. For example, to disable %UpdateStats and ignore the Truncate command, set DbFlags to 3 (setting bits one and two).

The default is value is 1.

See "Using PeopleCode in Application Engine Programs" (PeopleTools 8.53: Application Engine), "PeopleCode Built-in Functions and Language Constructs" (PeopleTools 8.53: PeopleCode Language Reference).

Suppress SQL Error

This option is not available through the PSADMIN interface, but it exists in the PSTOOLS section of the PSAPPSRV.CFG file for small, medium, and large configurations.

For security purposes, this option has a default value of 1 to prevent SQL error details from being displayed to users. Any SQL errors that occur don't display details, but refer users to consult the system

log. The details that were in the SQL message are written to the log file. This helps to prevent SQL injection vulnerabilities.

If you want SQL error details to be visible to users, set this property as follows:

```
Suppress SQL Error=0
```

Note: In developer configurations, the Suppress SQL Error option doesn't exist in PSAPPSRV.CFG, and the system assumes a value of 0.

Related Links

"Preventing SQL Injection" (PeopleTools 8.53: PeopleCode Developer's Guide)

Integration Broker Options

The following parameters apply to Integration Broker. All of the options for Integration Broker are discussed in detail in the Integration Broker PeopleBooks.

Related Links

[PeopleSoft Integration Broker Administration](#)

Min Message Size for Compression

Use this parameter to configure the threshold of a message before the system compresses the message.

Thread Pool Size

Set the thread pool size used by the SyncRequest PeopleCode event. The Minimum value is 1 and maximum allowable is 20.

Search

These options enable you to configure PeopleSoft search. These options are documented in detail in another section of this PeopleBook.

Note: If you do not specify a search configuration type, the system assumes the default configuration based on your operating system.

Related Links

[Configuring Verity Search Options](#)

Search Indexes

Use this option to specify the location of all the files pertaining to the search index. Index name is the same as the location. This option is documented in detail in another section of this PeopleBook.

See [Specifying the Verity Index Location](#), [Sharing Indexes Between Application Servers and PeopleSoft Process Scheduler](#).

PSRENSRV Options

PSRENSRV is a modified web server designed for real time event notification. The primary purpose of PSRENSRV is to publish events to the browser.

See "Understanding REN Servers" (PeopleTools 8.53: PeopleSoft MultiChannel Framework).

log-severity_level

This is the log severity level for the PSRENSRV process. Settings are Error, Warning, Notice or Debug. Default is Warning.

io_buffer_size

This is the TCP buffer size when serving content. This should not exceed a value of 65536.

default_http_port

This is the REN servers http port. The default value is 7180.

default_https_port

This is the REN servers https port. The default value is 7143.

default_auth_token

The fully qualified domain name of the server (network domain, not the PeopleSoft application server domain). This value should match the value of the web server's authentication token domain.

PSPPMRSRV Options

PSPPMRSRV servers subscribe to performance metrics published by the web service at the PPMI URL (specified in the Performance Monitor administration pages) and insert them into the database. If you select *Y* when you are asked whether you want Performance Collators configured, then the number of PSPPMRSRVs specified in Min Instances=1 will be started. Min and Max instances should be set to the same value, as new PSPPMRSRV servers *are not* spawned on demand.

Related Links

Performance Monitor

Min Instances

The number of servers started at boot time. This translates to the PSPPMSRV server's `-m` (min) parameter in the UBB file.

Max Instances

The maximum number of servers that can be started. This translates to the PSPPMSRV server's `-M` (max) parameter in the UBB file.

Select Server Process Options

After you enter all of the previous parameter values for the domain, PSADMIN prompts you for the following server process options. You can use these prompts to reduce the number of server processes that start when the domain boots. This, in turn, makes your configuration simpler while conserving system resources.

For instance, if you enter `n` for any of the following prompts, the corresponding server process (or a set of server processes) is not configured for the domain. If you enter `n` to all of the prompts, the domain will contain only the minimum required server processes.

Note: You can also enable and disable server processes using the Quick Configure Menu.

Do you want the Publish/Subscribe servers configured?

If you want the messaging server processes to be configured and booted, enter `y`. If you are not implementing the Integration Broker technology for a domain, enter `n`.

Note: In addition to setting this option, in Integration Broker you must also activate the domain on which the pub/sub server resides before you can use the pub/sub system.

See PeopleSoft Integration Broker Administration.

Move quick PSAPPSRV services into a second server (PSQCKSRV)?

Enter `n` if very few clients access the domain and concurrency is not an issue. Enter `y` to enable the PSQCKSRV in situations where concurrency and optimal transaction throughput are needed.

Move long-running queries into a second server (PSQRYSRV)?

If you want all user-generated queries to be initiated by PSQuery and handled by a dedicated server process, enable this option to improve overall performance.

Do you want JOLT configured?

The Jolt listener is required to support the PeopleSoft Internet Architecture by enabling transmission between the web server and the application server.

Do you want JRAD configured?

JRAD applies to specific configurations only. Accept the default unless you are attempting to configure JRAD for use with the Jolt internet relay.

Related Links

[Using Jolt Internet Relay](#)

Do you want WSL Configured?

Configures the Workstation Listener for Development Environment (Windows) workstation connections.

Do you want to enable PeopleCode Debugging (PSDBGSRV)?

Enter *y* to debug PeopleCode programs with the current domain.

Do you want Event Notification configured (PSRENSRV)?

Select *Y* to start the PSRENSRV servers.

See [PSRENSRV Options](#).

Do you want MCF Servers configured?

Select *Y* to start the Multi Channel Framework servers.

See [PeopleSoft MultiChannel Framework](#).

Do you want Performance Collators configured?

If the domain is servicing a Performance Monitor database, select *Y* to start the PSPPMRSRV servers.

See [PSPPMRSRV Options](#).

Do you want Analytic Servers configured?

Configures analytic servers to run in the domain to process Analytic Calculation Engine requests and to perform optimization processing.

Do you want Domains Gateway configured?

Enable this option if you are configuring a remote, or external, search server to which this domain will send search requests. That is, if you are configuring a Type-3 search option for an application server domain, you need to enable the domains gateway on the application server domain to a communication connection between the application server and its remote search domain.

Working with Oracle WebLogic

Understanding WebLogic

You use the Oracle WebLogic Server as a web server in the PeopleSoft Internet Architecture to deploy PeopleSoft applications. The PeopleSoft Internet Architecture installation on the WebLogic Server provides these primary server configuration options.

- Single server.

This domain configuration contains one server named PIA, and the entire PeopleSoft application is deployed to it. This configuration is intended for single user or very small scale, non-production environments.

- Multi server.

This domain configuration contains seven unique server definitions and a WebLogic cluster, and the PeopleSoft application is split across multiple server instances on a *single* machine. This configuration is intended for the production environment.

- Distributed managed server.

This option is an extension of the “Multi server” selection and installs the necessary files to boot managed servers spread across *multiple* machines. This option requires a “Multi server” installation to be performed to some other location that contains the configuration for this managed server.

Note: The topics related to WebLogic in this PeopleBook discuss PeopleSoft-specific subjects. Always refer to your Oracle WebLogic documentation for all general server administration information.

Working With the WebLogic Server Administration Console

The WebLogic Server Administration Console is the main utility that is used to administer and monitor the WebLogic server processes.

Note: For the sake of brevity, in this document the WebLogic Server Administration Console will be referred to as the Administration Console.

The Administration Console provides an interface to monitor and configure aspects of a PeopleSoft application from a web server perspective.

Access the Administration Console by pointing your browser to:

`http://<server>:<port>/console`

For example,

<http://server08.us.oracle.com:80/console>

Where the port is the port on which your PeopleSoft Internet Architecture domain is listening.

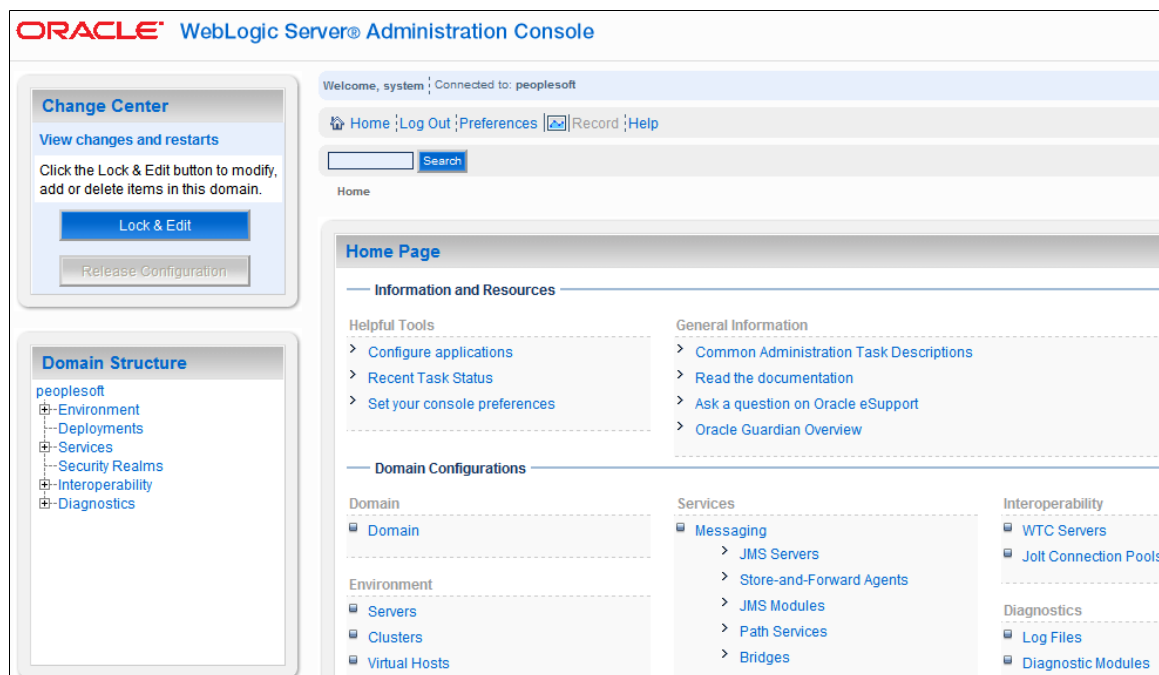
If it is a multi server installation the Administration Console from the Admin Server is always running on port 9999. For accessing the Administration Console in a multi server installation, the URL is:

<http://<server>:9999/console>

Before the Administration Console opens, you will be prompted for the WebLogic administration ID and password that you specified during the PeopleSoft Internet Architecture installation.

Image: Oracle WebLogic Server Administration Console

This example illustrates the fields and controls on the Oracle WebLogic Server Administration Console. You can find definitions for the fields and controls later on this page.



The Domain Structure section enables you to navigate to a variety of interfaces related to the elements of your PeopleSoft domain.

Note: Before making any configuration modifications or performing any administrative operations, such as stopping the server, always click Lock & Edit first.

Starting WebLogic

This section discusses how to:

- Start WebLogic on Microsoft Windows.
- Start WebLogic on UNIX.

Related Links

[Administering a WebLogic Server Life Cycle](#)

Starting WebLogic on Microsoft Windows

On Microsoft Windows, you can start WebLogic using:

- command prompt.
- Windows service.

Using the Command Prompt

Running WebLogic as a foreground process is beneficial if you need to monitor WebLogic in real time. To run WebLogic as a foreground process, enter one of the following commands at the command prompt in `PIA_HOME\webserv\domain\bin`.

Configuration Type	Command
Single server	<pre>startPIA.cmd</pre> <p>Or, you can use the PSADMIN PIA domain menu options.</p>
Multi server - domain admin server	<pre>startWebLogicAdmin.cmd</pre>
Multi server - managed server	<pre>startManagedWebLogic.cmd server</pre> <p>For example,</p> <pre>startManagedWebLogic.cmd PIA</pre>

Using the Windows Service

Benefits of running WebLogic as a Windows service include:

- WebLogic can automatically start when the Windows server boots.
- You can start and stop the service from a remote Windows machine.

To install the service, open the command prompt, and enter the appropriate command from your `PIA_HOME\webserv\domain\bin` directory:

Configuration Type	Command
Single server	<pre>installNTservicePIA.cmd</pre>
Multi server	<pre>InstallNTservice.cmd weblogic_server_ instance_name</pre> <p>For example:</p> <pre>installNTservice.cmd PIA</pre>

To start WebLogic as a Windows service, you can:

- Start the service named *WebLogicdomain-servername* (for example, peoplesoft-PIA) by using the Services utility in the Windows Control Panel.
- Start the service from a command prompt by entering the following command:

```
NET START peoplesoft-PIA
```

Note: If WebLogic fails to start as a service, try starting it as a foreground process to make sure there are no general issues related to the startup.

To uninstall the service, enter the following command:

```
uninstallNTservicePIA.cmd
```

Starting WebLogic on UNIX

To start PeopleSoft on UNIX execute the appropriate script in the WebLogic domain directory that the PeopleSoft install created, as in PIA_HOME/webserv/peoplesoft/bin.

Configuration Type	Command
Single server	startPIA.sh
Multi server - domain admin server	startWebLogicAdmin.sh
Multi server - managed server	startManagedWebLogic.sh server For example, startManagedWebLogic.sh PIA

When you run the above scripts, the server runs as background process.

Stopping WebLogic

You can stop Weblogic using:

- Administration Console
- command line

Related Links

[Administering a WebLogic Server Life Cycle](#)

Stopping WebLogic Using the Administration Console

To stop the PeopleSoft server:

1. In the left pane of the console, expand Environment, and select Servers.

2. In the Summary of Servers table, click the name of the server that you want to shut down.
3. Select the Control tab, and then the Start/Stop tab.
4. From the Shutdown drop-down menu, select one of the following options:
 - *When work completes*: This command initiates a graceful shutdown, which gives WebLogic Server subsystems time to complete application processing currently in progress.
 - *Force Shutdown now*: This command initiates a forced shutdown, in which the server instructs subsystems to immediately drop current requests.
5. Click Yes to confirm and shut down the server.

Note: If you shut down the Administration Server, the Administration Console is no longer active.

Stopping WebLogic Using the Command Line

To stop WebLogic using the command line, enter the appropriate command.

Configuration	Windows	UNIX
Single server	stopPIA.cmd	stopPIA.sh
Multi server (WebLogic Admin server)	stopWebLogic.cmd	stopWebLogic.sh
Multi server (Managed WebLogic server)	stopWebLogic.cmd <ManagedServerName>	StopWebLogic.sh <ManagedServer Name>

If WebLogic is running as a Windows service you can also stop the service in Windows Control Panel or using the NET STOP command.

Using WebLogic Server Administration Console to Monitor PeopleSoft Sessions

The Administration Console can display a list of established HTTP sessions for an instance of the WebLogic Server. Session Monitoring is automatically enabled for WebLogic. These instructions describe how to monitor the single server configuration of PIA in a test environment.

When in a production environment, note that a multi server configuration would be used to perform these steps to the server instance that you intend to monitor, such as PIA1 or PIA2, or both.

1. Start the PIA server.
2. Log on to PeopleSoft.

If possible, log on from different workstations using different PeopleSoft IDs. For the purpose of this test, do not log off.

3. Log on to the WebLogic Server Administrative Console.

4. Monitor established HTTP sessions for the PORTAL web application.

On the left, use the following navigation to view the list of established HTTP sessions for the PORTAL web application:

- a. Click Deployments, and view the deployment list in the right hand window.
- b. Click PeopleSoft.
- c. Select the Control tab.
- d. Select the PORTAL application module, where the context root of the module is '/'.
- e. Select the Monitoring tab.
- f. Select the Sessions tab.

Note: You can customize the list of fields that you want to monitor using the Customize this table link.

Image: WebLogic Server Administration Console – Monitoring Sessions

This example illustrates the fields and controls on the WebLogic Server Administration Console – Monitoring Sessions. You can find definitions for the fields and controls later on this page.

Settings for /

Overview Configuration Security Testing Monitoring

Web Applications Servlets Sessions PageFlows Workload

Use this page to view statistics about the sessions associated with this Web application.

Last Refreshed: Mar 24, 2009 2:33:40 PM

Customize this table

Servlet Sessions(Filtered - More Columns Exist)

Showing 1 to 2 of 2 Previous | Next

Context Root	Server	Creation Time	Time Last Accessed	Application	Monitoring ID
	PIA	Tue Mar 24 14:33:00 PDT 2009	Tue Mar 24 14:33:01 PDT 2009	peoplesoft	PSADMIN@141.144.74.106/QEDMO
	PIA	Tue Mar 24 13:30:56 PDT 2009	Tue Mar 24 14:31:19 PDT 2009	peoplesoft	QEDMO@141.144.171.211/QEDMO

Showing 1 to 2 of 2 Previous | Next

Note: An established HTTP session remains on the web server until the client logs out of PeopleSoft or until the user's HTTP session times out. Closing the browser does not log out a PeopleSoft user. As a result, when a user closes the browser without logging out of the PeopleSoft session, the corresponding HTTP session remains on the web server until it times out. HTTP session timeouts are controlled by the site's PeopleSoft web profile.

Related Links

[Tuning Performance and Monitoring Resources](#)

Setting Up Reverse Proxy Servers

This section provides an overview and discusses how to:

- Configure Microsoft Internet Information Server (IIS) as an RPS.
- Configure WebLogic as an RPS.
- Configure Oracle Sun Java System Web Server as an RPS.
- Configure Apache HTTP as an RPS.

Understanding Reverse Proxy Servers For PeopleSoft Implementations

This section describes how to proxy content to a single server configuration of PIA. In a production environment, a multi server configuration would be used to perform these steps to proxy content to your managed server instance of PIA, PIA1, PIA2, and so on.

Microsoft Internet Information Server (IIS) can be configured as a reverse proxy server (RPS) for one or more WebLogic Server instances. Multiple instances can either be independent instances, or grouped into a cluster. When using a reverse proxy, all URLs that are used to access your PeopleSoft application (even URLs that are stored in the database), need to point to the reverse proxy, and not to the WebLogic server.

Oracle only supports IIS 7.0 as an RPS on Windows Server 2008 x64 Standard Edition, or Enterprise Edition environments. These instructions are based on a logical separation of WebLogic Server and Microsoft IIS 7.0, where both web servers are installed on the same machine. If your configuration has WebLogic Server and Microsoft IIS on separate machines, you must perform these additional tasks:

- From the WebLogic server, copy `WL_HOME\server\plugin\win\x64\iisproxy.dll` to `c:\inetpub`, or to any other directory where an administrator has appropriate access to it on your Microsoft IIS server.
- From the WebLogic server, copy `WL_HOME\server\plugin\win\x64\iisforward.dll` to `c:\inetpub`, or to any other directory where an administrator has appropriate access to it on your Microsoft IIS server.
- In the procedures described below, change any reference to: `WL_HOME\server\plugin\win\x64\`, to `c:\inetpub\`, or to the other directory where you copied the two files mentioned above.

Note: In these instructions, make the appropriate adjustments for your server's architecture (32-bit or 64-bit).

Microsoft IIS 7.0, 1.1 plug-in is available for download in [My Oracle Support](#) and [Oracle Software Delivery Cloud](#).

See also, [Configuring Microsoft IIS 7.x as an RPS](#) on My Oracle Support.

Configuring Microsoft IIS 7.0 as an RPS

This section describes how to proxy content to a single server configuration of PIA. In a production environment, a multi server configuration would be used to perform these steps to proxy content to your managed server instance of PIA, PIA1, PIA2, and so on.

Microsoft Internet Information Server (IIS) can be configured as a reverse proxy server (RPS) for one or more WebLogic Server instances. Multiple instances can either be independent instances, or grouped into a cluster. When using a reverse proxy, all URLs that are used to access your PeopleSoft application (even URLs that are stored in the database), need to point to the reverse proxy, and not to the WebLogic server.

Oracle only supports IIS 7.0 as an RPS on Windows Server 2008 x64 Standard Edition, or Enterprise Edition environments. These instructions are based on a logical separation of WebLogic Server and Microsoft IIS 7.0, where both web servers are installed on the same machine. If your configuration has WebLogic Server and Microsoft IIS on separate machines, you must perform these additional tasks:

- From the WebLogic server, copy `WL_HOME\server\plugin\win\x64\iisproxy.dll` to `c:\inetpub`, or to any other directory where an administrator has appropriate access to it on your Microsoft IIS server.
- From the WebLogic server, copy `WL_HOME\server\plugin\win\x64\iisforward.dll` to `c:\inetpub`, or to any other directory where an administrator has appropriate access to it on your Microsoft IIS server.
- In the procedures described below, change any reference to: `WL_HOME\server\plugin\win\x64\`, to `c:\inetpub\`, or to the other directory where you copied the two files mentioned above.

Note: In these instructions, make the appropriate adjustments for your server's architecture (32-bit or 64-bit).

Configuring the Microsoft IIS 7.0 as an RPS

After you download the Microsoft IIS 7.0, set up the Microsoft IIS plug-in for IIS 7.0:

1. Create a web application in IIS Manager:
 - a. Right click Web Sites and select Add Web Site.
 - b. Enter the *web site name* of your web application.
 - c. Select the path of your *web application port* that can be any valid port number not currently in use.
 - d. Click OK button to create the web application.

If you can see the name of your application under Web Sites it means that your application has been created and started running.

2. Click the web application node under Web Sites to see all the settings related to the application.
3. Click Handler Mappings link to set the mappings to the handler for a particular MIME type.
4. Click StaticFile link and change the Request path field from `*` to `*.*`. Click OK.
5. Click the web application node under Web Sites on left panel and then click Add Script Map... under the Actions panel on right-hand side.
6. The Edit Script Map dialog box appears. Enter `*` for the Request path.
7. Browse to the `iisproxy.dll` file in Executable field. Name it *proxy*.
8. Click Request Restrictions... button and deselect the Invoke handler only if the request is mapped to check box.
9. Click OK to add this Handler mapping. Click Yes on the Add Script Map dialog box.
10. Click Root node of the IIS Manager tree and click ISAPI and CGI Restrictions. Select Allow unspecified ISAPI modules check box.

11. Create a file called *iisproxy.ini* with the following contents and place it in the directory with the plug-in:

```
WebLogicHost= @hostname@
WebLogicPort= @port@
ConnectRetrySecs=5
ConnectTimeoutSecs=25
Debug=ALL
DebugConfigInfo=ON
KeepAliveEnabled=true
WLLogFile=@Log file name@
SecureProxy=OFF
```

12. Open the Internet Explorer browser and enter *http://<hostname>:<port>*. You should be able to see the Medrec Sample Application from your Oracle WebLogic Server. If you want to run the plug-in in SSL mode, change the value of WeblogicPort to the SSL port of your application, and change the SecureProxy value to ON.

Configuring WebLogic as an RPS

This section discusses how to configure a WebLogic server as an RPS.

Note: Using WebLogic as a reverse proxy server is not recommended in a *production* environment.

Creating the RPS

To create an RPS, select *Multi Server Domain* as the configuration during the PeopleSoft Internet Architecture installation. The multi server domain installation automatically defines a server named "RPS" in addition to the main PIA server. The RPS is configured to be a reverse proxy server to other managed servers.

By default, the following settings are applied to the RPS:

Setting	Value
Name	RPS
HTTP Listen Port	8080
HTTPS Listen Port	8443
Default web application	HttpProxyServlet
Address of back-end WebLogic content server	The hostname of the machine from which the PIA setup was run, with the HTTP listen port specified during the PIA setup.

The default address specified for the back-end WebLogic content server assumes that it's the same machine as the one on which you're configuring the RPS, using the HttpProxyServlet application. There's no need to change this setting unless the content server is a different machine, or you enable load balancing with multiple content servers. If it's a different machine, you must change this setting to specify the correct content server. If you enable load balancing, you'll need to specify additional content servers.

Enabling Load Balancing

In addition to the *HttpProxyServlet* application, the PIA setup also defines an *HttpClusterServlet* application in your WebLogic configuration, which by default isn't active. The primary difference between the two applications is that for a given HTTP request, *HttpProxyServlet* can proxy content only from a single back-end content server, whereas *HttpClusterServlet* can proxy content from multiple back-end content servers, all of which serve the same content. This enables the RPS to load-balance the requests across a cluster of WebLogic servers.

You can configure the RPS for load balancing by changing the default web application from *HttpProxyServlet* to *HttpClusterServlet*, which becomes active as a result.

To change the default web application:

1. Start the WebLogic server.
2. Sign in to the Administration Console.
3. In the Domain Structure tree, click Deployments.
4. Click *HttpProxyServlet*, and select the Targets tab.
5. Clear the RPS Server check box, click *Apply*, and click *Activate Changes*.
6. In the Domain Structure tree, click Deployments.
7. Click *HttpClusterServlet*, and select the Targets tab.
8. Select the RPS Server check box, and click *Apply*.
9. Log out of the Administrative Console.

Specifying Back-End WebLogic Content Servers

You need to specify back-end WebLogic content servers only for the currently designated default web application (*HttpProxyServlet* or *HttpClusterServlet*).

To do this, you edit the appropriate web.xml configuration file directly.

Web Application	web.xml file modification
HttpProxyServlet	<p>You need to change the following settings only if the back-end WebLogic content server is on a different machine than the one where you're configuring the RPS.</p> <p>Edit the web.xml configuration file in <i>PIA_HOME</i>\webserv\<i>weblogic_domain</i>\applications\HttpProxyServlet\WEB-INF</p> <p>Modify the param-value elements for the <i>WebLogicHost</i> parameter and the <i>WebLogicPort</i> parameter to specify the hostname and HTTP listen port, respectively, of the back-end content server.</p>

Web Application	web.xml file modification
HttpClusterServlet	<p>Edit the web.xml configuration file in <i>PIA_HOME</i>\webserv\<i>weblogic_domain</i>\applications\HttpClusterServlet\WEB-INF.</p> <p>Modify the param-value element for the <i>WebLogicCluster</i> parameter to specify multiple back-end content servers separated by “ ” symbols, using the following format:</p> <pre>host1:http_port:https_port host2:http_port:https_port</pre>

Starting the RPS

To start the RPS, open a command prompt, change to *PIA_HOME*\webserv*weblogic_domain*, and launch the following commands:

1. Open a command prompt, and change to *PIA_HOME*\webserv*weblogic_domain*.
2. Start the WebLogic Admin Server.

```
startWebLogicAdmin
```

3. Start the RPS server.

```
startManagedWebLogic RPS
```

Note: You can also run the RPS as a service on Windows.

Related Links

[Starting WebLogic](#)

[Stopping WebLogic](#)

Configuring Oracle Sun Java System Web Server as an RPS

This section describes how to proxy content to a single server configuration of PIA. When in production, a multi server configuration would be used to perform these steps to proxy content to your managed server instance of PIA, PIA1, PIA2, and so on.

These instructions assume you have downloaded the Sun Java System Web Server from Oracle.

Installing the WebLogic Plug-In

Configuring the Oracle Sun Java System Web Server as an RPS involves installing a WebLogic plug-in on the Oracle Sun Java System Web Server. The plug-in comes in the form of a set of .DLL files for Windows systems and library files for UNIX systems.

To configure Oracle Sun Java System Web Server as an RPS:

1. Locate the appropriate WebLogic plug-in files within your WebLogic installation, which are "proxy*.dll" or "proxy*.so" or "proxy*.sl".

Note: If you are going to run your Sun web server on the *same* machine as WebLogic, it is recommended that you don't copy the plug-in files to your Sun web server.

The plug-in files are located in the following location:

WL_HOME/server/plugin/*operating system*/*architecture*

Where:

- *WL_HOME* represents the top level installation directory for your WebLogic server.
- *operating system* corresponds to the operating system, such as AIX, Linux, or Windows.
- *architecture* corresponds to the bit architecture of your operating system (32-bit, 64-bit).

2. Copy the appropriate WebLogic plug-in files to your Sun web server installation.

Copy the plug-in files to:

Sun_Home\operating system\plugins

Where:

- *Sun_Home* refers to the high-level installation location of your Oracle Sun Java System Web Server.
- *operating system* refers to the operating system the Oracle Sun Java System Web Server is installed on.

3. Define the NSAPI Module.

This step covers modifying the magnus.conf configuration file to reference the provided NSAPI module. Make sure to make a backup of it prior to modifying it.

Open the magnus.conf file located in:

C:*Sun_Home*\server\https-*machine_name*.com\config

Add the following lines to the bottom of the magnus.conf file. This instructs the Sun web server to load the native library as an NSAPI module. For *Sun_Home* and *drive*, substitute the actual location, including the drive letter of the NSAPI module you copied in at previous steps:

```
Init fn="load-modules" funcs="wl-proxy,wl-init" shlib=<drive>:<Sun_Home>
/plugins/proxy36.dll Init fn="wl-init"
```

If you skipped Step 1 because Sun web server and WebLogic are running on the same machine, update your configuration file similar to the following:

```
Init fn="load-modules" funcs="wl_proxy,wl_init" shlib="<drive>:/WebLogic_home
/wlserver_10.3/server/bin/proxy36.dll" Init fn="wl_init"
```

4. Define which requests to be handled by the plug-in.

The type of requests to be handled by the plug-in, and subsequently handed off to WebLogic, must be declared as part of an object definition in the `magnus.conf` file. A specific string in the URL, referred to as a *ppath*, can identify these requests.

To proxy all requests of a single PeopleSoft Internet Architecture site, such as *ps* (which would be accessed as `http://machine_name.com/ps/signon.html`), define the following object tag in the `magnus.conf` file. Define this and any other object tags directly following the default object tag.

```
<Object name="ps" ppath="*/ps/*">
Service fn=wl-proxy WebLogicHost=server1\
  WebLogicPort=7001
</Object>
```

The default object tag is generally several lines long and can be identified by `<Object name=default>...</Object>`.

To proxy additional sites, add subsequent object tags referencing the other site names:

```
<Object name="hr" ppath="*/hr/*">
Service fn=wl_proxy WebLogicHost=server1\
  WebLogicPort=7001
</Object>
```

To proxy all requests that are made to the Sun web server, create a single object tag named "peoplesoft" and set the `ppath` parameter to `*`.

5. Apply changes to the Sun web server.

With these settings saved, access the Server Manager, as in `http://localhost:8888`. When prompted, click the Apply button to update the Sun web server with your changes and restart it.

6. Make sure the PIA WebLogic Server is running.
7. Confirm the configuration.

To confirm the configuration between the WebLogic Server and the Sun web server, access PeopleSoft using the typical URL. For example,

`http://urserver.com/ps/signon.html`

If you are able to logon to PeopleSoft, your proxy plug-in installation and configuration was successful.

Working With the Oracle Sun Java System Web Server RPS Configuration

If you plan to proxy all requests for the PeopleSoft Internet Architecture through the Sun web server, you must also update any URLs that are defined in the PeopleSoft database to reference the proxy server (Sun web server), not the WebLogic server.

URL	Description
Portal	For the PeopleSoft portal, any content URLs that you have defined that directly reference PeopleSoft content (<code>psc</code> , <code>psp</code>) on the WebLogic server (meaning that the WebLogic server is referenced in the URL) must be updated to reference the Sun server in the URL

URL	Description
Integration Gateway	For the PeopleSoft integration gateway, any node definitions that directly reference an integration gateway on the WebLogic server must be updated to reference the Sun server in the URLs.
Report Repository	For the PeopleSoft report repository, any report node definitions that directly reference a report server on the WebLogic server must be updated to reference the Sun server in the URLs.
Miscellaneous	Any of your own definitions or objects that reference the URL of the WebLogic server must be updated to reference the Sun server in the URLs.

The magnus.conf file strictly enforces where text can be added. To avoid problems, follow these guidelines:

- Eliminate extraneous leading and trailing white space. If you must enter more characters than can be fit on one line, place a backslash (\) at the end of that line and continue typing on the following line. The backslash directly appends the end of the first line to the beginning of the following line.
- If a space is necessary between the words that end the preceding line and begin the next line, use one space, either at the end of the first line (before the backslash), or at the beginning of the second line.
- Do not split attributes across multiple lines.

The WebLogic online documentation contains a complete listing of WebLogic plug-in attributes and parameters.

Configuring for Deleting Attachments

If you intend to allow end users to delete attachments, and you are using the Oracle Sun Java System RPS, you need to make sure the access control lists for Oracle Sun Java System RPS are configured correctly.

To configure Oracle Sun Java System RPS for deleting attachments:

1. Start the Sun Java admin server and login to the administrative console.
2. Navigate to Configurations, and click the appropriate configuration.
3. Select Access Control, Access Control Lists, where you'll see the default and es-internal access control lists.
4. Modify the default access control list.
 - Click default.
 - For Allow anyone and Allow all, click delete rights.
 - Save.

5. Modify the es-internal access control list.
 - Click es-internal.
 - Click on allow anyone and enable delete rights.
 - Click on deny anyone and disable delete right.
 - Save.
6. Save and deploy the change so it is reflected in the instances.

Configuring Apache HTTP as an RPS

This section describes how to configure Apache to be a proxy for a single server configuration of PIA running on WebLogic. When in production, a multi server configuration would be used to perform these steps to proxy content to your managed server instance of PIA, PIA1, PIA2, and so on.

To configure Apache HTTP:

1. Download and install the Apache HTTP server.

See [Apache website](#).

2. Install the Apache HTTP server plug-in.

The installation of the Apache plug-in depends on whether you are installing the plug-in as a dynamic shared object (DSO) or a statically linked module. If you have downloaded the binary distribution of Apache, you will probably install the Apache plug-in from as a shared object. If you are in doubt as to which type, install the plug-in as a DSO. Exact instructions are available in the WebLogic documentation.

3. Specify the parameters that will be used by the Apache plug-in by defining them in an `IfModule` tag for WebLogic in the Apache `httpd.conf` file.

Add this tag in the `### Section 2: 'Main' server configuration section` of `httpd.conf`. For example, to configure the Apache to proxy all requests that it receives to a WebLogic server that is running on a machine named `crm.peoplesoft.com` and listening on port `7001`, you would define the following tag:

```
<IfModule mod_weblogic.c>
  WebLogicHost crm.peoplesoft.com
  WebLogicPort 7001
  MatchExpression /</IfModule>
```

To proxy requests to a cluster of WebLogic servers, replace the two attributes, `WebLogicHost` and `WebLogicPort`, with `WebLogicCluster`.

The syntax of the `WebLogicCluster` is `wlserver1:port,wlserver2:port`.

See The Red Paper posted on My Oracle Support: *Clustering and High Availability for PeopleTools*

If you specified an `AuthTokenDomain` during the PeopleSoft Internet Architecture installation, you must set the `WLCookieName` for the reverse proxy to that same value. To do so, add the `WLCookieName` attribute and set its value to `CookieName`, as specified on the WebLogic server in the PORTAL web application's `weblogic.xml` file in `PIA_HOME\webserv\peoplesoft\applications\peoplesoft\PORTAL.war\WEB-INF\weblogic.xml`.

4. Start the Apache HTTP server following the Apache usage instructions.
5. Start the WebLogic server.
6. Confirm the configuration between both the WebLogic server and Apache servers by signing onto the PeopleSoft system by using the typical URL. For example,

`http://Apachehost/ps/signon.html`.

If you can sign in to the PeopleSoft system, your installation and configuration was successful.

Setting The HTTP Session Timeout

PeopleSoft Internet Architecture does not use session timeout controls set on the web server.

The HTTP session timeout controls are accessible on the Security page of the web profile interface. At runtime, the session timeouts set in the current web profile override any HTTP session timeouts set on the web server.

See "Configuring Portal Security" (PeopleTools 8.53: Portal Technology).

Setting Authentication Failure Timeout

To limit the effectiveness of DOS attacks on failed authentications, you can use the `psft_failtimeout` Java option. Add this option in the `setEnv` script and assign a value in seconds. By setting the value to 60 seconds, for example, you override the default session timeout of 120 seconds (two minutes) when a user authentication fails or when a user is not yet authenticated.

For example,

```
SET JAVA_OPTIONS_WIN32=-server -Xms256m -Xmx256m -Dpsft_failtimeout=60  
-XX:MaxPermSize=128m -Xcomp
```

To determine the proper value for this property, you need to check the time in seconds that it takes to send an http(s) request from the browser to the web server and multiply the result by 2.

Enabling or Disabling HTTP Keep Alive

This section describes how to change HTTP Keep-Alive settings for a single server configuration of PIA. When in production, a multi server configuration would be used to perform these steps to your managed server instance of PIA, PIA1, PIA 2, and so on.

Keep-Alive, or more accurately termed "Persistent Connections" is a default feature of HTTP. Keep-Alive allows for the client (generally a web browser) and the web server to maintain open connections between requests for specified period of time. That time period is generally less than 60 seconds. The benefit of a persistent connection is that the client and the server do not need to perform the overhead of opening a new connection with every single request.

Enabling Keep-Alive is generally recommended, but in some situations it may introduce problems. Sporadic "The Page cannot be displayed" errors, for example, can be the result of a problem with Keep-Alive. In situations where Keep-Alive issues are suspected, disabling the web server Keep-Alive and performing tests will help to determine if the problem is indeed related to connection persistence.

To enable or disable Keep-Alive:

1. Start the PIA server.

See [Starting WebLogic](#).

See [Stopping WebLogic](#).

2. Log on to the Administrative Console.
3. Navigate to the server's HTTP settings page.
 - a. In the Domain Structure tree, click Environments.
 - b. Click the PIA server (or your custom named PeopleSoft server) from the window on the right.
 - c. Click the Protocols tab.
 - d. Click the HTTP tab.
4. Change the Keep-Alive settings.
 - To disable Keep-Alive: Deselect Enable Keepalives and click Save. With Keep-Alive disabled, Duration and HTTPS Duration are not used.
 - To enable Keep-Alive: Select Enable Keepalives and update Duration and HTTPS Duration values as needed for your system.

Click Save and Activate Changes.

5. Restart WebLogic Server.

Changing WebLogic User Passwords

The WebLogic domain built by the PIA install includes the following WebLogic IDs:

- system
- operator
- monitor

It is *highly* recommended to change the passwords for all of these user IDs on any production servers. You set the system password during the installation of the PeopleSoft Internet Architecture domain.

To change the passwords:

1. Start the PIA server.

See [Starting WebLogic](#).

See [Stopping WebLogic](#).

2. Log in to the Administrative Console.
3. Change a WebLogic Server user's password.
 - a. In the Domain Structure tree, click Security Realms.
 - b. In the Realms table, click myrealm.
 - c. Select the Users and Groups tab.
 - d. In the table of available users, click *system*.
 - e. Select the Passwords tab.
 - f. Enter and re-enter a new password for this user.
 - g. Click Save.
4. With the system user ID, modify the boot.properties file with the new values:
 - a. Open *PIA_HOME\websrv\<domainname>\servers\PIA\security\boot.properties*.
 - b. Remove the previous, encrypted system value, and enter the new value in clear text.

For example,


```
#Tue Mar 31 12:09:14 PDT 2009
password=newpassword
username={3DES}S43Sj+0Gar0\=
```
 - c. Save the file.
 - d. Reboot the server, and test the login.

Note: Rebooting the server should encrypt the clear text value you entered.

Note: When you run WebLogic as a Windows service, WebLogic uses the default ID of *operator* and password of *password*. Changing the password for the WebLogic ID that runs the Windows service requires an additional manual step of updating the setEnv.cmd, located in *PIA_HOME\websrv\peoplesoft\bin*. Set the WLS_PW environment variable to reflect the new password. Then, reinstall the Windows service by running the installNTservice command file located in *PIA_HOME\websrv\peoplesoft\bin*.

Note: If you use PSADMIN to manage a PIA domain, make sure to update the setEnv.cmd to reflect any password updates.

Implementing WebLogic SSL Keys and Certificates

This section provides an overview of Secure Sockets Layer (SSL) encryption with WebLogic and discusses how to:

- Obtain encryption keys.
- Prepare keys and certificates for the keystore.
- Import keys and certificates into the keystore.
- Configure WebLogic SSL encryption keys.

Understanding SSL Encryption with WebLogic

To use SSL encryption with WebLogic and the current PeopleTools release, the WebLogic keystore must contain the following appropriately configured encryption keys:

- The web server's private key.
- The web server's public key, digitally signed by a trusted *certificate authority* (CA).
- The digitally signed public key of the same CA that signed the web server's key.

A public key is transferred and stored as a data element in a digital certificate or a *certificate signing request* (CSR). You can obtain public keys from a variety of sources in several different formats.

When setting up SSL, you need to:

- ensure that the encryption keys are correctly formatted.
- install them in the keystore.
- configure them using the Administration Console.

Note: If you've already installed and configured a set of encryption keys for use with WebLogic in a previous PeopleTools release, you must migrate them as external files to the keystore within the current WebLogic version.

Obtaining Encryption Keys

If you already have a set of existing encryption keys configured as external files, you don't need to obtain new ones. To find the existing keys, refer to the documentation for the PeopleTools and WebLogic releases for which those keys were installed.

The following procedure describes how to obtain new encryption keys, using as an example the trial certificate available from Verisign.

To obtain new encryption keys:

1. At a command prompt, change to the following directory:

```
PIA_HOME\webserv\domain_name\piabin
```

Where *domain_name* is the name of the installed PeopleSoft Pure Internet Architecture domain for which you want to obtain encryption keys.

2. Enter the following command:

```
pskeymanager -create
```

Note: Pskeymanager is a script wrapper to Java's keytool, provided by PeopleSoft to manage the WebLogic keystore. For usage information, enter `pskeymanager -help`.

3. Follow the prompts and enter the requested information to create a new private key and a CSR for your web server.
 - Pskeymanager uses the keystore in `PIA_HOME\webserv\domain_name\piaconfig\keystore\pskey.`
 - Pskeymanager prompts you for an alias for the new keys, for example, `ServerABC`. This is the name you'll use to refer to the keys in the future.
 - Pskeymanager prompts you for distinguished name fields. Enter the appropriate values for your organization.
 - Pskeymanager prompts you for information about the CSR expiration date, key size, key algorithms, and the private key password. All of these fields have default values.

Pskeymanager creates the private key inside the keystore, and creates the CSR as a file called `ServerABC_certreq.txt` in the current directory. You use the CSR to obtain your signed public key certificate and a root certificate from a CA.

4. Decide which trusted CA you want to sign your web server's public key.

You can use any CA that's compatible with Sun's Java JKS standard, such as Verisign.

5. Open your CSR file in a text editor and copy its entire contents, including the first and last lines:

```
-----BEGIN NEW CERTIFICATE REQUEST-----
...
...
-----END NEW CERTIFICATE REQUEST-----
```

6. Access Verisign's test certificate enrollment site at <https://www.verisign.com/ssl/buy-ssl-certificates/free-ssl-certificate-trial/index.html>.

Verisign guides you through the CSR submission process, including:

- Accepting the Verisign license agreement.
- Entering your technical contact information, which includes the email address where Verisign can send your signed public key.
- Pasting your CSR contents in the provided text field.
- Verifying your CSR.
- Confirming and submitting your order.

Verisign also provides its own digitally signed public key in a certificate, which is known as a *trusted CA certificate*, *root certificate*, or *achain certificate*.

- Download the VeriSign test CA root certificate from <http://digitalid.verisign.com/cgi-bin/getcacert>.

When prompted, save `getcacert.cer` to `PIA_HOME\webserv\domain_name`.

Note: If you need to FTP your certificate to UNIX, you must FTP it in ASCII mode to `PS_HOME/webserv/domain_name`.

- Check your email.

Verisign digitally signs your web server's public key, then returns it to you in a certificate, called the *server certificate*. Following is an example of the contents of a server certificate:

```
-----BEGIN CERTIFICATE-----
DMICHDCACyCEAHSeRkM2guFL+6OvHr4AS0wDQYJKoZIhvcNAQEEBQAwwgakkxFjAP
AANVBAoTDVZlcm1TaWduLCBLbAMxRzBFBGNVBAStPnd3dy52ZXJpc2lnbi5jb20S
VcVwb3NpdG9yeS9UZXR0Q1ETIEluY29ycC4gQnkgUmVmLiBMaWFiLiBMVEQuMUyF
LIGEc3VyYW5jZXMgKEMpVRMxOSDFertdsfh67TIwNDAwMDAwMFOXTAwMTIxODIA
ONT1LVoweTELMaKGA1UERhMCMVVMxEzARBgNVBAgTCkNhbgGlm3JuaWExEzARBgNK
VBAUCOBzZWZzYW50b24BEzARBgNVBAoUC1Blb3BsZVNvZnQxZDAsBgNVBAsUC1BT
Eb3sZVVvb2xzMRUwEwADVQDFAXEQLJPV04xMTE0MDAwXDANBgkqhkiG9w0BAQET
SAALADBEAkeAucfM/GOQhdkk4Q0ZD5i1l4gp6WTYMc4IaReoCYkEAmDKAVcYzY3R
Mdbp4RC8SABd3bjjDOHcoCak9U6oSvL+HQIDAQABMA0GCsqGSIb3DQEBAUAA0EO
Arm3uf634Md0fqqNxAAL+e9rbY0ia/X48Axloi17+kLTVI1YPOp+Jy6Slp5iNIFC
DhskdDFH45AjSDAFhjruGHJK56SDFGqwg23SFRfgtjkjyu673424yGWE5Gw4576K
DsdDFG256EDHY45yTRH67i345314GQE356mjsdhjUwbtRhr43Gq3QEVe45341tS
YDY6d471DmQxDs9wGt1bkQ==
-----END CERTIFICATE-----
```

- Copy the entire certificate contents, and save it as a text file called `ServerABC-cert.pem` in `PIA_HOME\webserv\domain_name`.

Be sure to include the first and last lines.

Note: If you need to FTP your certificate to UNIX, you must FTP it in ASCII mode.

Note: Make backup copies of the server certificate and the root certificate before proceeding.

Preparing Keys and Certificates for the Keystore

Your encryption keys must be in *privacy enhanced mail* (PEM) format, which is Base64-encoded data. Base64 encoding uses only ASCII characters. A PEM-formatted key or certificate file has an extension of either `.pem` or `.cer`. If the file is in the binary *distinguished encoding rules* (DER) format, it has a `.der` extension. Use the `der2pem` Java utility to convert DER-formatted keys and certificates to PEM format.

For SSL to work, your WebLogic server must present its own public key to each client browser, along with the self-signed public key of a root CA that's also in the browser's keystore, as well as any keys necessary to establish a *chain of trust* between the two. All of these keys must be part of the same certificate file before you can import them into the WebLogic keystore.

If you generated the private key using `pskeymanager` on a WebLogic platform, it's automatically correctly formatted, password protected, and installed in the keystore with no additional steps required. However, if the private key was configured as an external file on an earlier WebLogic platform/version, you must properly format it and incorporate a password, before importing it into the current WebLogic keystore along with the public key certificates.

Converting DER Files to PEM Format

It's important to convert all DER-formatted key and certificate files to PEM format before you work with them further.

To convert DER-formatted key and certificate files to PEM format:

1. At a command prompt, change to the following directory:

```
PIA_HOME\webserv\domain_name
```

Where *domain_name* is the name of an installed PeopleSoft Pure Internet Architecture domain.

2. Enter the following command:

```
setenv.cmd
```

This sets the appropriate environment for Java commands.

3. For each DER-formatted key or certificate file, enter the following command:

```
java utils.der2pem filename.der
```

Make sure that you include the DER file's directory path. A new PEM file by the same name is created in the same location.

If you converted a private key file to PEM format, you must modify the header and footer to be compatible with WebLogic.

To modify the private key file header and footer:

1. Open the PEM-formatted private key file in a text editor.
2. Change the following line:

```
-----BEGIN CERTIFICATE-----
```

To this:

```
-----BEGIN RSA PRIVATE KEY-----
```

3. Change the following line:

```
-----END CERTIFICATE-----
```

To this:

```
-----END RSA PRIVATE KEY-----
```

4. Save and close the private key file.

Establishing the Server Certificate Chain of Trust

Your server certificate must contain, in addition to the web server's public key, any keys necessary to establish a chain of trust that culminates in the self-signed root certificate of a trusted root CA. That CA's root certificate must be in the keystore of any browser that's used to access your web server. Most browsers have an extensive set of trusted root certificates in their keystores.

First append the root certificate of the CA who issued your server certificate to the server certificate file. If you determine that the root certificate is not likely to be in your users' browsers, you must also append

to the certificate file a chain certificate that was issued to your CA by another CA, then a chain certificate issued to that CA, and so on, until you append a root certificate that was issued by a trusted CA to itself.

For example, if your server certificate file is `demo_cert.pem` and the CA's root certificate is `ca_cert.pem`, you can open `demo_cert.pem` in a text editor, then insert the contents of `ca_cert.pem` after adding a new line at the end of the file. Make sure that each certificate follows the previous one on the next line, as follows:

```

...
...
DosdDFG256EDHY45yTRH67i345314GQE356mjsdhhjuwbtrh43Gq3QEVe45341tS
YDY6d47lDmQxDs9wGt1bkQ==
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
DMICHDCCAcYCEAHSeRkM2guFL+6OvHr4AS0wDQYJKoZIhvcNAQEEBQAwwakxFjAP
...
...

```

The result is that `demo_cert.pem`, for example, now contains the data from both certificates.

Note: You can also use the `type` command in Windows or the `cat` command in UNIX to combine the certificate files.

Password Protecting the Private Key

Private keys inside the WebLogic keystore are password protected. You can't import an external private key file into the keystore without a password. If it isn't currently password protected, use the WebLogic `wlkeytool` utility to incorporate a password into the private key file.

To password-protect an external PEM-formatted private key file:

1. At a command prompt, change to the following directory:

```
WL_HOME\server\native\win\32
```

Where `WL_HOME` is the root directory where you installed WebLogic.

2. Enter the following command:

```
wlkeytool insecure_privatekey.pemsecure_privatekey.pem
```

Where `insecure_privatekey.pem` is the name of the original private key file, and `secure_privatekey.pem` is the name of the resulting password-protected private key file.

Note: Make sure that you include directory paths for the private key files.

The following message appears:

```
Enter password to unprotect private key:
```

3. Press `Enter`.

The following message appears:

```
Private key not PKCS8 encoded, trying RSA key
Private key file opened successfully
Enter password to protect private key :
```

4. Enter the password that you want to use for this key.

The following message appears:

```
Verify password to protect private key :
```

5. Enter the password again to confirm it.

The utility creates the password protected private key file that you specified. You import this key into the WebLogic keystore.

Importing Keys and Certificates Into the Keystore

Each WebLogic domain maintains its own keystore in *PIA_HOME*\webserv*domain_name*\keystore\pskey, and all servers within a domain can share the same keystore.

These utilities are available for importing keys and certificates into the keystore:

<i>Utility</i>	<i>When to Use</i>
pskeymanager	If you created the private key using the pskeymanager utility on a WebLogic platform, it's already installed in the keystore. You need only use pskeymanager to import your server certificate, which should contain your web server's signed public key, your trusted CA's root certificate, and any public keys necessary to establish a chain of trust between them.
ImportPrivateKey	If the private key was previously configured as an external file on an earlier WebLogic platform, you must import it into the WebLogic keystore along with the server certificate, using the <i>ImportPrivateKey</i> utility. The private key should be password-protected.

Using pskeymanager to Import the Server Certificate

To import the server certificate into the WebLogic keystore:

1. At a command prompt, change to the following directory:

```
PIA_HOME\webserv\domain_name\piabin
```

Where *domain_name* is the name of the installed PeopleSoft Pure Internet Architecture domain.

2. Enter the following command:

```
pskeymanager -import
```

Note: PeopleTools provides pskeymanager (a script wrapper to Java's keytool) to help manage the WebLogic keystore. For usage information, enter `pskeymanager -help`.

3. Follow the prompts and enter the requested information to create a new private key and a CSR for your web server.

Keep the following in mind:

- pskeymanager uses the keystore in *PIA_HOME*\webserv*domain_name*\piaconfig\keystore\pskey.

- pskeymanager prompts you for an alias for the server certificate, for example, *ServerABC*. This should be the same alias that you specified for the corresponding private key when you created it.
- pskeymanager prompts you for the name of the server certificate file, for example, *ServerABC-cert.pem*. Include the file path if necessary.

Using ImportPrivateKey to Import an External Private Key File with the Server Certificate

To import a password-protected private key and the server certificate into the WebLogic keystore:

1. At a command prompt, change to the following directory:

```
PIA_HOME\webserv\domain_name\bin
```

Where *domain_name* is the name of an installed PeopleSoft Pure Internet Architecture domain.

2. Enter the following command:

```
setEnv.cmd
```

This sets the appropriate environment for Java commands.

3. Enter the following command:

```
java utils.ImportPrivateKey keystore\pskey store_passprivatekey_aliasprivatekey=>
y_passservercert_fileprivatekey_file
```

The parameters for this command are as follows:

store_pass	Specify the password for the WebLogic pskey keystore. The default password is <i>password</i> .
privatekey_alias	Specify an alias for the private key. This is the name by which the key will be accessible inside the keystore.
privatekey_pass	Specify the password for the private key.
servercert_file	Specify the path and name of the server certificate file that includes the issuing CA's root certificate.
privatekey_file	Specify the path and name of the private key file.

Configuring WebLogic SSL Encryption Keys

This section describes how to configure the SSL encryption keys that you previously imported into the WebLogic keystore in *PIA_HOME\webserv\domain_name\keystore\pskey*, where *domain_name* is the name of an installed PeopleSoft Pure Internet Architecture domain.

The following procedure applies to a single server configuration of PIA. In a production environment, you would perform these steps for managed server instances of PIA, PIA1, RPS, and so on, in a multi-server domain configuration.

To configure WebLogic SSL encryption keys for the PIA server:

1. With the PIA server running, sign in to the Administration Console.
2. Access the keystore configuration pages.
 - a. In the Domain Structure tree, expand Environment, click Servers, and click PIA from the Servers list.
 - b. Select the Keystores tab.
 - c. To change the configuration section, click Lock & Edit.
 - d. Select *Custom Identity and Custom Trust* from the Keystores list.
3. Update the fields on the Configure Keystore Properties page as follows:

Field	Value
Custom Identity Key Store File Name	<i>keystore/pskey</i> This should be the relative path and name of the keystore into which you imported your SSL keys.
Custom Identity Key Store Type	<i>JKS</i> Don't change this value.
Custom Identity Key Store Passphrase	<i>password</i>
Confirm Custom Identity Key Store Passphrase	Same as the value of Custom Identity Key Store Pass Phrase.
Custom Trust Key Store File Name	<i>keystore/pskey</i> This should be the relative path and name of the keystore into which you imported your SSL keys.
Custom Trust Key Store Type	<i>JKS</i> Don't change this value.
Custom Trust Key Store Passphrase	<i>password</i>
Confirm Custom Trust Key Store Passphrase	Same as the value of Custom Trust Key Store Pass Phrase.

4. Click Save.
5. Access the SSL tab, and update the values on the SSL Private Key Settings as follows:

Field	Value
Private Key Alias	Specify a unique identifier, such as the web server's machine name. This should be the alias that you specified for this server's private key.

Field	Value
Passphrase	<i>password</i>
Confirm Passphrase	Same as the value of Passphrase.

6. Click Save and Activate Changes.

You *must* click the Activate Changes button to apply your changes. If you close your browser without clicking Activate Changes, your changes will be lost.

7. Restart the WebLogic PIA server.

Enabling TLS-Only on WebLogic

This section discusses:

- Configuring TLS-only on WebLogic Server.
- Configuring browsers for TLS.
- Configuring reverse proxy servers for TLS.

Configuring TLS-Only on WebLogic Server

The `weblogic.security.SSL.protocolVersion` command-line argument lets you specify which protocol is used for SSL connections.

To enable TLS-only on WebLogic Server:

1. Open `PIA_HOME/webserv/<Domain>/bin/setEnv.cmd` (Windows) or `setEnv.sh` (UNIX) in edit mode.
2. Append the following option to the `JAVA_OPTIONS` variable:

```
-Dweblogic.security.SSL.protocolVersion=TLS1
```

For example, on Windows in the `setEnv.cmd`, the updated `JAVA_OPTIONS_WIN` would appear as:

```
SET JAVA_OPTIONS_WIN=-jrockit -XnoOpt -XXnoJITInline -Xms512m -Xmx512m
-Dtoplink.xml.platform=oracle.toplink.platform.xml.jaxp.JAXPPlatform
-Dweblogic.security.SSL.protocolVersion=TLS1
```

3. Save the change.
4. Restart PeopleSoft Internet Architecture.

For more information about Supporting TLS1 protocol, refer to Oracle WebLogic documentation for specifying the version of SSL to be used.

See [Specifying the Version of the SSL Protocol](#).

Configuring Browsers for TLS

This section covers steps for configuring TLS on browsers.

Setting Up TLS on Microsoft Internet Explorer

To set up TLS on Internet Explorer:

1. Launch Internet Explorer.
2. Select Tools, Internet Options, and select the Advanced tab.
3. In the Settings box in the Security section, disable *Use SSL 3.0* and enable *Use TLS 1.0*.
4. Click OK and restart the browser.

Setting Up TLS on Mozilla Firefox

To set up TLS on Firefox:

1. Launch Firefox.
2. Select Tools, Options, click the Advanced icon, and select the Encryption tab.
3. In the Protocols group box, disable *Use SSL 3.0* and enable *Use TLS 1.0*.
4. Click OK and restart the browser.

Configuring Reverse Proxy Servers for TLS

It is strongly recommended that you access the vendor's documentation of the web server you are using for a reverse proxy server and use their instructions for setting up TLS.

Working With WebLogic Session Cookies

When a user signs in to a PeopleSoft application, the portal servlet generates a cookie containing the user's HTTP session ID, and sends it to the user's browser to maintain the state of the session. The name of the cookie is fixed for all users accessing that portal.

On a WebLogic portal, the session cookie's name is generated at install time based on the portal hostname and port number, which uniquely identify the portal within your PeopleSoft system. This name is stored in the portal's `weblogic.xml` file.

However, the cookie name must not start with a number, and it must not contain any periods. If your users are experiencing problems signing in to PeopleSoft applications at different URLs from the same browser session, make sure that the session cookie names at those sites are valid.

To ensure valid WebLogic session cookie names:

1. Shut down your WebLogic server.
2. Open the `weblogic.xml` file for your web server in a text editor.

You can find it in *PIA_HOME\webserv\domain_name\applications\peoplesoft\PORTAL\WEB-INF*.

3. Check the value of the session parameter called CookieName.

Ensure that the content of the param-value element doesn't start with a number or contain any periods. For example, the following session cookie name is invalid:

```
<session-param>
  <param-name>CookieName</param-name>
  <param-value>57.28.208.21-80-WebLogicSession</param-value>
</session-param>
```

You can replace the periods with dashes (-). Following is a valid version of the session cookie name:

```
<session-param>
  <param-name>CookieName</param-name>
  <param-value>c57-28-208-21-80-WebLogicSession</param-value>
</session-param>
```

4. Save and close the file.
5. Restart your WebLogic server.

Securing Servlets on WebLogic

This section describes how to restrict access to a web resource for a single server configuration of PIA. When in production, a multi server configuration would be used to perform these steps to your managed server instances of PIA, PIA1, PIA2, and so on. WebLogic Server provides an optional level of security to restrict access to resources on the web server.

The following steps describe how to restrict access to the PeopleSoft Portal servlet using a WebLogic ID and password. This, for example, could be applied to the report repository servlet.

To restrict access to a servlet:

1. Start the PIA server.
2. Open the Administration Console.
3. Change the security model to "Custom Roles And Policies":
 - a. In the Domain Structure section select Security Realms, myrealm.
 - b. Click Lock and Edit.
 - c. For the Security Model Default field, select Custom Roles And Policies option from the drop down list.
 - d. Click Save.
 - e. Click Activate Changes.
 - f. Select the default security model for the application to Custom Roles and Policies.
4. Enable security policy checks for web applications.

- a. Edit config.xml under *PIA_HOME*\webserv*<domain>*\config.

Note: Backup the file before you make any changes.

- b. Find the <app-deployment> section and add the following line in all of the <app-deployment> (or application deployment) sections:

```
<security-dd-model>CustomRolesAndPolicies</security-dd-model>
```

For example, one of the sections will look similar to the following:

```
<app-deployment>
  <name>peoplesoft</name>
  <target>PIA</target>
  <module-type>ear</module-type>
  <source-path>applications/peoplesoft</source-path>
  <deployment-order>1</deployment-order>
  <security-dd-model>CustomRolesAndPolicies</security-dd-model>
  <staging-mode>nostage</staging-mode>
</app-deployment>
```

- c. Save the file and restart the server so that the changes will take effect.
5. (Optional) Define the WebLogic users that you want to use.

If you want to use one of the provided WebLogic user accounts (system, operator, or monitor) you can skip this step. Otherwise, create a new WebLogic user account:

- a. Under Domain Structure, select Security Realms, myrealm.
- b. Click Users and Groups tab.
- c. Under the Users tab, click New to create a new user.
- d. Enter user name and password and click OK.

6. (Optional) Create a user group, and add user(s).

If you want to create a user group, add your users to that group and in the following steps select Caller is Member of group instead of User name of caller. To create a group:

- a. Under Domain Structure, select Security Realms, myrealm.
- b. Click the Users and Groups tab.
- c. Under the Groups tab, click New to create a new group.
- d. Enter the name of the group and description. Click OK.
- e. Select the Users tab, select your new user, and then click the Groups tab.
- f. Move the appropriate group from the Available box to the Chosen box.

7. Define a security policy for the PeopleSoft Portal web application.

To restrict access to the Portal web application, perform the following in the navigation window on the left:

- a. Under Domain Structure, select Deployments, and select peoplesoft from the list of applications.
 - b. Under Overview tab, select the portal module which appears as “/” in the Modules and Components table.
 - c. Click the Security tab.
 - d. Define a new URL pattern for this web module. Select New and enter the URL pattern as “/*” or specify the URL requiring authentication, and click OK.
 - e. Select the URL pattern that you just created from the table and enter a security policy for this URL pattern.
 - f. Add the condition to give access to the particular user you want to have access to this URL or any other conditions by clicking Add Conditions.
 - g. To restrict access to a specific user, select the policy condition of User name of caller, click Add, and when prompted specify the user name. Repeat this step for additional users, groups, or access times. For access times, the server's local time is used.
 - h. Click Finish and go back to the policy page and click Save.
- This action does not require a server reboot.
8. Test the configuration.

Test your new security policy by accessing the URL you defined. If the security policy is active, you'll be prompted to sign in using a user ID that you added.

Adjusting the JVM Heap Size

Java options, such as JVM heap size and VM mode, used by the WebLogic server are stored in your WebLogic domain's setEnv script, stored in:

```
PIA_HOME\webserv\peoplesoft\bin
```

These options are specified in the script using the `JAVA_OPTIONS_OSplatform` environment variable. If you need to adjust any of the Java options, including changing the JVM heap size, you must manually edit the script.

The Microsoft Windows setEnv.cmd script contains the following default setting:

```
SET JAVA_OPTIONS_WIN=-jrockit -XnoOpt -Xms512m -Xmx512m
-Dtoplink.xml.platform=oracle.toplink.platform.xml.jaxp.JAXPPlatform
```

The UNIX standard setEnv.sh script contains the following default settings for supported Linux and UNIX platforms:

```
JAVA_OPTIONS_AIX="-Xms128m -Xmx256m -Dtoplink.xml.platform
=oracle.toplink.platform.xml.jaxp.JAXPPlatform
```

```
-Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0"

JAVA_OPTIONS_HPUX="-Xms256m -Xmx256m -XX:MaxPermSize=256m
-Dtoplink.xml.platform=oracle.toplink.platform.xml.jaxp.JAXPPlatform
-Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0"

JAVA_OPTIONS_LINUX="-jrockit -XnoOpt -Xms512m -Xmx512m
-Dtoplink.xml.platform=oracle.toplink.platform.xml.jaxp.JAXPPlatform
-Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0"

JAVA_OPTIONS_SOLARIS="-Xms256m -Xmx256m -XX:MaxPermSize=256m
-Dtoplink.xml.platform=oracle.toplink.platform.xml.jaxp.JAXPPlatform
-Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0"
```

You modify the `Xms` parameter to adjust minimum heap size, and modify the `Xmx` parameter to adjust maximum heap size. For most operating systems, the minimum and maximum values of the heap size are equal, which is recommended for better performance.

In a multi-server domain, the platform-specific versions of the `JAVA_OPTIONS` environment variable that appear in the `setEnv` script apply only to managed servers. The administration server doesn't use any of these variables, but it assumes default JVM heap size values of "`-Xms256m -Xmx256m`".

To adjust the JVM heap size for the administration server, add the environment variable `JAVA_OPTIONS_ADMINSERVER` following the last entry for `JAVA_OPTIONS_OSplatform`, and set it to your required minimum and maximum values, for example:

```
JAVA_OPTIONS_ADMINSERVER="-Xms128m -Xmx256m"
```

Note: If you're running WebLogic as a Microsoft Windows service and you modify `setEnv.cmd`, you must reinstall the service.

Related Links

[Managing JVM Heap Size](#)

Determining the Service Pack Level

A summary of installed products and their versions and service pack levels is maintained in the `BEA_HOME\registry.xml` file. However, to confirm version information, it's more accurate to check the WebLogic log. A failed service pack install may be indicated in the log, but not found at runtime.

This section discusses how to:

- Check the WebLogic log
- Query WebLogic

Checking Version Information in the WebLogic Log

The WebLogic log is located in:

```
PIA_HOME\webserv\peoplesoft\servers\\logs\weblogic_server_weblogic.log
```

For version information, look for an entry similar to:

```
###<Apr 13, 2011 3:29:46 PM PDT> <Info> <WebLogicServer> <RTDC79601VMC> <>
<[ACTIVE] ExecuteThread: '0' for queue: 'weblogic.kernel.Default (self-tuning)'\>
```



```
<> <> <> <1302733786761> <BEA-000000> <WebLogic Server "PIA" version:
WebLogic Server 10.3.4.0  Fri Dec 17 20:47:33 PST 2010 1384255
Copyright (c) 1995, 2009, Oracle and/or its affiliates. All rights reserved.>
###<Apr 13, 2011 3:29:47 PM PDT> <Notice> <Log Management> <RTDC79601VMC> <>
<[ACTIVE] ExecuteThread: '0' for queue: 'weblogic.kernel.Default (self-tuning)''>
<> <> <> <1302733787089> <BEA-170019> <The server log file
D:\853-801\websevr\AftrOn1034\servers\PIA\logs\PIA_weblogic.log is opened.
All server side log events will be written to this file.>
```

Checking Version Information at the Command Line

To query a *running* WebLogic for version information from the command line, submit the following arguments to setEnv.

Windows:

While the WebLogic sever is up and running, from the command line, run setEnv and the following command in sequence.

```
C:\PIA853\websevr\peoplesoft\bin>setenv.cmd
C:\PIA853\websevr\peoplesoft\bin>java weblogic.Admin
-url t3://localhost:<portnumber> -username <username> -password <password> VERSION
WebLogic Server 10.3.4.0  Fri Dec 17 20:47:33 PST 2010 1384255
```

UNIX:

While the WebLogic sever is up and running, from the shell command window, source the setEnv.sh file and issue the following command in sequence. For example,

```
rtdc60007stdb:$ ./setEnv.sh
rtdc60007stdb:$ java weblogic.Admin
-url t3://localhost:80 -username <username> -password <password> VERSION
```

Enabling HTTP Access Log

This section describes how to change HTTP logging for a single server configuration of PIA. When in production, a multi server configuration would be used to perform these steps to your managed server instance of PIA or PIA1, PIA 2, and so on.

To enable or disable HTTP access log:

1. Make sure the PIA server is running.

See [Starting WebLogic](#).

See [Stopping WebLogic](#).

2. Log on to the Administrative Console.
3. Open Server's Logging configuration page.
 - a. In the Domain Structure tree, expand Environment, and clickServers.
 - b. Click PIA (or your custom server name) in the Servers list.
 - c. Select the Logging tab, and select the HTTP tab.

4. Enable HTTP access logging.
 - a. Click the Lock&Edit button.
 - b. Select the HTTP access log file enabled check box to turn on the access.log.
 - c. Modify the Log file name field if desired.
 - d. Click Save andActivate Changes.

5. Restart the WebLogic Server.

Chapter 8

Working with IBM WebSphere

Understanding WebSphere Application Server within Your PeopleSoft Implementation

This section discusses:

- Deploying PeopleSoft Applications with WebSphere.
- Using the Integrated Solution Console.
- WebSphere Application Server profiles.
- IBM HTTP Server.

Oracle only documents IBM WebSphere where it relates specifically to the Oracle product which supports it. For all general IBM WebSphere information see [WebSphere Application Server Information Center](#).

Deploying PeopleSoft Applications With WebSphere

The IBM WebSphere Application Server (WAS) is a J2EE application server that PeopleTools uses as a web server to deploy PeopleSoft applications, as an alternative to Oracle WebLogic Server. To install and use WebSphere, install it using the detailed instructions in the PeopleTools installation documentation.

Note: WebSphere Base edition and Network Deployment edition (ND) are packaged and installed together. PeopleSoft applications use the Network Deployment edition for application deployment.

Note: Oracle *does not* provide IBM WebSphere Application Server. PeopleTools sites a certified version on My Oracle Support, Certifications, and you need to contact IBM to acquire the supported software.

Using The Integrated Solutions Console

To view and configure WebSphere settings, you use the web-based administrative console named Integrated Solutions Console, which is based on the Integrated Solutions Console (ISC) framework, providing a consistent and integrated capability for administering IBM software. The ISC enables you to access settings related to key areas of server administration, including:

- servers
- applications
- security
- services
- environment

- users and groups
- system administration
- monitoring and tuning

To access ISC:

1. Enter the following URL:

`http://WAS_Hostname:9061/ibm/console`

You can find the port number for the ISC associated with your Pure Internet Architecture domain by checking the Administrative console port value in the AboutThisProfile file in `D:\PIA_HOME\webserv\<<PIA_Domain>\logs`. For example:

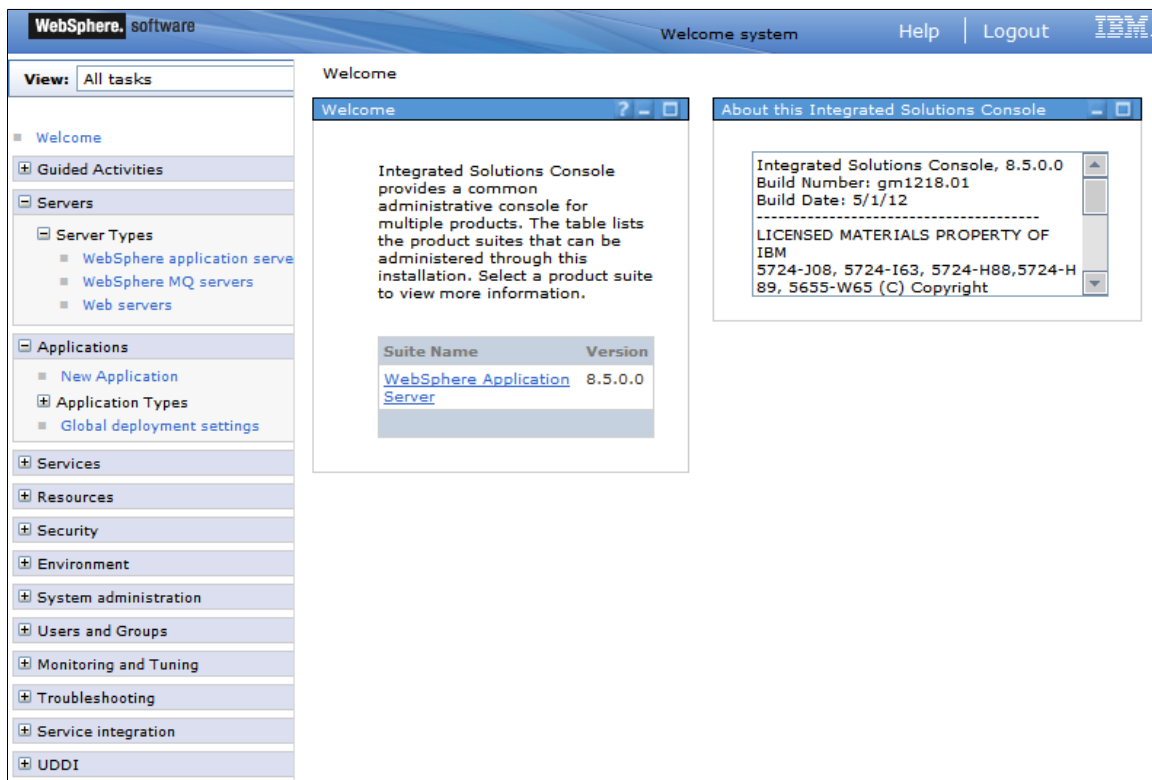
```
Administrative console port: 9061
```

2. On the ISC login page, enter the appropriate credentials in the User ID field and clickLogin.

Use the user ID and password you specified as the administrative ID when you installed the PeopleSoft domain. If you have enabled Global Security, add the appropriate user ID and password. If Global Security is not enabled, add any user name in the User ID field (or leave it blank).

Image: Integration Solutions Console

This example illustrates the fields and controls on the Integration Solutions Console.



If you have more than one WebSphere server profile installed on a machine, the ISC port number for each installation is different. For each PeopleSoft domain you install, IBM creates the required profile(s).

A key benefit of the ISC is that it allows administrators to create a custom navigation list of tasks performed more frequently, using the View, My Tasks feature in the upper, left-hand corner. It also provides an additional set of system management, monitoring and tuning features.

Note: You can perform numerous administrative tasks using ISC. This guide discusses those that pertain most specifically to deploying PeopleSoft applications. It is assumed that you are familiar with the topics contained in the IBM WebSphere documentation.

Note: In this document, we refer to the ISC as the "administrative console."

WebSphere Application Server Profiles

Different WebSphere Application Server environments are defined using *profiles*. A profile defines the runtime environment (JVM) for web applications. A profile includes:

- files and settings that the server processes require during runtime.
- a common set of shared WebSphere product binaries.
- a unique directory name containing required files and subdirectories.

PeopleSoft Internet Architecture makes use of the *application server* environment type to deploy the PeopleSoft applications. When you run the PeopleSoft Internet Architecture installation, one of the elements that the system creates are the necessary application server profiles. The number of application server profiles created varies depending on whether you elected to implement a single server or multi server installation.

Note: The Application Server profile location differs from the default location where the PIA installation process creates server profiles.

Single Server Installation

If you specified the default application name of "peoplesoft" at install time, an application server profile with the same name gets created in *PIA_HOME*\websrv\<>peoplesoft<. When the PIA install creates an application server profile, it creates a default server names "server1" (single JVM process) and all of the PeopleSoft web modules are deployed to this single server.

You can view the PeopleSoft modules, such as PORTAL, PSEMHUB, PSIGW, and so on, in the single server profile directory structure.

PIA_HOME\websrv\<>peoplesoft<\installedApps\peoplesoftCell\peoplesoft.ear

Multi Server Installation

If you specified the default application name of "peoplesoft" at install time, the application server profile is created under *PIA_HOME*\websrv. In the case of a multi server installation, the profile contains the following separate servers.

Profile	Description
server1_peoplesoft	Hosts the PORTAL, PSIGW, and other web modules used for PeopleSoft online transactions.
PSEMHUB_peoplesoft	Hosts the PSEMHUB web module used by the PeopleSoft Environment Management Hub.

IBM HTTP Server

The IBM HTTP Server is the IBM version of the Apache HTTP Server. It is a separate installation required for IHS and IHS plug-ins. You may use the IBM HTTP Server as a reverse proxy server within your PeopleSoft implementation.

For more information see your PeopleTools installation guide, “Installing Web Server Products”.

Starting and Stopping WebSphere Application Servers

By default, all of the servers of a WebSphere instance are stopped when you install PIA. You need to start the server in order to access the PeopleSoft components.

To start and stop WebSphere servers, use the delivered .BAT and .SH files located in:

PIA_HOME\webserv\<<profilename>\bin

Note: During the server startup, the Dynamic runtime provisioning option is enabled by default. This option is used to enable only the necessary containers (servlet container, EJB container, and so on) when a server starts.

Starting the WebSphere Server

To start the WebSphere server, enter the following command:

Operating System	Command
Windows	startServer.bat <i>server_name</i> -profileName <i>profilename</i>
UNIX	startServer.sh <i>server_name</i> -profileName <i>profilename</i>

Stopping the WebSphere Server

To start the WebSphere server, enter the following command:

Operating System	Command
Windows	stopServer.bat <i>server_name</i> -profileName <i>profilename</i>
UNIX	stopServer.sh <i>server_name</i> -profileName <i>profilename</i>

Managing WebSphere Server with PSADMIN

You can also start and stop PIA domains running on IBM WebSphere using the PIA menus in PSADMIN.

See [Using the Web \(PIA\) Server Menu](#).

Configuring Reverse Proxy Servers For WebSphere

This section contains an overview and discusses how to:

- Configure an IBM HTTP server as a reverse proxy server.
- Configure Microsoft IIS as a reverse proxy server.
- Configure Oracle Sun Java System Web Server as a reverse proxy server.

Understanding Reverse Proxy Servers With IBM WebSphere

Using reverse proxy servers adds an additional, protective layer between your application and the internet or your end users. A reverse proxy server receives user requests, and sends them to a back end content server, usually behind a firewall. The back end server, in this case your PeopleSoft web server, remains unknown to the user.

For your PeopleSoft implementation, you can configure reverse proxy servers for WebSphere on the following web servers:

- IBM HTTP Server.
- Microsoft Internet Information Server (IIS).
- Oracle Sun Java System Web Server.

The communication between your web server and your reverse proxy server is configured using delivered plug-ins. You must install the web server software before you can install a plug-in for the web server. You can install the web server plug-in by itself on a machine where WebSphere Application Server ND has been installed but the plug-in has not. You can also install a plug-in on a remote machine where the HTTP proxy server is already installed.

Web Server Plug-In

Web server plug-ins enable the web server to communicate requests for dynamic content, such as servlets, to the application server. A web server plug-in is associated with each web server definition. The configuration file (plugin-cfg.xml) that is generated for each plug-in is based on the applications that are routed through the associated web server.

WebSphere RPS Plug-in

The RPS plug-in is used to forward HTTP requests from the proxy server to the PeopleSoft web server. The RPS plug-in provides:

- XML-based configuration file.
- Standard protocol recognized by firewall products.

- Security using HTTPS, replacing proprietary Open Servlet Engine (OSE) over Secure Sockets Layer (SSL).

Configuring IBM HTTP Server as a Reverse Proxy Server

To configure the IBM HTTP Server for use as a reverse proxy server, you use the IHS plug-in.

Before you perform the following steps, you need to install the following items:

- IBM HTTP Server.
- Web server plug-ins.

See *PeopleTools 8.53 Installation* for your database platform, *Installing Web Server Products*, "Installing IBM HTTP Server and Web Server Plug-ins"

To configure IHS for reverse proxy:

1. Start WebSphere server and open the ISC window.
2. Navigate to Environment, Virtual Hosts, pia_host, Host Aliases.

The PeopleSoft application is deployed on a virtual host called "pia_host".

3. Create new entries for the required ports.

For example:

Hostname = *, Port =10001 (for web server port)

Hostname = *, Port =10002 (for HTTP Administration Server port)

Hostname = *, Port =10043 (for SSL port assigned to IHS)

4. In a multi server environment, repeat the steps 2. and 3. for the other virtual hosts and "psemhub_host".
5. Click Apply and save the settings to "master".

This updates <PIA_HOME>/webserv/profile_name/config/cells/node_name/virtualhosts.xml

6. From the WebSphere Plug-ins installation, copy the configureWebserverDefinition script from the <Plugin_Install_Root>/bin to the directory <PS_HOME>/webserv/profile_name/bin and run it.

This creates the web server definition in WebSphere server.

7. Generate the plugin-cfg.xml by selecting the web server definition in Servers, Web servers.
8. Copy the plugin-cfg.xml from <PS_HOME>/profile_name/config/cells/cell_name/nodes/node_name/servers/WebserverDefinition to <Plugin_Install_Root>/config/WebserverDefinition so that IHS can communicate with WebSphere directly and access the PeopleSoft application.
9. Restart the WebSphere server, IBM HTTP Server, and IBM HTTP Administration Server.
10. Verify accessing the PeopleSoft application using the IHS HTTP port.

Note: In scenarios where the system needs to process large amounts of data, a page may become stuck in the processing status. While the change is reflected on the database, it cannot be viewed on a page until another session is started. This problem can be resolved by increasing the `ServerIOTimeout` from 60 seconds to 600 seconds, for example, in the `Plugin-cfg.xml` file located in `<PLUGIN_HOME>/config/webserver1`.

Configuring Microsoft IIS as a Reverse Proxy Server

This section discusses how to configure Microsoft IIS as a reverse proxy server for WebSphere. Before you perform these steps, the following items need to be installed:

- Microsoft IIS.
- Web server plug-ins.

Microsoft IIS Installation should have the following components already installed in order for the RPS setup to be successful:

- IIS Management Compatibility
- IIS Management Console
- IIS Scripting Tools
- IIS WMI Compatibility
- IIS Metabase compatibility
- ISAPI Extensions
- ISAPI Filters

See http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.nd.doc/info/ae/ae/tins_webplugins.html.

See *PeopleTools Installation*, Installing Web Server Products, "Installing IBM HTTP Server 7.0 and Web Server Plug-ins"

See your Microsoft IIS documentation.

Installing WebSphere Web Server Plugin for Microsoft IIS

Before installing the plugin, create a new IIS website according to the instructions given in the IBM WebSphere "Installing Web Server Plug-ins" documentation. Alternately, you can use the Default website if it is available. The Plugins install wizard prompts for the location where the WebSphere plugins for IIS need to be installed and also for the location where WebSphere is installed. The wizard also prompts for the name of the web server definition to be created.

Note: The WebSphere version may come with 32-bit and 64-bit web server plugins. The 32-bit plugin installer is in CD1 and 64-bit plugin installer is in CD2. Install the appropriate plugin based on your machine architecture and the mode (32-bit or 64-bit) in which your IIS is running.

Configuring the Plugin on Microsoft IIS

To configure the plugin:

1. From the WebSphere plugin installation, copy the `configureWebserverDefinition` script from the `<Plugin_Install_Root>/bin` to the directory `<PS_HOME>/webserv/profile_name/bin`.
2. Start the WebSphere server and run the `configureWebserverDefinition` script from the PIA profile location.

This creates the web server definition in the WebSphere server.

3. In the Administrative Console select Environment, Virtual hosts, and add new Host Alias entries for the IIS web server HTTP port into following virtual hosts:
 - `pia_host`
 - `psemhub_host`
4. Select Servers, Web servers, and select the server definition and generate the `plugin-cfg.xml`.
5. Copy the `plugin-cfg.xml` from `<PS_HOME>/profile_name/config/cells/cell_name/nodes/node_name/servers/WebserverDefinition` to `<Plugin_Install_Root>/config/WebserverDefinition`.

This enables the IIS web server to communicate with WebSphere directly and access the PeopleSoft application.

6. Restart the Microsoft IIS web server, and start the IIS website that you have created.
7. Restart the WebSphere server and login to the Administrative Console.
8. (Optional) Enable Administrative Console to manage the IIS web server.

This will show the IIS web server as running. This step is needed only if you want to manage your IIS web server from the WebSphere Administrative Console.

- a. In the Administrative Console select Servers, Server Type, Web servers.
 - b. Click the web server definition to create one for the IIS web server.
 - c. Enter the port number for the IIS web server.
9. Access the PeopleSoft application through the IIS web server HTTP port:

`http://<hostname>:<IIS_HTTPPort>/ps/signon.html`

Note: If you have a Windows 64-bit machine, IIS runs in 64-bit mode by default and requires 64-bit WebSphere plugins to be installed. However, if you have installed 32-bit WebSphere plugins then you can set IIS to run in 32-bit mode by executing the following script at the command prompt:

```
CSCRIPT %SYSTEMDRIVE%\Inetpub\AdminScripts\adsutil.vbs SET W3SVC/
AppPools/Enable32bitAppOnWin64 1
```

Restart the IIS web server following the execution of this script.

Configuring Oracle Sun Java System Web Server as a Reverse Proxy Server

The following steps discuss how to configure the Oracle Sun Java System Web Server as a Reverse Proxy Server. Before you perform these steps, the following items need to be installed:

- Oracle Sun Java System Web Server.
- Web server plug-ins.

Installing WebSphere Plugin for Oracle Sun Java System Web Server

The Plugins install wizard prompts for the location where WebSphere and Oracle Sun Java System Web Server are installed. The wizard also prompts for the name of the web server definition to be created, the location of the obj.conf file, and the location magnus.conf file. These two configuration files can be found under the Oracle Sun Java System Web Server installation in the following directory:

```
$SunJava_Home/https-<hostname>/config
```

Configuring the Plugin

To configure the plugin:

1. From the WebSphere Plugins installation, copy the configureWebserverDefinition script from the <Plugin_Install_Root>/bin to the directory <PS_HOME>/webserv/profile_name/bin.
2. Start the WebSphere server and run the configureWebserverDefinition script from PIA profile location.

This creates the web server definition in WebSphere server.

3. In the Administrative Console, select Environment, Virtual Host, and add the new Host Alias entries for Oracle Sun Java System Web Server HTTP port into following virtual hosts:

- pia_host
- psemhub_host

4. Select Servers, Web servers and generate the plugin-cfg.xml by selecting the web server definition.
5. Copy the plugin-cfg.xml from <PS_HOME>/profile_name/config/cells/cell_name/nodes/node_name/servers/WebserverDefinition to <Plugin_Install_Root>/config/WebserverDefinition,

This enables the Oracle Sun Java System Web Server to communicate with WebSphere directly and access the PeopleSoft application.

6. Note the following entries in the Obj.conf file.

```
After the <Object name=default> tag

Service fn="as_handler"
AddLog fn="as_term"
```

7. Note the following entries in the Magnus.conf file.

UNIX:

```
Init fn="load-modules"
    funcs="as_init,as_handler,as_term"
    shlib="/opt/IBM/WebSphere/Plugins/bin/libns61_http.so"

Init fn="as_init"
    bootstrap.properties="/opt/IBM/WebSphere/Plugins/config/webserver1
/plugin-cfg.xml"
```

Windows:

```
Init fn="load-modules"
    funcs="as_init,as_handler,as_term"
    shlib="C:\IBM\WebSphere\Plugins\bin\ns41_http.dll"

Init fn="as_init"
    bootstrap.properties="C:\IBM\WebSphere\Plugins\config\webserver1
\plugin-cfg.xml"
```

8. Restart the Oracle Sun Java System Web Server.
9. Restart the WebSphere web server.
10. Access the PeoplesSoft application through the Oracle Sun Java System Web Server HTTP port.

`http://<hostname>:<SunJava_HTTPPort>/ps/signon.html`

11. (Optional) Enable Administrative Console to manage the Oracle Sun Java System Web Server.

This will show the Oracle Sun Java System Web Server as running. This step is needed only if you want to manage your Oracle Sun Java System Web Server from the WebSphere Administrative Console.

- a. In the Administrative Console select Servers, Server Type, Web servers.
- b. Click the web server definition to create one for Oracle Sun Java System Web Server.
- c. Enter the port number for the Oracle Sun Java System Web Server's.

Configuring for Deleting Attachments

If you intend to allow end users to delete attachments, and you are using the Oracle Sun Java System RPS, you need to make sure the access control lists for Oracle Sun Java System RPS are configured correctly.

To configure Oracle Sun Java System RPS for deleting attachments:

1. Start the Sun Java admin server and login to the administrative console.
2. Navigate to Configurations, and click the appropriate configuration.
3. Select Access Control, Access Control Lists, where you'll see the default and es-internal access control lists.
4. Modify the default access control list.
 - Click default.

- For Allow anyone and Allow all, click delete rights.
 - Save.
5. Modify the es-internal access control list.
 - Click es-internal.
 - Click on allow anyone and enable delete rights.
 - Click on deny anyone and disable delete right.
 - Save.
 6. Save and deploy the change so it is reflected in the instances.

Setting Up SSL For WebSphere

This section provides an overview and discusses how to:

- Generate a certificate using pskeymanager.
- Configure the WebSphere container to support SSL.

Understanding WebSphere Key Stores

WebSphere manages keys in key store files. There are two types of files:

- key stores
- trust stores

These store types are very similar, however the trust store contains only trusted signers. The Certificate Authority (CA) certificates and other signing certificates are kept in a trust store. Personal certificates with private keys are stored in a key store.

The pskeymanager utility is a PeopleTools wrapper to Java's keytool, used to manage the predefined WebSphere keystore located in the following directory:

```
PIA_HOME\webserv\profilename\piabin\pskeymanager
```

Generating a Certificate Using pskeymanager

Use the following steps to generate a self-signed certificate for the web container.

To generate a certificate using pskeymanager:

1. At a command prompt, change to the WebSphere domain directory, for example:

```
PIA_HOME \webserv\profilename\piabin
```

2. Create a new private key and certificate request for your server.

a. Run the following command:

```
pskeymanager.cmd -create
```

b. Follow the prompts and specify the required information for creating a certificate, such as alias, common name, organizational unit, location, and so on.

c. Make sure a Certificate Signing Request (CSR) file named *alias_certreq.txt* was created.

You submit this data to a CA for obtaining a public key that you can load into your key store.

3. Decide which CA you wish to use.

You may use an CA that is compatible with Sun's Java JKS standard.

As an example, the following steps indicate how to submit the CSR that you generated to Verisign to obtain a trial certificate.

4. Submit your CSR to a CA.

For example, access Symantec's site at:

<https://www.verisign.com/products-services/index.html>

When prompted, copy and paste the contents of your CSR, provide all necessary contact information, and submit the request.

5. Check your email for the certificate sent from the CA.

The certificate from the CA should look similar to the following:

```
-----BEGIN CERTIFICATE-----
DMICHDCCAcYCEAHSeRkm2guFL+6OvHr4AS0wDQYJKoZIhvcNAQEEBQAwwgkxvFjAP
AANVBAoTDVZ1cm1TaWduLCBLbAMxRzBFBgNVBAsTPnd3dy52ZXJpc2lnbi5jb20S
VcVwb3NpdG9yeS9UZXR0Q1ETIEluY29ycC4gQnkgUmVmLiBMaWFiLiBMVEQuMUYF
LIGEc3VyYW5jZXMGKEMpVRMxOSDFertdsfh67TIwNDAwMDAwMFoXDTAwMTIxODIA
ONT1LVoweTELMaKGA1UERhMCMVVMxEzARBgNVBAGTCkNhbGlm3JuaWEExEzARBgNK
VBAUCOBsZWFzYW50b24BEzARBgNVBAoUC1Blb3BsZVNvZnQxWDASBgNVBAsUC1BT
Eb3sZVVvb2xzMRUwEwADVQQDFAxEQ1JPV04xMTE0MDAwXDANBggkqhkiG9w0BAQET
SAALADBEAkeEAucfM/GOQhdkk4Q0ZD5i1l4gp6WTYMc4IaReoCYkeAmDKAVcYzY3R
Mdbp4RC8SABd3bjjDOHcoCak9U6oSvL+HQIDAQABMA0GCSqGSIb3DQEBAUAA0EO
Arm3uf634Md0fqqNxxhAL+e9rbY0ia/X48Axloi17+kLTVI1YPOp+Jy6S1p5iNIFC
DhskdDFH45AjSDAFhjruGHJK56SDFGqwq23SFRfgtjkjyu673424yGWE5Gw4576K
DsdDFG256EDHY45yTRH67i345314GQE356mjsdhhjUwBtrh43Gq3QEVE45341tS
YDY6d471DmQxDs9wGt1bkQ==
-----END CERTIFICATE-----
```

6. Copy the entire certificate, including --BEGIN CERTIFICATE-- and --END CERTIFICATE--, and save it as a file named *webservname-cert.pem*.

Note: To save the file, don't use a word processor that inserts formatting or control characters.

Note: If you need to FTP your certificate to UNIX, you must FTP it in ASCII mode.

7. Download the CA root certificate:

For example, if downloading the Verisign\Symantec trial root CA certificate.

- a. Download the Trial Root CA certificate from:

https://www.verisign.com/support/verisign-intermediate-ca/Trial_Secure_Server_Root/index.html

- b. From the specified link, click Select All, and copy the contents of the certificate into the `verisignRootCA.cer` file and save it to your WebSphere domain directory.

- c. Download VeriSign's Trial Intermediate CA certificate from:

<https://www.verisign.com/support/verisign-intermediate-ca/trial-secure-server-intermediate/index.html>

- d. Click Select All and copy the contents of the certificate into the `verisignInterCA.cer` file, and save it into your WebSphere domain directory. You can also append the contents of this Trail Intermediate CA certificate to the Root CA certificate file `verisignRootCA.cer`.

Note: If you need to FTP your certificate to UNIX, you must FTP it in ASCII mode to your WebSphere domain directory.

8. Import the CA's certificates into your key store.

To import the CA's public certificate into your key store, run:

```
pskeymanager.cmd -import -trustcacerts
```

For example, when prompted for an alias, specify the appropriate name to store the CA as, for example `VerisignTrialCA`. This name is only an alias for this certificate.

When prompted for the certificate file to import, specify the root certificate, such as `verisignRootCA.cer` file.

If any other certificates (such as the Verisign Intermediate certificate) are saved into a different file, run the command to import that certificate also.

9. Import your certificate into your keystore.

To import your public certificate into your keystore, run the following command from the command prompt

```
pskeymanager.cmd -import
```

When prompted for an alias, specify the same alias you did when you created your private key and certificate request.

When prompted for the certificate file to import, specify your certificate file, `webservername-cert.pem`.

Configuring the WebSphere Container to Support SSL

To complete the SSL configuration, the web container must be modified to use the self-signed certificates you created.

To set up WebSphere Container SSL:

1. Start ISC, and select Security, SSL certificate and key management, Manage endpoint security configurations.
2. On the Local Topology tab, expand the Inbound tree, and click on the appropriate node, as in *peoplesoftNode*.
3. In the Related Items list on the right, click *Key stores and certificates*.
4. In the resource table, click *NodeDefaultKeyStore* in the Name column.
5. In the Additional Properties list on the right, click *Personal certificates*.
6. Click Import.
7. On the General Properties page, select the Key store file radio button, and complete the following:
 - a. In the Key file name field, enter the fully qualified path to the keystore file containing the certificate to import.

For example,

C:\PSHCM\webserv\peoplesoft\piabin\pskeymanager
 - b. From the Type dropdown list, select *JKS*.
 - c. In the Key file password, enter the password you specified when creating *pskey*.
 - d. Click Get Key File Aliases.

The system searches the key store and should populate the Certificate alias to import list.
 - e. If you want to use a new alias, enter a new value in the Imported certificate alias field, otherwise leave it empty.
8. Click Apply and OK.
9. Save the configuration in the Administrative Console.

Note: To configure Outbound SSL, repeat the same steps within the Outbound tree.

Enabling TLS-Only on WebSphere

Transport Layer Security (TLS) protocol is an improvement on the SSL v3 protocol. This section discusses:

- Configuring WebSphere for TLS.
- Configuring browsers for TLS.
- Testing TLS.
- Configuring Reverse Proxy Servers for TLS.

Configuring WebSphere for TLS

To enable TLS-only on WebSphere:

1. Login to ISC (<http://host:adminport/ibm/console>).
2. Under the Security menu, select SSL certificate and key management, SSL configurations, NodeDefaultSSLSettings, Quality of protection (QoP) settings.
3. Change the Protocol value to *TLS* or *TLSv1*.

This ensures that WebSphere server will accept only TLS connections. That is, when the web server acts as a server (inbound) or as client (outbound) the SSL connections will be established through the TLS protocol. When testing from a browser make sure to check the browser settings to initiate TLS handshakes only.

Configuring Browsers for TLS

This section covers steps for configuring TLS on browsers.

Setting Up TLS on Microsoft Internet Explorer

To set up TLS on Internet Explorer:

1. Launch Internet Explorer.
2. Select Tools, Internet Options, and select the Advanced tab.
3. In the Settings box in the Security section, disable *Use SSL 3.0* and enable *Use TLS 1.0*.
4. Click OK and restart the browser.

Setting Up TLS on Mozilla Firefox

To set up TLS on Firefox:

1. Launch Firefox.
2. Select Tools, Options, click the Advanced icon, and select the Encryption tab.
3. In the Protocols group box, disable *Use SSL 3.0* and enable *Use TLS 1.0*.
4. Click OK and restart the browser.

Testing TLS

After setting TLS for WebSphere and browsers, the TLS communication can be verified by logging in to the PeopleSoft application through WebSphere's default SSL port (HTTPS).

For example:

```
https://<host_name>:<https_port>/<PIA site>/signon.html
```

You can find the HTTPS port in the WebSphere Administrative Console, by selecting Servers, Application Server, server1, ports. Find the port corresponding to the entry WC_defaulthost_secure.

Configuring Reverse Proxy Servers for TLS

It is strongly recommended that you access the vendor's documentation of the web server you are using for a reverse proxy server and use their instructions for setting up TLS.

Securing The Administrative Console and Applications For WebSphere

This section provides an overview, and discusses how to:

- Secure the Administrative Console.
- Secure applications (servlets).

Understanding WebSphere Security

This section discusses:

- Registries and repositories.
- Security Domains.

Registries and Repositories

With the IBM WebSphere Application Server, a user registry, or repository, authenticates a user and retrieves information about users and groups to perform security-related functions, including authentication and authorization. WebSphere makes access control decisions based on the information stored in the user registry or repository.

WebSphere supports multiple types of registries and repositories, including:

- Local operating system registry.
- Standalone lightweight directory access protocol (LDAP) registry.
- A standalone custom registry.
- Federated repositories.

Security Domains

WebSphere allows "fine-grained security configuration" enabling security to be configured at the cell, node, server, or cluster level. Also, with WebSphere, you can configure application security separately for administrative security within a cell environment. Administrative security can be enabled through Global Security, while application security can be enabled by creating a new Security Domain and customizing the security configurations specific to the domain.

You configure separate applications to use different security configurations by assigning the servers, cluster, or Service Integration Buses hosting the applications or servlets to appropriate security domains.

Note: For a user to configure multiple security domains, they must be assigned to the administrator role.

For example, administration can be configured to use a federated repository while the applications can be configured to use an LDAP registry. In previous versions of Websphere, administrative and user applications use global security attributes by default, where a user registry defined in global security authenticates users for every application in the cell. Using multiple security domains provides more flexibility and simpler configurations.

This section illustrates how you can configure application security separately by creating a new security domain and assign it to the server level scope. The Administrative Console is secured with global security settings. The primary admin user belongs to the Local Operating System registry realm. The Report Repository Servlet of the PeopleTools application is secured by providing access to a user from another realm.

Note: For simplicity, in this example, both realms are Local Operating System realms. Realms may also be implemented as a Federated Repository or LDAP Registry.

Securing the Administrative Console

This section discusses how to secure the Administrative Console for the profile associated with your PeopleSoft application. It is assumed that you have already set up user names and passwords on the host machine.

Note: The Administrative Console is secure by default in the current release. This section applies if someone at your site has disabled the Administrative Console security, and you want to reapply it.

For Windows, the user should be in the Administrators group. On UNIX, use the root user.

This user information (user ID and password) will be used during configuration and authentication.

Configuring Administrative Security

To configure Administrative Security:

1. Open the Administrative Console of the profile hosting the PeopleSoft application.
2. Select Security, Global Security.
3. In the User account repository group box, set Available realm definitions to Local operating system, and click Configure.
4. In the Primary administrative user name field, enter a valid user name.

In this example "ansrivat11" will be the admin user ID. This value is the name of a user with administrative privileges that is defined in the registry. This user name is used to access the Administrative Console or used by wsadmin. On UNIX this should be "root" user.

5. Select the Server User Identity that is stored in the repository radio button, and enter the same user name as mentioned above for Server user ID or administrative user on version WAS 6.x node field, and enter the user's password that corresponds to server identity.
6. In the Custom properties table, click New, and enter these values:

Name: *com.ibm.websphere.registry.UseRegistry*

Value: *local*

7. Click Apply and Save.
8. Select Security, Global Security and click the Administrative user roles link, ensure that the user ID you just specified as the Primary administrative user name appears assigned to the Primary administrative user name role.
9. Return to Security, Global security and set these options:
 - a. Select the Enable administrative security check box.
 - b. Select the Enable application security checkbox only if you want the application security to have the same global security configuration as the admin user. In this example, we illustrate the application security enabled separately within a separate domain.
 - c. Deselect the Java 2 security options. PeopleSoft does not adhere completely to Java 2 security.
 - d. Click Apply.

Configuring SSL

If you are using the default SSL configuration, extract the signer certificate from the WebSphere default Trust Store. If you have set up a customized SSL configuration extract the signer certificate corresponding to that configuration.

To configure SSL:

1. In the Administrative Console select Security, SSL certificates and key management.
2. Under Related Items, click Key stores and certificates.
3. Click NodeDefaultTrustStore.
4. Click Signer certificates.
5. Select Root signer and click Extract.
6. Enter a unique path and filename for the signer, such as c:\temp\rootsigner.arm.
7. Click OK.
8. Add the exported signer certificate to trust.p12 in the <PS_HOME>/webserv/<profile name>/etc directory to enable SSL connectivity.
 - a. Open the key management utility, iKeyman, for that product version.
 - b. Start ikeyman.bat or ikeyman.sh from <PS_HOME>/webserv/<profile name>/bin.
 - c. Select Key Database File, Open.
 - d. Open <PS_HOME>/webserv/<profilename>/etc/trust.p12.
 - e. Enter WebAS for the password. (Do not change this password.)
 - f. Select Add and enter the path to the signer certificate you extracted.

9. Log out from the Administrative Console and stop the web server.

Modifying soap.client.props File

In any text editor, open the soap.client.props file located in *PIA_HOME*\webserv\<<profile name>\properties. Set securityEnabled to true, and specify the appropriate user ID and password for loginUserID and loginPassword. For example,

```
com.ibm.SOAP.securityEnabled=true

com.ibm.SOAP.loginUserid=<user ID>
com.ibm.SOAP.loginPassword=<password>
```

Use the user ID and password used to access the Administrative Console. If you want to encode the password in this file then run the PropFilePasswordEncoder script located in the folder PS_HOME/webserv/<profileName>/bin. For example,

```
PropFilePasswordEncoder.sh PS_HOME/webserv/<profileName>
/properties/soap.client.props com.ibm.SOAP.loginPassword
```

Testing Administrative Console Security

Successfully starting the WebSphere Application Server and successfully logging into the Administrative Console verifies that Admin Security is enabled.

To test Administrative Console security:

1. Start the server.
2. Launch the Administrative Console.
3. When prompted to log in, submit the user ID and password you configured previously.
4. To ensure no PeopleSoft elements have been changed, sign on to PIA as well.

Configuring Application Security

This section describes how to configure Application Security for WebSphere. In the context of PeopleSoft, the PeopleSoft servlets (PORTAL, Report Repository, and so on) run in the PeopleSoft application. In this discussion, we secure the access to the Report Repository servlet in the PeopleSoft application as the example.

Modifying and Deploying the Delivered EAR file

In this procedure, you modify the delivered peoplesoft.ear file and deploy it using the PeopleSoft install program.

To modify and deploy peoplesoft.ear:

1. Locate the PeopleSoft ear file and move it to a temporary directory.
 - a. Copy the PeopleSoft application ear file (peoplesoft.ear) from PS_HOME\setup\PsmPPIAInstall\archives folder into a temporary directory and extract it. In this discussion, C:\temp is used.
2. Modify the application.xml file and add the <security_role> element as shown below.
 - a. Open C:\temp\META-INF\application.xml.

- b. Add the security section shown below:

```
<application>
...
...
...
<module>
<web>
<web-uri> helloportletapp.war</web-uri>
<context-root /helloportletapp</context-root>
</web>
</module>
<security-role>
<description>Role for SchedulerTransfer Servlet</description>
<role-name>SchedulerTransferRole</role-name>
</security-role>
</application>
```

- c. Save and close the file.

3. Modify the PORTAL web.xml file.

- a. Extract C:\temp\PORTAL.war to C:\temp\PORTAL directory.

- b. Open C:\temp\PORTAL\WEB-INF\web.xml.

- c. Add the <security-constraint> and <security-role> element after the <welcome-file-list> element, and before the </web-app> element. In this case, the SchedulerTransferRole is mapped to the resource Report Repository servlet.

```
...
</welcome-file-list>
<security-constraint>
<web-resource-collection>
<web-resource-name>SchedulerTransferWebResource</web-resource-name>
<description>SchedulerTransferWebResourceDescription</description>
<url-pattern>/SchedulerTransfer/*</url-pattern>
<http-method>GET</http-method>
<http-method>POST</http-method>
</web-resource-collection>
<auth-constraint>
<description></description>
<role-name>SchedulerTransferRole</role-name>
</auth-constraint>
</security-constraint>
<security-role>
<description></description>
<role-name>SchedulerTransferRole</role-name>
</security-role>
</web-app>
```

- d. Save and close the file.

4. Recreate PORTAL.WAR.

- a. Repackage PORTAL.war by running the following command

```
cd C:\temp\PORTAL
jar -cvf ..\PORTAL.war *
```

- b. Delete the C:\temp\PORTAL directory.

5. Repackage the peoplesoft.ear file by running the following command in the temporary directory.

```
jar -cvf ..\peoplesoft.ear *
```

6. Copy the recreated ear file into PS_HOME\setup\PsmPPIAInstall\archives.
7. Deploy the recreated ear file onto WebSphere using the PeopleSoft Internet Architecture installation program.

Create a New Security Domain For Application Security

Creating a separate security domain for applications enables you to assign tailored security attributes to the application.

To create a security domain:

1. Select Security, Security domains, and click New.
2. Provide a name and description for the security domain, and click OK.
3. Define the scope.

For this example, the scope is server1, but other implementations can set the scope at Cell, Clusters, Nodes, and so on.

4. Under Security Attributes, expand Application Security and select the Customize for this domain radio button, and select the Enable application security check box.
5. Expand User Realm, and select Customize for this domain, select a Realm type, and click Configure.

In this example, Local operating system is used, but any of the realm types can be used. The realm selected will need to contain the users who are authorized to access the resource, which in this case is the Report Repository servlet.

For this example, when configuring the realm, we use appRealm for the Provide a realm name field. In the Custom properties table we use com.ibm.websphere.registry.UserRegistry (Name) and local (Value).

Mapping Security Roles to User Groups

To map security roles to user groups:

1. Select Applications, Application Types, WebSphere enterprise applications, and click peoplesoft.
2. Click Security role to user/group mapping.
3. Select SchedulerTransferRole and click on Map Users.
4. Select the User realm as the one created for the application, in this case it is the "appRealm".
5. Click the Search button to display all the available user IDs for this realm.

For this example, we select the user ID ansrivat12 for the SchedulerTransferRole. This ensures that only the user ansrivat12 is authorized to access the Report Repository servlet.

Testing Application Security

Restart the web server. Test the security of Report Repository servlet (as we configured it in this example) by using the following URL:

```
http(s)://<hostname>:<PIA http(s) port>/SchedulerTransfer/ps
```

You should be prompted for user ID and Password dialog. If authentication is successful, you should be shown the servlet output.

Setting HTTP Session Timeout

HTTP session timeout controls are accessible on the Security page of the web profile interface. PeopleSoft Internet Architecture ignores any session timeout control set on the web server. At run time, the session timeouts set in the web profile override any HTTP session timeouts set at the web server level.

Setting Authentication Failure Timeout

To limit the effectiveness of DOS attacks on failed authentications, you can use the `psft_failtimeout` Java option. Add this option in the `setEnv` script and assign a value in seconds. By setting the value to 60 seconds, for example, you override the default session timeout of 120 seconds (two minutes) when a user authentication fails or when a user is not yet authenticated.

For example,

```
SET JAVA_OPTIONS_WIN32=-server -Xms<nnn>m -Xmx<nnn>m -Dpsft_failtimeout=60
-XX:MaxPermSize=<nnn>m -Xcomp
```

To determine the proper value for this property, you need to check the time in seconds that it takes to send an `http(s)` request from the browser to the web server and multiply the result by 2.

Working With JVM Heap Size

Adjusting the JVM heap size settings can improve performance in some situations, however, the default settings are typically adequate for most situations.

To access the JVM heap size properties:

1. In the Administrative Console, select Servers, Server Types, WebSphere application servers, and click on your server in the resource list.
2. Select the Configuration tab, and in the Server Infrastructure section expand Java and Process Management, and click Process definition.
3. In the Additional Properties, click Java Virtual Machine.

The JVM heap size properties are:

Property	Description
Initial heap size	The initial heap available to the JVM. If blank, the system assumes the default value of 50 MB.
Maximum heap size	The maximum heap available to the JVM. If blank, the system assumes the default value of 256 MB. IBM recommends that if you determine that garbage collection occurs more than desired, increase the Maximum heap size value.

Working with Logging and Tracing Options

In addition to the logging and tracing options PeopleTools provides, WebSphere also offers a variety of tracing options.

See your IBM WebSphere Application Server and Administrative Console documentation.

Enabling HTTP Access and HTTP Error Logging

To enable HTTP access and error logging:

1. In the Administrative Console, select Servers, Server Types, WebSphere application servers, server1, and click the NCSA access and HTTP error logging link.
2. Enable HTTP access logging by selecting the options Enable logging service at server start-up and Enable access logging.

The HTTP access logs will be written to `PIA_HOME/webserv/profileName/logs/server1/http_access.log`.

3. Enable HTTP error logging by selecting Enable error logging.

The HTTP error logs will be written to `PIA_HOME/webserv/profileName/logs/server1/http_error.log`.

Enabling General Logging and Tracing

To access WebSphere logging and tracing options:

1. In the Administrative Console, select Servers, Server Types, WebSphere application servers, and clicking on your server in the resource list.
2. In the Troubleshooting section, click Logging and tracing.

The diagnostic trace, JVM log, and process log output files are directed to a variable, `${SERVER_LOG_ROOT}`, which refers to the following location on a typical PeopleSoft installation:

`PIA_HOME\webserv\profileName\logs\server`

The IBM Service log output file is directed to a variable, `{LOG_ROOT}`, which refers to the following location on a typical PeopleSoft installation:

PIA_HOME\webserv\profilename\logs

Chapter 9

Configuring Verity Search and Building Verity Search Indexes

Understanding Verity Search Indexes

This section provides an overview of search indexes and discusses:

- Types of indexes.
- Components of the search architecture.
- Index building.
- Search index limitations.
- User search strategies.

Important! The Verity search engine is not supported for use with the PeopleSoft Search Framework. If you intend to configure the PeopleSoft Search Framework and any of the search features based on the PeopleSoft Search Framework, such as Application Search or Keyword Search, you must install Oracle Secure Enterprise Search.

Overview of Verity Search Indexes

A search index is a collection of files that is used during a search to quickly find documents of interest. The process of creating the search index is also called building the search index. The set of files that make up the index is a *collection*. This collection contains a list of words in the indexed documents, an internal documents table containing document field information, and logical pointers to the actual document files.

Fields contain metadata about a document. For example, Author and Title might be fields in an index. *VdkVgwKey* is a special field that identifies each document and is unique to all of the documents in the collection.

The document table is a relational table with one row for each document and columns of fields. Every index can be modified by defining a set of fields for it.

In PeopleSoft search implementations, every search index has a home location where all of the files pertaining to that index are located. This directory is the home directory of the index and is typically located at *PS_CFG_HOME/data/search/INDEXNAME*. Under this directory is another directory named for the database to which the application server or the Process Scheduler is connected. The actual collection files reside in this database directory.

Every search index can be modified by changing the configuration files that are associated with the index. These configuration files are known as *style* files and reside in the style directory. A typical configuration of style files define fields for a particular index.

Types of Verity Indexes

PeopleSoft software supports three types of search indexes:

- Record-based indexes.
- HTTP spider indexes.
- File system indexes.

Record-Based Indexes

Record-based indexes are used to create indexes of data in PeopleSoft tables. For example, if the PeopleSoft application has a catalog record that has two fields (Description and PartID), you can create a record-based index to index the contents of the Description and PartID fields. Once the index is created, you can use the PeopleCode search application programming interface (API) to search this index.

HTTP Spider Indexes

HTTP spider indexes index a web repository by accessing the documents from a web server. You typically specify the starting uniform resource locator (URL). Then the indexer walks through all documents by following the document links and indexes the documents in that repository. You can control to what depth the indexer should traverse.

File System Indexes

File system indexes are similar to HTTP spider indexes, except that the repository that is indexed is a file system. You typically specify the path to the folder or directory. Then the indexer indexes all documents within that folder. HTTP spider indexes and file system indexes are sometimes collectively referred to as *spider* indexes. The indexer recognizes a wide variety of document formats, such as Word or Excel documents. Any document that is an unknown format will be skipped by the indexer.

Components of the Verity and PeopleSoft Search Architecture

PeopleSoft search architecture uses two main technologies: that provided by the PeopleSoft Portal and that provided by Verity. They are connected by the PeopleSoft search API.

PeopleSoft Portal Technologies

The PeopleSoft Portal search technology contains the following components:

- Search input field.
Captures a query string that is entered by users in the portal header.
- Search API.
Passes the query string that is captured in the search input field to the Verity search engine.
- Portal Registry API.
Applies security to filter the search results.
- Portal registry.

Contains a repository of content references that can be searched.

- Search results page.

Formats and displays search results for the user.

- Search options.

Enables users to personalize search behavior and results.

Note: By default, the PeopleSoft search performs case-insensitive searches.

Verity Technologies

The basic items of the Verity architecture that are incorporated in the PeopleSoft Portal search architecture are:

- Verity collection.

This is the set of files forming a search index. When a user performs a search, the search is conducted against the Verity collection. You can create and maintain your own collections with the Search Design and Search Administration PeopleTools.

- BIF file.

This is an intermediate file that is created in the process of building a Verity collection. The BIF file is a text file that is used to specify the documents to be submitted to a collection. It contains a unique key, the document size (in bytes), field names and values, and the document location in the file system.

- XML file.

This is another intermediate file that is created in the process of building a Verity collection. The XML file is a text file named *indexname.xml* that contains all of the information from the documents that are searchable but not returned in the results list. This information is stored in zones. Zones are specific regions of a document to which searches can be limited.

- Style files.

These files describe a set of configuration options that are used to create the indexes that are associated with a collection.

- mkvdk.

This Verity command-line tool is used to:

- Index a collection.
- Insert new documents into a collection.
- Perform simple maintenance tasks, like purging and deleting a collection.
- Control indexing behavior and performance.

PeopleSoft Verity Search Utilities

To create and administer search indexes for use with PeopleSoft software, use the PeopleTools utilities under PeopleTools, Search Engine. The utilities enable you to administer indexes and to create file system, spider, and record-based indexes.

Building Verity Indexes

For both HTTP spider and file system indexes, options are available to include or exclude certain documents based on file types and Multipurpose Internet Mail Extensions (MIME) types. The index building procedure is different for record-based indexes and the spider indexes. Typically, the index building procedure is carried out from an Application Engine job that is scheduled by using the process scheduler.

The steps for building record-based indexes are:

1. The data from the application tables is read and two files called *indexname.xml* and *indexname.bif* are created.

indexname.xml contains one XML record for each document that needs to be indexed. The XML record contains all of the data that needs to be indexed. *indexname.bif* contains field information, the VdkVgwKey document, and offsets to denote the start and end of each document in the XML file.

2. The XML and the bulk insert file (BIF) files are typically generated through PeopleCode and reside in the home location of the index. The Verity utility, *mkvdk*, is called, passing in the BIF file as the argument to build the index.

The steps for building spider indexes are:

1. The Verity utility, *vspider*, is called.

The *vspider* utility takes a number of arguments, but the most important ones are the starting URL or directory to spider and the number of links to follow.

2. The *vspider* utility walks through all of the documents in the repository and builds the index.

Verity Search Index Limitations

PeopleSoft Verity search index limitations include:

- Verity does not run on IBM z/OS.

For the most current support information for Verity (and any product), always check the Certifications tab on My Oracle Support. Review the operating system certifications for Verity, as well as the deployment options (Type-1, Type-2, Type-3).

- Verity collections must reside on the PeopleSoft application server or be accessible from it through a shared drive.

Satisfying this requirement can take several forms, depending on the application server's operating system. On Microsoft Windows, this could be a network drive. On UNIX, this could be an NFS-mounted drive.

- Verity collections are most efficient if you index large groups of data, rather than indexing one or two documents at a time.

Small updates degrade the index and require that you run the Verity cleanup utility.

- Style files are located in the style subdirectory of the index.

To make style changes, apply them to the files in this directory.

- You can have only one language per collection.

Additionally, a number of Verity search index features are limited to certain maximum values, as follows:

Feature	Limitation
Wildcards	Wildcard auto-expansion is limited to 16,000 matches.
Number of collections	The maximum number of physical collections that can be searched at one time is 128.
Documents per collection	The maximum number of documents allowed per collection is 16 million, subject to disk space availability.
Fields per collection	The maximum number of fields allowed per collection is 250.
Field length	The maximum length of any field is 32 kilobytes. Note: The actual number of characters that translates to depends on the character set being used.
Field value length in bulk files	The maximum length of a field value in a bulk file is 32 kilobytes. Note: The actual number of characters that translates to depends on the character set being used.
Zones per document	The number of zones allowed per document is unlimited.
Characters in path	The maximum path size allowed is 256 characters.
Maximum documents with sort specification	The maximum number of documents that are returned when a sort specification is applied is 16,000.
Sort fields per search	The maximum number of fields that can be included in a sort specification is 16.

Refer to the Verity documentation for details about these features.

User Search Strategies for Verity

A user submits a search request by entering a search string into the search input form field in the portal header. The “<form action=...>” element in the portal header is generated at runtime to link to a PeopleSoft Internet Architecture page, and a Java script submits the form. The query string is passed to the Search API as a parameter named PortalSearchQuery to find matching results. Those results are filtered for security through PeopleCode by the Portal Registry API. The search results page echoes the

original query string and displays a list of content references that match the request. If the user clicks the Go button but does not enter a search query, the search results page displays without any results.

The search results page performs the following steps:

- Changes the case of the entered text to all uppercase characters.

By default, the Verity search engine searches for all mixed-case variations when a query string is entered in all lowercase or in all uppercase. However, search queries that are entered in mixed-case automatically become case sensitive. (For example, a query on *Apple* behaves as if the user had specified *Apple*, which would find only the precise string *Apple*, while a query on *apple* finds *APPLE*, *Apple*, and *apple*.) But the portal makes one important change: It changes the case of the query string to all uppercase, prohibiting users from truly executing case-sensitive searches. This avoids situations where mixed-case searches would otherwise return no results. On the search results page, however, the original case is echoed back to the user.

- Formats the query string to pass to the Search API.

This includes filtering out expired and hidden content reference, and content references that are not valid yet.

- Calls the Search API.

This returns the query results.

- Calls the Portal Registry API.

This is done to apply security filtering to the results. Security is applied in PeopleCode by checking the Authorized property.

- Formats and displays search results.

This completes the user's search request.

Note: End users must enter a double backslash (\\) when they need to submit a search request containing a backslash in the text. Using a single backslash (\), will cause undesired results.

Configuring Verity Search Options

This section contains an overview and discusses how to:

- Configure search to run natively within the application server (Type-1).
- Configure search to run as a separate process managed by the application server (Type-2).
- Configure a separate Search Server (Type-3).

Understanding PeopleSoft Search Configurations

PeopleSoft offers these configuration options for enabling PeopleSoft search:

- Type-1: Verity running within the application server domain.

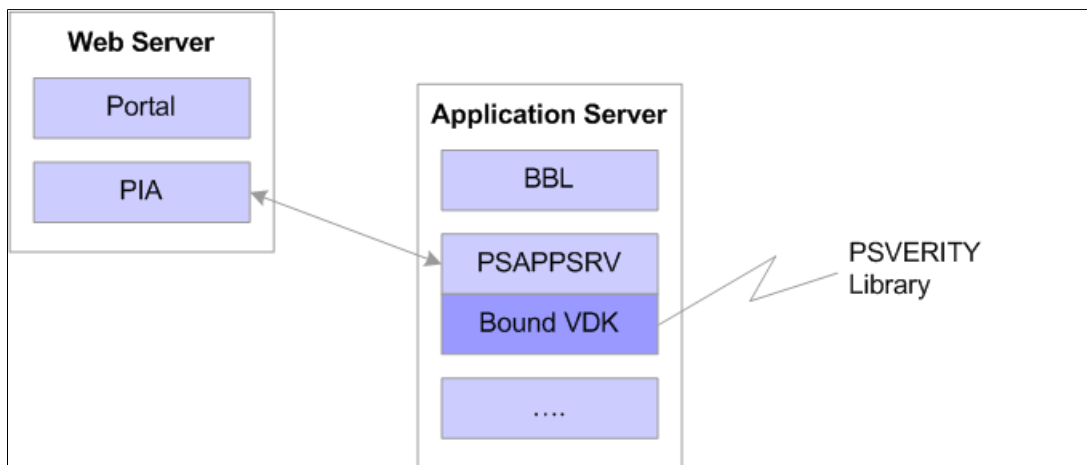
- Type-2: Verity running within a separate process managed by the application server.
- Type-3: Verity running within a separate search server.

Note: In some cases, the operating system determines which search configuration options can be used. *Always* refer to the *PeopleSoft Hardware and Software Requirements* guide, the Certifications area on Metalink, or customer support for the most recent support information.

Type-1: Verity Running within the Application Server Domain

Image: Type 1 search configuration: VDK bound to PSAPPSRV processes the request with the PSVERITY library

In this configuration, Verity runs within the application server. Its libraries are linked to the application server. For example, as shown in the following diagram, the Verity VDK is bound to the PSAPPSRV server process. When a search request is submitted, the VDK bound to PSAPPSRV processes the request with the PSVERITY library.



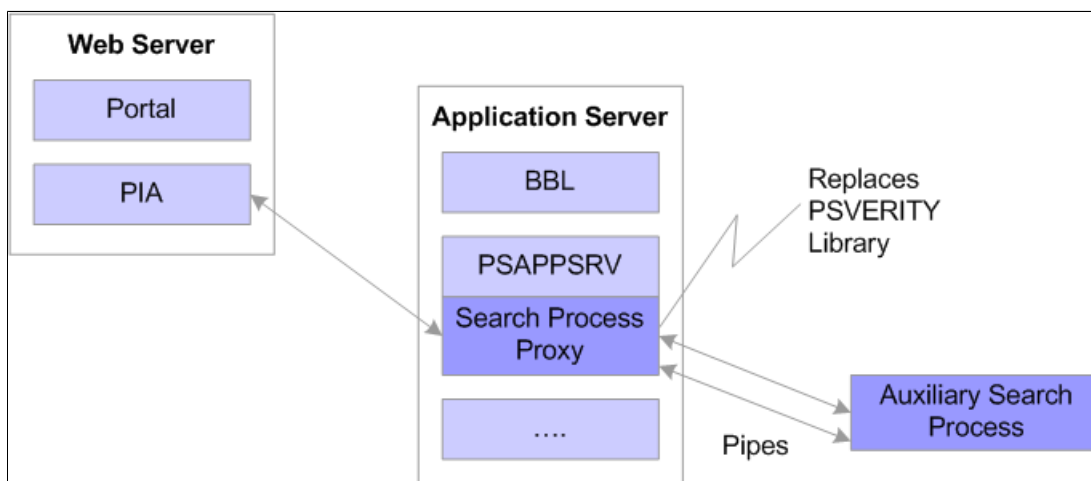
Note: This configuration has been used in PeopleSoft applications in all previous releases of the PeopleSoft Internet Architecture.

Type-2: Verity Running as a Separate Process Managed by the Application Server Domain

Having Verity run as a separately managed process enables application server domains running within the 64-bit framework to interoperate with Verity running within the 32-bit framework.

Image: Type 2 search configuration: The spawned process hosts the VDK processing on behalf of the application server domain

The following diagram illustrates that in this configuration, when the first search request is submitted, the PSAPPSRV server process spawns an auxiliary process to run within the application server domain. The spawned process hosts the VDK processing on behalf of the application server domain. A proxy search library within the application server routes search requests from the PeopleSoft Internet Architecture to the auxiliary search process.

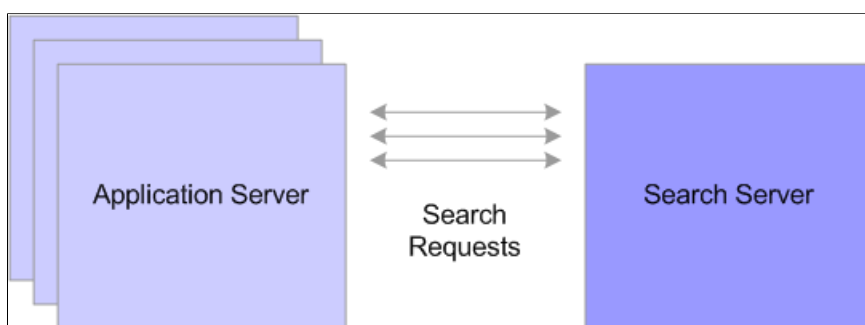


The proxy search library and the auxiliary search process transmit data using efficient system resources (anonymous pipes). Having both processes running on the same computer can reduce performance degradation introduced from the extra communication layer of the network in a Type-3 configuration.

Type-3: Verity Running within a Separate Search Server

Image: Type 3 search configuration: application server domain sends search requests to the search server on a separate physical machine

The following diagram illustrates an example that you can implement the Type-3 search configuration to centralize the configuration of the search features as well as the maintenance and storage of search indexes.



In this configuration, the application server domain routes search requests to the search domain running on a remote search server. Multiple application server domains may use the same search server to execute

search requests. In this scenario, the application server domain is the "client," submitting search requests and the search domain is the "server," processing requests and returning results.

Note: Tuxedo must be installed on both the application server machine and search server machine.

Configuring Verity Search to run within the Application Server (Type-1)

This configuration requires the application server to be installed as outlined in the PeopleTools Installation guide for your platform. This installation process installs the required application server and Verity software.

In the Search section of PSADMIN, enter 1 for the Deployment Type parameter.

```
Values for config section - Search
  Deployment Type=1
  Application Server Port=
  Remote Search Server Credentials=
```

Note: If you do not assign a value to the Deployment Type parameter, the system assumes the default configuration for your operating system.

Configuring Verity Search to Run as a Separate Process (Type-2)

This configuration requires the application server to be installed as outlined in the PeopleTools Installation guide for your platform. This installation process installs the required application server and Verity software.

In the Search section of PSADMIN, enter 2 for the Deployment Type parameter.

```
Values for config section - Search
  Deployment Type=2
  Application Server Port=
  Remote Search Server Credentials=
```

Note: If you do not assign a value to the Deployment Option parameter, the system assumes the default configuration for your operating system.

Configuring a Separate Verity Search Server (Type-3)

Setting up a remote search server to process requests for application server domains requires you to complete configuration steps on the:

- search server.
- application server(s).

Configuring the Search Server Domain

To configure a separate search server:

1. Ensure the environment is set up correctly.

Configuring a search domain is comparable to creating an application server domain on an application server. You need to make sure:

- Tuxedo is installed locally.
 - PS_HOME is available (locally or remotely)
 - PS_CFG_HOME is set correctly on the search server machine.
2. Launch PSADMIN, and select Search Server from the PeopleSoft Server Administration menu.
 3. On the PeopleSoft Search Server Administration menu select 2) *Create a domain*, and enter a name for the search domain.

```

-----
PeopleSoft Search Server Administration
-----

  1) Administer a domain
  2) Create a domain
  3) Delete a domain
  q) Quit

Command to execute (1-3, q) : 2

Please enter name of domain to create :SAMPLE

```

4. Select 1) *search*, for a configuration template.

```

Configuration templates:

  1) search

Select config template number: 1

```

5. When prompted to configure the search domain and change any configuration values, enter *y* to indicate "yes."
6. In the [Startup] section, add the information required for the search domain to connect to the application database.

The values entered should be identical to the connect information in any application server domain connecting to the same database.

Note: The search domain must connect to the same database as the application servers sending requests to the search domain.

7. In the [Database Options] section, select the same options you use for other application server domains in your environment.
8. In the [Domain Settings] section, select the same options you use for other application server domains in your environment, including, for example, Add To Path for specifying database driver locations.

Note: Make note of the unique Domain ID value. It is required when configuring the application server domains using the search server.

9. Modify the options in the [PSSRCHSRV] section.

Min Instances, Max Instances, Service Timeout

These parameters operate the same as PSAPPSRV.

See [PSAPPSRV Options](#).

Search Server Port

Enter the port address the search domain will monitor for search requests.

The default is *7778*.

Application Server Credentials

Enter a list of application server domains that will be using the search domain. The application servers need to be identified by Domain ID, Server ID, and port in the following format.

```
<Domain ID>|<Server ID>:<port>
```

Note: Server ID can be an IP Address or a hostname.

When multiple domains use the same search server, separate the entries by a comma (.). For example, the following illustrates how to enter two different domains running on two different servers.

```
APPDOM1|appsrv_computer1:7777,APPDOM2|appsrv_computer2:7777
```

Note: The Domain ID value can be found in the [Domain Settings] section of PSADMIN.

Configuring an Application Server Domain to use a remote Search Server

Once you have a remote search domain configured, you then need to modify each application server domain that will use that search server to process search requests.

To configure an application server domain to use a remote search server:

1. Launch PSADMIN, and initiate the configuration interface for the desired application server domain.
2. Modify the [Search] section.

Deployment Type

Enter *3* to indicate Type-3 configuration.

Application Server Port

Enter the port number on which the application server domain will "listen" for responses from the search domain. Make sure this value is the same port number you specified in the search domain in the Application Server Credentials parameter.

Remote Search Server Credentials

Specify the search server domain that will be used by the application server domain. The search server needs to be identified by Domain ID, Server ID, and port in the following format.

```
<Domain ID>|<Server ID>:<port>
```

- When prompted to configure Domains Gateway (External Search Server) indicate *y* for "yes."

Note: The Domains Gateway can also be enabled in the Quick Configure menu.

Setting Up Failover Search Domains

To provide high availability and to compensate for the possibility of issues with the network, a server machine, a search domain, or simply having to shut down a search domain for maintenance, you can configure failover search domains. In these situations, the unavailability of a primary search domain does not affect end users.

Failover search domains only process search requests when the primary search domain is unavailable. If the primary search domain is unavailable, the system seamlessly routes search requests to the next search domain specified in the failover string sequence.

For example, assume an application server domain has the following search domains specified in the search domain failover string in this order:

- SRCH_PRIMARY
- SRCH_FAILOVER1
- SRCH_FAILOVER2

If SRCH_PRIMARY is unavailable, the system checks to see if SRCH_FAILOVER1 is available, and if so, begins routing search requests to SRCH_FAILOVER1. If SRCH_FAILOVER1 is not available, then the system checks the availability of SRCH_FAILOVER2, and so on. When the primary search domain becomes available again, the system begins routing search requests to that search domain.

To set up failover search domains:

1. Install and configure the number of failover search domains you require.

It is recommended that each failover search domain reside on a separate server machine for optimal failover coverage.

2. For each failover search server domain, specify the complete list of application server domains that could potentially use the search domain for failover coverage.

Use PSADMIN, or edit the PSSRCHSRV.CFG manually. Specify the application server domains using the Application Server Credentials parameter in the PSSRCHSRV section.

3. For each application server domain using a particular set of search server domains, modify the Remote Search Server Credentials parameter in the [Search] section to include the connect information for each search domain, with the primary search domain appearing first and a comma (,) separating multiple values.

```
<Domain ID>|<Server ID>:<port>,<Domain ID>|<Server ID>:<port>
```

For example,

```
Remote Search Server Credentials=SRCH_PRIMARY|ts-sun04:7778,
SRCH_FAILOVER1|ts-sun05:7778,SRCH_FAILOVER2|ts-sun06:7778
```

Verity Search Server Administration

While the administrative tasks associated with search servers are similar to your application server or Process Scheduler administration, keep the following items in mind when managing search servers.

Working with Search Domains in PSADMIN

When administering search server domains, you use a subset of PSADMIN menu options.

```
-----
PeopleSoft Search Domain Administration
-----
      Domain Name: search01

1) Boot this domain
2) Domain shutdown menu
3) Domain status menu
4) Configure this domain
5) TUXEDO command line (tmadmin)
6) Edit configuration/log files menu
7) Clean IPC resources of this domain
8) Domain Gateway TUXEDO command line (dmadmin)
q) Quit
```

Command to execute (1-8, q) :

Using these menus is similar to the menus for an application server domain, except that items that are not applicable do not appear. For example, there are no menu options for purging cache, preloading cache, or setting up messaging servers because they do not apply in the context of search servers.

For search servers, the following options differ slightly from application server domain options:

Boot this domain

For application server domains, you have options to boot a domain in serial or parallel mode. Because the number of server processes within a search domain are typically fewer than a large domain, the option of a parallel boot to save time is unnecessary. With search domains, you are not presented with boot options, and the domain boots in serial mode.

Domain Gateway TUXEDO command line (dmadmin)

The dmadmin is similar to the tmadmin interface. dmadmin is an interactive command interpreter used for the administration of domain gateway groups defined for a particular Tuxedo application.

Locating Logging Information in Type-3 Search Configurations

The system writes logging information to these files:

- TUXLOG for both the application server and search server domain.
- APPSRV_MMDD.LOG for the application server domain.
- SRCHSRV_MMDD.LOG for the search server domain.

Monitoring Domain Gateway Connections

The domain gateway is a subcomponent of a Tuxedo domain that allows it to communicate with another domain through the network. The domain gateway ensures that the application server and search server

domains are successfully connected and able to transmit data. An application server domain and search domain can start independent of one another and do not report any obvious signs of being successfully connected when they start.

When working with search domains and troubleshooting Domain Gateway issues:

- Ensure that the domain gateway is enabled. Check the Tuxedo logs of both the application server and the search server. Both logs should indicate the gateway connection.
- Check machine and port configuration. Failure to connect, or connections with numerous disconnections can be caused by incorrect port and machine address information or another machine using the same port. Use canonical names if you are using a non-numerical IP address.

Use the `dmadmin` command line interface to monitor a Domain Gateway connection between a local and remote domain. Access this interface through the PeopleSoft Search Domain Administration menu in PSADMIN. The following commands can be helpful when working with search domains.

Command	Description
<pre>pd -d <LOCAL DOMAIN ACCESSPOINT_ID></pre>	<p>Use the <code>pd</code> (print domain) command to confirm whether or not the application server and search server domains are connecting and transmitting data. Confirm successful connection by viewing the 'Connected domains' list.</p> <p>The <code><LOCAL DOMAIN ACCESSPOINT_ID></code> is formed by prepending "SS" to the domain ID. For example, if the domain ID is <code>SRCHSERV</code> the value of <code><LOCAL DOMAIN ACCESSPOINT_ID></code> is <code>SS_SRCHSERV</code>.</p> <p>The following is sample output:</p> <pre>pd -d SS_SRCHSERV Local domain :SS_SRCHSERV Connected domains: Domainid: APPDOM1</pre> <p>If the search domain is not connected to the application server domain you will see output similar to this:</p> <pre>pd -d SS_SRCHSERV Local domain :SS_SRCHSERV Connected domains: Disconnected domains being retried: Domainid: APPDOM1</pre> <p>This examples show only one application server domain and one search domain. In reality, multiple application server domains would connect to one search domain. The <code>pd</code> command lists the status of each of the application server domains connected to a search domain.</p>
<pre>pstats -d <LOCAL DOMAIN ACCESSPOINT_ID></pre>	<p>Use the <code>pstats</code> command to extract monitoring statistics from the Tuxedo MIB regarding the domain gateway connection. This can help to identify the amount of requests being processed for application server domain clients.</p>
<pre>h</pre>	<p>Displays and describes all <code>dmadmin</code> commands.</p>

See Oracle Tuxedo documentation for complete `dmadmin` documentation.

Building Search Indices

For a search server (Type-3 configuration), a Process Scheduler deployed on the search machine should be used for indexing. Because Verity libraries may be available only on the search machine, and because any index would be used by the search server on the search machine, it is recommended to build the indices on the search machine to avoid having to relocate indexes from other machines. A recommended

approach is to deploy a Process Scheduler server along side the search server and specify that Process Scheduler server for generating indexes (PeopleTools, Search Engine, Administration, Schedule).

Note: For building search indexes on a Type-3 configuration, it is strongly recommended to use the PSNT Process Scheduler Server running on the same server machine as the Type-3 configuration.

If Verity is not supported on the operating system where your production application server domains run, another option is deploying an application server domain along side the search server. This application server would be accessible through its own web server instance possibly on a different port than the production application server. This provides access to an application server with Verity support, allowing the creation of indices interactively. Also, the indices would be created where the search server can locate them.

Working with Verity Indexes

This section provides overviews of common controls and supported MIME types, and discusses how to:

- Open existing collections.
- Create new collections.

Understanding Common Verity Controls

The following controls appear on the pages that are used for designing record-based, file system, or HTTP spider indexes.

Index	Shows the name of the index that you opened or the name that you gave the index on the Add New Value page.
Build Index	Invokes the collection build program. Before clicking this button, select all of the appropriate options for the collection.
Test Index	After building an index, click to test that the build program assembled the index properly. The Test Index page contains a single text field with a query button. Enter text to search for in the collection and click the [?] button to submit the query. The results return a list of the keys that are stored by Verity in the collection.
Show Logs	View the log files that are produced by the collection build program during execution. This is used mainly for troubleshooting.
Append to Verity Command Line	This control is for PeopleSoft internal use only.

Understanding Supported MIME Types

The following list contains the supported document MIME types. Any document that is not one of these types is ignored during the indexing process.

- application/msword
- application/wordperfect5.1
- application/x-ms-excel
- application/x-ms-powerpoint
- application/x-ms-works
- application/postscript
- application/rtf
- application/x-lotus-amipro
- application/x-lotus-123
- application/x-ms-wordpc
- application/x-corel-wordperfect
- application/x-wordprocessor
- application/x-spreadsheet
- application/x-presentation
- application/x-graphics
- application/x-keyview
- application/x-ms-write
- application/pdf
- application/x-executable
- message/rfc822
- message/news
- text/html
- text/sgml
- text/xml
- text/ascii
- text/enriched
- text/richtext
- text/tab-separated-values
- text/plain

- text/x-empty
- image/gif
- application/x-verity

Opening Existing Verity Collections

To open an existing collection:

1. Select PeopleTools, Search Engine.
2. From the available menus, select the type of collection that you want to open, as in record-based indexes, file system indexes, or HTTP spider indexes.
3. On the Find an Existing Value tab, use the Search for drop-down list box to select the appropriate criteria (begins with or contains).
4. In the edit box to the right, enter the character string that reflects the appropriate begins with or contains criteria.
5. Click Search.

Creating New Verity Collections

To create a new collection:

1. Select PeopleTools, Search Engine.
2. From the available menus, select the type of collection that you want to create, as in record-based indexes, file system indexes, or HTTP spider indexes.
3. Select the Add a New Value page.
4. Enter a name for the collection.
5. Click Add.
6. Specify the appropriate attributes for the collection as described in the following sections.
7. Save your work.

Note: You cannot create indexes of the same name even if they are of different types; for example, record, HTTP, or file.

8. Build the index.

Building Record-Based Verity Indexes

The record-based index extracts data from database tables and inserts the data into BIF and XML files, which are then indexed by Verity. The individual creating the index chooses the records (tables) to be indexed.

Note: The record-based index supports only data that is stored in PeopleSoft databases.

This section discusses how to:

- Modify record-based index properties.
- Add subrecords to search indexes.

Modifying Record-Based Verity Index Properties

Select PeopleTools, Search Engine, Record-Based Indexes to access the Design a Search Index page.

Image: Design a Search Index page (Record-Based)

This example illustrates the fields and controls on the Design a Search Index page (Record-Based). You can find definitions for the fields and controls later on this page.

Primary Record
Subrecords

Design a Search Index

Index: TEST

Index Location: NEW

Key returned in search results:

<pairs/>
▼

[Edit Key](#)

Build Index

[Test Index](#)
[Show Logs](#)

Parent Data Record

***Record (Table) Name:** 🔍

WHERE clause to append:

WHERE

Fields

***How to zone the index:** One zone ▼ [Click here for help with the Fields columns](#)

Fields Included in the Index

Record	Field Name	Verity Field	Word Index	Has attachment
1		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Append to Verity Command Line:

Parent Data Record

Record (Table Name)

Enter tables, views, or a PeopleSoft view that contains data. To combine the data from multiple PeopleSoft tables, to create a view on those tables and specify the name of that view here.

WHERE clause to append

Fine-tune the data that you receive by entering a Structured Query Language (SQL) WHERE clause.

Key returned in search results

Use to synthesize the VdkVgwKey, which supports an XML-like syntax enabling you to modify the tag that is returned by Verity.

You have the following options:

- `<pairs/>`: Inserts a string of NAME=VALUE;. One such pair is returned for each key of the record.
- `<row/>`: Inserts the record keys in a SQL-like syntax.
- `<field fieldname='MYFIELD'/>`: Inserts the value of MYFIELD if it exists in the record.
- `<sql stmt='SQL STATEMENT'/>`: Inserts the value that is returned by the SQL statement. The system accepts only the first row that is returned, and PeopleSoft software does not support SQL statements returning more than one column.

Edit Key

Click to access the page where you can change the results that are returned by the Key returned in search results functionality.

Fields

How to Zone the Index

One Zone: Select to put all of the data into one zone. With this option, the collection builds more quickly but the application can't restrict searches to the portions of the index that come from a particular field.

Field Zones: Select to create one zone for each PeopleSoft field on the record. Applications can specify that they want to access that particular zone in their searches.

Field Name

After you specify a record name, the fields in that record appear in this grid. Select the following options for each field in the record: Verity Field, Word Index, or Has Attachment (each option is explained in the following sections).

Verity Field

Select if the PeopleSoft field should be indexed as a Verity field.

In general, PeopleSoft fields that contain a lot of descriptive text, such as description fields, should be indexed as word indexes (See the following definition) and PeopleSoft fields that contain metadata about what is being indexed (such as ProductID) should be indexed as Verity fields.

Word Index

Select if this PeopleSoft field should be indexed as a word index. See the preceding Verity Field definition for guidelines on defining a PeopleSoft field as a Verity field versus defining it as a word index.

Has attachment

Enables you to index attachments that are referenced in the field as uniform resource identifiers (URIs). Refer to the PeopleCode Developer's Guide for a description of file attachments. If this field contains the URL to an attachment, select this check box. The indexer downloads the attachment and indexes it as part of the document. This item is enabled only if the corresponding PeopleSoft field contains character data, because numeric fields cannot contain URLs.

To use this field, you need a record that is designed with this feature in mind. In the record, each row has a text field that contains a URI or an empty string.

The text must be a valid File Transfer Protocol (FTP) URI (including the login and password string) of the following form:

- ftp://user:pass@host/path/to/filename.doc.
- A valid record URI of the form record://RECORDNAME/path/to/file.doc.
- A string of the form <urlid name="A_URLID"/>/path/to/file.doc.

The third form references an entry in the URL table (Utilities, Administration, URLs). If the URL ID that is named in the name attribute is valid, the entire URI is rewritten with the part in brackets replaced by the actual URI.

For example, if A_URLID is equal to ftp://anonymous:user@resumes.peoplesoft.com, the entire string in the previous example becomes ftp://anonymous:user@resumes.peoplesoft.com/path/to/file.doc and is treated like any other FTP URI.

Rows of data with empty strings in the URI field are ignored with no error.

If the string is one of these three valid URI forms and a document can be retrieved at that URI, the document is indexed with the same key as the rest of the row of data and is searchable.

To add subrecords to the index, select the Subrecords tab, and insert the child records that you want to include in the index.

Adding Subrecords to Verity Search Indexes

Select PeopleTools, Search Engine, Record-Based Indexes, Subrecords.

To index more than one record as a single document, the records must be hierarchically related. For example, the record that is specified on the previous page must be a parent of all the others. Formally, this means that the keys of each subrecord named must be a superset of the keys of the parent record. The parent record is the one that you specify in the Record (Table Name) field on the Primary Record page.

To add subrecords to an index:

1. Create and save the index definition.
2. Select PeopleTools, Search Engine, Record-Based Indexes, Subrecords.
3. Click the Add a new row button to insert the names of the records that are children of the parent record that is defined on the Primary Record page.

On the Primary Record page, the fields of the child record are added to the Fields grid. When you build the index, data from the child records whose keys match the row in the parent record is included as part of the parent record. When an end user searches for data that is found in the child record, the system returns a reference (VdkVgwKey) for the parent record.

Building File System (Spider) Verity Indexes

You can index file systems that are local to the application server. This refers to any file system on the physical server on which your application server domain runs, and it also refers to any drives that are accessible from the application server machine. File systems might include file servers, report repositories, and so on.

The index is compiled by using `vspider`. The program descends into the directory structure recursively and indexes the file types that you've selected to be indexed. It indexes only files that Verity supports for collections.

This section discusses how to:

- Set file system options
- Define what to index

Setting File System Options

Select PeopleTools, Search Engine, Filesystem Indexes to access the Design a Search Index page.

Image: Design a Search Index page (File system)

This example illustrates the fields and controls on the Design a Search Index page (File system). You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Design a Search Index' page for a file system. At the top, there are two tabs: 'Filesystem Options' (selected) and 'What to Index'. Below the tabs, the title 'Design a Search Index' is displayed. Underneath, there are two labels: 'Index: TEST' and 'Index Location: NEW'. To the right of these labels is a yellow 'Build Index' button, and further right are links for 'Test Index' and 'Show Logs'. Below this is a table with a dark blue header. The header row contains 'Filesystem paths', 'Customize | Find | View All | [grid icon] | First [left arrow] 1 of 1 [right arrow] Last', and a search icon. The table has two columns: 'List local filesystem paths to spider' and 'Remap Path to This URL'. Below the table, there is a text area labeled 'Append to Verity Command Line:'.

List local filesystem paths to spider

Specify the network file system path that contains the documents to index. Ensure that the local application server has the proper access to the file systems that you include in the list.

For Microsoft Windows, this means the drive mappings must be set up from the applications server. For UNIX, this means the correct network file system (NFS) mappings must be set on the application server.

To add a system path to the list, click the plus button. To remove a file system, click the minus button.

Remap Path to This URL

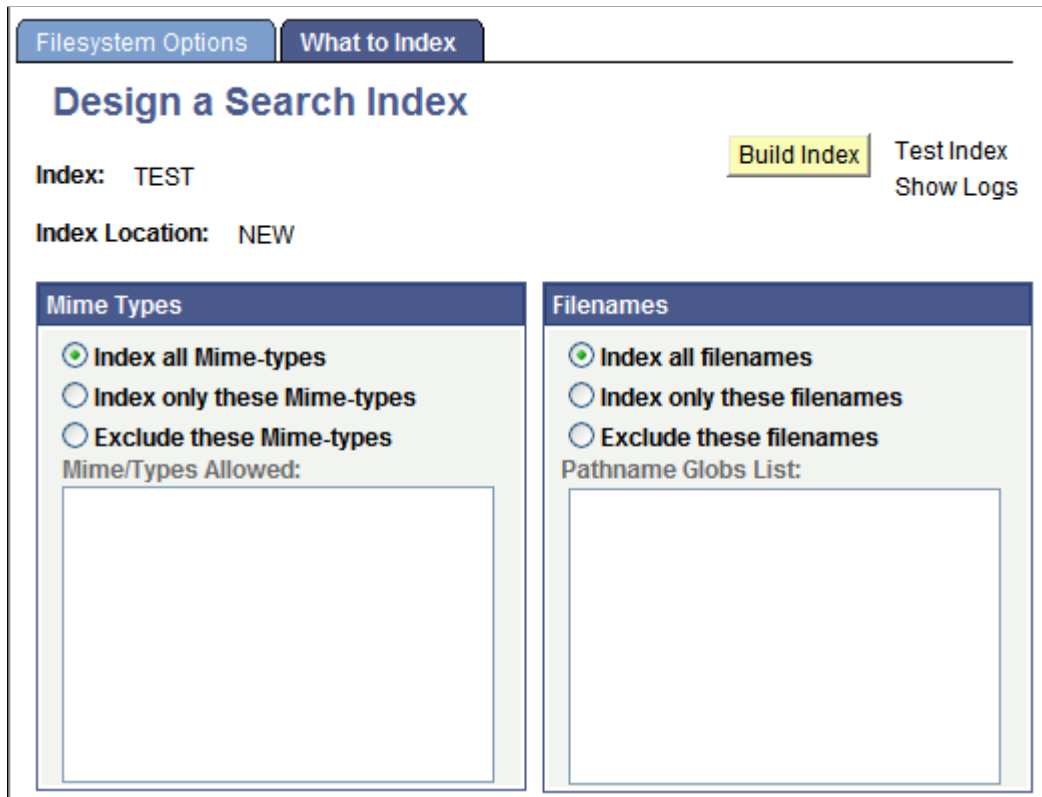
Do not use.

Defining What to Index

Select PeopleTools, Search Engine, Filesystem Indexes, What to Index to access the What to Index page.

Image: What to Index page

This example illustrates the fields and controls on the What to Index page. You can find definitions for the fields and controls later on this page.



MIME Types

Index all Mime-types

Select to index all MIME types on a website.

Index only these Mime-types

Select to index only a certain MIME type, and specify the file type in the MIME/Types Allowed list box. Separate multiple MIME types with a space.

Exclude these Mime-types

Select to exclude a set of MIME types, and specify the MIME types to exclude. Separate multiple MIME types with a space.

MIME/Types Allowed

Add a list of MIME types, separated by spaces, if you selected Index only these Mime-types or Exclude these Mime-types.

Filenames

Index all filenames

Select to index all file types.

Index only these filenames

Select to index only a certain file type, and specify the file type in the Pathname Globs List list box.

Exclude these filenames

Select to exclude a set of file types, such as temporary files, but to index all others. Also specify the file types to exclude.

Pathname Globs List

Add the files that you want to incorporate into your index. Separate the entries with spaces. You can use wildcard characters (*) to denote a string and “?” to denote a single character. For example, the string '*.doc 19??.excel' means select all files that end with the “.doc” suffix and Microsoft Excel files that start with 19, followed by 2 characters.

Building HTTP Spider Verity Indexes

HTTP spider indexes are similar to the indexes that the spider functionality compiles for the file system index. When using the spider index on a website, vspider starts at the home page of the site and then follows each link on that page to the next level of the site. For each page at the next level, vspider follows each link on each page. After following a link, vspider indexes all of the data on the target page.

You can specify as many websites as you want, and you can configure the depth, or number of layers of links, that vspider follows into a website and index.

This section discusses how to:

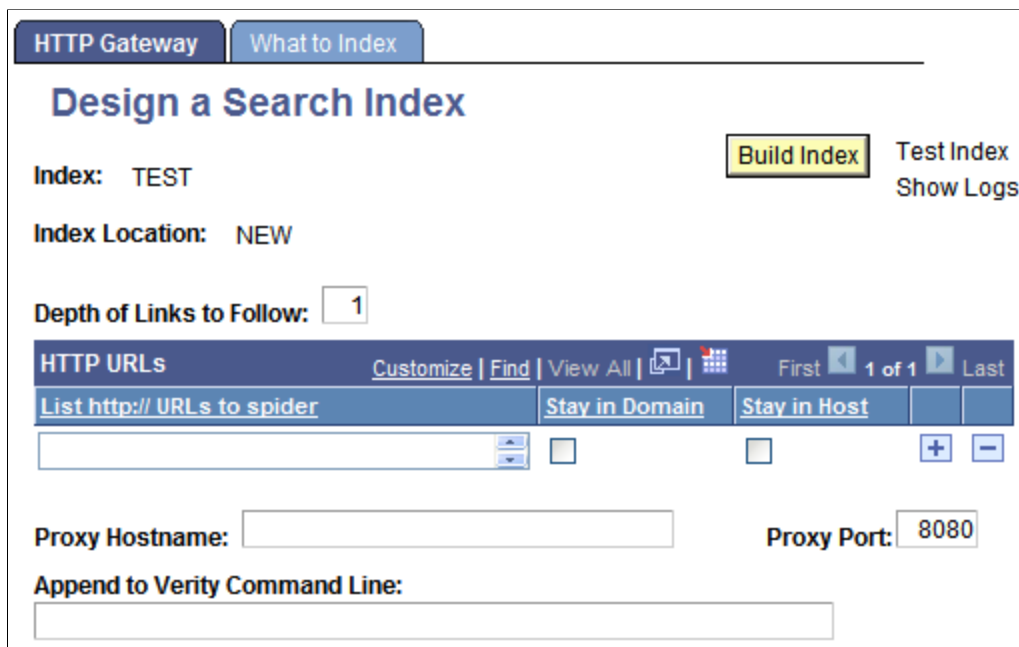
- Define HTTP gateway settings.
- Define what to index.

Defining HTTP Gateway Settings

Select PeopleTools, Search Engine, HTTP Spider Indexes to access the HTTP Gateway page.

Image: HTTP Gateway page

This example illustrates the fields and controls on the HTTP Gateway page. You can find definitions for the fields and controls later on this page.



Depth of Links to Follow

Set the level of detail that you want to index within a certain site. If you enter *1*, vspider starts at the homepage and follows each link on that page and indexes all of the data on the target pages. Then it stops. If you enter *2*, vspider follows the links on the previous pages and indexes one more level into the website.

As you increase the number, the number of links that vspider follows increases geometrically. Do not set this value too high, because it can impact performance negatively. You should not need to set this value higher than 10.

List http://URLs to spider

Click the plus button to add multiple URLs to spider. Click the minus button to remove a URL from the list. If you forget to include the *http://* (scheme) portion of the URL, the system automatically includes it.

URLs should contain only the alphanumeric characters as specified in RFC 1738. Any special character must be encoded. For example, encode a space character as *%20*, and encode a *<* as *%3c*. Additional examples are available.

See [Uniform Resource Locators \(URL\)](#).

Stay in Domain

Select to limit spidering to a single domain. For example, suppose that you are spidering *www.peoplesoft.com* and you

select this option. If a link points to a site outside the PeopleSoft domain (as in yahoo.com), the collection ignores the link.

Stay in Host

Select to further limit spidering within a single server. If you select this option, the collection contains references to content only on the current web server or host. Links to content on other web servers within the domain are ignored. For example, if you are spidering www.peoplesoft.com and you select this option, you can index documents on www.peoplesoft.com, but not on www1.peoplesoft.com.

Proxy Hostname and Proxy Port

Enter a host and port for vspider to use. Enter the same settings that you would use in your web browser if you need a proxy to access the internet.

Defining What to Index

Select PeopleTools, Search Engine, HTTP Spider Indexes, What to Index. The fields on this page are documented in a previous section.

See [Defining What to Index](#).

Administering Verity Search Indexes

After you design and build your search indexes, the Search Administration interface enables you to schedule when and how frequently the indexes must be rebuilt. An important aspect of maintaining the collections involves scheduling PeopleSoft Process Scheduler jobs that, on a regular basis, rebuild the collection completely or incrementally update the index. Search index administration also includes deleting old indexes and building indexes to support additional languages.

This section discusses how to:

- Specify the index location.
- Administer the search index.
- Edit properties.
- Schedule administration.
- Share indexes between application servers and PeopleSoft Process Scheduler.

Specifying the Verity Index Location

By default, search index files are located in:

PS_CFG_HOME/data/search/indexname/db_name/language_code

You can store indexes in different locations, but you need to specify the custom location in the CFG file for an application server, Process Scheduler, or search server. Use the [Search Indexes] section in the

PSAPPSRV.CFG, PSPRCS.CFG, or PSSRCHSRV.CFG files to specify alternate search index locations and multiple locations, if necessary.

Note: This procedure assumes that you've already used the Search Index Designer to define, build, and store the search indexes that you will specify in the CFG file.

Note: You must edit the CFG file manually to include the locations. You do not add search index locations through PSADMIN.

To add a search index location:

1. Open the CFG file.
2. Locate the Search Indexes configuration section.

For example:

```
[Search Indexes]
;=====
; Search index settings
;=====
: Search indexes can be given alternate locations if there is an entry here.
; Entries look like: IndexName=fs location (ie EMPLOYEE=c:\temp)
```

3. Add an entry for each search index location that you want to specify by using the following syntax:

index_name=location

For example, to specify the location for search INDEX_A and INDEX_B, your entries would look similar to the following:

```
[Search Indexes]
;=====
; Search index settings
;=====
: Search indexes can be given alternate locations if there is an entry here.
; Entries look like: IndexName=fs location (ie EMPLOYEE=c:\temp)
INDEX_A=c:\temp
INDEX_B=n:\search
```

Note: Make sure that your entries are not commented out with a semicolon (;) appearing before them.

Note: For the Process Scheduler configuration file, PSPRCS.CFG, include the same location as specified in the application server configuration file.

Note: When specifying the index to be generated in a custom location, the directory structure the system builds within the custom location will be slightly different from that built in the default location. The directory structure within custom index locations will not have a directory for the database name.

4. Save the CFG file.

Administering the Verity Search Index

Select PeopleTools, Search Engine, Administration to access the Search Index Admin page.

Image: Search Index Admin page

This example illustrates the fields and controls on the Search Index Admin page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Search Index Admin' page with the title 'Delete, Modify and Schedule Builds for Indexes'. Below the title is a table titled 'Defined Search Indexes' with the following columns: Index, Index Location, Edit Properties, and Schedule. There is one row with index '1 TEST' and location 'C:\Documents and Settings\admin\psft\pt8.50-901-R1\data\search\TEST\QEDMO'. Below the table is a 'Process Monitor' link and a warning message: 'Deleting the Index Definition also removes the actual collections stored in the filesystem, if any have been built.' A yellow button labeled 'Delete checked indexes' is also visible.

Defined Search Indexes				
	Index	Index Location	Edit Properties	Schedule
<input type="checkbox"/>	1 TEST	C:\Documents and Settings\admin\psft\pt8.50-901-R1\data\search\TEST\QEDMO	Edit Properties	Schedule

[Process Monitor](#)

Deleting the Index Definition also removes the actual collections stored in the filesystem, if any have been built.

[Delete checked indexes](#)

Index	Displays the name of the index so that you can identify specific indexes. To select an index, select the check box to the left of the index name.
Index Location	Displays the current location of the index.
Edit Properties	Click to access the interface for changing the index location and to build indexes to support additional languages.
Schedule	Click to access the interface for scheduling the program that maintains your collection.
Delete checked Indexes	If you have selected indexes to be deleted, click this button to remove them from the system. The deletion process deletes the index definition and the collections that are stored in the file system.

Note: If you attempt to delete a scheduled index, you may see SQL errors on IBM DB2 UDB or Sybase database platforms.

Editing Properties

Select PeopleTools, Search Engine, Administration, Edit Properties.

Index Location	Displays the current location of the index.
Language Code	Select the language for which you want to build an index.

Language to Map	Currently disabled.
Build	After you add the additional indexes, click to create the indexes.

Note: Style files are located in the style subdirectory of the index. To make style changes, apply them to the files in this directory.

Scheduling Administration

Select PeopleTools, Search Engine, Administration, Schedule.

Image: Scheduling search builds

This example illustrates the fields and controls on the Scheduling search builds. You can find definitions for the fields and controls later on this page.



Add a new Recurrence Definition

In PeopleSoft Process Scheduler, you define run recurrence definitions that enable you to schedule jobs to run at regular intervals, such as monthly, weekly, daily, and so on. The more current you keep the collections, the more accurate your search results will be.

Type of Build

Rebuild: Select to drop the existing collection and rebuild a new collection. This applies to all types of collections.

Increment: Use only for the spider indexes. For record-based indexes, only the *Rebuild* option is available.

Run Recurrence Name

Select the appropriate run recurrence definition for the collection maintenance requirements.

Server Name

Specify the PeopleSoft Process Scheduler server on which you want the build program to run. The PeopleSoft Process Scheduler system must be installed and configured before you can schedule the collection build program to run as a job.

Sharing Indexes Between Application Servers and PeopleSoft Process Scheduler

The index files reside on a file system at the home location and must be accessible to all application servers and process schedulers that will manipulate the index. An application server uses the index for searching while the process scheduler invokes an Application Engine program that builds the indexes. Therefore, if you are running a process scheduler on a different machine than the application server, ensure that the index files are accessible to both. You can do this three ways:

- Make a Microsoft Windows shared drive or NFS file system available for the index.
Specify the index location in both the application server and the process scheduler to point to the shared directory.
- Run an instance of the process scheduler on the application server host and schedule only the building of indexes on this process scheduler.
Because the process scheduler and the application servers are running on the same host, they create and read files from the same location.
- Use an external program such as FTP or Secure Copy (SCP) to copy all of the files and directories in the index home location from the process scheduler host (after the index has been built) to the application server host so that they are available for searching.

Modifying the VdkVgwKey Key

To make the VdkVgwKey more readable and easier to parse, use the following XML-like syntax:

```
<field fieldname='MYFIELD' />
<row />
<pairs />
<sql stmt="SELECT 'Y' FROM PS_INSTALLATION" />
```

- Fieldname and the SQL statement support single and double quotes, as well as no quotes at all (in which case only the first word is considered part of the option).
Using double quotes for the SQL statement is recommended.
- The SQL statement must return only one column.
Multiple rows are ignored. Trying to return more than one column results in a collection-build-time error.
- Currently, the only tag style that is supported is <tag/> with the slash (/) at the end.
- The VdkVgwKey can include any amount of literal text interspersed with the tags.
This text is copied into the VdkVgwKey that goes into the BIF file, unmodified.
- Field names are automatically set in uppercase.

Building Verity-Based Registry Search Indexes

This section discusses how to:

- Associate keywords with content references.
- Build search indexes for the portal registry.

Associating Keywords with Content References

Access the Content Ref Administration page (PeopleTools, Portal, Structure and Content Click the Edit link for the appropriate content reference.)

Image: Content Ref Administration page: Content Reference Attributes group box

The following examples illustrates how keywords can be associated with a content reference on the Content Ref Administration page.

The screenshot shows a web form titled "Content Reference Attributes". It contains the following elements:

- Name:** A text input field containing the text "KEYWORD".
- Label:** An empty text input field.
- Attribute value:** A larger text input field containing the text "skill".
- Translate:** A checkbox labeled "Translate" which is currently unchecked.

You can add additional keywords to a content reference by adding portal KEYWORD or KEYWORD_LNG attributes. Specific portal features, such as the Pagelet Search page, make use of these keywords when searching for definitions.

Note: While the KEYWORD attribute is also included in Verity-based search indexes, it is not included in indexes generated by the PeopleSoft Search Framework.

Building Search Indexes for the Portal Registry

Access the Build Search Index page (PeopleTools, Portal, Build Registry Search Index).


Important! The Build Search Index page is used for building Verity-based portal registry indexes only and does not build a portal registry index for the PeopleSoft Search Framework.

You can build a search index from the portal administration pages and schedule the build process using PeopleSoft Process Scheduler. Each time that you build the search index, the index is entirely rebuilt, and any previous entries are overwritten.

Note: Only content references can be indexed.

Image: Building the portal registry search index

The process for building a search index for the portal registry is fairly easy because data in the portal registry is already in a known format and this data is used to build the search index. When you build a search index from the portal administration pages, an Application Engine program calls a C++ program, PORTAL_INDEX, which reads from the portal registry. The C++ program then generates the bulk insert files to create the search index. The following illustration represents this process:



To build a portal registry search index:

1. Select or add a run control ID.

You can use any run control ID; however, you may want to use a portal name to associate the run control ID with the registry for which you are building the search index.

2. Select the All Installed Languages check box to build the index for all installed languages.
3. Select the language code for which you want to build the search index.

The default language code is the base language.

4. Click Run to launch the build process.

PeopleSoft Process Scheduler initiates an Application Engine program named PORTAL_INDEX. Click the Process Monitor link to view the status of the index build. Click Report Monitor to review output.

Related Links

[Understanding Verity Search Indexes](#)

Using PeopleSoft Configuration Manager

Understanding PeopleSoft Configuration Manager

PeopleSoft Configuration Manager simplifies Windows workstation administration by enabling you to adjust PeopleSoft registry settings from one central location. It contains a variety of controls that let you set up Windows workstations running development and administrative applications, such as Application Designer and Data Mover.

You can set up one workstation to reflect the environment at your site, and then export the configuration file, which can be shared among all the workstations at your site. You can also define separate profiles for connecting to different PeopleSoft databases. PeopleSoft configuration parameters are grouped on the Configuration Manager tabs according to the function, feature, or tool that they control.

Note: The changes you make within PeopleSoft Configuration Manager do not take effect until the next time a user signs on to PeopleSoft.

Note: When referencing install locations in your environment using Configuration Manager, keep in mind that some files may be located in PS_APP_HOME or PS_CUST_HOME instead of PS_HOME, depending on how you've implemented your PeopleSoft application.

Note: PeopleSoft supports a number of versions of UNIX and Linux in addition to Microsoft Windows. Throughout this discussion, we make reference to operating system configuration requirements. Where necessary, this topic refers to specific operating systems by name (Solaris, HP/UX, Linux, and so forth). However, for simplicity the word UNIX refers to all UNIX-like operating systems, including Linux.

Related Links

[Setting Up the PeopleTools Development Environment](#)

Common Elements in PeopleSoft Configuration Manager

OK	Saves your settings and exits PeopleSoft Configuration Manager.
Cancel	Closes PeopleSoft Configuration Manager without saving any changes that you have made.
Apply	Saves your changes without exiting.

Starting PeopleSoft Configuration Manager

This section discusses:

- Launching Configuration Manager.
- Managing PeopleTools 8.4x and PeopleTools 8.5x Configuration Manager settings.

Launching Configuration Manager

You can start PeopleSoft Configuration Manager by one of two methods:

- Double-click the Configuration Manager shortcut in your PeopleSoft program group.
- At a command prompt, enter:

```
PS_HOME\bin\client\winx86\pscfg.exe
```

Important! Due to changes in the registry structure beginning with PeopleTools 8.50, if you intend to run PeopleTools 8.50 and later versions of Configuration Manager and previous versions of Configuration Manager (PeopleTools 8.49, 8.48, and so on), on the same development workstation, pre-PeopleTools 8.50 values stored in the registry will be deleted.

Important! Certain PeopleSoft utilities require setting an environment variable, PS_SERVER_CFG, to run properly. However, PeopleSoft Configuration Manager isn't compatible with PS_SERVER_CFG being set. Before you start Configuration Manager, you must ensure that PS_SERVER_CFG is not set. You can make this convenient by using a DOS batch file to unset PS_SERVER_CFG, launch Configuration Manager, then after Configuration Manager exits, reset PS_SERVER_CFG to its previous value.

Managing PeopleTools 8.4x and PeopleTools 8.5x Configuration Manager Settings

PeopleTools 8.4x and PeopleTools 8.5x releases manage registry settings differently. This is due, in part, so that PeopleTools complies with security policies introduced with Microsoft Windows Vista and Microsoft Windows 7.

PeopleTools 8.4x Configuration Manager uses both HKEY_CURRENT_USER and HKEY_LOCAL_MACHINE to store various settings. PeopleTools 8.5x Configuration Manager stores all settings under HKEY_CURRENT_USER. For example, in PeopleTools 8.4x, profile settings are stored in HKEY_LOCAL_MACHINE, whereas, in PeopleTools 8.5x profile settings are stored in HKEY_CURENT_USER.

Due to these changes in the registry structure, the first run of the PeopleTools 8.5x Configuration Manager will delete any PeopleTools 8.4x configuration settings stored in the registry in the HKEY_CURRENT_USER hive. PeopleTools 8.4x and PeopleTools 8.5x registry settings can coexist, but the PeopleTools 8.4x settings must be preserved prior to the first run of the PeopleTools 8.5x Configuration Manager so that they can be restored afterwards.

To preserve PeopleTools 8.4x registry settings:

1. Use the Registry Editor to export and save the contents of HKEY_CURRENT_USER\Software\PeopleSoft\PeopleTools\Release8.40.
2. Run PeopleTools 8.5x Configuration Manager, and modify settings as needed.
3. Import the .reg file containing the PeopleTools 8.4x exported registry contents.

After completing these steps, PeopleTools 8.4x and PeopleTools 8.5x will then use the same registry values with the exception of the Profiles registry settings. PeopleTools 8.4x versions will use the Profiles settings under HKEY_LOCAL_MACHINE. PeopleTools 8.5x versions will use the Profiles settings under HKEY_CURRENT_USER.

Note: If the PeopleTools 8.4x settings are not saved or they are deleted when the PeopleTools 8.5x Configuration Manager runs, you can recreate the PeopleTools 8.4x settings by running the PeopleTools 8.4x Configuration Manager and reentering the required settings.

Specifying Startup Settings

Select the Startup tab.

Use the Startup tab to customize the default values that appear on the signon screen.

Signon Defaults

Database Type

Select the database type to appear as a default on the PeopleSoft Signon dialog box. Select *Application Server* to sign in to an application server instead of a database. To enable users to change their database type selection in the signon dialog box, you must select the Database Type option in the User Can Override group.

Note: When you select *Application Server* from the Database Type drop-down list, the Server Name and Database Name fields are disabled. The system obtains these values from the application server.

Application Server Name

If you selected *Application Server* from the Database Type drop-down list, specify the application server's name in this field. You must have already configured your application server and registered it on the Profile tab.

Server Name

Enter the name of the default database server. This parameter is only enabled for Informix, Sybase, and Microsoft SQL Server, and refers to the instance to which the user connects.

For Informix, enter the server name in lowercase.

Database Name

Enter a default database name. You can choose any valid PeopleSoft database name. As with the database type, you must select the appropriate option in the User Can Override group.

to enable users to override the default database name selection when they sign in.

User ID

Specify the default user ID to sign in to PeopleSoft.

You can use the user ID in conjunction with a PSUSER module containing a user-defined sign-in process. The PSUSER code, if present, can evaluate and modify the user ID value before you attempt to sign in to the selected database.

Connect ID and Connect Password

PeopleSoft uses the connect ID for the initial connection to the database. Use the Connect Password field to define a default connect ID password.

Note: The connect ID edit box must contain a value, or the user can't sign in to the system in a two-tier environment.

See "PeopleSoft Authorization IDs" (PeopleTools 8.53: Security Administration).

See *PeopleTools Installation Guide for Your Database Platform*.

Numeric Keypad - Enter Key Tabs to Next Field

In Microsoft Windows applications, pressing the Enter key in a dialog box selects the default action button. For example, in the PeopleSoft Signon dialog box, pressing Enter is the same as clicking the OK button. Selecting the Numeric keypad check box overrides this default behavior for the Enter key on the numeric keypad; instead of selecting the action button, pressing the Enter key moves the cursor to the next field in the dialog box.

Note: This check box affects the Enter key on the numeric keypad, but not the Enter key on the main keyboard.

User Can Override

Some PeopleSoft sites use multiple database types and names. Using the check boxes in the User Can Override group box, you can enable users to enter a database type, database name, or user ID other than the default provided at sign-in. In most cases, you use these controls to prevent users from attempting to sign in onto any database other than the default.

Database Type

When selected, users can choose a database other than the default. Selecting this check box selects the Database Name and User ID options automatically. You cannot clear Database Name or User ID without first clearing Database Type. When configuring a workstation to connect in both two-tier and three-tier, you must select this box. The user needs to specify a two-tier or three-tier connection from the PeopleSoft Signon dialog box.

Database Name

When selected, the User ID check box is automatically selected, although you can clear it. To clear Database Name, you must clear the Database Type check box.

User ID Select to enable users to users override only the user ID submitted at when they sign in. You cannot clear User ID if Database Type is selected.

Cache Files Directory

Enter a parent directory to hold your client cache file directories. The directory you specify must be writable for a user with Standard User privileges. For example, %USERPROFILE%\psft\pt\client\cache.

Cache files store database object information locally and are automatically downloaded the first time you open a PeopleSoft database object. They are also downloaded automatically if the master copy of the object on the database server has changed. For each PeopleSoft database to which you connect, the system creates a single child cache files directory to store the cache files for that database.

Clicking Purge Cache Directories brings up a dialog box with your existing cache file directories.

You can select a single directory and click Delete, or you can click Delete All to remove all directories. If a cache file directory is missing (after you delete it), the system automatically rebuilds it the next time that cache files are downloaded. After you delete the appropriate cache directory, click Close to return to the Startup tab.

Specifying Display Settings

Select the Display tab.

Use the Display tab to configure the appearance of the PeopleSoft graphical user interface. For instance, you can adjust page width and height to fit in with the other elements on your desktop.

Note: These settings affect the interface display for PeopleTools Windows applications used for development and administration, such as Application Designer and Data Mover. These settings have no effect on PeopleSoft applications displayed to end users on browsers.

Language

In the Language drop-down list box, specify which language you want to display. The default setting is US English.

Note: You select from the languages that PeopleSoft delivers. Although you can implement applications to appear in other languages, you cannot switch to custom languages using PeopleSoft Configuration Manager. Switch to these languages by manually changing the registry setting.

Page Display

You can adjust window display size or the height and width.

Display Size, Width, and Height Specify display size in pixels. This setting affects the default size of the PeopleSoft window as displayed in the corresponding Width and Height fields. Select from:

- *640 X 480*: The default window size is 640 pixels by 448 pixels.

- *800 X 600*: The default window size is 800 pixels by 576 pixels.
- *1024 X 768*: The default window size is 1024 pixels by 744 pixels.
- *Custom*: You can manually set the default window size by specifying width and height values.

Note: Changing these parameters does not affect open windows. If either value is either blank or zero, the values are reset to 640 by 480 pixel resolution.

Page Sizing

Use this field to specify how interfaces that were designed for a different-size window should be displayed. Select from:

- *CLIP*: Interface controls are always displayed in their normal size. If a page is too large for the window, the page information is clipped along the right and bottom edges of the window. Use scroll bars to view the remainder of the page.
- *SCALE*: Windows are scaled to fit the allotted space as necessary. For example, if your display size is set to 640 by 480 pixels, and you open a window designed to display in an 800 by 600 pixel window, the page controls are scaled down so that all page information appears. Conversely, if you open a page designed for 640 by 480 pixel resolution in a larger window, the interface controls are scaled to fill the window completely.

Show Page in Navigator

Select to see the navigator tree view and the page view at the same time.

Highlight Popup Menu Fields

Select to highlight fields with associated pop-up menus. The box is clear by default. In most cases, it's a good idea to indicate which fields contain pop-up menus. Pop-up menus are indicated by a black rectangle surrounding the perimeter of a page control.

Show Database Name

Select to display the name of current database in the status bar at the bottom of a window. This feature is useful if you are running multiple instances of PeopleTools.

Note: The database name may be abbreviated to fit on the screen.

Font

Use the Font options to configure the way that text appears on the windows.

Click the Font button to bring up a standard font selection pop-up menu.

Business Process Display

Select from:

- *On*: The navigator appears with each menu group that you open.
- *Off*: You must open the navigator manually.
- *First*: The navigator appears on the first instance of PeopleSoft only. Subsequent instances do not display the navigator.

Specifying Crystal Report, Business Interlink, and JDeveloper Settings

Select the Crystal/Bus.Interlink/JDeveloper tab.

Crystal Options

If you have Crystal Reports installed on a workstation, the Crystal executables path is populated automatically. If Crystal Reports is installed on a network drive, use this field to reflect the location of the Crystal Reports executables. For example, you might enter *n:\hr920\bin\client\winx86\crystal*.

Use the Default Crystal Reports field to specify the default location of reports. If this setting does not apply to your site's Crystal Reports implementation, leave this field blank. Depending on your application installation, keep in mind that you may need to specify your PS_APP_HOME location.

When you select Use Trace during execution, Crystal Reports writes the trace statements to a log file that you specify in the Trace File field.

Business Interlink Driver Options

In the Business Interlink Directory box, enter the complete path to the directory that contains the drivers that PeopleSoft Business Interlinks uses to communicate with external systems.

Note: PeopleSoft Business Interlink is a deprecated product. These options currently exist for upgrade compatibility and transition.

JDeveloper Home Directory

Oracle JDeveloper is used when developing transformations using Oracle XSL Mapper. You need to specify the high-level installation directory as well as the appropriate Classpath string. These settings are discussed in detail in the Integration Broker documentation.

Related Links

"Configuring Crystal Reports" (PeopleTools 8.53: Crystal Reports for PeopleSoft)

"Developing Transforms Using Oracle XSL Mapper" (PeopleTools 8.53: PeopleSoft Integration Broker)

Specifying Trace Settings

Select the Trace tab.

Use the Trace tab to select tracing options for various parts of the PeopleTools system, such as SQL statements, PeopleCode, and Application Engine. If you work on tuning your PeopleSoft system and improving online performance, familiarize yourself with this tab. When you update the Trace tab, the new settings take effect the next time you launch PeopleTools.

Note: The Trace tab in PeopleSoft Configuration Manager traces only Microsoft Windows client (two-tier) interactions. Use these settings only when you require tracing on the client.

You can override some of the trace options on this tab from the Trace SQL and Trace PeopleCode pages in PIA.

See [Configuring SQL Trace](#), [Configuring PeopleCode Trace](#).

SQL Informational Trace

Select this check box to trace information messages from the Runstats command in DB2 UDB for z/OS executed as a result of an %UpdateStats meta-SQL command.

PeopleTools Trace File

The default filename for the PeopleTools trace file is DBG1.TMP. The system writes the file to the following directories:

- In Microsoft Windows: %TEMP% directory
- In UNIX: \$PS_HOME/log/dbname

Important! The PeopleTools trace file stores elapsed times for PeopleCode and SQL events to a precision of one microsecond (six decimal places). However, due to limitations of the operating system, Windows precision is actually in milliseconds (three decimal places), so the last three digits in a Windows trace will always be zero. Elapsed times in UNIX are accurate to one microsecond.

To specify a different PeopleTools trace file:

1. Click the button on the right side of the PeopleTools Trace File edit box.
A standard Open dialog box appears.
2. Navigate to and select the new trace file.
3. Click Open.

The PeopleTools Trace File field displays the path and filename.

Related Links

"Setting Options in PeopleSoft Configuration Manager" (PeopleTools 8.53: Application Engine)

Specifying Workflow Settings

Select the Workflow tab.

Use the Workflow tab to specify the options and locations related to the PeopleSoft Workflow implementation at your site.

Maximum Worklist Instances	Enter a number to limit the number of worklist instances or entries that appear when viewing worklists. The default value is <i>250</i> . If you do not want any rows returned, leave the field blank.
SMTP Server	Specify the SMTP settings for email routings.

Related Links

Workflow Technology

Specifying Remote Call/AE Settings

Select the Remote Call/AE tab.

Some PeopleSoft applications use the Tuxedo Remote Call feature, which invokes data-intensive transactions, such as COBOL processes, on a remote server.

Timeout	Enter the amount of time after which Remote Call terminates the child COBOL process. The default is <i>50</i> seconds.
Redirect Output	Select to specify whether the standard out or standard error of the child COBOL process is directed to a file. This check box is clear by default.
Support COBOL Animation	Select to save the COBOL input file so that you can reuse it with COBOL animator. This check box is clear by default.
Normal, Minimized, andHidden	Specify how the window state of the child COBOL process appears on the desktop. <ul style="list-style-type: none"> • Select Normal to have the window state appear like a DOS window on the desktop. • Select Minimized to have the window state appear as an icon on the task bar. • Select Hidden to have the window state run unseen in the background.
Disable DB Stats (disable database statistics)	Select to turn off the %UpdateStats meta SQL construct. This setting applies to Application Engine programs.

See "Meta-SQL" (PeopleTools 8.53: Application Engine).

Configuring Developer Workstations

Select the Client Setup tab.

As part of the PeopleSoft installation process, you need to configure developer workstations (also called the PeopleTools development environment) to run successfully with your PeopleSoft system. You use developer workstations for development and administrative tasks that require the use of a Windows workstation. Such tasks include, updating record definitions with Application Designer and running scripts using Data Mover. Development workstations can access the PeopleSoft system using both two-tier and three-tier connections.

Use the Client Setup tab to configure developer workstations and invoke the Client Setup process. Although this tab is specifically for developer settings, all of the PeopleSoft Configuration Manager settings may affect developers, especially the Startup tab and the Process Scheduler tab for the default profile.

See *PeopleTools Hardware and Software Requirements*

Shortcut Links

Here are the various shortcut links:

Application Designer	Adds a shortcut for the main PeopleTools development environment.
Configuration Manager	Adds a shortcut for PeopleSoft Configuration Manager, which enables you to edit registry settings relevant to PeopleSoft.
Data Mover	Adds a shortcut to launch PeopleSoft Data Mover.
Uninstall Workstation	Adds a shortcut for Uninstall Workstation, which uninstalls the most recent client setup.
PeopleTools RPT Converter	Adds a shortcut to a standalone program that converts RPT files from the format PeopleSoft used in previous releases to the currently supported Crystal format. You only need to run this program if you are upgrading from previous versions of PeopleTools.
nVISION	Adds a menu item for PS/nVision to the PeopleSoft 8 menu group in the Microsoft Windows Start menu.

Note: Back up RPT files before you run the converter program, which significantly alters them.

Install Workstation

Select the Install Workstation check box to run the Client Setup process. Only select the check box after specifying all the appropriate selections on all PeopleSoft Configuration Manager tabs. If you do not select this box, the Client Setup process will not run.

After you select this check box, click either OK or Apply.

Enable Dirty-Read for 2Tier Query

(Applies only to DB2 systems.) Select this option if you need to run "dirty read" queries, through PeopleSoft Query during a two-tier connection.

See PeopleSoft Query.

Related Links

[Setting Up the PeopleTools Development Environment](#)

Importing and Exporting Environment Settings

Select the Import/Export tab.

Use this tab to export, or save to file, the specified environment settings, and to import previously exported settings. This feature is useful when you plan to configure multiple workstations with similar settings.

Export to a File

Click to write current configuration settings to a file. A Save dialog box appears. Note the file name that you give the configuration file.

Note: Click Apply before you export a file. This ensures that the exported configuration file reflects the current settings.

Import from a File

Click to import previously saved configurations on another workstation. Importing a configuration file overrides all the current environment settings on the machine that you import to.

When you click this button, an Open dialog box appears. Navigate to the directory containing the appropriate configuration file, select the file, and click Open.

Warning! In addition to overwriting environment settings, this function also overwrites all existing settings made in Application Designer.

Configuring User Profiles

This section discusses how to:

- Define a profile.
- Specify databases and application servers.
- Configure process scheduler.
- Configure nVision.
- Specify common settings.

Note: The term "user profiles" is used here to refer to user configurations for a workstation. These profiles are not to be confused with PeopleSoft security user profiles.

Defining a Profile

Select the Profile tab.

Use this tab to define one or more user profiles, each of which specifies connection parameters and file location information for a PeopleSoft installation.

Many PeopleSoft installations include multiple databases. For example, there may be one database for tracking financial information, such as expense reports, and another database for human resources processes, such as benefits enrollment. Each of these databases has its own set of supporting files, SQR reports, COBOL processes, and so on. PeopleTools locates these files by referring to the Microsoft Windows registry. By defining multiple profiles, you can tell PeopleTools to use different directory paths depending on the database.

When you first open PeopleSoft Configuration Manager, the Profile tab displays a single profile named Default. To set the parameters for this profile, make sure that it's selected, and click the Edit button. The Edit Profile dialog box appears.

Each workstation must have a default profile, which is used when the user signs in to a database or application server that isn't listed in any profile. If the workstation requires only one set of profile settings, you can use the default profile. You can also set up multiple PeopleSoft Configuration Manager profiles. The profiles are set for Microsoft Windows workstations and are specific to each user and not shared by all workstation users.

Note: You can use profiles to easily switch between applications.

Specifying Databases and Application Servers

From the Edit Profile dialog box, select the Database/Application Servers tab.

Use this tab to specify the configured databases and application servers associated with this profile. When a user enters one of these databases or application servers in the PeopleSoft Signon dialog box, PeopleTools uses the registry settings associated with this profile.

Note: You can assign multiple databases and application servers to a single profile. However, each database and application server must be assigned to only one profile. If you try to add a database to a second profile, PeopleSoft Configuration Manager asks you if you want to remove it from the previous profile and add it to the current one.

Note: Before you enter a database or application server on this tab, you should have already installed and configured it as documented in the PeopleSoft installation documentation for your database platform.

Application Server Name

Enter a name for an application server that you have configured. This name will appear in the drop-down list box on the PeopleSoft Signon dialog box. Choose a name that's intuitive for your site.

Note: Application server names cannot exceed 24 characters.

Machine Name or IP Address

Enter the IP address or the resolvable server name of the application server you specified in the Application Server Name field. You specified the IP address in the [Workstation Listener] section of your PSAPPSRV.CFG file when you installed your PeopleSoft application server. For example, you could enter *207.135.65.20* or *rsp-hp32*.

Port Number

Enter the port number for the application server that you specified in the Application Server Name field. You specified the port number when you installed and configured the application server using PSADMIN. A port number is an arbitrary number between 0 and 9999 that is determined by site specifications.

TUXEDO Connect String

Use this field to support dynamic load balancing. You can specify a free-form connect string that allows a client to connect to another application server in case another is either down or being used to full capacity.

Note: The Tuxedo connect string cannot exceed 1000 characters.

When configuring load balancing, you might choose from the following approaches:

- Round robin load balancing.

With this approach, you specify multiple application servers, and each client picks a server randomly. This approach assumes that application server will receive an equal number of connections. To specify round robin load balancing, use the following syntax for the connect string:

(//IP address 1:port 1//IP address 2:port 2//IP address n:port n)

You can specify the IP address using either dotted notation or by using the server's DNS name. Either way, the slashes (//) preceding the IP address are required.

If the selected application server is unavailable, your connection attempt fails, and the system does not try to connect you to the other application servers defined within the parentheses.

Spaces are not allowed in any part of the connection string. The system automatically removes embedded spaces before storing the value in the registry.

- Round robin with failover.

With this approach, you define a failover connection string. Use the following syntax:

(//IP address 1:port 1//IP address 2:port 2),(//IP address 3:port 3)

If the application server selected from the first group of parentheses (IP addresses 1 and 2) is unavailable, the system automatically attempts to connect to an application server defined in the second group (IP address 3). If that application server fails, the system attempts to connect to the next group to the right, sequentially.

If multiple application servers are defined within any group, the system round-robins between them. If the selected application server fails, the system attempts to connect to the next application server to the right, if any. The following list shows three examples of connect strings that use this approach:

- (//sp-ibm01:8000//sp-ibm02:8000),(//sp-nt01:8000)
- (//208.136.78.88:8000//208.136.78.88:8050//208.136.78.88:8080)
- (//sp-sun01:8000),(//sp-sun02:8000),(//sp-sun03:8000)

Domain Connection Password

Specify for the password for establishing a connection to the domain. The domain configuration contains the DomainConnectionPwd parameter, which provides an additional layer of security for the domain by preventing connections from any unauthorized clients. Three-tier workstation connections must submit the same password as the specified DomainConnectionPwd value.

For a three-tier Windows workstation connection, you enter the password in the Domain Connection Password field.

Set and Delete Buttons

When you click Set, your application server information is displayed in the grid at the top of the dialog box. You can enter a new application server name and set up a different server if you like.

Note: The settings in the grid are not saved until you click Apply or OK. If you click Cancel without first clicking Apply or OK, you lose all the information in the grid.

To remove an application server configuration, select its application server name in the grid and click Delete.

Configuring Process Scheduler

Access the Process Scheduler tab.

Use this tab to specify the directories that are associated with PeopleSoft Process Scheduler jobs, such as SQR and COBOL directories.

General

PeopleSoft Home Directory

Enter your high-level PeopleSoft directory, such as *N:\HR920*, or *N:\PT853*. This value is referred to as your PS_HOME.

PeopleSoft Apps Home Directory

If you have installed your PeopleSoft application into a directory other than PS_HOME, specify that location here, such as *N:\HR920*. This value is referred to as your PS_APP_HOME.

Note: If the PS_APP_HOME value is not set and the PS_HOME value is set, then the system assumes that PS_HOME is the location of your PS_APP_HOME. That is, the system assumes you have installed your PeopleSoft application into the same directory as PeopleTools in PS_HOME. Changes need to be added and saved to the Process Scheduler tab for the system to implement this.

PeopleSoft Customizations Home Directory

If you have created custom files or modified delivered files, you can store them in PS_CUST_HOME to simplify system resource sharing as well as upgrades and updates.

See [Working with PS_CUST_HOME](#) for more information.

Crystal Reports

Enter the file path to \CRWRTPATH, where Crystal Reports sends your reports.

Output Directory

(Optional) Enter the directory used with the Output Destination field when scheduling a PeopleSoft Process Scheduler request.

Log Directory

Enter the directory for SQR, COBOL, and PeopleSoft Process Scheduler log files.

Temp Directory

Enter the path to your temporary directory, for example, C:\TEMP. This directory stores log files and other output files.

Database Drivers

Enter the path to the directory where your database drivers reside.

Note: The PeopleTools Development Environment running on the Windows workstation (Application Designer, Data Mover, and so on) currently requires 32-bit connectivity software, while the connectivity software for the application server is 64-bit. Make sure to point to the correct connectivity software for your Windows workstations. See *PeopleTools 8.53 Installation* for your database platform, for more information.

Word Executables Directory

Enter the directory containing Microsoft Word executable.

Redirect Output

Select to redirect on-screen COBOL Display statements to a log file. (If this check box is clear, you see the onscreen messages only.) Sending the messages to a log file is useful for debugging purposes. The log file is created in the %TEMP%\PS_HOME\DBNAME directory. In addition to the output generated by COBOL Display statements, the log file contains errors generated by the COBOL runtime system.

Note: To use the Application Engine debug feature, clear Redirect Output.

Application Engine

Debug

Select to enable the Application Engine command-line debugger.

Warning! Select the Debug check box only when you are testing and troubleshooting client-side processes. If you select Debug and submit a process request to the server, the process hangs, waiting for a user command.

Disable Restart

Select to disable the Application Engine restart feature, which lets you restart an abnormally terminated Application Engine program. When selected, Application Engine programs start from the beginning. This option is useful during debugging. Do not select it in a production environment.

SQR

SQR Executables

Enter the path to the directory where SQR executables reside.

SQR Flags

Enter the SQR parameters that PeopleSoft Process Scheduler should pass on the command line to the SQR executables. The following SQR flags are required for launching SQR reports:

- -i: Specifies the path to SQC files.
- -m: Specifies the path to the ALLMAXES.MAX file.
- -f: Specifies the output path.
- -o: Directs log messages to the specified file.
- -ZIF: Sets full path to the and name of the SQR initialization file, SQR.INI.

PSSQR Path

Enter the directory paths that the SQR executable should use to locate SQR reports. For example, *C:\PT8.53\SQR;C:\Apps\sqr*.

COBOL

COBOL Executables

Enter the path to \CBLBIN, where COBOL executables reside.

Note: When specifying COBOL locations, consider the location of your PS_APP_HOME.

Configuring nVision

Access the nVision tab.

Use this tab to specify where PS/nVision should look for files and how it should operate. PeopleSoft Query Link, the feature that enables you to send PeopleSoft Query output to a spreadsheet, also uses these settings.

Space between Query Columns

This parameter sets the number of blank Microsoft Excel characters that PeopleSoft Query Link places between query output columns. To eliminate column spacing, set Space between Query Columns to zero.

Directory Paths

Specify the locations of directories associated with PS/nVision jobs.

Note: When specifying nVision directory locations, consider the location of your PS_APP_HOME.

Customization Macros	Enter the directory containing macros for PS/nVision and PeopleSoft Query Link. It is usually <i>PS_HOME\excel</i> .
Report Layouts	Enter the location of PS/nVision layout fields.
Drilldown Layouts	Enter the location of PS/nVision drilldown files, for example, <i>c:\user\nvision\layout\drilldn</i> .
Report Instance	Enter the directory in which PS/nVision places report instances; for example, <i>c:\user\nvision\instance</i> .
Query Templates	Enter the directory to look for the QUERY.XLT file, which defines the Microsoft Excel styles used to format output. The default is <i>PS_HOME\excel</i> .
Style Sheets	Enter the directory where the NVSUSER style wizard locates nPlosion style sheets.
Trace Level	<p>Indicate whether you want PS/nVision to generate independent trace log files of two-tier activity, and at what level, for each nVision process. Select one of the following values:</p> <ul style="list-style-type: none"> • <i>0</i>: Disable tracing. This is the default value. • <i>1</i>: Generate basic high level information. <p>This setting can be used to check whether nVision has successfully launched and is able to connect to Excel and process the request. Some of the key entries in a level 1 trace log are:</p> <ul style="list-style-type: none"> • Command Line Arguments. • Trace Level. • Excel Pid. • Run Control Name. • Report Id. • Business Unit. • Drill Layout.

- Report Id.
 - Instance Name.
- 2: Generate level 1 tracing plus high level code flow.
 - 3: Generate level 2 tracing plus runtime SQL statements.
 - 4: Generate level 3 tracing plus most function calls and output values.

Use this setting to identify problems that are intermittent and hard to predict.

The trace log files are generated in the c:\temp directory, named with the format psnvs_*process_id*.nvt, for example, psnvs_1024.nvt. You can view these log files in a text editor.

See "Viewing Process Detail Actions" (PeopleTools 8.53: PeopleSoft Process Scheduler).

Note: Extensive tracing will affect PS/nVision performance. Two-tier log files aren't automatically purged by PS/nVision. Users must manually delete them from the temp directory to save disk space.

Related Links

PS/nVision

Specifying Common Settings

Select the Common tab.

Sybase Packet Size

Specify a TCP packet size. The minimum value is 512 and the maximum value is 65538. The default packet size is 512. If you change the packet size, make sure to make the corresponding changes to the Sybase server. See the material on Sybase administration and tuning on the PeopleSoft Customer Connection website, as well as your Sybase documentation.

See Your Sybase reference manuals.

Application Designer Image Conversion

When you upgrade to newer version of PeopleTools, you'll need to convert images to a new format, which may require more storage space. If the images exceed the record size limit of your platform, you can shrink the images to conform to this limit.

Convert and Shrink Images to Platform Limit

Select to convert and shrink images to fit your selected database platform limit, as shown in the Image Size Limit field.

- Convert and Shrink Images to Image Size Limit** If you are upgrading to a different database platform, select this option and specify the correct value in the Image Size Limit field.
- Don't Convert, but Shrink Images to Image Size Limit** Select for images that have already been converted, but need to be converted so they meet the platform size limits.

Data Mover Directories

You can control several PeopleSoft Data Mover environment variables through PeopleSoft Configuration Manager.

Note: When specifying Data Mover directory locations, consider the location of your PS_APP_HOME.

- Input Directory** Enter the directory where PeopleSoft Data Mover should search for its input data files. If no path is specified for SET INPUT in the Data Mover script, Data Mover searches directories for the database file in the following order.
1. Specified output directory.
 2. C:\TEMP.
- Output Directory** Enter the directory where PeopleSoft Data Mover scripts and .DAT files will be created. Data Mover must have write access to the location.
- The default is *PS_CFG_HOME\data*.
- Log Directory** Enter the location of PeopleSoft Data Mover log files. This location must allow Data Mover write access in the case of a read-only PS_HOME configuration.
- The default is C:\Documents and Settings\admin\Local Settings\Temp.

Related Links

"Understanding PeopleSoft Data Mover" (PeopleTools 8.53: Data Management)

Specifying Command Line Options

In addition to its GUI interface, PeopleSoft Configuration Manager offers command line options. Syntax for PeopleSoft Configuration Manager command line options is as follows:

```
pscfg -command
```

For example:

```
pscfg -import:n:\config\hr900.cfg
```

Import File

To import configuration settings from a named file, enter `-import: filename`.

Export File

To export the current configuration settings, enter `-export: filename`.

Run Client Setup

To run the Client Setup process, enter `-setup`.

Note: You must use the `-setup` command in conjunction with the `-import` command if you are setting up a new workstation.

Run Client Setup Without Displaying Messages

To run the Client Setup process without displaying messages or dialog boxes, enter `-quiet`.

Note: Output messages are written to a log file called `%temp%\PSINSTAL.LOG`.

Install MSS DSN

To install MSS DSN, enter `-dsn`.

Note: For Microsoft SQL Server, the server name value is used to automatically create your ODBC data source name.

Uninstall Workstation

To clear the PeopleSoft settings from the registry or uninstall the PeopleSoft workstation, enter `-clean`.

The `-clean` command removes the following items from the workstation:

- PeopleSoft registry settings.
- All cache files from the current `\CACHE` directory.
- Shortcut links.
- PeopleSoft program group.

Make sure that removing all of these items is acceptable before issuing the `-clean` command.

Help

To view PeopleSoft Configuration Manager command-line options online, enter `-help` or a question mark (`?`).

Setting Up the PeopleTools Development Environment

This section provides overviews of the PeopleTools development environment and the client setup process and discusses how to:

- Verify *PS_HOME* access.
- Verify connectivity.
- Verify supporting applications.
- Use the Configuration Manager pages.
- Run the Client Setup process.

Understanding the PeopleTools Development Environment

Most user workstations are equipped with supported web browsers, but with no special PeopleSoft software installed. The traditional Microsoft Windows client is supported for application developer and administrator use. The PeopleTools development environment runs on a supported version of Windows.

This topic describes how to configure these Windows-based clients using PeopleSoft Configuration Manager. As before, such clients can connect to the PeopleSoft database directly using client connectivity software (a two-tier connection), or through a PeopleSoft application server (a three-tier connection).

Understanding the Client Setup Process

Before running the Client Setup process, create all the profiles you need.

The Client Setup process installs a PeopleSoft program group on the workstation.

If the Install Workstation check box on the Client Setup tab is selected, these Client Setup functions are performed when you click **OK** or **Apply** from PeopleSoft Configuration Manager.

See [Configuring Developer Workstations](#).

Note: Any files installed by the Client Setup process on the workstation from the file server use the paths specified in the default profile.

Verifying PS_HOME Access

To use the PeopleTools development environment, each workstation must have access to the file server *PS_HOME* directory (the high-level directory where PeopleSoft executables were installed) and have a drive mapped to the directory. Workstation users must have read access to the *PS_HOME* directory.

Verifying Connectivity

Database connectivity is required on all Microsoft Windows-based clients that make two-tier connections to the database. A two-tier connection is required if any of the following is true:

- You sign in to the Application Designer in two-tier.

- You run PeopleSoft Data Mover scripts.
- You run COBOL and SQR batch processes on the client.

Note: PeopleSoft Windows-based development tools, such as Application Designer, are 32-bit applications in the current release and require 32-bit connectivity software.

Verify Supporting Applications

Supporting applications must be installed on any Microsoft Windows-based client on which batch processes are run locally.

SQR

On Microsoft Windows-based clients, you can install SQR locally, or you can map to a copy installed on the file server. Because SQR does not require local registry settings, you can execute SQR from any Windows-based client once SQR has been installed to a shared directory. Installing SQR locally results in improved performance; over a slow network connection, the improvement is significant.

Crystal Reports

Optionally install Crystal Reports on Microsoft Windows-based two-tier clients. As with SQR, you can install Crystal Reports locally, or you can map to a copy installed on the file server. Because Crystal Reports does not require local registry settings, you can run Crystal Reports from any two-tier client once it has been installed to a shared directory. Installing Crystal Reports locally results in improved performance; over a slow network connection, the improvement is significant.

Crystal Reports requires that you install the PeopleSoft ODBC driver on the workstation where Crystal Reports processes run.

Microsoft Office

Install Microsoft Office on any two-tier client that runs PS/nVision or Microsoft Word batch processes. Microsoft Office must be installed locally, because it requires registry settings.

Using the Configuration Manager

PeopleSoft Configuration Manager enables you to manage registry settings on a Windows workstation that control a variety of interface options, such as signon defaults, display settings, and environment profiles.

Related Links

[Understanding PeopleSoft Configuration Manager](#)

Running the Client Setup Process

To run the Client Setup process:

1. In Configuration Manager, select the Client Setup tab.

2. In the Group Title text box, enter the name of the program group for the icons you want on the client workstation.
3. Select check boxes to create shortcut links for PeopleSoft applications that you want to access from the workstation.

When you run the Client Setup process, it removes existing shortcuts in the PeopleSoft 8 program group and installs shortcuts for the applications that you have selected. If you later want to install or uninstall shortcuts, you can always run Client Setup again.

4. Select the Install Workstation check box.

Client Setup runs when you click Apply or OK in PeopleSoft Configuration Manager. If this check box is not selected, the Client Setup process creates or updates settings in the registry, but it doesn't set up the PeopleSoft 8 program group or install local DLLs.

5. Click Apply to run the Client Setup process and apply other PeopleSoft Configuration Manager settings.
6. To view a list of the files installed and actions taken by the Client Setup process, open the psinstal.log file in your Temp directory.

Related Links

[Configuring Developer Workstations](#)

Chapter 11

Using PeopleTools Utilities

Understanding the PeopleTools Utilities

As you work with the PeopleSoft system, you find that there are some administrative tasks that you only need to perform occasionally. These tasks include such things as maintaining error messages and setting DDL model defaults. The PeopleTools Utilities menu is where you find tools for accomplishing some of these more infrequent tasks.

The documentation of the utilities matches the menu structure of the Utilities interface. For example, the PeopleTools Options utility is under the Administration menu in the Utilities interface; therefore, the documentation for PeopleTools Options is in the Using Administration Utilities section in this topic. Also, in many cases this book refers to other PeopleBooks for the detailed documentation of a utility.

Using the System Information Page

This section provides an overview of the system information page and discusses how to view the system information page.

With the combination of accessing PeopleSoft applications with a browser, single signon between databases, and the PeopleSoft Portal, users and system administrators need a quick tool to provide orientation information and information regarding the current environment. For this reason, PeopleSoft provides the system information page.

Understanding the System Information Page

With single-signon and the portal, it may not be apparent to all end users just exactly what databases or applications they are currently accessing. Viewing environment information can help end users orient themselves.

In most cases, the administrators use the system help page to aid in troubleshooting. If a user has trouble accessing a particular application, the system administrator can instruct the user to provide the system information that appears in this page so that the administrator can immediately identify the current application server, database, software version, operating system, and so on.

Viewing the System Information Page

To view the System Information help page, you press the CTRL+J hotkey while a PeopleSoft page is active. The following example illustrates the type of information that appears.

Image: System Information page

This example illustrates the fields and controls on the System Information page. You can find definitions for the fields and controls later on this page.

Browser	IE/7.0
Operating System	WINXP
Browser Compression	ON (gzip)
Tools Release	8.50-901-R1
Application Release	PeopleTools 8.50.00.000
Service Pack	0
Page	PMN_PRCSLIST
Component	PROCESSMONITOR
Menu	PROCESSMONITOR
User ID	PTDOCGS
Database Name	QEDMO
Database Type	ORACLE
Application Server	//BUFFY:9211
Component Buffer Size (KB)	70
continue	

To return to the previous page, click continue.

The following table briefly describes each item:

Item	Description
Browser	The browser version and type, such as Internet Explorer.
Operating System	The operating system that runs on the computer on which the browser is running. For example, this refers to the operating system of the end user's workstation or the operating system running on a kiosk machine. It does not refer to the operating system that runs on the application server, web server, or database server.

<i>Item</i>	<i>Description</i>
Browser Compression	<p>Indicates if browser compression is enabled in the Compress Responses field on the General page of the current web profile. Values are:</p> <ul style="list-style-type: none"> • ON: The flag is on in the web server configuration and the page is compressed. The compression type is either gzip or zip. • OFF: The page is not compressed because the flag is cleared in the web profile. • OFF (not supported): The page is not compressed because the browser doesn't support compression, however the flag is turned on in the web profile.
Tools Release	The version of PeopleTools that is currently installed at the site. For example, PeopleTools 8.5, 8.51.01, and so on.
Application Release	The version of PeopleSoft applications that are currently installed at the site.
Service Pack	Typically, updates to PeopleSoft applications arrive in the form of a service pack. This item shows the current service pack that is applied to the applications.
Page	The current page that the user is accessing.
Component	The component to which the current page belongs.
Menu	The name of the menu under which the component appears.
User ID	The user ID of the user that is currently accessing PeopleSoft.
Database Name	The name of the database that the user is currently performing a transaction in.
Database Type	The type of the current database, as in Microsoft, Oracle, DB2 UDB, and so on.
Application Server	The domain name server name or Internet Protocol (IP) address and the JSL port number.
Component Buffer Size (KB)	<p>The component buffer size, which reflects the data buffer size, not including metadata,, such as the record definition or component definition. This metric is the same metric also displayed by the Performance Monitor.</p> <hr/> <p>Note: The Performance Monitor does not need to be configured for this value to be populated.</p> <hr/>

You enable and disable the System Information page using the Show Connection Information check box on the Debugging page of the current web profile.

Related Links

"Configuring Trace and Debug Options" (PeopleTools 8.53: Portal Technology)

Using Administration Utilities

This section discusses:

- PeopleTools Options.
- Message Catalog.
- Spell Check System Dictionary.
- Translate Values.
- Load Application Server Cache.
- Tablespace Utilities.
- Tablespace Management.
- DDL Model Defaults.
- Strings Table.
- XML Link Function Registry.
- Merchant Integration Utilities.
- TableSet IDs.
- Record Group.
- TableSet Control.
- Convert Panels to Pages.
- Update Utilities.
- Remote Database Connection.
- URL Maintenance.
- Copy File Attachments.
- Query Monitor.
- Sync ID Utilities.
- Gather Utility.

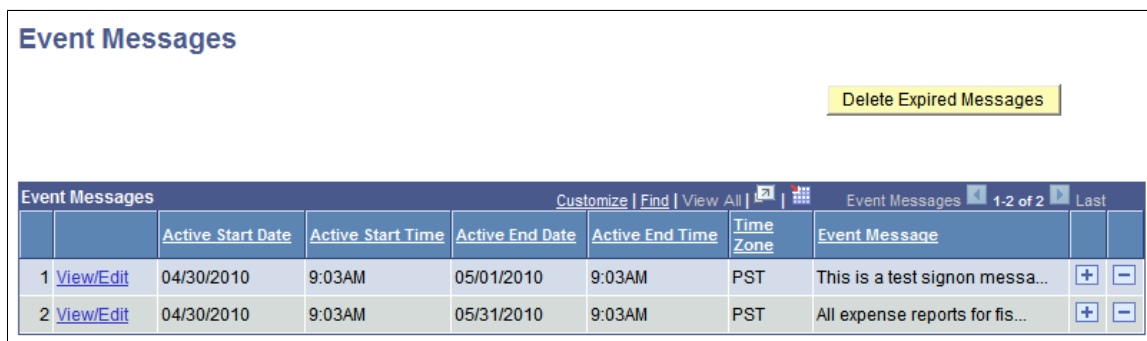
Signon Event Message

Select PeopleTools, Utilities, Administration, SignOn Event Message to access the Event Messages page.

Use this page to define messages to display to end users when they sign on to the PeopleSoft system. As long as the signon message is active, each time a user signs on to PeopleSoft, the event message will appear after the user clicks Sign In on the signon page.

Image: Event Messages page

This example illustrates the fields and controls on the Event Messages page. You can find definitions for the fields and controls later on this page.



View/Edit

Click to view or modify a signon message.

Delete Expired Messages

Click to delete messages that have an End Time and End Date value earlier than current time.

Note: To delete individual messages at any time, use the Delete row (-) button.

Note: Only users with the role of PeopleSoft System Administrator will have access to the pages used to view and define signon event messages.

Adding Signon Event Messages

To add a new signon event message:

1. Select PeopleTools, Utilities, Administration, SignOn Event Message.
2. Add a new row to the Event Messages grid.

Defining Signon Event Messages

Use the Message Definition dialog box to define signon event messages.

Click View/Edit in the Event Messages grid to display the Message Definition page.

Image: Message Definition page

This example illustrates the fields and controls on the Message Definition page. You can find definitions for the fields and controls later on this page.

Message Definition

Event Message

Reminder: All servers will be down for maintenance Saturday, July 24th from 6 AM to Noon PST.

Message Active

Time Zone: PST

Start Time: <input type="text" value="9:14PM"/>	Start Date: <input type="text" value="07/20/2011"/>
End Time: <input type="text" value="9:14PM"/>	End Date: <input type="text" value="07/21/2011"/>

Event Message

Compose the signon message in this edit box. The system will display to the user all text entered into this edit box. When creating messages, consider:

- Message size limit is 32 KB.
- Only *plain* text is supported (not HTML, XML, Rich Text, or text including links, bold, underline, and so on).
- Translation of the message is the responsibility of the system administrator creating the message. The system does not translate the message depending on the language of the user. For example, in Quebec, Canada, you might consider displaying an English and French version of the message.
- The message will be seen by all users signing on to the database through a browser, regardless of role, department, time zone, and so on. As such, care should be taken to ensure the message has broad relevance.

Note: The message will not be displayed to Windows Development Environment users.

Message Active

Specify the duration for which the message will be active by setting a start time and date and an end time and date. The system only displays messages where the current time is after the start time and date and before the end time and date.

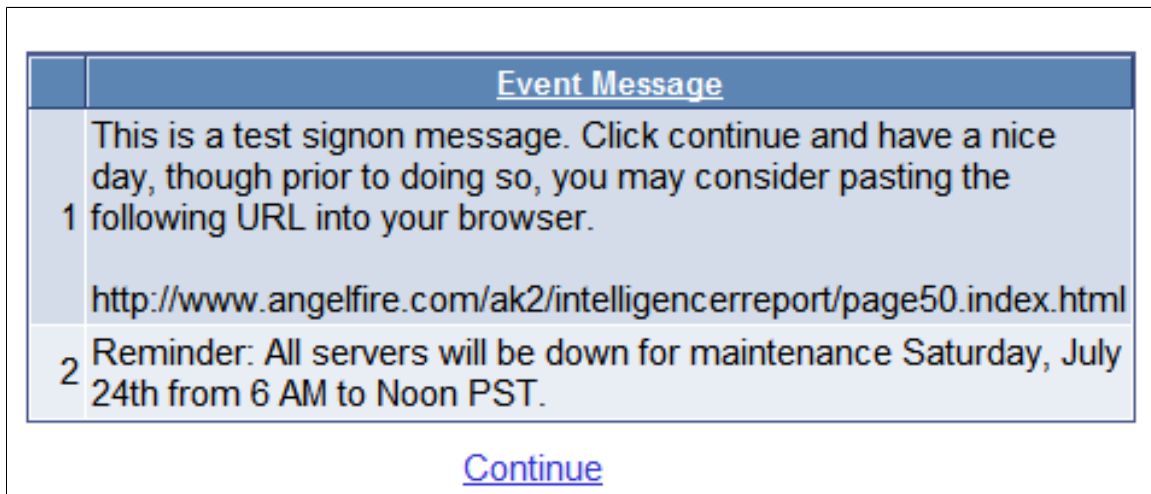
Note: The Time Zone is read-only, and displays the time zone in which the database runs. It does not necessarily reflect the time zone that would be applicable to the text in the message. If there are any specifics related to the time zone that users need to be aware of, that should be communicated in the event message text.

Working With Signon Event Messages

The Event Messages grid can contain multiple signon event message definitions. However, only those messages that are currently active will be displayed to the user.

Image: Event Message display page

This example illustrates the fields and controls on the Event Message display page. You can find definitions for the fields and controls later on this page.



Multiple signon event messages are displayed to the user by start time and date, with the earliest appearing first.

After reading the message, the user clicks Continue, and accesses the application as usual.

Important! The event message display uses the component PT_EVENTDISPLAY. If you intend to use signon event messages, all users of your system need to be able to access the PT_EVENTDISPLAY component. By default, the Permission List PTPT1000 ships providing access to PT_EVENTDISPLAY. If you use custom permission lists, make sure that all users have access to the PT_EVENTDISPLAY component through their permission lists.

PeopleTools Options

Select PeopleTools, Utilities, Administration, PeopleTools Options to access the PeopleTools Options page. Use this page to set a number of options that affect multiple PeopleTools and applications, such as language options and change control settings.

Image: PeopleTools Options (1 of 2)

This example illustrates the fields and controls on the PeopleTools Options (1 of 2). You can find definitions for the fields and controls later on this page.

PeopleTools Options

Environment Long Name Environment Short Name

System Type Undefined Database Right Align Field Labels

Language Settings

Language Code English *Sort Order Option Binary Sorting

Translation Change Last Update

<p>Background Disconnect Interval <input type="text" value="30"/></p> <p><input type="checkbox"/> Multi-Company Organization</p> <p><input checked="" type="checkbox"/> Multi-Currency</p> <p><input checked="" type="checkbox"/> Use Business Unit in nVision</p> <p><input checked="" type="checkbox"/> Use Secure Rep Rqst in nVision</p> <p><input type="checkbox"/> Multiple Jobs Allowed</p> <p><input checked="" type="checkbox"/> Allow DB Optimizer Trace</p> <p><input checked="" type="checkbox"/> Grant Access</p> <p><input checked="" type="checkbox"/> Platform Compatibility Mode</p> <p><input type="checkbox"/> Allow NT batch when CCSID<>37</p> <p><input type="checkbox"/> Save Error is Fatal</p> <p><input type="checkbox"/> Set Focus on Save Button</p> <p>*Case Insensitive Searching On - CaseSensitive Default Off <input type="button" value="v"/></p> <p>Style Sheet Name <input type="text" value="PSSTYLEDEF"/> <input type="button" value="m"/></p> <p>Branding Application Package <input type="text" value="PT_BRANDING"/></p> <p>Branding Application Class <input type="text" value="BrandingBase"/></p>	<p>Temp Table Instances (Total) <input type="checkbox"/></p> <p>Temp Table Instances (Online) <input type="checkbox"/></p> <p>*Maximum Message Size <input type="text" value="10,000,000"/></p> <p>Maximum number of Segments <input type="text" value="10"/></p> <p>Base Time Zone PST <input type="button" value="m"/></p> <p>Last Help Context # Used <input type="text" value="100222"/></p> <p>*Data Field Length Checking Others <input type="button" value="v"/></p> <p>*Maximum Attachment Chunk Size <input type="text" value="28,000"/></p> <p>Upgrade Project Commit Limit <input type="text" value="50"/></p> <p>*Enable Switch User All <input type="button" value="v"/></p> <p>Max rows in search results <input type="text" value="0"/></p> <p>Default rows in search results <input type="text" value="300"/></p> <p>Default rows in SES results <input type="text" value="50"/></p>
--	---

Tree Manager Options

Use Tree Update Reservation

Max Tree Inactivity Period,min

Image: PeopleTools Options (2 of 2)

This example illustrates the fields and controls on the PeopleTools Options (2 of 2). You can find definitions for the fields and controls later on this page.

Environment Long Name and Environment Short Name

Enter a long name and a short name for the current PeopleSoft environment. PeopleSoft software update tools use this information to identify the database when searching for updates. For example, enter *HR Demo Environment* for the long name, and *HR Demo DB* for the short name.

System Type

Select an appropriate system type from the dropdown list, for example, *Demo Database*. This information helps to further identify the current environment for the purpose of searching for and applying software updates.

Right Align Field Labels

By default, the system uses left label alignment when adding new fields to pages in Application Designer. To make all new page field labels use right label alignment (at design time), set the Right Align Field Labels check box on the PeopleTools Options page.

See "Setting Page Field Properties for Controls" (PeopleTools 8.53: PeopleSoft Application Designer Developer's Guide).

Language Settings

Language Code

The base language of an application is the application's primary language, normally the language that is used most commonly throughout the enterprise. A database can have only one base language. All other language translations that are stored in the database are referred to as nonbase languages (or sometimes as foreign languages).

You can't change the Language Code setting on this page. This field is for display purposes only. To change the base language, use the SWAP_BASE_LANGUAGE Data Mover command.

The Language Code field box identifies the database's base language.

Translations Change Last Update

If you select the Translations Change Last Update check box, and you use the PeopleTools translate utilities to translate objects, the system updates the Last Updated information of the translated object to the date/time/userid of the translation. If it's turned off, then the date/time/userid of the object does not change when it's translated.

Note: This only applies when you're using the page-based PeopleTools translation utilities; the Translation Workbench always updates the last updated information.

Sort Order Option

Select the sort order that is appropriate for the site.

See the Global Technology PeopleBook for descriptions of the options.

General Options**Background Disconnect Interval**

The value in seconds that you enter here acts as the default for Security Administrator profiles.

Multi-Company Organization

Turn on Multi-Company Organization if more than one company makes up the organization.

This option affects how Application Processor displays company-related fields in search dialogs and pages. See the HRMS documentation for more details.

Multi-Currency

The Multi-Currency setting is a systemwide switch that enables automatic formatting of currency amount fields that have associated currency control fields. Another function of this setting is to globally display currency control fields. If you turn off this option, automatic formatting based on currency control fields is no longer active and all currency control fields are thus hidden.

When the Multi-Currency setting is on, it also validates user-entered currency data against the currency's defined decimal precision. This validation causes the system to issue an error if a user attempts to enter a decimal precision that is greater than that which is allowed by the currency code definition.

Under most circumstances, leave Multi-Currency selected.

Use Business Unit in nVision

Deselect the Use Business Unit in nVision option if you're using an HRMS database. Otherwise, select it.

Use Secure Rep Rqst in nVision

Select this check box if you want the report request in nVision to be secure. The default setting is selected.

Multiple Jobs Allowed

Selecting Multiple Jobs Allowed enables HRMS systems to support employees holding concurrent jobs with more than one set of enrollments.

This option affects how Application Processor displays employee-record-number-related fields in search dialogs and pages. See the HRMS documentation for more details.

Allow DB Optimizer Trace

Typically, you turn on this trace only during periods in which you are collecting detailed performance metrics. When you are not tuning your performance, the DB Optimizer trace should be turned off.

Grant Access

When adding a new user by using PeopleTools Security, the system automatically grants the new user select-level access to a set of PeopleTools SQL tables used for security. If you are using a SQL security package and do not want PeopleTools Security to perform any SQL grants, turn off Grant Access.

Platform Compatibility Mode

Enables you to add the capability to set a database compatibility mode as an overall database setting, forcing developers to create applications by using all platforms as the least common denominator. This option enables developers, who create applications for multi-platform deployment, to catch platform-specific issues at design time rather than during testing.

Note: This option is used mainly by PeopleSoft development teams that need to develop applications to run on all supported database platforms. To support numerous database platforms, PeopleSoft needs to have a tablespace for each physical table record definition.

If platform compatibility is enabled for a database, the system forces developers to enter a tablespace name when saving a record definition regardless of the current platform. If this option is disabled, you are only prompted for a tablespace name if you are developing on a platform that utilizes tablespaces. This prevents table record definitions being added to the database without a tablespace name.

Allow NT batch when CCSID <>37

Enables you to override non-z/OS COBOL batch restrictions. If the DB2 z/OS database's CCSID is NOT 37, PeopleSoft blocks batch COBOL from running against z/OS Databases on Windows unless you choose this override.

Note: Even if you choose this override, if you use %BINARYSORT() in the COBOL, the system issues an error on Windows. RemoteCall COBOL can run on Windows and UNIX regardless of this option setting, even if CCSID is NOT 37, but the system issues an error.

Save Error is Fatal

Select this option when you have non-repeatable PeopleCode logic in your application's SavePreChange or Workflow. In previous releases, PeopleSoft applications were coded to assume that errors during save are always fatal, but the current PeopleTools release no longer behaves this way. Use this option

to ensure predictable behavior with your application without having to modify your older application code.

This check box is cleared by default. If you get an error during save processing, the transaction continues and you're allowed to attempt to save again. When this option is selected, if you get an error during save processing the transaction is aborted and all changes are lost. This applies to errors that occur between and including the SavePreChange event to the SavePostChange event. It also includes the component processor save processing. It doesn't include errors from the SaveEdit event.

For example, suppose you have some calculations that occur in SavePreChange which are based on the buffers and also modify the buffers. If there's an error during the save and you attempt to save again, the calculations are repeated, but this time based on the buffers that were already modified by the first time the calculations were done. Therefore the second time the calculations are done they will be incorrect, which could lead to incorrect data being saved to the database. In this case you would want to turn on the Save Error is Fatal option, because a fatal error on save is more desirable than incorrect data being put into the database.

Set Focus on Save

If selected, focus is set on the Save button when a user saves a component. If not selected, focus is set on the first control on the page that can assume focus when a user saves a component. By default, the option is not selected.

Note: This setting has a system-wide effect.

Temp Table Instances (Total):

The value that you specify in the Temp Table Instances (Total) edit box controls the total number of physical temporary table instances that Application Designer creates for a temporary table record definition when you perform the Build process.

This value indicates the total number of undedicated temporary table instances. The maximum number of temporary table instances that you can specify is 99.

Temp Table Instances (Online)

Enter the available online instance values. When you invoke a process online, PeopleTools randomly allocates a single temporary table instance number to programX for all of its dedicated temp table needs. The higher the number of online instances that is defined, the less likely it is for two online processes to get the same value.

Maximum Message Size

Use this parameter to set the maximum message size being handled by Integration Broker.

There is a practical limit to how large a message can be. This limit can be determined by a variety of factors, including network speed, available memory, processing load, and so on.

Base Time Zone

Although you can display time data a number of different ways, PeopleSoft databases store all times relative to a system-wide base time zone. You can adjust the display of the time that an end user sees using the Use Local Time Zone (LTZONE) setting in PeopleTools, Personalizations.

This base time zone is the one that the database server uses. In order for PeopleSoft to properly manage time data, the system needs to know which time zone that is. Set the Base Time Zone to the time zone that the database server's clock uses.

Note: After changing this setting, reboot any application servers that are connected to the database. It is critical for the correct operation of the system that this time zone match the time zone in which the database is operating. Any discrepancy in the base time zone as defined in this page and the time zone in which the database system is operating leads to inaccurate time processing.

Last Help Context # Used

This field is no longer used.

Data Field Length Checking

Normally, field length validation is based on the number of characters that are allowed in a field. For example, a field defined as CHAR(10) in Application Designer holds ten characters, regardless of which characters you enter. In a Unicode database, double-byte characters, such as those found in Japanese, are counted the same as single-byte characters, such as those found in the Latin alphabet.

If you create a non-Unicode database, the field length in Application Designer represents the number of bytes that are permitted in the field, not the number of characters. When the non-Unicode database uses a single-byte character set (SBCS), you can only enter single-byte characters, so the number of characters and the number of bytes are the same. However, because double-byte character sets (DBCS) typically allow a mix of single- and double-byte characters, the number of characters that are allowed in a field in a non-Unicode DBCS database varies. This is true for both shifting and non-shifting double-byte character sets.

For example, if a user enters ten Japanese characters into a field that is defined as CHAR(10) in Application Designer, this string needs 20 bytes of storage in a nonshifting double-byte character set and 22 bytes of storage in a shifting double-byte character set. This ten-character input fails insertion into both these databases.

Use the Data Field Length Checking option to ensure field length validation appropriate to the database's character set. Values are DB2 MBCS, MBCS, and Others.

Choose Others if you are using a Unicode-encoded database or a non-Unicode single-byte character set database. This prevents

special field length checking. As discussed above, these types of databases do not require such checking.

Choose DB2 MBCS if you are running a Japanese database on the DB2 UDB for z/OS platform. This enables field length checking based on a shifting DBCS character set.

Choose MBCS if you are running a non-Unicode Japanese database on any other platform. This enables field length checking based on a nonshifting DBCS character set.

The non-Unicode DBCS settings are specifically oriented towards Japanese language installations, as Japanese is the only language that PeopleSoft supports in a non-Unicode DBCS encoding. All other languages requiring double-byte character sets are only supported by PeopleSoft by using Unicode encoded databases.

Maximum Attachment Chunk Size

Controls the size of the file attachment chunks that are stored in the database. This allows you to upload a file attachment that is larger than could be stored in a single row of the database. The system divides the data into chunks to store in different rows of the database table.

The default is 28,000 bytes, and maximum attachment chunk size allowed by PeopleTools is 50 MB. Application Designer data types control how the system manages "long data."

Note: The Maximum Attachment Chunk Size determines the size of the chunks *stored* in the database. This setting does not affect the maximum size of the file chunks that can be *read* from that table. The maximum size of the file chunks that the system reads from the database table is determined by the Maximum Length property defined for that field in Application Designer.

Note: When working with file attachments it is imperative that you understand how the file processing PeopleCode constructs behave, such as the AllowLargeChunks parameter, used with the AddAttachment built-in function.

Note: Sybase has a 32 KB limitation for long character fields.

See "Specifying Long Character Field Attributes" (PeopleTools 8.53: PeopleSoft Application Designer Developer's Guide).

See "Functions by Category" (PeopleTools 8.53: PeopleCode Language Reference).

See "Understanding the File Attachment Functions" (PeopleTools 8.53: PeopleCode Developer's Guide).

Upgrade Project Commit Limit

Sets the limit on how many rows can be modified by an upgrade project before the system issues a COMMIT statement.

Enable Switch User

The Enable Switch User option enables you to limit the users who can change identities in a PeopleSoft system. The feature applies only when accessing PeopleSoft using a browser; it has no effect on two-tier or three-tier connections.

Most sites have no reason for individual users to change their PeopleSoft system identity during a session. For those sites that do require this capability, the number of users who need to switch to another user profile typically is fairly small.

For example, the users who can switch identities are usually limited to the system's GUEST user (if the system has one) and, perhaps, a few system administrators or power users.

Options for Enable Switch User are:

- *All*: Use this value to indicate that all users are permitted to change their identities during their browser session. This is the default value.
- *None*: Use this value to indicate that no user is permitted to change identities during a browser session
- *Some*: Use this value to indicate that some, selected users are permitted to change their identities during a browser session. When you select this value, the Allow Switch User checkbox appears on the General page of the user profile definition in the PeopleTools Security interface. Select Allow Switch User to indicate that the individual user is permitted to change identities within a PeopleSoft session. User profiles that have the Allow Switch User option unselected will be immediately logged out when the system detects a change in the user's identity.

Once you specify which users can switch identities, if an identity change is attempted (switch user) by an unauthorized user, the system:

- Logs that user off the system, immediately.
- Displays the following message on the signout page: "Illegal Identity switch has been detected by the System. Please re-login."
- Writes an entry in the web server logs indicating that an illegal switch user was attempted.

Note: Assume that the user profile to which a user switches is the “destination” user profile, and the user profile from which a user switches is the “source” user profile. If a user is allowed to switch identities during a session, it is allowed by the source user profile, not the destination user profile. For example if you allow UserA to switch to UserB, it is UserA's user profile that must have the Allow Switch User option selected. The setting for UserB's profile is not relevant in determining whether a switch user from UserA to UserB is allowed. This feature does not control the user profile to which a user can switch; it only addresses which users are permitted to switch identities during a browser session.

See "Specifying User Profile Attributes" (PeopleTools 8.53: Security Administration).

Case Insensitive Searching

Enables you to provide case-insensitive searching for the PeopleSoft search records when searching for PeopleSoft definitions. The options are:

- *Off* : Enables case sensitive searching and sets it to "on" *without* displaying the Case Sensitive check box on the search page.
- *On - Case Sensitive Default Off*: (Default) Displays the Case Sensitive check box on the search page unselected, providing the user the option of enabling case-sensitive searching for a particular search by selecting the Case Sensitive check box.
- *On - Case Sensitive Default On* : Displays the Case Sensitive check box on the search page selected, providing the user the option of switching to *case-insensitive* searching for a particular search by *deselecting* the Case Sensitive check box.

Note: This option is not associated with the Verity search technology.

Style Sheet Name

Set the style sheet, which controls the look and feel. You can set the individual style sheets in Application Designer, and these override the general style sheet for the application, which is set here. The default style sheet for PeopleSoft applications is PSSTYLEDEF_SWAN and for PeopleTools it is PSSTYLEDEF.

Branding Application Package

Specifies the application package that contains the branding application classes to generate the portal headers, footers and menu pagelet icons. The default is the standard PeopleTools branding, PT_Branding. For PeopleSoft Interaction Hub, a different branding application package is specified.

Max rows in search results

Sets the absolute maximum number of rows, system-wide, an end user can specify to be returned on a standard search page. The default value is 0, which makes the Default rows in search results value the absolute maximum.

If this value is greater than 0 (enabled), then the following will appear at the top of the search page:

"Maximum number of rows to return (up to <Max rows in search results>):"

End users specify a custom number of rows on the search results page in the edit box appearing after this line.

When the user clicks Search, if the key values supplied retrieve more than the <Max rows in search results> value, then a message will appear above the search results similar to:

"Only the first 300 results of a possible 386 can be displayed."

In this case, 300 is the value of the Default rows in search results edit box, unless the user has changed it. The 386 is the actual number of results the search key values would have returned. The system runs a separate query to report this value.

This has performance implications, so system administrators should consider whether to enable this feature. To prevent end users from specifying any custom number of rows returned by the search, set Max rows in search results to 0. If disabled, the number of rows returned by a search is the value from the Default rows in search results setting, and cannot be modified by the end user.

Note: When setting this value to add flexibility to your end users search experience, consider the performance implications of setting this number too high for your system. Perform adequate testing and adjust this value to suit the requirements of your system.

Default rows in search results

Sets the default number of rows, system-wide, to be returned on a standard search page. The default value is 300.

If an end user needs to increase the number of rows returned on a search page, they can specify a number up to the maximum, specified by the Max rows in search results value.

Default grids to scrollable

Sets all grids displayed to the end user as scrollable grids, by default.

Branding Application Class

The main branding application class that generates header, footer, and menu pagelet icons. The default is the standard PeopleTools branding, BrandingBase. For PeopleSoft Interaction Hub, a different branding application class from a different branding application package is used. It generates

different header, footers, and menu pagelet icons dynamically, based on the user role or security.

Help Options

FI Help URL

This setting applies only to the Windows environment (such as Application Designer) when the user presses F1 or selects Help, PeopleBooks Help while in PeopleTools.

The F1 Help URL can direct users to any location on the web, such as a custom help system or the website for the company's help desk. It can be a fully qualified uniform resource locator (URL), which is passed literally to the browser, or it can contain one or both of these system variables.

`%CONTEXT_ID%` is the object name or context ID of the currently displaying page or dialog box.

`%LANG_CD%` is the three-letter language code for the user's preferred language.

The PeopleBooks context sensitive help system requires that you enter a URL with a specific format, as follows:

```
http://helpwebserver:port/productline
/flsearch.htm?ContextID=%CONTEXT_ID%&LangCD=%LANG_CD
%
```

For example:

```
http://myhelpwebserver:8080/htmldoc/flsearch.htm?
ContextID=%CONTEXT_ID%&LangCD=%LANG_CD%
```

Specify the URL that is needed to link to the correct location in your HTML PeopleBooks. When users click the Help button, the appropriate context-sensitive PeopleSoft documentation should appear. To remove the help link, leave this value blank, and users won't see a Help link on the application page. Construct the URL like this: `http://helpwebserver:port/productline/flsearch.htm? ContextID=%CONTEXT_ID%&LangCD=%LANG_CD%` For example: `http://myhelpwebserver:8080/htmldoc/flsearch.htm?ContextID=%CONTEXT_ID%&LangCD=%LANG_CD%`

If using the hosted PeopleBooks site on `www.oracle.com`, the URL you use is different. For example:

```
http://www.oracle.com/pls/topic/lookup?id=%CONTEXT_
ID%&ctx=pt853
```

The hosted PeopleBooks are available at:

```
http://www.oracle.com/pls/psft/hompage
```

See the product documentation for *PeopleTools 8.53 Installation* for your database platform.

Ctrl-F1 Help URL

This setting only applies to the Windows environment (such as Application Designer).

The Ctrl+F1 URL allows you to provide an alternate location for help. For example, you may set the main F1 Help URL to the PeopleBook and the Ctrl+F1 for the company's help site.

WSRP Display Mode

This option determines how WSRP-enabled content appears for users of remote WSRP portals that consume PeopleSoft WSRP content.

See "Setting WSRP Display Mode" (PeopleTools 8.53: Portal Technology).

Database Encryption Algorithm

(Applies to Oracle databases only.) Enables you to select the encryption algorithm the system uses to encrypt database fields to additional security.

See "Implementing Oracle Transparent Data Encryption" (PeopleTools 8.53: Data Management).

Message Catalog

Select PeopleTools, Utilities, Administration, Message Catalog to access the Message Catalog page.

Image: Message Catalog page

This example illustrates the fields and controls on the Message Catalog page. You can find definitions for the fields and controls later on this page.

The screenshot displays the Message Catalog interface. At the top, the title 'Message Catalog' is shown. Below it, there are input fields for 'Message Set Number' (value: 2), 'Description' (value: PeopleCode), and 'Short Description' (value: PeopleCode). A 'Messages' table is visible, with the first message selected. The selected message details include: 'Last Update Timestamp: 06/04/1999 8:24AM', '*Message Number: 1', '*Severity: Error', '*Message Text: Invalid use of !.', and 'Description: The exclamation symbol (!) can only be used as a token within PeopleCode as part of a not-equal operator (!=). Either make the operator a proper not-equal expression or delete the'.

You add and maintain system messages by using the Message Catalog page. PeopleSoft error messages are stored in the Message Catalog, and organized by message set number. Each message set consists of a category of messages, ranging from PeopleTools Message Bar Items and PeopleCode Runtime Messages to PeopleSoft Payroll and PeopleSoft General Ledger application messages.

Message Set Number	Identifies the message set.
Description	The Message Set Description is a reference that is used on reports and pages for easy identification.
Short Description	The Message Set Short Description is a reference that is used on reports and pages for easy identification.
Message Number	Each message set consists of one or more rows of messages that are identified by a message number.
Severity	<p>You assign each message a severity, which determines how the message appears and how the component processor responds after the user acknowledges message. The severity levels are:</p> <p>Cancel: This severity should be reserved for the most severe of messages, as when a critical error occurs and the process must be aborted or a machine needs to be shut down. To indicate how rarely this severity level is appropriate, of all PeopleTools messages only five or so have a severity level of Cancel. In almost all cases, you use one of the other severity levels.</p> <p>Error: Processing stopped, and data cannot be saved until the error is corrected.</p> <p>Message: This is an informational message and processing continues normally.</p> <p>Warning: User can decide to either stop or continue processing despite the error.</p>
Message Text	In the Message Text edit box, you see the message text. Any reference to the characters %n, as in %1 or %2, is replaced by parameter values that the system provides.
Explanation	The Explanation text provides a more in-depth explanation of why the message is generated and how to fix the problem. This text appears below the Message Text when the message appears.

PeopleTools uses some messages, but the applications use the other messages, which get called by the Error, Warning, Message Box, MsgGet, and MsgGetText built-in PeopleCode functions.

Note: You can create messages and message sets to support new or customized functionality in the system. You can also edit the messages that PeopleSoft delivers if required. In both of these cases, remember that Oracle reserves all message sets up to 20,000 (that is, 1 to 19,999). If you add a message set or edit a message set with a number that is less than 20,000, it may be overwritten in future upgrades.

To add a message set:

1. Select Utilities, Administration, Message Catalog, and on the search page click Add New Value.
2. Enter the value of the new Message Set Number and click OK.
3. Enter a description and short description of the type of messages that this message set contains.

Try to group the messages logically. For instance, create one message set for the new budgeting application and a different one for the customized billing pages.

4. Add messages.
5. Save your work.

To add a message:

1. Open the desired message set.
2. In the Message Catalog page, click the plus sign button to add a new row.

The Message Number value is automatically set to the next unassigned number in the message set.

3. Select a Severity level, enter message text and a detailed explanation.
4. Save your work.

Spell Check System Dictionary

PeopleSoft PeopleTools provides personal and system-level dictionaries. End users and system administrators can add words to the dictionary for use with the spell check feature. Typically, system administrators add words to the system-level dictionary that are used company-wide; end users add additional role-specific terminology to their personal dictionaries.

Select PeopleTools, Utilities, Administration, System Dictionary to access the system-level dictionary.

Select the All Languages page to enter words that are valid across all languages. Select the Language Specific page for those words that are valid to a specific language:

To add words to the system dictionary by language:

1. Select Spell Check System Dictionary, Language Specific.
2. Select the desired language from the Spell Check Language drop-down list box.
3. Select Session to add a word to the current session's spell check dictionary. After saving this word, the language field refreshes to the current spell check language.
4. Enter the word (maximum 40 characters) that is to be added in the Spell Check Word field.
5. Save your changes.

Case Sensitivity for Spell Check

The words that you add to your personal dictionary are case-sensitive and are validated by the following rules:

1. If the added word is all lower case, such as worklist, then the following are considered valid:
 - Exact match, all lower case (worklist).
 - All uppercase (WORKLIST).

- Initial capitals (Worklist), regardless of its position in the sentence. Mixed case (WorkList) is considered incorrect.
2. If the added word is all uppercase, such as CRM (customer relationship management), then only an exact match is valid.
 3. If the added word is in initial capitals, such as California, then only an exact match and all upper case (CALIFORNIA) are considered valid.
 4. If the added word contains an embedded capital letter, such as PeopleSoft, then only an exact match is valid. Therefore, if case is not relevant to the validity of the word, use all lower case.

Table Structure for Word Storage

System and personal words are stored in the database in the PSSCWORDDEFN table with the following fields:

- SCOPRID indicates whether a word is a system word or a user's personal word.
- SCLANG stores the dictionary language for which the word is considered valid. If the system administrator chooses to store the word for all languages, this field is left blank.
- SCWORD stores the actual word, with a maximum length of 40 characters.
- SCNEGWORDFLG is a flag used to determine if a word is negative (incorrect) or not. Values can be 'N' or 'Y'. PeopleSoft does not currently use this feature, so the value should always be set to 'N'.

To load values in bulk into PSSCWORDDEFN:

1. Using the method of your choice (as in a SQL script), issue SQL similar to the following:

```
insert into PSSCWORDDEFN (SCOPRID, SCLANG, SCWORD, SCNEGWORDFLG)
values ('SYSTEM', 'SC00', 'nnn', 'N')
```

Note: For each word you want to add to the library, you need a separate insert command, and the value 'nnn' will be changed in each of those insert statements to be the next value in the list of words you want to add.

2. Add a value (any value) to the Language Specific tab and click Save.

This alerts the runtime system to update the cached version of the PSSCWORDDEFN table.

Note: In the current release, the maximum number of rows in the PSSCWORDDEFN table should not exceed 2,850.

Translate Values

You use the Translate Values interface to maintain the values in the translate table. If it's allowed by site security administrators, power users can now learn to add their own pick lists (translate values) to an application:

Select PeopleTools, Utilities, Administration, Translate Values to access the Maintain Translate Values page.

Image: Maintain Translate Values page

This example illustrates the fields and controls on the Maintain Translate Values page. You can find definitions for the fields and controls later on this page.

Use this page to maintain translate values

Field Name: ANALYSIS_PLATFORM Length: 3

Maintain Translate Values						
Customize Find View All First 1-5 of 5 Last						
General Last Updated						
	*Value	Effective Date	Status	*Long Name	*Short Name	
1	ESS	01/01/1900	Active	Hyperion Essbase	Essbase	+ -
2	IAS	01/01/1900	Inactive	PS/ROLAP	PS/ROLAP	+ -
3	MOS	01/01/1900	Inactive	Microsoft OLAP Services	MstfOLAP	+ -
4	PPL	01/01/1900	Active	Cognos PowerPlay	PowerPlay	+ -
5	STS	01/01/1900	Active	Generic Star Schema	StarSchema	+ -

Value

Enter the value for the translate selection.

Effective Date

Specify a date for the value to become active.

Note: If you are adding a second row for the same translate value, you must enter a unique effective date.

Status

Specify whether the value is active or not.

Long Name

Enter a long description for identification. There is a 30-character limit.

Short Name

Enter a shorter description for identification. There is a 10-character limit.

Related Links

"Using the Translate Table" (PeopleTools 8.53: PeopleSoft Application Designer Developer's Guide)

Load Application Server Cache

The Load Application Server Cache page enables you to run Application Engine programs that loads the *shared* file cache or the database cache for application server domains.

Select PeopleTools, Utilities, Administration, Load Application Server Cache to access the Load Application Server Cache page.

Image: Load Application Server Cache page

This example illustrates the fields and controls on the Load Application Server Cache page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Load Application Server Cache' page. At the top, there is a 'Run Loadcache' tab. Below it, the page title 'Load Application Server Cache' is displayed. The 'Run Control ID' is 'LOADCACHE_TEST'. There are links for 'Report Manager' and 'Process Monitor', and a 'Run' button. An 'Output Directory' field is present, along with a checked 'Database Cache' checkbox. A table titled 'Select Languages to Load' is shown with the following data:

Language Code	Load Language
DUT	<input type="checkbox"/>
ENG	<input checked="" type="checkbox"/>

Below the table are 'Select All' and 'Clear All' buttons.

You need to run a cache loading program described in this section only if you intend to implement shared file caching on the application server or database caching.

To load *non-shared* file cache or database cache with specific cache projects, you use a separate process, which is documented elsewhere in this guide.

See [Configuring an Application Server Domain to Preload Cache](#).

See [ServerCacheMode](#).

Understanding Application Server Caching

Each server process running within a domain has these caching features:

Cache Type	Description
Memory cache	Stores PeopleTools definition metadata to improve system performance. Memory cache is always enabled for each server process in a domain.
File cache	Stores PeopleTools definition metadata, such as the metadata for pages, locally on the application server. File caching is controlled by the EnableServerCaching parameter in PSAPPSRV.CFG.

Cache Type	Description
Database cache	Stores PeopleTools definition metadata, such as the metadata for pages, in database tables. Database caching is controlled by the EnableDBCACHE parameter in PSAPPSRV.CFG. Any domain that has EnableDBCACHE enabled and connects to the database uses the cache information stored in the database.

With file caching, you have the option of implementing shared or non-shared file caching.

File Cache Option	Description
Shared	With shared file caching, all of the server processes within a single domain share the same file cache location.
Non-shared	With non-shared file caching, each server processes within a single domain maintains its own file cache. For example, with non-shared caching configured, each PSAPPSRV server process uses its own file cache.

Oracle provides cache loading programs designed specifically to load cache prior to users accessing the system, which you run from the Load Application Server Cache page. The cache loading programs retrieve all of the PeopleTools definition metadata from the database and builds either file cache that you can deploy locally on an application server or it populates the database cache tables, depending on what type of cache you are implementing.

Running a cache loading program is the equivalent of having a user access every page in the system once so that all the metadata associated with each page is retrieved from the database and stored in the local cache. With all the metadata already stored in cache, the end users experience more predictable performance.

With cache already loaded, end users don't have to wait for the system to retrieve a definition from the database and cache it, each time a definition is accessed for the first time. The system retrieves all of the required objects from the cache as opposed to querying the database for each request. This provides significant performance benefits for first-time and large transactions.

If you elect to run a cache loading program, consider the following items:

- You need to run one of the cache loading programs at least once. Once the cache is loaded, the cache loading programs download only new or changed metadata to the cache.

Note: Items that are outdated in the cache are marked invalid but are not deleted. This includes design time changes, upgrades, patches, and so on. However, the system uses only the most current cache definitions.

- The first time that you run a cache loading program, it can take a significant amount of time to complete, for example 2 to 30 hours. The duration of the program run depends on a number of factors, including the number of active languages selected, the size of the database, and the power of the server machine. Subsequent program runs complete in less time if there is already valid cache data in the target cache directory or the database cache tables. The programs are designed to update only the changed objects.

- You can copy the output of a cache loading program to other machines (for shared file caching).

Note: Shared file cache output is not portable to different operating systems. For instance, if you generate the cached metadata on a Windows server, you can't copy the cache files to a UNIX server.

Warning! If for any reason you update PSSTATUS.LASTREFRESHDTM, the system marks all items in the shared cache as invalid, and you will need to rerun one of the cache loading programs again.

Oracle provides these options for loading your cache:

Preloading Cache Option	Process Scheduler Process Type	Description
Serial cache loading	Process	Serial cache loading refers to running an Application Engine program named LOADCACHE. The LOADCACHE program, adds each applicable cache object to the cache, sequentially.
Parallel cache loading	PSJob	<p>Parallel cache loading is intended to reduce the time required to build the cache. Parallel cache loading splits the work of the LOADCACHE program into two separate Application Engine programs that run simultaneously within the PSJob PLCACHE.</p> <p>The PLCACHE job runs a basic setup program, and then runs these programs:</p> <ul style="list-style-type: none"> LCACHE_INDEP: Loads all the cacheable objects that are independent of other cached objects. LCACHE_DEP: Loads all the cacheable objects that are dependent on other cached objects. Objects that are dependent on other cached objects are those objects that, when loaded into cache, cause instances of other object to be loaded as well.

Loading Shared File Cache or Database Cache

To load shared file cache or database cache:

- Make sure that the database that the application server runs against produces a clean SYSAUDIT report.

If SYSAUDIT is not clean, the cache loading programs can fail.

- Access the Load Application Server Cache page, entering an appropriate Run Control ID.

Select PeopleTools, Utilities, Administration, Load Application Server Cache to access the Load Application Server Cache page.

3. (For shared file cache) In the Output Directory field, enter a valid location.

Note: In almost all situations, the system ignores any custom value entered in the Output Directory and generates output into the default cache-loading output directory: *PS_CFG_HOME\appserv\procs\ProcessScheduler_domain\CACHE\CACHE\stage\stage*. However, if a cache loading program runs through a PSNT Process Scheduler server running on a remote drive, *and* the application server cache directory and Process Scheduler cache directory are located on the local drive, then the value entered in the Output Directory field supersedes the default output directory.

4. (For database caching) Select the Database Cache check box.

This disables the output directory field. With database caching enabled, the cache loading program inserts the cache directly into the database cache table, PSOBJCACHE.

5. Select the languages to load.

If not all possible languages are required to be loaded for a domain, you can improve performance of the program run by selecting only those languages that are required.

6. Click Run.

The Process Scheduler Request page appears.

7. From the Server Name dropdown list, select the name of the server to run the cache-loading process.

8. From the Process List, select *Serial LoadCache* to run the single LOADCACHE Application Engine program, or select *Parallel LoadCache* to run the PLCACHE PSJob to run multiple Application Engine jobs in parallel.

9. (For Parallel LoadCache) Ensure that the Process Scheduler server definition is set to run the LOADCACHE process category and that the Max Concurrent value is set appropriately.

Parallel LoadCache is assigned to the LOADCACHE process category. It will run only on servers that are configured to run the LOADCACHE process category.

For the LOADCACHE process category, set the Max Concurrent value greater than 0. The recommended value is 2.

If you modify the server definition, restart the server for your changes to take effect.

Note: When running the cache loading programs, only one server at a time should have the MaxConcurrent value set greater than 0. This guarantees that the processes running within the Parallel LoadCache job will run on the same server. If processes within the job run on different servers, this could result in each server having an incomplete set of generated cache files.

See "Creating Server Definitions" (PeopleTools 8.53: PeopleSoft Process Scheduler).

10. Click OK to launch the cache loading program.

After you invoke the process, you can use the Report Manager and Process Monitor links to monitor the progress.

The first time that you run the program, the process may take numerous hours to complete (less time for the Parallel LoadCache option).

11. When the cache loading program completes, verify the existence of the cache contents.

For shared caching, the cache loading program creates the shared file cache in the following directory:

- Windows: *PS_CFG_HOME*\appserv\prcs*ProcessScheduler_domain*\CACHE\CACHE\stage\stage
- For UNIX/Linux: *PS_CFG_HOME*/appserv/prcs/*ProcessScheduler_domain*/CACHE/STAGE/stage

Where *ProcessScheduler_domain* is the Process Scheduler domain in which you ran the process.

If database caching is enabled, the program loads the cache directly into the database. To verify the existence of the database cache, run the following SQL statement using your SQL editor:

```
SELECT * FROM PSOBJCACHE;
```

12. Once you've determined that the cache has been successfully loaded, shut down the application server domain(s), and modify the domain configuration to recognize the generated cache.

For shared file caching, edit the PSAPPSRV.CFG file, and enable shared caching, by uncommenting and setting `ServerCacheMode=1` for all appropriate domains.

For database caching, edit the PSAPPSRV.CFG file, and enable database caching, by uncommenting and setting `EnableDBCACHE=Y` for all appropriate domains.

13. (For shared file cache) Create a \SHARE directory within the CACHE directory for the appropriate application server domains, and copy the output of the load caching program into the \SHARE directory.

For example, copy the output of the load caching program to *PS_CFG_HOME*\appserv*domain*\CACHE\SHARE directory for the appropriate application server domain(s).

Note: UNIX and Linux systems expect an uppercase \SHARE directory. For Windows, the \SHARE directory can be mixed case.

14. Boot (or reboot) the application server domain(s).

Working with Database Cache

Database caching provides another implementation option. Determining which caching option works best for your site requires analysis and testing. For many cases, file caching will continue to be the preferred method.

The benefits of database caching include:

- Cache loading only needs to be run once for all domains connecting to a database as they will all share the same database cache.
- New domains can immediately take advantage of the database cache once `EnableDBCACHE` is set in the domain configuration. That is, no additional cache loading programs need to be run or additional output that needs to be copied.

- Eliminates administrative overhead of managing disk space and cache files on the file system.

Note: The database cache stores the cached metadata as a BLOB in a single table for the most efficient retrieval.

Note: The database cache is generated specific to the system's byte-order (little-endian or big-endian). The database cache is shared by all systems that have the same byte-order.

Maintaining the Versions of Cached Definitions

PeopleTools uses version numbers to keep cached objects synchronized with the data in the database. When a system undergoes change (such as during custom configurations, updates, or upgrades) versions can get out of sync, and they need to be reset.

Use the VERSION Application Engine program to maintain the definition version synchronization in your system. VERSION can run in these modes.

Mode	Invoked From	Run Control ID	Description
Reset All Versions	Command line only	Must start with RESETVERSIONS. For example, the run control ID could be RESETVERSIONS, RESETVERSIONS2, RESETVERSIONS3, and so on. Note: The run control ID is not case sensitive.	Resets all versions to 1. Before running in this mode, shutdown all servers and client connections. After running, regenerate cache.
Reset Differences Only	Command line or Process Scheduler	Any run control ID, <i>not</i> beginning with RESETVERSIONS or REPORTONLY.	Synchronizes only the definition versions where discrepancies exist. Servers can remain running and clients can remain connected. Cache does not need to be regenerated.
Report Only	Command line or Process Scheduler	Must begin with REPORTONLY. For example, the run control ID could be REPORTONLY, REPORTONLY2, REPORTONLY3, and so on. Note: The run control ID is not case sensitive.	Generates a report indicating where version discrepancies exist, but does not update any versions. Servers can remain running and clients can remain connected. Cache does not need to be regenerated.

When running VERSION in any mode, the system generates a report and writes it to %PS_SERVDIR%\log_output.

When running VERSION through Process Scheduler, keep these items in mind:

- Invoke VERSION from PeopleTools > Application Engine > Request AE.
- Create the appropriate run control ID (or use an existing run control ID), such as REPORTONLY.
- Set Process Frequency to *Always* so that the run control ID can be used more than once.

Related Links

[Configuring an Application Server Domain to Preload Cache](#)

"Using the Command Line to Invoke Application Engine Programs" (PeopleTools 8.53: Application Engine)

Tablespace Utilities

Select PeopleTools, Utilities, Administration, Tablespace Utilities to access the Tablespace Utilities page.

To comply with requirements for DB2 UDB for z/OS, the Tablespace Utility now includes both tablespace name and database names when you define a tablespace using the Tablespace Management page. Use the Add/Delete/Rename Tablespaces page to change the list of tablespace and database names.

Image: Add/Delete/Rename Tablespaces

This example illustrates the fields and controls on the Add/Delete/Rename Tablespaces. You can find definitions for the fields and controls later on this page.

Add/Delete/Rename Tablespaces							
Tablespaces Defined in the Database							
Customize Find View All First 1-7 of 14 Last							
	*Tablespace Name	Database Name	Type	DB2 Unix Type	Comment		
1	PSIMAGE	PSPTDMO	Regular	DMS		+	-
2	PSIMGR	PSPTDMO	Regular	DMS		+	-
3	PTAMSG	PSPTDMO	Regular	DMS		+	-
4	PTAPP	PSPTDMO	Regular	DMS		+	-
5	PTAPPE	PSPTDMO	Regular	DMS		+	-
6	PTAUDIT	PSPTDMO	Regular	DMS		+	-
7	PTLOCK	PSPTDMO	Regular	DMS		+	-

Tablespace Name Enter the name of the Tablespace that you want to add.

Database Name Enter the database name into which you want to add the space.

Type Select the type of Tablespace.

Use the plus/minus buttons to add and delete tablespaces.

Tablespace Management

Select PeopleTools, Utilities, Administration, Tablespace Management to access the Tablespace Management component

These pages enable you to modify the tablespace definition.

Tablespace Defn Page

This page shows the identification values for the tablespace.

Tablespace List Page

This page is where you add records to a particular tablespace. Use the plus and minus buttons to add and delete rows from the list.

Tablespace DDL Page

This page enables you to view and override DDL parameters if needed. View the default DDL in the Default Tablespace DDL list. You override specific parameters, if needed, in the Override Tablespace DDL list. Enter the parameter that you want to override in the Parameter Name column, and enter the override value in the Override column.

DDL Model Defaults

Select PeopleTools, Utilities, Administration, DDL Model Defaults to access the DDL Model Defaults page.

This page is used to view and edit the DDL for creating tablespaces, indexes and tables. Any changes that you make here are global.

Image: DDL Model Defaults

This example illustrates the fields and controls on the DDL Model Defaults. You can find definitions for the fields and controls later on this page.

DDL Model Defaults

Platform ID: 2 Oracle Copy...

Sizing Set: 0

DDL Find | View All First 1 of 5 Last

Statement Type: Table + -

*Model SQL: `CREATE TABLE [TBNAME] ([TBCOLLIST]) TABLESPACE [TBSPCNAME]
STORAGE (INITIAL **INIT** NEXT **NEXT** MAXEXTENTS **MAXEXT**
PCTINCREASE **PCT**) PCTFREE **PCTFREE** PCTUSED`

Parameter Count: 6

Parameters		Customize Find View All First 1-3 of 6 Last	
DDL Parm	DDL Parameter Value		
INIT	40000	+ -	
MAXEXT	UNLIMITED	+ -	
NEXT	100000	+ -	

Platform ID

Identify the type of platform of the current database type.

Sizing Set

Specify multiple Sizing Sets if needed. Sizing Sets are a way to maintain multiple versions of the DDL Model statements for a particular database platform. For example, you could have one sizing set to be used during a development phase, when tables only have test data, and you could have separate sizing set to be used during production, when tables have much more data.

Copy

Copies information from one sizing set to another.

Statement Type

Indicates the type of statement that's entered in the Model SQL edit box. Values for this field can be *Table*, *Index*, and *Tablespace*.

Model SQL

This field displays the model SQL statements, which you can edit. Valid statements are CREATE TABLE, CREATE INDEX, CREATE TABLESPACE, and a platform-specific statement for updating statistics.

Some platforms have all the statements, some do not.

Parameter Count

The Parameter Count is calculated based on how many nonblank DDL Parm rows that you define.

DDL Parm

Enter the DDL Parm (parameter) to which you want to assign a default.

DDL Parameter Value

The DDL Parameter Value field is where you override a DDL parameter's default value with your own for the selected statement type.

Using the DDL Model Defaults page, you can maintain DDL model statements and default parameters for Data Mover and Application Designer when DDL (PSBUILD.SQL, for example) is generated during the build process.

Using this utility, you can:

- Scroll through all the statement types and platforms that are defined in the PSDDLMODEL table.
- Change the default values or overrides for the parameters in the DDL model statements.
- Add, delete, or change DDL parameters and values.

The Platform IDs are as follows:

<i>Number</i>	<i>Platform</i>
0	SQLBase (no longer supported).
1	DB2.
2	Oracle.
3	Informix.
4	DB2/Unix.
5	Allbase (no longer supported).
6	Sybase.
7	Microsoft.
8	DB2/400 (no longer supported).

Note: There is no validation performed on the Model SQL statement, the DDL Parm syntax, or the relationship between the statement and the parameters.

BOE Integration Administration

This page enables you to manage your BusinessObjects Enterprise installation.

See "Crystal Reports for PeopleSoft Overview" (PeopleTools 8.53: Crystal Reports for PeopleSoft).

Strings Table

Select PeopleTools, Utilities, Administration, Strings Table to access the Strings Table page.

Image: Strings Table page

This example illustrates the fields and controls on the Strings Table page. You can find definitions for the fields and controls later on this page.

Strings Table

Program ID: QEGBR09 Report Language: English

String List Find | View All First 1-7 of 22 Last

*String Source	*String ID	Default Label	String Text	Width
Text	100%-AREA	<input type="checkbox"/> <input type="checkbox"/>	100% Area Chart	15
Text	100%-BAR	<input type="checkbox"/> <input type="checkbox"/>	100% Bar Chart	14
Text	AREA	<input type="checkbox"/> <input type="checkbox"/>	Area Chart	10
Text	AUGUST	<input type="checkbox"/> <input type="checkbox"/>	August	6
Text	BAR	<input type="checkbox"/> <input type="checkbox"/>	Bar Chart	9
Text	FLOATING-BAR	<input type="checkbox"/> <input type="checkbox"/>	Floating Bar Chart	18
Text	HIGH-LOW-CLOSE	<input type="checkbox"/> <input type="checkbox"/>	High-Low-Close	14

String Source

Options are:

RFT Long: Select if you want the long description of the field to be displayed in the column heading as set in Application Designer.

RFT Short: Select if you want the short description of the field as set in the Application Designer to be displayed in the column heading.

Text: Select to enter a custom column heading for the report.

String ID

Use the browse button to select the string ID that is to be used for the column heading in the SQR report.

Default Label

The default label is enabled if you select the RFT Long or RFT Short string source, otherwise, the check box is disabled.

Remember that fields can have multiple labels. Select the Default Label option to ensure that the default label is used. If you do not use the field's default label, you must select which of the field's labels to use using the label properties button.

String Text

Enter the text for the custom column heading, This is the text that is displayed if you set the string source to Text.

Width

The default value is the current width of the string that you enter or select. Be sure to update the width based on the actual space that is available on the report layout to avoid limiting a translator to an artificially short length, which is likely to degrade the quality of the translation.

Lookup Exclusion

Enter the record name for the tables you want to exclude from the auto lookup feature. A prompt or lookup button opens a lookup page in the user's browser populated with up to 300 available values for that field. The user can then either select the desired value or refine their search further. For extremely large tables, the system administrator has the option of excluding that table from auto prompting by adding the table to this list.

Related Links

"Prompt Fields" (PeopleTools 8.53: PeopleSoft Application Designer Developer's Guide)

XML Link Function Registry

The XML Link Function Registry is used exclusively in conjunction with the XML Link technology, which is associated with PeopleSoft Business Interlinks.

Note: PeopleSoft Business Interlinks is a deprecated product. This option currently exists for upgrade compatibility and transition.

Merchant Integration Utilities

There are two utilities that are related to the Merchant Integration technology that are provided for upgrade support only: Merchant Categories and Merchant Profile.

Refer to PeopleSoft documentation from previous releases for information regarding these utilities. These utilities are not intended for any new development purposes.

TableSet IDs

Select PeopleTools, Utilities, Administration, TableSet IDs to access the Tableset ID page.

Use this utility to create Set IDs. Before doing this:

- Add the SETID field (as a key field) to the record definition for that table.
- Define a Set Control Field as the field controlling the assignment of table sets.

SetID

Enter the set ID as defined in the record definition.

Descriptions/Comments

Add any descriptions and comments that are necessary for identification and internal documentation.

Record Group

Used to group record definitions for the tables that you want to share, as well as any dependent record definitions.

Select PeopleTools, Utilities, Administration, Record Group to access the Record Group page.

Image: Record Group page

This example illustrates the fields and controls on the Record Group page. You can find definitions for the fields and controls later on this page.

Description

The Record Group ID description should provide enough information to encompass a category of related tables, not just the table that you are specifically sharing.

Short Description

Enter a short description.

Force Use of Default SetID

This overrides alternate setIDs that are entered so that the default is used.

Record (Table) Name

This prompt list comes from a SQL view of record definitions that are defined with that Set Control Field that aren't already associated with a record group.

Record Description

Automatically populated when the Record (Table) Name is selected.

TableSet Control

The following pages are used to control table sets.

Record Group Page

Select PeopleTools, Utilities, Administration, TableSet Control to access the Record Group page.

Used to define which record groups use which table set.

Default SetID

This is the setID that the system uses as you add additional record definition groups to be shared within this tableset.

SetID

Although this database is set up to share only one accounting-related record group, you may have multiple record groups to which you assign default unique Set IDs.

Tree Page

Select PeopleTools, Utilities, Administration, TableSet Control, Tree to access the Tree page.

Used to share Trees as well as tables and views.

Default SetID

The Default setID that you assign to this field value automatically appears. If you create another tableset for sharing trees, you can change this value.

Tree Name

Use the browse button to select from a list of only the tree definitions that are defined with the same Set Control Field.

SetID

Use the browse button to select the appropriate SetID.

Convert Panels to Pages

The following pages are used to convert panels that are used in previous PeopleSoft Windows applications to pages that are used for browser access.

Scope Page

Select PeopleTools, Utilities, Administration, Convert Panels to Pages to access the Scope page.

This utility helps you update panels from PeopleTools 7.5x versions to reflect the pages that are used for the PeopleSoft Internet Architecture.

Project List

Insert projects, containing panels that you want to convert, into this scroll. In addition, if you use the Apply Panel Group Defaults option, any panel group that is contained in projects in this scroll are processed. Note that exceptions may be defined see the task titled, Project Exceptions.

Page List

Insert panels that you want to convert to pages into this scroll.

Project Exceptions

If you want to ensure that a group of panels or panel groups is never processed for conversion, you can insert them into an application upgrade project and insert the project name in this scroll.

Page Exceptions

Panels that are inserted into this scroll are not be processed.

See PeopleSoft ugrade documentation.

documentation for more information.

Options Page

Select PeopleTools, Utilities, Administration, Convert Panels to Pages, Options to access the Options page.

Specify the options for the conversion process:

Convert Scrolls to Scroll Areas

If you select this option, scroll-to-scroll area conversions take place for panels with scroll bars. If this is unchecked, no scroll-to-scroll area conversion takes place.

Convert Scroll Action Buttons to Scroll Areas

Some scroll bars may exist with scroll action buttons that are already defined. This option determines whether these scrolls should be converted or ignored. If they are converted, the scroll action buttons are removed before the scroll bar is converted to a scroll area.

If you select this option, scrolls with scroll action buttons are converted. If this options is not checked, scrolls with scroll action buttons are ignored.

Panels with Level 1 Scrolls

If you select this option, panels with level 1 scrolls are processed for scroll conversion.

Panels with Level 2 Scrolls

If you select this option, panels with level 2 scrolls are processed for scroll conversion.

Panels with Level 3 Scrolls

If you select this option, panels with level 3 scrolls are processed for scroll conversion.

Convert Level 1 Scrolls

If you select this option, level 1 scrolls are converted to scroll areas.

Convert Level 2 Scrolls

If you select this option, level 2 scrolls are converted to scroll areas.

Convert Level 3 Scrolls

If you select this option, level 3 scrolls are converted to scroll areas.

Max # Scrolls

This parameter is a general scroll count limit for scroll conversion processing. For example, if this is set to 5, any panel with more than five scrolls that are not invisible is ignored. This is a simple way of eliminating complex panels from automatic scroll conversion.

Apply Specific Page Size

This option is used to define whether a specific size should be assigned to a panel. If you select this option, the panel size that is defined in the drop-down list box is applied to the panel. If this is unchecked, no changes are made to the panel size.

Note: Note. When you select a specific panel size, the panel size is applied to standard panels only (secondary panels and subpanels are not sized automatically).

Apply Default Style Sheet

If you select this option, the style sheet that is associated with a panel is updated with a blank value, so that the panel's style sheet appears by default from PSOPTIONS.STYLESHEETNAME ('PSSTYLEDEF').

Apply Frame/Horz/GrpBox Styles

If you select this option, the conversion process looks for frames, group boxes, and horizontal rules that have no styles associated with them, and that appear to be associated with a specific scroll area by virtue of their position within a scroll area. It then assigns level-specific styles, based on the occurs level of the scroll area.

Convert Frames to Horizontal

Horizontal lines were a new page object for PeopleSoft 8.

This option applies only for applications upgraded from releases previous to PeopleSoft 8. If you select this option, the conversion process looks for frames on the panel with upper and lower coordinates less than nine grid units apart. These frames are then converted to horizontal lines.

Delete All Frames

If you select this option, the process removes all frames on the converted panel.

Note: If Convert Frames to Horizontal and Delete All Frames are both checked, the conversion from frame to horizontal takes place first, then any remaining frames are deleted.

Turn On Grid 'Odd/Even Style'

This applies to grids that are on a panel being converted. If you select this option, the conversion process determines if grids on the panel have their 'Odd/Even Style' turned on. If it is not turned on, the conversion process turns on this option.

Turn On 'Show Prompt Button'

This option applies to edit box fields that are not invisible and are not display-only. If you select this option, the conversion process turns on the Show Prompt Button option for edit box fields that have it turned off.

Apply Component Defaults

Used to apply standard defaults to component definitions. The defaults that are set are dependent on the Use characteristics of the component. See Application Designer, Component Properties/Use and Component Properties/Internet tabs.

Turn Off 'Show Grid Lines'

Turns off the Show Grid Lines option for grids that have it checked on.

Language Code

Enables you to convert panels whose language code differs from that in PSOPTIONS. Select a language code from the drop-down list box.

Update Utilities

The Update utilities enable you to keep track of the PeopleSoft updates that you apply to the database.

Updates By Release Label

The release label refers to the official release name, such as PeopleTools 8.53.00

Updates By Update ID

The update ID refers to the patch or project name that you apply to the system. The update ID is typically the report ID for a TPRD incident.

Remote Database Connection

Use the Remote Database Connection page to set up remote databases for use with the Remote Data Access (RDA) feature. Select PeopleTools, Utilities, Administration, Remote Database Connection to access the Remote Database Access Management page.

Name	Enter the name of the remote database connection.
Database Type	Available types are Microsoft, DB2 (z/OS), DB2/UNIX, Sybase, Informix, Oracle, and Sybase.
Description	Enter a description of the remote database.
Server	Enter the server name where the remote database resides.
Database	Enter the remote database name.
Local Connect	One connection must be defined as the Local Connect for the current PeopleSoft instance (the local database). Check this to specify which database is the local.
DB Server Port	This value is automatically populated with a default value that is based on the database type. You may need to change this value depending upon the database server configuration.
User ID	Enter the user ID that is needed to connect to the remote database.
Password	Enter the password that is associated with the user ID.
Test Connection	Select this to test the remote database connection.
Connection Type	For Oracle database type only. TNS Names or Specific. TNS Names represent a preconfigured file (tnsnames.ora) that consists of previously defined database connection information. Enter Specific if you want to set up a database that does not already have a TNS entry defined.
TNS Entry	For Oracle database type only.
Inf Svr Name	For Informix database type only.

Security in Remote Databases

To ensure security and limit the risk of unauthorized access to databases, follow these recommendations:

- The remote system's database administrator should create a user with read-only access to the tables that may be accessed by other systems using PeopleSoft's RDA.

Use this restricted user ID and password in configuring a source RDA node.

- The local system's database administrator should create a user with insert/update access to the RDA destination tables only.

Use this restricted user ID and password in configuring the target RDA node.

URL Maintenance

Select PeopleTools, Utilities, Administration, URLs to access the URL Maintenance page.

Use the URL Table to store URL addresses and to simplify specifying and updating URLs. URLs that are saved here can be referenced from page controls such as a push button or link. The associated URL can be either an internet or intranet link.

Image: URL Maintenance page

This example illustrates the fields and controls on the URL Maintenance page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'URL Maintenance' page with the following fields and values:

- URL Identifier:** QE_NT4
- *Description:** QE_NT4
- *URL:** https://ptsec01.peoplesoft.com:7002/servlets/iclientservlet/psNT4/?cmd=start&
- Comments:** (Empty text area)

Description

Users can search for URLs by description.

URL

Enter the entire URL.

Note: For the file attachment functionality, in specifying the URL for the FTP server, the FTP server's machine name can be more than 30 characters, but the length of the full URL is limited to 120 characters. The URL format for file attachments stored in the database is *record://recordname*, as in *record://ABSENCE_HIST_DB*.

Comments

This field can be used to make notations and comments and is not displayed elsewhere.

Adding New URLs

To add a new URL entry in the URL table:

1. Click the Add a New Value link.

A new page appears, prompting you to enter the URL Identifier. Enter the name that you want to use to identify the new URL address.

2. Select Add.
3. Enter the Description, URL, and Comment, if any.
4. Select Save.

You must save the page before you can add another URL, or update or display existing URL addresses.

5. Select Add to add another URL.

Viewing and Updating the URL Table

To update or display the URL table:

1. From the URL Maintenance search page, click Search.
2. Select the URL Identifier link that you want to update from the Search Results.
3. Make changes to the page and save.

Specifying URL Properties

You access the URL Properties page by clicking the URL Properties link on the URL maintenance page.

Image: URL Properties page

This example illustrates the fields and controls on the URL Properties page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'URL Properties' page. At the top, there is a navigation bar with 'Customize | Find | [2] | [grid icon] | First 1-5 of 5 Last'. Below this is a table with the following data:

	*Property Name	Property Value		
1	AUHTYPE	3 - ANY	+	-
2	PRIVATEKEY	id_rsa	+	-
3	PUBLICKEY	id_rsa.pub	+	-
4	USER	tsawyer	+	-
5	PASSWORD	3QwPKYQEt1ktIY0fQNHVPQ==	+	-

Below the table is a section titled 'Password Encryption'. It contains two input fields: 'Password' and 'Confirm Password', both filled with dots. Below these fields is a yellow 'Encrypt' button. To the right of the button is an 'Encrypted Password' field containing the text '3QwPKYQEt1ktIY0fQNHVPQ=='. The 'Encrypted Password' field is highlighted in blue in the original image.

You create the URLID using the URL Maintenance page, but you assign various properties to the URLID for the following protocol types using the URL Properties page.

- FTP
- FTPS
- SFTP
- HTTP
- HTTPS

The properties you assign depend on the protocol and the type of security you are implementing.

The Password Encryption box enables you to enter and confirm any password value you intend to assign to a PASSWORD property. Click Encrypt to encrypt the password, then copy the results into the edit box for PASSWORD.

Note: If you intend to implement secure FTP, SFTP, HTTP, or HTTPS, Oracle recommends you understand the details and behavior of the attachment PeopleCode constructs, such as AddAttachment.

See "Understanding the File Attachment Functions" (PeopleTools 8.53: PeopleCode Developer's Guide).

To assign properties to a URL:

1. On the URL Maintenance page, add a URL value containing *ftp://*, *ftps://*, *sftp://*, *http://*, or *https://* in the URL field.
2. Click Save.
3. Click the URL Properties link.
4. Set the URL properties for the protocol and your security preference.
5. Click Save.

The available URL properties for each protocol appear in the following table:

Protocol	Properties	Description
FTP/S	ACTIVEMODE	To enable active mode, add the ACTIVEMODE property to the URL and set it to <i>Y</i> . The default FTP connection mode is extended passive mode.
	ACTIVEPORTOPTION	This property can be used along with ACTIVEMODE. When active mode is enabled, you can use ACTIVEPORTOPTION to specify the IP address and port on which the FTP server can be accessed. This is useful when the server is behind a firewall. By default, ACTIVEPORTOPTION uses the default IP address of your system. If you want to use a particular IP address, set the ACTIVEPORTOPTION value to either the full IP address, a host name to resolve to an IP address, or a local network interface name. You can also specify a port range. For example: <i>10.176.147.111:10000-13000</i>

Protocol	Properties	Description
	CERTALIAS	<p>(Certificate Alias) The Certificate Alias must be an alias name of a certificate stored in the database (using the PeopleTools Digital Certificates page).</p> <hr/> <p>Note: Currently, only PEM certificates are supported for FTPS.</p> <hr/>
	ENABLEEPRT	<p>This option can be used only with Active Mode. If Active Mode is enabled and ENABLEEPRT is set to <i>N</i>, then the system will use a PORT (IPv4) Active Mode connection. By default, ENABLEEPRT is Y, if Active Mode is set to Y.</p>
	FILE_EXT_LIST	<p>Select the name of a file extension list created using the File Extension List page.</p> <p>A file extension list enables you to restrict file attachment features to work only with appropriate file types, and reject those file types that you believe may be suspicious, potentially harmful, or not appropriate for a particular call to a file-processing built-in PeopleCode function.</p> <p>See File Extension List.</p>

Protocol	Properties	Description
	KEYSTOREPASSWORD	<p>Required for FTPS and HTTPS repositories.</p> <p>For attachments transferred from the PeopleSoft system to the FTPS or HTTPS repository, the system retrieves the key pair for the client certificate from the digital certificate store and writes them to a file in PKCS12 format with password protection. The value of this property will be used as the password for the PKCS12 file.</p> <p>The PKCS12 file enables connection and file transfer, and it exists only temporarily in <PS_SERVDIR>\files \<CERT ALIAS NAME> for the duration of the file transfer.</p> <p>The system deletes the file after the file transfer transaction.</p> <hr/> <p>Note: If the system fails to delete the certificate alias file, a message will be written to the application sever log. The maximum number of files that can exist at any time is equal to the total number of FTPS and HTTPS URL identifiers defined in the system.</p> <hr/>
	EXTENDEDPASSIVEMODE	<p>Enables you to control whether extended passive mode (EPSV) will be used by FTP.</p> <p>To enable select 1, to disable select 0.</p> <p>EPSV is used by default. That is, by default, this value is considered to be 1.</p> <p>If the client fails to connect to the server with EPSV, then the system will try passive mode (PASV). To use PASV only, add EXTENDEDPASSIVEMODE to the URL Properties and set it to 0.</p>
	PASSWORD	<p>The password associated with the USER property, which identifies the FTP user ID.</p>

Protocol	Properties	Description
	SSLUSAGELEVEL	<p><i>0 - No SSL:</i> No SSL will be used.</p> <p><i>1 - Try SSL:</i> Try using SSL, but proceed as normal otherwise.</p> <p><i>2 - Control:</i> Require SSL for the control connection.</p> <p><i>3 - SSL Only:</i> (Default) Require SSL for all communication.</p>
	USER	The FTP user ID used for authentication when accessing the FTP site.
	VERIFYHOST	<p><i>0:</i> Do not verify the server for host name.</p> <p><i>1:</i> Checks if there exists any value in the common name field in the server certificate. Does not verify if it matches with what the client specifies.</p> <p><i>2:</i> (Default) Checks for a match with the host name in the URL with the common name or Subject Alternate field in the server certificate.</p>
	VERIFYPEER	<p><i>False:</i> Do not verify the peer.</p> <p><i>True:</i> (Default) Verify the peer by authenticating the certificate sent by the server.</p>
SFTP	AUTHTYPE	<p>Select the authentication type:</p> <ul style="list-style-type: none"> • Public Key • Password • Any

Protocol	Properties	Description
	FILE_EXT_LIST	<p>Select the name of a file extension list created using the File Extension List page.</p> <p>A file extension list enables you to restrict file attachment features to work only with appropriate file types, and reject those file types that you believe may be suspicious, potentially harmful, or not appropriate for a particular call to a file-processing built-in PeopleCode function.</p> <p>See File Extension List.</p>
	PASSWORD	Specify the user password. You can enter the password in the Password Encryption box, click Encrypt, then copy the encrypted value to the Password property.
	PASSWORDKEY	Enter the password for the private key.
	PRIVATEKEY	Select the private key.
	PUBLICKEY	Select the public key.

Protocol	Properties	Description
	SSHKEYALIAS	<p>Select the SSH certificate saved to the database using the PeopleTools Security, Digital Certificates page (PeopleTools, Security, Security Objects, Digital Certificates). The SSH certificate added through the Digital Certificates page contains both the public and private key data, identified by the Alias column value on the Digital Certificates page.</p> <p>If using the SSHKEYALIAS URL property, the Property Value prompt displays only the list of SSH certificates that have been added to the Digital Certificates page.</p> <p>If you have added the SSH certificate using the Digital Certificates page, and you have assigned an SSH certificate to the SSHKEYALIAS URL property, the system ignores the PUBLICKEY and PRIVATEKEY properties, regardless if they refer to valid key files in the file system.</p> <p>If you provided a password (or passphrase) when generating your SSH certificate, specify that value using the PASSWORDKEY URL property.</p> <p>See "Configuring Digital Certificates" (PeopleTools 8.53: Security Administration)</p>
	USER	Specify the user ID to be authenticated.
HTTP	FILE_EXT_LIST	<p>Select the name of a file extension list created using the File Extension List page.</p> <p>A file extension list enables you to restrict file attachment features to work only with appropriate file types, and reject those file types that you believe may be suspicious, potentially harmful, or not appropriate for a particular call to a file-processing built-in PeopleCode function.</p> <p>See File Extension List.</p>

Protocol	Properties	Description
	PASSWORD	Specify the user password. You can enter the password in the Password Encryption box, click Encrypt, then copy the encrypted value to the Password property.
	PROXY	Specify the proxy server (if used).
	PROXYPORT	Specify the proxy server port (if used)
	USEAUTHTOKEN	<p>Uses the PS_TOKEN authentication cookie, if single signon is configured between source and target destination.</p> <hr/> <p>Note: This must be set to <i>Y</i> if a PeopleSoft web server is used for the HTTP repository.</p> <hr/>
	USER	<p>Specify the user ID to be authenticated.</p> <hr/> <p>Note: USER must be set if a PeopleSoft web server is used for the HTTP repository.</p> <hr/>
HTTPS	CERTALIAS	<p>(Certificate Alias) The Certificate Alias must be an alias name of a certificate stored in the database (using the PeopleTools Digital Certificates page).</p> <hr/> <p>Note: Currently, only PEM certificates are supported for FTPS.</p> <hr/>

Protocol	Properties	Description
	KEYSTOREPASSWORD	<p>Required for FTPS and HTTPS repositories.</p> <p>For attachments transferred from the PeopleSoft system to the FTPS or HTTPS repository, the system retrieves the key pair for the client certificate from the digital certificate store and writes them to a file in PKCS12 format with password protection. The value of this property will be used as the password for the PKCS12 file.</p> <p>The PKCS12 file enables connection and file transfer, and it exists only temporarily in <PS_SERVDIR>\files \<CERT ALIAS NAME> for the duration of the file transfer.</p> <p>The system deletes the file after the file transfer transaction.</p> <hr/> <p>Note: If the system fails to delete the certificate alias file, a message will be written to the application sever log. The maximum number of files that can exist at any time is equal to the total number of FTPS and HTTPS URL identifiers defined in the system.</p> <hr/>
	FILE_EXT_LIST	<p>Select the name of a file extension list created using the File Extension List page.</p> <p>A file extension list enables you to restrict file attachment features to work only with appropriate file types, and reject those file types that you believe may be suspicious, potentially harmful, or not appropriate for a particular call to a file-processing built-in PeopleCode function.</p> <p>See File Extension List.</p>
	PASSWORD	<p>Specify the user password. You can enter the password in the Password Encryption box, click Encrypt, then copy the encrypted value to the Password property.</p>

Protocol	Properties	Description
	PROXY	Specify the proxy server (if used).
	PROXYPORT	Specify the proxy server port (if used).
	SSLUSAGELEVEL	<p><i>0 - No SSL:</i> No SSL will be used.</p> <p><i>1 - Try SSL:</i> Try using SSL, but proceed as normal otherwise.</p> <p><i>2 - Control:</i> Require SSL for the control connection.</p> <p><i>3 - SSL Only:</i> (Default) Require SSL for all communication.</p>
	USEAUTHTOKEN	<p>Uses the PS_TOKEN authentication cookie, if single signon is configured between source and target destination.</p> <hr/> <p>Note: This must be set to <i>Y</i> if a PeopleSoft web server is used for the HTTP repository.</p> <hr/>
	USER	<p>Specify the user ID to be authenticated.</p> <hr/> <p>Note: USER must be set if a PeopleSoft web server is used for the HTTP repository.</p> <hr/>
	VERIFYHOST	<p><i>0:</i> Do not verify the server for host name.</p> <p><i>1:</i> Checks if there exists any value in the common name field in the server certificate. Does not verify if it matches with what the client specifies.</p> <p><i>2:</i> (Default) Checks for a match with the host name in the URL with the common name or Subject Alternate field in the server certificate.</p>
	VERIFYPEER	<p><i>False:</i> Do not verify the peer.</p> <p><i>True:</i> (Default) Verify the peer by authenticating the certificate sent by the server.</p>

Working with FTP/FTPS

FTP over SSL (FTPS) can be achieved in two ways: implicit and explicit. Implicit SSL uses `ftps://` and explicit uses `ftp://`.

To use PASV only, set URL property `EXTENDEDPASSIVEMODE` to `0`.

Working with SFTP

With SFTP, these authentication methods are supported:

- password.
- key authentication.

Note: With key authentication, you have two choices: storing the public and private keys as files within the file system, or storing the public and private keys within an SSH certificate added using the Digital Certificates page.

To set up password authentication create a URL identifier with an SFTP URL value, such as `sftp://10.145.641.122/`, then set these URL properties:

- `AUTHTYPE: 2 - PASSWORD`
- `USER: user ID`
- `PASSWORD: encrypted user password`

For public key authentication *with key files stored in the file system*, you can copy the public key file to `<PS_SERVDIR>\sshkeys\public` and the private key file to `<PS_SERVDIR>\sshkeys\private`. Alternatively, you can copy the key files to a different location and define an environment variable `PS_SSHKEYPATH`.

Note: `PS_SSHKEYPATH` can replace only `PS_SERVDIR`. The following directory structure must exist in your custom location, `.\sshkeys\public` and `.\sshkeys\private`. For example, `<PS_SSHKEYPATH>\sshkeys\public`.

If storing the SSH key data *in the file system*, set these URL properties appropriately:

- `AUTHTYPE: (1 - PUBLIC KEY)`
- `USER`
- `PRIVATEKEY`
- `PUBLICKEY`
- `PASSWORDKEY`

For public key authentication *with an SSH certificate stored in the PeopleTools keystore*, save the private key and public key data to the database using the Digital Certificates page (PeopleTools, Security, Security Objects, Digital Certificates page). If your SSH certificate is stored *in the database* by way of the Digital Certificates page, set these URL properties appropriately:

- `AUTHTYPE: (1 - PUBLIC KEY)`

- USER
- SSHKEYALIAS
- PASSWORDKEY

Note: Oracle recommends storing the SSH public and private key data in the database by creating an SSH certificate stored in the PeopleTools keystore for increased security and ease of maintenance. If the private key and public key files are stored in the file system, Oracle recommends implementing a password-protected private key.

For more information on the Digital Certificates page, see "Configuring Digital Certificates" (PeopleTools 8.53: Security Administration).

Working with HTTP/HTTPS

HTTP/S file transfer requires the http/https repository available (file storage location). The repository can either be on a PeopleSoft web server or non-PeopleSoft server. The PeopleSoft Internet Architecture provides a handler servlet (psfiletransfer) to handle the requests. If you use a non-PeopleSoft web server you need to create a handler for non-PeopleSoft servers/repositories.

Once the repository is created, the configuration.properties file for the appropriate PeopleSoft site needs to be modified to include the HttpRepository property. For example:

```
HttpRepositoryPath=D:/HttpRepository
```

Note: Use a forward slash (/).

To set up HTTP:

After creating a URL identifier with an http:// URL (as in, http://<PIA URL>/psfiletransfer/<sitename>/), select these URL properties:

- USEAUTHTOKEN: Y
- USER: *user_ID*

Set the Authentication Option field on the default local node to *Password*.

To set up HTTPS:

Setup the target web server for HTTPS, and add the Root CA and Client Certificate to Digital Certificate store.

Create an URL Identifier using *https://*, and add these URL properties:

- CERTALIAS: *your certificate alias*
- SSLUSAGELEVEL: *3-SSL Only*
- VERIFYHOST: 2
- VERIFYPEER: Y
- USER: *user ID*

- USEAUTHTOKEN: *Y*

You can have the HTTP repository on a separate PeopleSoft web server too. For example, PIAsystem1 can have the URL identifier pointing to a URL on PIAsystem2. This requires following to be in place:

- The user specified in the URL properties must exist in both the systems with the proper permissions.
- The default local nodes on both systems should have Authentication Option set to *Password* and should be pinged successfully.
- The default local nodes of each system are added to single-signon list for each database.

Note: To upload non-ASCII file names to an HTTP Repository, the application server needs to be booted in the native language (such as Japanese or Chinese). If the application server is booted in English, file names with non-ASCII characters will not be uploaded to the HTTP Repository.

See "Selecting Character Sets" (PeopleTools 8.53: Global Technology) for more information on how character sets on various tiers of the PeopleSoft system can affect file attachment processing.

Specifying URLs When Attaching Files

When attaching files, whether you can pass a string URL or only use the URL identifier depends on the protocol and repository type being used. When attaching files, keep these guidelines in mind:

- Using the URL Identifier is mandatory for FTPS, SFTP, HTTP, and HTTPS.
- For FTP, Record, and File, you can use the URL Identifier or use a URL string (and not reference the stored URL Identifier).

For example, with plain FTP you could use a URL string, such as:

```
ftp://user:password@site/
```

Or, if a URL Identifier has been created, you can reference that in one of these formats:

- *ftp://userid:password@ftpurl/*: where the FTP user and FTP password are provided in clear text, but the URL information comes from the URL Identifier.
- *ftp://userid@ftpurl/*: where the FTP user is provided in clear text, but the FTP password and the URL information comes from the URL Identifier.
- *ftp://ftpurl/*: where all the required information comes from the URL Identifier.

File Extension List

Select PeopleTools, Utilities, Administration, File Extension List to access the File Extension List page.

A file extension list contains a list of file types, identified by file extension, that the system either accepts or rejects during specific file attachment operations. Once you have created a file extension list you specify it as the value for the FILE_EXT_LIST property of a URL object using the URL Maintenance page. This enables you to only upload the appropriate file types at runtime, and reject those file types you think may be suspicious, potentially harmful, or not appropriate for a particular call to a file-processing built-in PeopleCode function.

A file extension list applies only when performing these file attachment operations:

- Uploading files (AddAttachment, MAddAttachment PeopleCode)
- Viewing files (ViewAttachment PeopleCode)
- Detaching files (DetachAttachment PeopleCode)

Note: File extension lists can be assigned only to a URL object by way of the FILE_EXT_LIST property, and they apply only to these file attachment operations.

See [URL Maintenance](#).

See "Understanding the File Attachment Functions" (PeopleTools 8.53: PeopleCode Developer's Guide).

You can create different file extension lists to restrict certain file types, depending on the needs of the application. For example, for an application that users access from outside your firewall, such as a job applicant interface, you might want to impose strict restrictions on the accepted file types to avoid allowing potentially harmful file types into your system. For an internal-facing application, you may not require such strict requirements for file attachments.

Image: File Extension List page

This example illustrates the fields and controls on the File Extension List page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'File Extension List' configuration page. It includes a header 'File Extension List' and a sub-header 'File Extension List'. The form contains the following fields and controls:

- List Name:** JOB_APPLICANT_EXT_LIST
- Description:** Accepted Resume or CV Types
- Extension List Type:** A group of radio buttons with options: Inactive, Absolute (selected), Relative, and Unlimited.
- Standard Extensions:** A link to view standard extensions.
- Contained Extensions:** A table with columns for File Extension, Status, and actions (+, -).

Contained Extensions		Personalize	Find	First	1-2 of 2	Last
	*File Extension	*Status				
1	.PDF	Accept				
2	.TXT	Accept				

List Name

Displays the name you entered for the file extension list in the Add New Value tab. File extension list names cannot exceed 30 characters.

Description

Enter any descriptive information to help indicate how this list is to be used.

Extension List Type

Inactive

When set to *Inactive*, the system ignores the file extension list. Any file extension entries in the Contained Extensions grid are ignored at runtime for an Inactive list type. From a development

perspective, it may be useful to keep a working list of file extensions, but keep the list itself inactive until it is fully tested.

Note: When troubleshooting or testing, setting a file extension list to *Inactive*, immediately disables any restrictions imposed by that file extension list.

Absolute

With an *Absolute* file extension list, you can create a list in which all the acceptable filename extensions are explicitly enumerated and have a status of *Accept*. That is, *only* the file extensions that appear in the Contained Extensions grid and have the Status set to *Accept* will be accepted for that file attachment operation at runtime. An Absolute file extension list is most appropriate for situations where you wish to accept a relatively small number of file types.

Any file extension set to *Reject* will not be accepted at runtime. However, *not* including such a file extension in your absolute list will have precisely the same effect.

Relative

The Relative file extension list type is based on the Standard Extensions list, and provides a flexible means to specify a finite number of files that are acceptable, as well as enabling you to extend the Standard Extensions list, if needed. A Relative file extension list is most appropriate for situations where you wish to reject a relatively small number of file types.

After reviewing the Standard Extensions list, add specific file extensions that you *do not* want to authorize for file attachment operations to the Contained Extensions grid, and set the Status value to *Reject*. By doing so, you indicate to the system that you authorize all of the extensions on the Standard Extensions list, *except* for those you've added to the Contained Extensions grid and set to *Reject*.

For a Relative file extension list type, you can use the *Accept* status value to effectively extend the current list of Standard Extensions for this specific relative file extension list. For example, assume that at your site you use a product that uses a file extension that does not appear in the Standard Extensions list. By adding that file extension to the Contained Extensions grid and setting its Status value to *Accept*, the PeopleSoft system will accept that file type for file attachment operations involving this specific relative file extension list.

Note: PeopleTools does not permanently add the extension to the Standard Extensions list when it is set to *Accept* for a relative file extension list type. That is, when you view the Standard Extensions list after setting the Status to *Accept* for a particular file extension, it will not appear in the list and will not affect the contents of the Standard Extensions list when it is accessed at runtime by other file extension lists you use. However, at runtime, the system will recognize that file extension as acceptable for this specific relative file extension list.

Unlimited

Do not use the Unlimited extension list type in the current release. This type of list is intended for possible use with future functionality.

Standard Extensions

Click to display a read-only list of standard file extensions that are known and potentially acceptable to the PeopleSoft system. This list comprises the standard list of file types that a PeopleSoft web server will handle. The displayed Standard Extensions list is provided for informational purposes to help you complete your file extension lists by indicating the set of file extensions that the PeopleSoft system recognizes.

To include additional file extensions to the Standard Extensions list, refer to the Relative file extension list type description.

Contained Extensions

File Extension

Enter the file extensions you want to accept or reject. Add one extension per row, with each extension not exceeding ten characters.

The system saves the value of an individual file extension as all upper case, regardless of the case you enter, and prepends a period ('.') to the file extension unless its first character is already a period.

Note: Files having no file extension may be indicated collectively by specifying a file extension of a period (".").

Status

For a specified file extension of type Absolute or Relative, you have these Status options:

- *Accept:* Indicates that the system will accept files having the specified extension.
- *Reject:* Indicates that the system will reject files having the specified extension.

Bypassing Restrict-by-File-Type Functionality for System Administrators

System administrators managing the file attachment features of your PeopleSoft system can bypass restrictions otherwise imposed by filename extension lists by using a delivered security role, *?PTFX Ignore Download Ext Lists?*.

During a call to ViewAttachment or DetachAttachment, if the user profile initiating the transaction is associated with the *?PTFX Ignore Download Ext Lists?* role, then the system ignores file download restrictions specified by any filename extension list.

Note: The *PTFX Ignore Download Ext Lists* role should be associated only with the appropriate user profile(s).

Note: The *PTFX Ignore Download Ext Lists* role contains no permission lists, however, its association with a user profile triggers internal PeopleTools security rules.

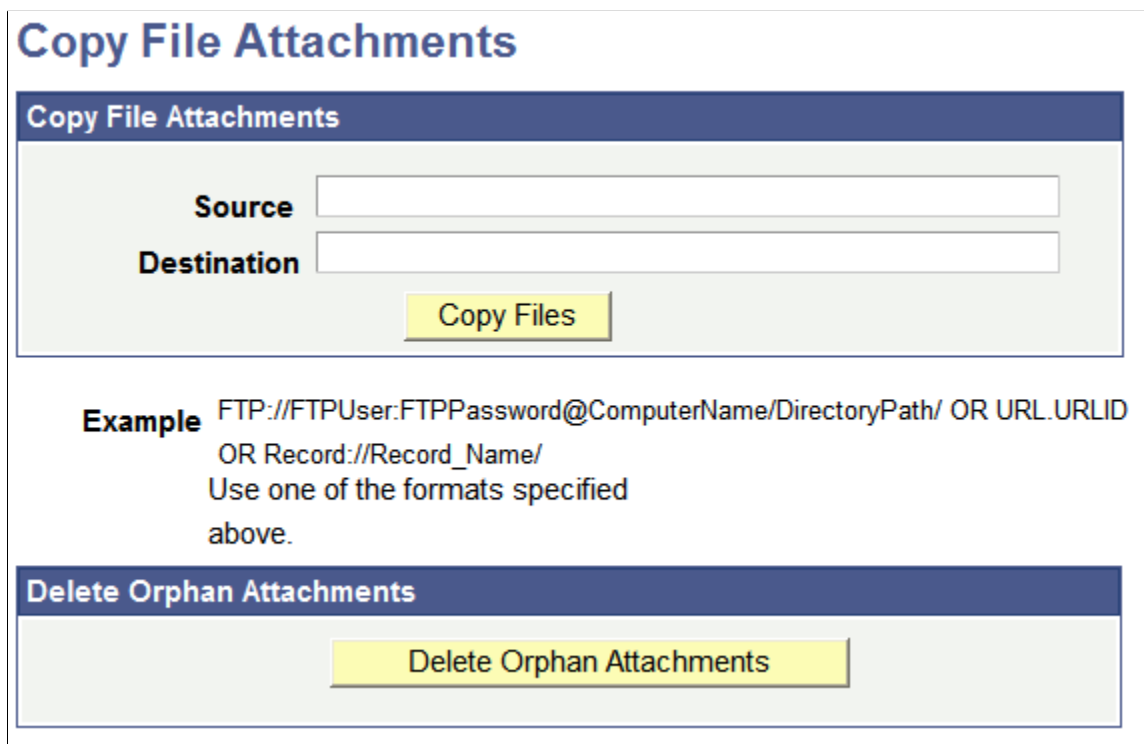
Copy File Attachments

Select PeopleTools, Utilities, Administration, Copy File Attachments to access the Copy File Attachments page.

This page enables you to manage your attachment archives. You can transfer file attachment archives from one location (source) to a new location (destination), and you can perform general housecleaning by deleting orphaned attachments.

Image: Copy File Attachments page

Copy file attachments from one location to another as well as purge system of orphaned file attachments



Copy File Attachments

Source	Enter the URL that corresponds to the current (source) location of the file attachment archive to be copied.
Destination	Enter the URL that corresponds to the destination location of the file attachment archive to be copied.
Copy Files	<p>Invokes the PeopleCode function (CopyAttachment) that copies the file attachment archive from the specified source storage location to the specified destination storage location.</p> <p>For example, you can copy from the FTP server to a database, a database to the FTP Server, and so on.</p>

Note: For the file attachment functionality, in specifying the URL for the FTP server, the FTP server's machine name can be more than 30 characters. The length of the full URL is limited to 120 characters.

Prior to copying files to a database storage location, you must first prepare the target record to store the attachments, which is discussed in the PeopleCode Developer's Guide.

See "Understanding the File Attachment Functions" (PeopleTools 8.53: PeopleCode Developer's Guide).

Example	Displays the sample formats used for FTP, URL, or Record attachment transfers. These sample formats apply to both the Source and Destination edit boxes.
----------------	--

Deleting Orphan Attachments

The accumulation of file attachments can consume a significant amount of space in your database. Therefore, on a regular basis, it is recommended that you make sure that orphaned file attachments are deleted from the storage location, reclaiming disk space being used unnecessarily. *Orphan file attachments* are stored file attachments for which there is no longer a corresponding file reference in a valid file reference table.

Click the Delete Orphan Attachments button to complete this task. This button invokes the CleanAttachments PeopleCode function.

Warning! There is no way to roll back changes made by the CleanAttachments function. Review the CleanAttachments documentation in the PeopleCode Language Reference PeopleBook describing the behavior of CleanAttachments in order to appropriately anticipate how it will behave. Oracle suggests that you perform a database backup before deleting orphan attachments.

Note: When large numbers of orphaned file attachments may exist, Oracle recommends using the delivered Application Engine program, CLEANATT84, rather than the Delete Orphan Attachments button to remove the orphaned files. This avoids potential time-out issues.

See "CleanAttachments" (PeopleTools 8.53: PeopleCode Language Reference)

Query Administration

System administrators can use Query Administration to monitor query performance and usage. Some of the conditions that you can monitor include average runtime, number of times run, and the dates last run. Using a predefined search, you can also select queries to review and report on.

Related Links

"Understanding Query Administration" (PeopleTools 8.53: PeopleSoft Query)

Sync ID Utilities

The Sync ID Utilities are used exclusively with PeopleSoft Mobile Applications technology.

Important! PeopleSoft Mobile Agent is a deprecated product. These pages currently exist for backward compatibility only.

nVision Report Request Admin

See "Securing and Sharing Report Requests in PIA" (PeopleTools 8.53: PS/nVision).

Analytic Server Administration

See "Administering Analytic Servers" (PeopleTools 8.53: Analytic Calculation Engine).

Upgrade Conversion

Defines upgrade drivers, providing details regarding Application Engine program, section, group, and calling sequence. If you need to specify any of these values, your upgrade documentation will provide the details.

See your upgrade documentation for more information.

Analytic Model Viewer

See "Viewing Analytic Model Properties" (PeopleTools 8.53: Analytic Calculation Engine).

Analytic Instance Load/Unload

See "Loading and Unloading Analytic Instances" (PeopleTools 8.53: Analytic Calculation Engine).

Analytic Instance Create/Del/Copy/

See "Creating, Deleting, and Copying Analytic Instances" (PeopleTools 8.53: Analytic Calculation Engine).

Pre-Load Cache Utilities

See [Configuring an Application Server Domain to Preload Cache](#).

Gather Utility

The Gather utility facilitates communications between PeopleSoft and the customer on technical questions or issues. The Global Support Center (GSC) directs the customer to the Gather Utility when problems arise. Customers can also use a self-service website to run this utility and send in relevant information about their problems or issues.

Using a simple command line interface, the Gather utility is a small Java application that can run on any platform to collect various files from the following environments:

- Application Server.
- Web Server.
- Any additional files that the user chooses (SQL Trace files, PeopleCode Trace Files, and so on).

The collected files are placed in a single jar file with `psft.jar` as the default name, in the temp directory. Subsequently, these files are sent to PeopleSoft.

Note: For this utility to work, the supported version of Java (JRE) must be installed on the target machine.

Getting Started

The following files reside in the starting directory:

- `Gather.class`: The main Java class file
- `Helper.class`: This class file is called by `Gather.class`
- `Runnit.bat`: A MS-DOS batch file that is used by Windows users.

UNIX users have to run the Gather utility manually.

- `Vars.sh`: a UNIX shell script.

Gather calls this automatically if the UNIX operating system is detected.

Windows Users

The following steps are used for Windows:

1. Make sure that you have the `PS_HOME` environment variable set.
This saves the user from having to type it in.
2. Go to `PS_HOME\utility`.
3. Type `runnit`.
4. Follow the directions that are on the screen.

UNIX Users

Use the following steps for UNIX:

1. At a command prompt, run the following command where PeopleSoft is installed:

```
../psconfig.sh
```

2. Go to the *PS_HOME/utility* directory.
3. Change permissions for all files:

```
chmod 777 *.*
```

4. Enter the following to start the utility:

```
java -cp .:$CLASSPATH Gather
```

Note: UNIX is case-sensitive – Gather is spelled with a capital G.

5. Follow the instructions that are on the screen.

Environmental Data

On Windows, both the set and netstat commands are invoked with the results copied to a file that is collected. On UNIX, the same thing is done with the env command.

Application Server Data

The following files are collected from the Application Server:

- PSAPPSRV.CFG
- PSAPPSRV.UBB
- LOGS/*.*—this usually includes all application serv/tuxedo logs, dump, and replay files.

This includes all subdirectories under LOGS .

Web Server Data

The gather utility collects numerous files (log files, configuration files, and so on) from each of the supported web servers. If an analyst only asks for a specific file, send that, but make sure to keep the other collected files in case they are needed.

Additional Files

There is always a need to include files that are not on the above list. These can include PeopleCode Trace files, SQL Trace files, SQL output, and so forth. The command line interface allows you to specify any file that you want to be included in the jar file.

QAS Administration

See "Using QAS Administration" (PeopleTools 8.53: Reporting Web Services).

Oracle Resource Management

See "Working With Oracle Consumer Groups" (PeopleTools 8.53: Data Management).

Using Audit Utilities

This section covers the utilities that are used for auditing the system's integrity.

This section discusses how to:

- Use the Record Cross Reference component.
- Perform a system audit.
- Perform database level auditing.

Using the Record Cross Reference Component

Select PeopleTools, Utilities, Audit, Record Cross Reference.

You use the Record Cross Reference component (XREF_PANEL_01) to view where a record is used throughout the application. There are two pages in this page group:

- Pages, Views, Search Records.
- Prompts, Defaults, PeopleCode.

Pages, Views, Search Records

This is a read-only page that shows which Projects, Menus, Pages, and Objects reference a particular record:

Image: Pages, Views, Search Records page

This example illustrates the fields and controls on the Pages, Views, Search Records page. You can find definitions for the fields and controls later on this page.

The screenshot displays the 'Pages, Views, Search Records' page for the record 'QEDERVD_ED_SVCS'. The page is divided into three main sections:

- Record:** QEDERVD_ED_SVCS
- Referenced in Project(s):** A table showing 4 projects: QEDM084, QEDM084, QEDMOALL, and QEDM084.
- Used as a Search Record on:** A table with columns 'Menu Name', 'Item Name', and 'Component'.
- Referenced on Page(s):** A table showing 1 page: QEORDER_TABLE_WF.

Prompts, Defaults, PeopleCode

On the Prompts, Defaults, PeopleCode page, the group boxes list the components that refer to the record.

Image: Prompts, Defaults, PeopleCode page

This example illustrates the fields and controls on the Prompts, Defaults, PeopleCode page. You can find definitions for the fields and controls later on this page.

Record: QEDERVD_ED_SVCS

Used as an Edit Table on: Customize | Find | View All | First 1 of 1 Last

Base Record	Field Name

Used as a Default Table in: Customize | Find | View All | First 1 of 1 Last

Base Record	Field Name

PeopleCode with Fields from this Record: Customize | Find | View All | First 1-8 of 8 Last

PeopleCode Reference Name	PeopleCode Fieldname	PeopleCode Recname	PeopleCode Type
QE_TOTAL_PRICE	QE_PRICE	QEORDER_DTL	FieldChange
QE_GRAND_TOTAL	QE_PRICE	QEORDER_DTL	FieldChange
QE_GRAND_TOTAL	QE_PRICE	QEORDER_DTL	RowInit
QE_GRAND_TOTAL	QE_QTY	QEORDER_DTL	FieldChange
QE_GRAND_TOTAL	QE_PRICE	QEORDER_DTL	RowDelete
QE_TOTAL_PRICE	QE_QTY	QEORDER_DTL	FieldChange
QE_TOTAL_PRICE	QE_PRICE	QEORDER_DTL	RowDelete
QE_TOTAL_PRICE	QE_PRICE	QEORDER_DTL	RowInit

PeopleCode referring to this: Customize | Find | View All | First 1 of 1 Last

PeopleCode Recname	PeopleCode Fieldname	PeopleCode Type

Used as an Edit Table on

Lists pages that use the record as edit table.

Used as a Default Table in

Lists pages that use the record as a default table.

PeopleCode with Fields from this Record

Shows where fields from this record are used in PeopleCode.

PeopleCode referring to this

Shows all PeopleCode that references this record.

Performing a System Audit

The System Audit (SYSAUDIT) utility is documented in the Data Management PeopleBook.

Related Links

"Running SYSAUDIT" (PeopleTools 8.53: Data Management)

Performing Database Level Auditing

This utility is used to support database level auditing features, and is documented in the Data Management PeopleBook.

Related Links

"Understanding Database Level Auditing" (PeopleTools 8.53: Data Management)

Using Debug Utilities

This section discusses how to:

- Use the PeopleTools Test Utilities page.
- Use the Trace PeopleCode utility.
- Use the Trace SQL utility.

Note: The Trace Page / Trace Panel page is no longer actively used or maintained.

Using the PeopleTools Test Utilities Page

Select PeopleTools, Utilities, Debug, PeopleTools Test Utilities to access the PeopleTools Test Utilities page:

Image: PeopleTools Test Utilities

This example illustrates the fields and controls on the PeopleTools Test Utilities. You can find definitions for the fields and controls later on this page.

The screenshot displays the 'PeopleTools Test Utilities' page with the following sections:

- Remote Call Test:** A yellow button labeled 'Test'.
- Interlink Test:** A yellow button labeled 'Interlink Test'.
- PeopleCode/Java Test:** A section containing two buttons: 'Test 1' (under 'Derived Class File') and 'Test 2' (under 'External Class File'). To the right of these buttons is a large empty rectangular box.
- File Attachment Test:** A section with a text input field labeled 'URL Id or String URL:'. Below the field is an example: 'Example: FTP://YourFTPUser:YourFTPPassword@YourComputerName/YourDirectoryPath OR URL.YourURLID'. At the bottom of this section is a yellow button labeled 'Attach'.

Remote Call Test

You use the Remote Call Test button to test the Remote Call configuration.

Interlink Test

Used to test Business Interlinks. (Business Interlinks is a deprecated product).

PeopleCode/Java Test

The Derived Class File button tests Java-PeopleCode integration. It tests to see that Java is being executed correctly through PeopleCode.

The External Class File button tests Java PeopleCode integration.

File Attachment Test

This enables you to test the file attachment PeopleCode functions with your file storage location. Enter the full path and password for the test file or a URL identifier (`URL.URL_ID`).

For example:

ftp://user01:password@ftpserver.oracle.com/myfiles

or

URL.FILEDB

Click Attach to attach the file. Once the file is attached, you can View, Detach, and Delete the file.

The file attachment PeopleCode functions and their usage within PeopleSoft applications are documented in detail in the PeopleCode Developer's Guide.

When testing file attachment features, keep these usage stipulations in mind:

- Using a URL Identifier is mandatory only for FTPS, SFTP, HTTP, and HTTPS attachment protocols.
- FTP, Record, and File attachment protocols, can use the URL Identifier, *or* they can be specified as a plain string.

Note: The File Attachment Test does not allow .JSP files to be uploaded or downloaded.

Note: PeopleTools supports uploading, downloading, and deleting empty file attachments. An empty file refers to a file that exists, but has a size of zero bytes.

Note: The File Attachment Test utility enables only a single file to be uploaded at a time, because the underlying PeopleCode function is AddAttachment. PeopleTools supports uploading multiple files at a time by using the MAddAttachment PeopleCode function.

Note: For the file attachment functionality, in specifying the URL for the FTP server, the FTP server's machine name can be more than 30 characters. The length of the full URL is limited to 120 characters.

See "Understanding the File Attachment Functions" (PeopleTools 8.53: PeopleCode Developer's Guide)[URL Maintenance](#).

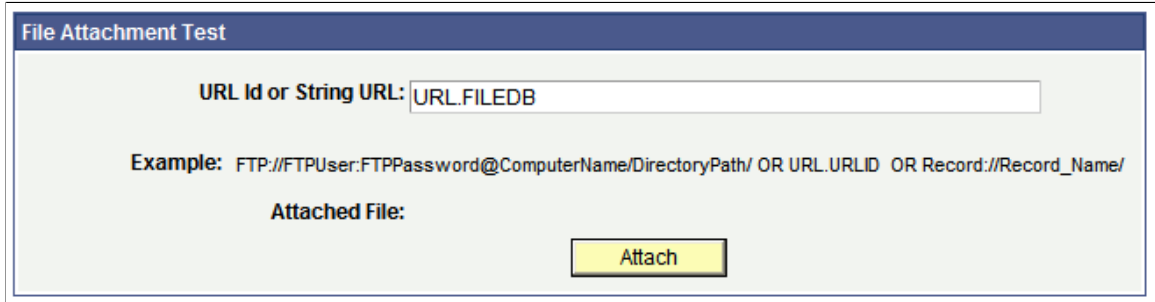
Example: File Attachment Test

The following shows the process of the file attachment test:

1. Enter a string URL or a URLID defined in the URL Maintenance utility, and click Attach.

Image: File Attachment Test example: Entering URL ID

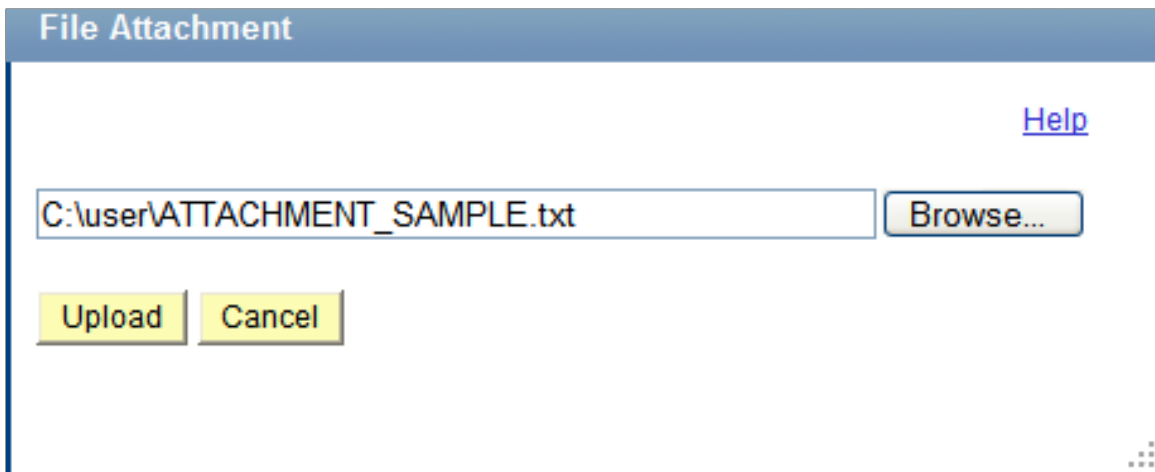
This example illustrates the fields and controls on the File Attachment Test example: Entering URL ID. You can find definitions for the fields and controls later on this page.



2. Browse to the location of the file to attach, and click Upload

Image: File Attachment dialog box

This example illustrates the fields and controls on the File Attachment dialog box. You can find definitions for the fields and controls later on this page.



3. Confirm that the file has been stored in the appropriate repository (attached).

Image: File Attachment Test: Confirming file attachment

This example illustrates the fields and controls on the File Attachment Test: Confirming file attachment. You can find definitions for the fields and controls later on this page.

Note: Once the file has been uploaded, you can View, Delete, or Detach the file.

Related Links

"Understanding the File Attachment Functions" (PeopleTools 8.53: PeopleCode Developer's Guide)
[URL Maintenance](#)

Replay Appserver Crash

See [Configuring the Application Server to Handle Cache Files and Replay Files](#).

Using the Trace PeopleCode Utility

The Trace PeopleCode utility is discussed elsewhere in this PeopleBook.

Related Links

[Setting Up the PeopleCode Debugger](#)
[Configuring PeopleCode Trace](#)

Using the Trace SQL Utility

The Trace SQL utility is discussed elsewhere in this PeopleBook.

Related Links

[Configuring SQL Trace](#)

Using International Utilities

The following sections cover the utilities that you use in globalization efforts.

This section discusses how to:

- Set international preferences.
- Set process field size.
- Administer time zones.
- Manage languages.

Setting International Preferences

Select PeopleTools, Utilities, International, Preferences.

Used to override the language that you select when you sign in to the database.

Language Preference

Use the International Preferences page to temporarily change the session's language preference that was specified during signon. This change lasts until you exit the PeopleSoft session or change the language preference again. Only languages that are enabled on the Languages page are available for selection.

Related Links

"Changing the Session Language While Signed In" (PeopleTools 8.53: Global Technology)

Setting Process Field Size

Select PeopleTools, Utilities, International, Process Field Size.

If you process currency values that require large numbers, such as Italian lira, that require fields longer than those that are included in the standard application, you can use the International Field Size page to expand amount fields throughout the application.

After you create or select a run control ID, set the appropriate lengths for a list of fields, then click the Run button to launch the batch program that performs the field size changes.

Field Name

Use the Browse button to select the field name.

Current Field Size

This is a read-only field indicating the current field size as stored in PSDBFIELDS.

Field Size - International

Enter the field size to expand (or contract) the field size for foreign fields.

Related Links

"Resizing Currency Fields by Using the International Field Size Utility" (PeopleTools 8.53: Global Technology)

Administering Time Zones

This utility is extensively documented in the PeopleTools Global Technology PeopleBook.

Related Links

"Understanding Time Zones" (PeopleTools 8.53: Global Technology)

Managing Languages

Select PeopleTools, Utilities, International, Languages to access the Manage Installed Languages page.

Image: Manage Installed Languages page

This example illustrates the fields and controls on the Manage Installed Languages page. You can find definitions for the fields and controls later on this page.

Manage Installed Languages										
Language										
*Language Code	Language	Enabled	Directionality	*ISO Locale	*Default Character Set	*Verify Locale Mapping	Spell Check Language	*Windows Character Set	*Verify Character Set	
1 ARA	Arabic	<input type="checkbox"/>	Right-to-Left	ar	ISO_8859-6	uni	US and UK English	CP1256	UTF-8	
2 BUL	Bulgarian	<input type="checkbox"/>	Left-to-Right	bg	ISO_8859-5	uni	US and UK English	CP1251	UTF-8	
3 CFR	Canadian French	<input type="checkbox"/>	Left-to-Right	fr-ca	ISO_8859-1	frenchv	French	CP1252	CP1252	
4 CRO	Croatian	<input type="checkbox"/>	Left-to-Right	hr	ISO_8859-2	uni	US and UK English	CP1250	UTF-8	
5 CZE	Czech	<input type="checkbox"/>	Left-to-Right	cs	ISO_8859-2	uni	Czech	CP1250	UTF-8	
6 DAN	Danish	<input type="checkbox"/>	Left-to-Right	da	ISO_8859-1	danishv	Danish	CP1252	CP1252	
7 DUT	Dutch	<input checked="" type="checkbox"/>	Left-to-Right	nl	ISO_8859-1	dutchv	Dutch	CP1252	CP1252	
8 ENG	English	<input checked="" type="checkbox"/>	Left-to-Right	en	ISO_8859-1	englishv	US and UK English	CP1252	CP1252	
9 ESP	Spanish	<input type="checkbox"/>	Left-to-Right	es	ISO_8859-1	spanishv	Spanish	CP1252	CP1252	
10 FIN	Finnish	<input type="checkbox"/>	Left-to-Right	fi	ISO_8859-1	finnishv	Finnish	CP1252	CP1252	
11 FRA	French	<input type="checkbox"/>	Left-to-Right	fr	ISO_8859-1	frenchv	French	CP1252	CP1252	
12 GER	German	<input type="checkbox"/>	Left-to-Right	de	ISO_8859-1	germanv	German (new)	CP1252	CP1252	
13 GRK	Greek	<input type="checkbox"/>	Left-to-Right	el	ISO_8859-7	uni	Greek	CP1253	UTF-8	
14 HEB	Hebrew	<input type="checkbox"/>	Right-to-Left	he	ISO_8859-8	uni	US and UK English	CP1255	UTF-8	
15 HUN	Hungarian	<input type="checkbox"/>	Left-to-Right	hu	ISO_8859-2	uni	Hungarian	CP1250	UTF-8	
16 ITA	Italian	<input type="checkbox"/>	Left-to-Right	it	ISO_8859-1	italianv	Italian	CP1252	CP1252	
17 JPN	Japanese	<input type="checkbox"/>	Left-to-Right	ja	Shift_JIS	japanv	US and UK English	CP932	Shift_JIS	
18 KOR	Korean	<input type="checkbox"/>	Left-to-Right	ko	CP949	koreav	US and UK English	CP949	CP949	
19 MAY	Bahasa Malay	<input type="checkbox"/>	Left-to-Right	ms	ISO_8859-1	uni	US and UK English	CP1252	UTF-8	

Use this page as a central utility to manage language information for the currently enabled languages.

Language Code

Use the search prompt to select the PeopleSoft language code from the PSXLATITEM table. The language description appears to the right of the code field.

Enabled

When you select this check box, PeopleSoft Internet Architecture enables you to log in with the language.

ISO Locale

Use the search prompt to select the ISO locale code from the PSLOCALEDEFN table. Consists of an ISO 639 language code, optionally followed by an ISO 3166 country code.

Default Character Set

Use the search prompt to select the character set from the PSCHARSETS table. Determines the default encoding for input and output files.

Verity Locale Mapping

Select the Verity locale code from the PSVERITYLOCALE table. Determines the locale to use for building search collections and searching data.

Spell Check Language	Select the spell check language from the PSXLATITEM table. This enables you to select the language of the spell check dictionary that is associated with a given language code.
Windows Character Set	Select the Microsoft codepage that is associated with the given language. This defines the codepage to use with certain Microsoft applications.
Verity Character Set	Select the character set that the Verity engine uses for its internal encoding in the given language. You should not modify the value in this field under normal circumstances.

Related Links

"Managing Languages in the PSLANGUAGES Table" (PeopleTools 8.53: Global Technology)

Using Optimization Utilities

The Optimization utilities are documented extensively in the Optimization Framework PeopleBook.

Related Links

PeopleSoft Optimization Framework

Using PeopleSoft Ping

The PeopleSoft Ping utility collects timestamps by sending a specific page to different tiers of the PeopleSoft system, starting at the browser, then going to the web server, the application server, the database and back. The timestamps that are collected are total time elapsed for the round trip, and arrival and departure time at each of the tiers.

To use the PeopleSoft Ping feature, select PeopleTools, Utilities, PeopleSoft Ping.

Image: PeopleSoft Ping page

This example illustrates the fields and controls on the PeopleSoft Ping page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'PeopleSoft Ping' interface. At the top, there is a 'Test Case Identifier' field containing 'PTDOCTEST' and a stopwatch icon. Below it, 'Repeat Time Interval (Seconds)' is set to '20' and 'Counter' is '3'. There are 'Run' and 'Stop' buttons. Performance metrics are displayed in green boxes: Total Time (0.578), Browser Time (0.36), WebServer Time (0.03), AppServer Time (0.095), and Database Time (0.093). A 'Sample Data' section includes a 'View All' link and a table with columns for Seq#, Sequence Text, and Description. A sub-table shows detailed data for sequence 10001, with columns for sequence number and components A through E.

Seq#	Sequence Text	Description
10001	010001	Test 010001

Sample Data	Seq#	11	101	102	103	104	105	A	B	C	D	E
10001	11	101	102	103	104	105	A	B	C	D	E	
10001	12	101	102	103	104	105	A	B	C	D	E	
10001	13	101	102	103	104	105	A	B	C	D	E	
10001	14	101	102	103	104	105	A	B	C	D	E	
10001	15	101	102	103	104	105	A	B	C	D	E	
10001	16	101	102	103	104	105	A	B	C	D	E	
10001	17	101	102	103	104	105	A	B	C	D	E	
10001	18	101	102	103	104	105	A	B	C	D	E	
10001	19	101	102	103	104	105	A	B	C	D	E	
10001	20	101	102	103	104	105	A	B	C	D	E	

Enter a Test Case Identifier to uniquely group each set of metrics. For Repeat Time Interval , enter an increment for the ping to run. To avoid creating unnecessary traffic and overhead to the PeopleSoft system, set the Repeat Time Interval to a relatively high value, such as 600 to 1800 seconds, during normal operations. You may need to increase the Session Timeout value accordingly.

PeopleSoft Ping Chart

Select PeopleTools, Utilities, PeopleSoft Ping, PeopleSoft Ping Chart.

PeopleSoft Ping includes a charting utility to zoom in to a specific time interval from the ping test.

You can change the displayed time interval to a subset of the full ping test period. Edit the start time and end time values, and click Redraw to refresh the chart display with the new time interval.

Click Query Viewer to query the database for the ping data. A new browser window opens, displaying the ping data for the full test period in a table.

PeopleSoft Ping Delete

To delete a ping page test case:

1. Select PeopleTools, Utilities, PeopleSoft Ping, PeopleSoft Ping Delete.

The Delete page lists the current test case identifiers.

2. Select the check box next to the identifier(s) you want to delete.
3. Click Delete.

PeopleSoft Ping Options

Select PeopleTools, Utilities, PeopleSoft Ping, PeopleSoft Ping Options.

The PeopleSoft Ping Options page enables you to set targets for each tier as well as overall completion time. If the ping process exceeds your targets, this affects the color-coding on the PeopleSoft Ping interface, using green, yellow, and red. Green is any time under the yellow and red targets, yellow is any time over the yellow target yet under red, and red is anything exceeding the red target. The metrics you enter depend on the typical conditions and expectations at your site.

Chapter 12

Tracing, Logging, and Debugging

Working with PeopleSoft Server Process Logs

For each PeopleSoft server process, such as PSAPPSRV, the system writes an individual log file to your specified logs directory. Depending on the logging level, these log files can contain varying amounts of information. Reviewing these log files regularly can help you to identify potential issues or system trends, and they are a valuable source of information when troubleshooting. To understand the log information, it is helpful to be able to identify the various log fields that the system writes to an individual line in the log file.

A typical log file contains these log fields:

[Server Process]:[Operating System Process ID] [Service Request Number] [Timestamp] [Log Level] [Message]

The following table provides examples of content for each trace field.

Server Process:PID	Service Request	Timestamp *(YYYY-MM-DDTHH-MM-SS-sss)	Log Level	Message
PSAPPSRV.5648	(993)	[2012-12-12T11:54:23.778]	(3)	Authenticate User not needed. New User Id/Password same as current.
PSAPPSRV.5648	(1000)	[2012-12-12T11:54:26.198]	(0)	Recycling server after 1000 services
PSAPPSRV.28708	(0)	[2012-12-12T11:54:35.514]	(0)	Cache Directory being used: D:\PT_SERVER\appserv \T1B85301\CACHE\PSAPPSRV_1\
PSAPPSRV.28708	(0)	[2012-12-12T11:54:38.339]	(0)	Server started

* Where YYYY = year, MM = month, DD = day, HH = hour, MM = minute, SS = seconds, and sss = milliseconds.

You can select how you prefer to separate the trace fields using the [Domain Settings] LogFieldSeparator setting.

Related Links

[Domain Settings](#)

[Trace Options](#)

[Specifying Trace Settings](#)

Setting Up the PeopleCode Debugger

This section discusses how to:

- Debug for a two-tier connection.
- Debug for a three-tier connection.
- Use the PeopleCode Debugger.

Note: PeopleCode debugging is not supported on z/OS.

You can debug the PeopleCode program configurations of a two-tier connection to the database or a three-tier connection to the database.

Note: When you debug PeopleCode with an application server, Application Designer should be run in three-tier mode. PeopleCode debugging by using a two-tier PSIDE and an application server is not supported on multi-homed (multiple Internet Protocol address) workstations.

Note: Application Designer (PSIDE.EXE) will not connect to the PeopleCode Debugger through a firewall using Network Address Translation (NAT).

Debugging for a Two-Tier Connection

Debugging in two-tier connections involves connecting directly to the database, not through the application server. Use this method to debug two-tier Windows applications.

Note: By default, the port number that the PeopleCode debugger uses is 9500. Unless this port number is being used by another application, you do not need to alter any environment settings, and after you sign on to the database, you are able to debug PeopleCode.

If you need to change the PeopleCode Debugger port settings, complete the following procedure.

To change the default PSDBGSRV listener port number:

1. Open PeopleSoft Configuration Manager.
2. Select the Trace tab.
3. Locate the PeopleCode Debugger section, and make sure that the default value for the Local PSDBGSRV Listener Port is suitable for the system.

For example, make sure that no other applications are configured to listen on the default port number (9500). If so, you must assign a port number that is not being used.

Note: If you're using a personal firewall, you must configure it to enable data packets to flow through the PSDBGSRV listener port. If you can't configure your firewall appropriately, you must shut it down while performing PeopleCode debugging.

Debugging for a Three-Tier Connection

Use three-tier debugging to debug three-tier Windows applications and PeopleSoft Internet Architecture (PIA) applications. For three-tier debugging, use PSADMIN to ensure that the following items are set:

- The appropriate PSDBGSRV listener port is specified in the PeopleCode Debugger section of PSADMIN.
- At least two PSAPPSRV processes are configured to boot in the domain with the service timeout parameter set to zero.
- Enter y for yes at the Enable PSDBGSRV Server Process prompt at the end of the PSADMIN interface.

Debugging on a Multi-Homed System

If you're debugging on a multi-homed (multiple IP address) system, you must explicitly specify an IP address in the Workstation Listener section of the PSADMIN configuration, rather than %PS_MACH %. The address you specify must be one by which the application server identifies the machine on which you're doing the debugging. This ensures that the workstation listener monitors requests from the correct location.

See [Workstation Listener Options](#).

Setting the PSDBGSRV Listener Port

In the PeopleCode Debugger section of PSADMIN make sure that the value that is assigned to the PSDBGSRV listener port is not already in use by another application or listener on the application server. The default value is 9500. If the default is not acceptable, assign a suitable value to the parameter. If it is acceptable, no changes are required.

For example,

```
Values for config section - PeopleCode Debugger
  PSDBGSRV Listener Port=9500
```

```
Do you want to change any values (y/n)? [n]:
```

Consider the following when debugging PeopleCode:

- If multiple application server domains are running on a single, physical machine, each domain needs to use different debugging port numbers.

Otherwise, there is contention for the PSDBGSRV listener port value. This is the same principle that requires each application server domain on a server to have unique workstation listener port numbers.

- When you are not debugging, turn off (set to 0) the Enable Debugging parameter.

The debugging mode results in an unavoidable amount of overhead, which can degrade performance.

- Regarding performance, do not perform debugging on a production domain.

Debugging should be performed on a designated testing domain only.

Enabling Multiple PSAPPSRV Server Processes

The minimum requirements for PeopleCode debugging are:

- Two PSAPPSRV server processes are configured to boot in the domain.
- The Service Timeout value in the PSAPPSRV configuration section must be set to 0.

For the debugger to work, it has to run in parallel with the application that it's debugging. Suppose that the domain has only one PSAPPSRV server process running. In this case, the PSAPPSRV can process the requests of only one component at a time, and therefore debugging is not possible. Debugging involves two items, the debugger (PSDBGSRV) and the PSAPPSRV server process that is running the application PeopleCode.

Provided that you have two PSAPPSRV server processes configured; one PSAPPSRV handles the debugger program, while the other handles the application that you're stepping through with the debugger. In this case, the two programs run in parallel, which enables interactive debugging.

The configuration templates that PeopleSoft delivers all have at least two PSAPPSRV processes. However, if you are using a custom template, make sure that you configure the domain to start two PSAPPSRV processes prior to debugging. To do this, in PSADMIN set the Min Instances parameter in the PSAPPSRV section to 2.

The following example shows a sample PSAPPSRV section properly configured for debugging PeopleCode:

```
Values for config section - PSAPPSRV
  Min Instances=2
  Max Instances=2
  Service Timeout=0
  Recycle Count=0
  Allowed Consec Service Failures=0
  Max Fetch Size=5000
```

Do you want to change any values (y/n)? [n]:

When configuring the PeopleCode debugger:

- PeopleSoft recommends using the Developer configuration template because this template, by default, provides two PSAPPSRV server processes and has service timeout set to zero.
- PeopleSoft recommends using a simple configuration where you are assured that the server that Application Designer connects to is the same server that the application you are debugging is running on.

Note: If you do not set the settings for PSAPPSRV correctly (at least two PSAPPSRV processes), PSADMIN automatically sets these values to comply with the minimum requirements when you enable PeopleCode Debugging (as discussed in the next section).

Note: When you enable the PeopleCode Debugger (PSDBGSRV), service timeout settings for server processes are set to zero (0) by the system, overriding any previous settings you may have made in PSADMIN. For example, if the service timeout settings for the PSAPPSRV service process are set to 300 prior to enabling the debugger, after you enable the debugger, the service timeout value will be zero (0). After using the debugger, you need to reset your service timeout settings to the desired values. Before enabling the debugger, it is recommended that you make a backup of your configuration file or make note of your service timeout settings.

Requesting a PSDBGSRV Server Process

After you specify the settings by using PSADMIN, the system prompts you with a series of options, such as setting up messaging server processes, enabling Jolt, and so on.

When you're prompted to enable the PSDBGSRV, enter y. Y appears in the Developer template by default.

Using the PeopleCode Debugger

After the system is configured properly, using the PeopleCode debugger is just a matter of signing on to the PeopleSoft system and entering the PeopleCode Debugger mode in Application Designer.

Note: You must use a unique user ID when you're performing PeopleCode debugging, as opposed to using a shared user ID, such as those that PeopleSoft delivers, for example QEDMO, PS, or VP1. Shared IDs are likely to be used by others that are connecting to the same test database, which can affect debugging.

Configuring PeopleCode Trace

Select PeopleTools, Utilities, Debug, Trace PeopleCode to access the Trace PeopleCode page.

You use this page to change the PeopleCode tracing options while online. This page does not affect trace options that are set in PeopleSoft Configuration Manager. Use Trace PeopleCode to create a file displaying information about PeopleCode programs processed from the time that you start the trace.

Trace Evaluator Instructions	Select to show a line-by-line trace of the program
List Evaluator Program	Select to show the code of the PeopleCode program.
Show Assignments to Variables	Select to show variable assignments.
Show Fetched Values	Select to show values that are from PeopleCode Fetch call.
Show Stack	Select to display the PeopleCode evaluator's stack after each PeopleCode (internal) instruction.
Trace Start of Programs	Select to show the starting and ending points of the program.
Trace External Function Calls	Select to show calls to application written functions.
Trace Internal Function Calls	Select to show the calls to PeopleTools built-in function calls.
Show Parameter Values	Select to show function parameter values.
Show Return Parameter Values	Select to show function return parameter values.
Show Each	Select to trace each statement in the program.

Note: The Trace PeopleCode Utility decreases system performance because of the overhead that occurs during the monitoring and recording of all PeopleCode actions.

The check boxes on this page correspond to the options on the Trace tab in Configuration Manager. However, the selections that appear on this page do not necessarily reflect those that are made in Configuration Manager. While the Configuration Manager settings are stored in the Windows registry and used at each signon, the settings in the Utilities page only apply to the current online session, and, once set, they override the Configuration Manager's settings.

The benefit of using this page to control PeopleCode tracing is that you can turn it on and off without having to restart PeopleTools, and without resetting the Configuration Manager settings. Keep in mind, though, your selections are not enabled until you save the page.

To enable/disable PeopleCode tracing while on line

1. Select PeopleTools, Utilities, Debug, Trace PeopleCode.

The Trace PeopleCode page appears.

2. Select/deselect the desired Options.
3. Save the page.

If you selected any of the check boxes, the system starts writing to the trace file.

Related Links

[Specifying Trace Settings](#)

"SetTracePC" (PeopleTools 8.53: PeopleCode Language Reference)

Configuring SQL Trace

Select PeopleTools, Utilities, Debug, Trace SQL to access the Trace SQL page.

You use this page to change the SQL tracing options while you're online. Your Configuration Manager settings are not affected:

Trace SQL Statement	Select to show the SQL statement.
Trace SQL Bind	Select to show bind values for SQL statements that have parameter markers.
Trace SQL Cursor	Select to show connect, disconnect, commit and rollback calls.
Trace SQL Fetch	Select to show fetch call for Select Statement.
Trace SQL API	Select to show other API calls (Execute, Describe, and so on.)
Trace SQL Set Select Buffer	Select to show Binds for Select columns.
Trace SQL -- Database Level	Select to specify low-level tracing at the database API (ODBC, ct-lib, and so on.)
Trace SQL -- Manager Level	Select to show calls for Cache calls.

The check boxes on the Trace SQL page correspond to options on the Trace tab in the Configuration Manager. However, the selections that appear on this page do not necessarily reflect those that are made in the Configuration Manager. The displayed page selections are not enabled until you save the page.

To enable or disable SQL tracing while online:

1. Select or deselect the desired trace options.
2. Save the page.

If you select any of the check boxes, the system starts writing to the trace file.

Related Links

[Specifying Trace Settings](#)

"SetTraceSQL" (PeopleTools 8.53: PeopleCode Language Reference)

Enabling IDDA Logging

This section contains an overview and discusses how to:

- Enable IDDA logging.
- Work with IDDA functional categories.
- Configure logging options.
- Viewing log results.

Understanding IDDA Logging

System administrators typically use numerous monitoring and logging utilities to diagnose system issues and internet applications. Such logging utilities include:

- TCPMON
- ieHttpHeaders
- Access log
- Heap dump
- Thread dump

All of these utilities provide different information, but each helps an administrator gain insight and detailed information related to specific system behavior. PeopleSoft provides a variety of logging and tracing mechanisms as well, enabling you to gather vital information at various levels of the architecture, including database server, application server, web server, and so on.

The PeopleSoft Instrumented Development Diagnostic Aid (IDDA) logger, enables you to gather specific information about various areas within the PeopleSoft Internet Architecture and PeopleSoft Interaction Hub, including:

- PeopleSoft Internet Architecture processing.
- Integration Broker.
- Reporting, Report Repository.
- Portal.
- Caching.
- File processing.
- Security, authentication.
- Performance Monitor.
- WSRP.
- Jolt.

Typically, administrators only run IDDA traces when instructed to do so by Oracle support contacts, looking for specific information. However, it can also be a useful troubleshooting tool.

Note: When ever tracing or logging is enabled, you should always expect a certain degree of performance degradation as the system incorporates the overhead involved with logging. Disable logging when you finish troubleshooting.

Enabling IDDA Logging

To enable IDDA logging:

1. Select PeopleTools, Web Profile, Web Profile Configuration, and open the current web profile.
2. Select the Custom Properties page.
3. Add a new row, and enter these values:

Column	Value
Property Name	<i>IDDA</i>
Validation Type	<i>Number</i>
Property Value	The sum of the bit values of the functional area(s) you want to log. For example, if you wanted to log PIA (1) and Portal (8), you enter 9.

4. Click Save.
5. Restart the PeopleSoft site.

Working with IDDA Functional Categories

The IDDA logger, gathers information for a wide range of technology within the PeopleSoft Internet Architecture. Depending on the needs of your troubleshooting, you can select different areas to log.

The different areas of technology are referred to as functional categories in the context of the IDDA logger. Each functional category is assigned a bit value (in a 32-bit space). When you enable IDDA logging, you enter specific bit values *or* the sum of the bit values of different areas.

The IDDA functional categories are:

Bit Value	Functional Category
1	PeopleSoft Internet Architecture
2	Integration Broker
4	Report repository
8	Portal
16	Web server caching
32	File processing (attachments)
64	Authentication
128	Performance Monitor
256	Web Services for Remote Portlets (WSRP)
512	Jolt

The type of information included in the log message differs per category. The log message is free format and can display any information that helps troubleshooting for that particular area, such as error codes, exception messages, and so on.

Configuring Logging Options

Once you've enabled IDDA logging, you can modify configuration options in the logging.properties file, which you can find in these locations:

Web Server	Location
Oracle WebLogic	<i>PS_HOME</i> /webserv/ <i>domain name</i> /applications/peoplesoft
IBM WebSphere	<i>PS_HOME</i> /webserv/ <i>profile name</i> /installedApps/ <i>node&cell name</i> /peoplesoft.ear/

The relevant configuration properties are:

Property	Description
.level	<p>Sets the global logging level.</p> <ul style="list-style-type: none"> • OFF: Disables all IDDA logging. • SEVERE: Displays server issues, such as premature termination and failure. • WARNING: Displays less severe issues, such as configuration issues. • INFO: Displays basic operational information, such as starting and stopping. • FINE, FINER, FINEST: Displays internal non-critical informational messages. • ALL: Enables all logging levels. <p>Default value is <i>INFO</i>.</p> <p>WARNING and SEVERE messages are always logged, unless IDDA logging is set to OFF.</p> <p>INFO or FINE, FINER, FINEST messages are only logged if the IDDA value is greater than zero.</p>
java.util.logging.FileHandler.pattern	<p>Sets the naming style for the log file and the output directory location.</p> <p>Default value is: <i>./servers/PIA/logs/PIA_servlets%u.log</i></p> <p>If you are running a multi-server, distributed environment (for clustering and failover purposes), the default value for <code>java.util.logging.FileHandler.pattern</code> is not applicable. You must set this property to point to a valid location in order to collect logging messages for a particular server.</p> <p>Logging output is tab-delimited.</p>
java.util.logging.FileHandler.limit	<p>Limits the size of the output file in bytes. When set to 0, there is no size limit.</p> <p>Default value is <i>0</i>.</p>
java.util.logging.FileHandler.count	<p>Sets the total number of log files. Default value is <i>5</i>.</p> <p>If the value is greater than 1, the system writes to a rotating set of log files. When the file reaches a given size limit or the web server restarts, the system ends writing to the current file and begins writing to a new file. The system names each file in the sequence it is saved by adding "0", "1", "2", (and so on) into the base filename.</p>

Viewing IDDA Logging Output

The output log files contain:

- Machine header information

- Log message information

Working with Machine Header Information

Each log file displays the following machine environment information at the top of each log file.

Header Entry	Description
Timestamp	Date, time, time zone.
PeopleTools Release	PeopleTools version number.
os.name	Operating system.
os.version	Operating system version.
os.arch	Operating system architecture (such as x86 for Windows servers).
java.version	Java version number.
java.vendor	Java vendor.
java.vm.info	Mode of the Java virtual machine (JVM), such as "compiled mode."
java.home	Java installation directory.
user.dir	The value of the user.dir property.
java.class.path	CLASSPATH setting on the server.
.level	Logging level, such as INFO, SEVERE, WARNING, and so on.
java.util.logging.FileHandler.pattern	Logging output directory and file naming convention.
Trace	Current IDDA functional group(s) being logged (integer reflecting the sum of the bit values assigned to each group).

Working with Log Message Information

Each entry in the log file contains this information.

Log Message Data	Description
Timestamp	Date, Time and Time zone. For example, 2/7/09 4:00:39 PM PST

Log Message Data	Description
Sequence	Tracks the sequential order of the messages. Starting at 1, the system increments by 1 per each log message.
Thread ID	Java thread ID.
Logging group	Bit value representing the functional grouping, as in 1 for PeopleSoft Internet Architecture.
Source class	Class name of where the message is logged. For example, <code>psft.pt8.auth.PSAuthenticator</code>
Source method	Method name of where the message is logged. For example, <code>SetCookie</code>
Log message	The actual log message.

Viewing Log Contents

The log file is a tab-delimited text file and can be opened in any standard text editor, such as Notepad or Textpad. Because the output is tab-delimited, you can also use a spreadsheet application, such as Microsoft Excel, for more efficient analysis. For example, viewing the output within a spreadsheet enables you to apply filters to columns and only view specific log messages, which can be helpful with large files.

Working with Jolt Configuration Options

Configuring Jolt Failover and Load Balancing

This section discusses how to:

- Configure weighted load balancing.
- Configure Jolt failover.

Configuring Weighted Load Balancing

With weighted load balancing, you can set the “weight” of the load, or amount of requests, being directed to a particular server. Weight values are integers 1–10, with 1 being low and 10 being a heavy load.

Servers that can handle extra work can take heavy loads, while servers that are either less powerful or are being used in other capacities can take lower loads. You specify weighted load balancing by modifying the server values in the `pserver` property in the PeopleSoft Internet Architecture configuration.properties file, using the following format.

```
pserver=Host1:Port1#Wt,Host2:Port2#Wt
```

For example,

```
pserver=appserver1:9000#3,appserver2:9010#1
```

In this case, `appserver1` would receive 3x more requests than `appserver2`.

Configuring Jolt Failover

You can also specify strict failover assignments with weighted load balancing, with the following options:

- Strict failover with weighted backup.
- Strict failover with sequential backup.

You add the failover string within curly brackets at the end of the server entry.

```
pserver=<host>:<port>#wt{failover servers}
```

With the failover string, you can set weighted backup by separating failover server with a comma (,).

```
pserver=Host1:Port1#Wt{Host3:Port3#Wt,Host4:Port4#Wt},Host2:Port2#Wt
```

In this case, Host 3 and Host 4 are failover servers when Host 1 is down. You can assign weighted load balancing to the backup servers just as you would a primary server.

You can also set a sequential backup with your failover string. To set sequential backup, you separate multiple backup servers using a semicolon (;).

```
psserver=Host1:Port1#Wt{Host3:Port3;Host4:Port4},Host2:Port2#Wt
```

In this case, the system assigns Host 4 the requests when both Hosts 1 and 3 are down.

Configuring Dynamic Reloading of the Application Server Connection String

The `DynamicConfigReload` property enables the PIA domain to reload the `psserver` string (application server connect string) dynamically. That is, after modifying the `psserver` string to include new servers, replace existing servers, or remove existing servers from the string, you do not need to restart the PIA domain. With `DynamicConfigReload` enabled, the system automatically detects when the `psserver` string has been updated and ensures the PeopleSoft site maintains seamless connections with the application server domains listed in the `psserver` string.

To enable dynamic `psserver` reloading, set `DynamicConfigReload` to 1. For example:

```
DynamicConfigReload=1
```

To disable dynamic `psserver` reloading, set the property to 0.

If dynamic `psserver` reloading is disabled, then you need to restart the PIA domain manually for it to recognize any changes made to the `psserver` string.

Configuring Jolt Session Pooling

Jolt session pooling is enabled by default. Jolt session pooling enables web server connections to be shared between user sessions, which reduces the usage of system resources, such as threads and file descriptors.

You control session pooling by modifying the `joltPooling` property in the `configuration.properties` file per site.

```
joltPooling=true
```

To enable Jolt session pooling, set the property value to *true*, and to disable Jolt session pooling set the property value to *false*.

Configuring Parallel Pagelet Loading

You can use the `parallelLoading` property in the `configuration.properties` file to control how your portal loads pagelets. Pagelets can be loaded sequentially or in parallel.

Loading pagelets sequentially refers to a single-threaded process, downloading multiple pagelets one at a time. Loading pagelets in parallel refers a multi-threaded process where multiple pagelets can be loaded at the same time. To enable parallel pagelet loading, set the `parallelLoading` property to *true*.

For example:

```
parallelLoading=true
```

Note: To load pagelets in parallel, Jolt pooling must be enabled.

To disable parallel loading of pagelets, set the `parallelLoading` property to false.

Configuring Domain Connection Password

The `DomainConnectionPwd` setting in the [Security] section of the application server domain, sets the domain connection password for all PeopleSoft Internet Architecture Jolt connections to that application server domain. It is not a required password setting, but *if set*, for the Jolt session to connect successfully to the application server domain, the `DomainConnectionPwd` property value in the configuration.properties file needs to match the `DomainConnectionPwd` parameter set for the application server domain.

You can encrypt the password value using the `pscipher` utility. For example:

```
DomainConnectionPwd={V1.1}6O5vpwGxd5o=
```

Related Links

[DomainConnectionPwd](#)

Using Jolt Internet Relay

This section discusses:


- Jolt Internet Relay architecture.
- A Jolt Internet Relay example.
- Implementation considerations.

Note: Jolt Internet Relay was used in previous releases at some sites for specific security requirements. Currently, it is not a widely used feature and should only be implemented in specific circumstances or by suggestion of Oracle support staff.

Jolt Internet Relay Architecture

Image: Jolt Internet Relay Architecture showing the Jolt Relay sending messages through a firewall to the Jolt Relay Adapter

Jolt Internet Relay provides an environment in which the PeopleSoft web server and application server can be further decoupled. This provides greater security at sites where security is of utmost importance. Jolt Internet Relay routes messages from a Jolt client to a Jolt Server Listener (JSL) or Jolt Server Handler (JSH), and eliminates the need for the JSL, JSH, and Tuxedo application to run on the same machine as the web server. Communication takes place between the JRLY and JRAD elements rather than between the Jolt client and JSL/JSH processes. Traditionally an application server domain opens between 2 and 6 ports for such communications. The use of Jolt relay restricts this to one port per domain on the application server machine. This enables an administrator to open just one port on the application server machine. The following diagram illustrates this feature:



Jolt Internet Relay consists of two elements: Jolt Relay (JRLY) and Jolt Relay Adapter (JRAD). It's important to understand the difference between these two elements.

JRLY consists of a standalone program and configuration file; the program runs on the same machine as the web server. JRLY receives Jolt messages from a PeopleSoft web application and routes those messages to JRAD on the application server. It receives the Jolt message through one port, the LISTEN port, and connects to the JRAD by using another port, the CONNECT port. JRLY is sometimes referred to as a front-end relay.

JRAD runs on the same machine as the application server. It's configured on the application server domain as part of the PeopleSoft PSADMIN domain configuration procedure. JRAD listens for JRLY messages on its LISTEN port and transfers the message to the JSL or JSH. JRAD is sometimes referred to as a back-end relay.

Note: Implementing Jolt Relay can impede performance. Always perform testing with typical production system load to ensure it will meet your service level requirements.

Implementation Considerations

Keep the following points in mind as you configure the Jolt Internet Relay components:

- The jrly binary and its corresponding jrly.config file must exist in the same directory. To start multiple Jolt Relays on a machine, copy the jrly binary and jrly.config into each subdirectory, modify the parameters in the jrly.config file, and start Jolt Relay. On Windows, you can define multiple Jolt Relay services on a machine.
- You can start the JRLY process before or after you start JRAD. The JRLY attempts to connect to JRAD on the client request. If the JRLY is unable to connect to the JRAD, the client is denied access and disconnected. The connection will be retried upon the first use of PeopleTools..
- If you're installing Jolt Internet Relay on UNIX and anticipate a large number of concurrent connected clients, increase the file descriptors limit before running the JRLY executable.
- At runtime, if you get the following message:

```
[Fri JUN1 6 20:25:11 1997] JRLY:accept():accept failed, err no: 23, strerror: =>
File table overflow
```

PeopleSoft recommends that you increase the MaxUSERS kernel parameter and regenerate the kernel.

- If you're unable to connect, make sure that you check the following items:
 - Port numbers do not match.

Print out the jrly.cfg file and the psappsrv.cfg file and compare the port numbers that you specified.
 - Make sure that the application server is running.
 - Make sure that JRLY is running.
- Jolt Internet Relay can be installed on an intermediate machine rather than the web server machine if necessary. This extra level of indirection can cause performance degradation.
- Make sure that JRAD is running on the application server and that you configure JRAD using PSADMIN.

Configuring JRLY

Configuring JRLY is identical on UNIX and Windows.

To configure JRLY, navigate to *TUXDIR\udataobj\jolt\relay* and open jrly.config in a text editor.

Important! On UNIX, you can edit this configuration file by using VI or an equivalent editor. However, on Windows, you must edit the file using an editor that preserves the file's UNIX line feeds. WordPad is valid for this purpose, but Notepad is not.

Modify the parameters in the configuration file to reflect the site specifications, as follows:

Parameter	Description
LOGDIR	<p><i>LOGDIR</i> specifies the directory where JRJLY creates access and error log files. This directory must exist; the JRJLY program does not start if it can't find this directory. The path that you specify for LOGDIR should be an absolute path (starting from / on UNIX systems, starting from a drive letter on Windows systems). The JRJLY accepts relative path names, but LOGDIR is relative to the directory from which the JRJLY program is started, unless you specify it as an absolute.</p>
ACCESS_LOG	<p><i>ACCESS_LOG</i> specifies the name of the file where JRJLY records access information. This log file is created in <i>LOGDIR</i>. If the log file already exists, the most recent information is appended to it.</p> <p>This parameter can be any valid file name. Everything after the equals sign (=) to the end of the line is considered as part of the file name, but leading and trailing blanks are ignored.</p> <hr/> <p>Note: If the JRJLY program can't create this file or open it for appending, the program exits.</p>
ERROR_LOG	<p><i>ERROR_LOG</i> specifies the name of the file where JRJLY records error information. This file follows all the rules that apply to the <i>ACCESS_LOG</i> parameter. JRJLY_error_log is created in /tmp.</p>
LISTEN	<p><i>LISTEN</i> specifies the host and port on the current machine (that is, the machine where you're installing Jolt Relay). JRJLY listens for client connections. The following formats are acceptable:</p> <pre>LISTEN=192.9.100.100:9000 LISTEN=//192.9.100.100:9000 LISTEN=sp-ibm02:9000 LISTEN=//sp-ibm02:9000</pre> <p>Specify the port number in decimal; it must match the port number that is specified by the <i>pserver</i> parameter in the configuration.properties file for the PIA web application.</p> <hr/> <p>Note: If a machine has multiple network interfaces, you should use the IP address notation, because specifying the hostname could be ambiguous (the result is OS dependent). If the JRJLY program can't establish a network listening end-point at the host and port specified, it prints an error and exits.</p> <hr/> <p>The hostname that's specified for this parameter must be the name of the host on which the program is running.</p> <hr/> <p>Note: You can create multiple configuration files to run multiple instances of JRJLY. Each configuration file must specify a different port number for this parameter.</p>

Parameter	Description
CONNECT	<p><i>CONNECT</i> specifies the location of the JRAD machine and process port on the application server machine to which the JRLY program connects. A JRLY program communicates only with a single JRAD. The address you specify for this parameter must match the JRAD listener address that's on the application server machine (check the PSAPPSRV.CFG file in <i>PS_CFG_HOME/appserv/domain</i>). The JRAD doesn't have to be running when you start the JRLY program. JRLY attempts to connect to the JRAD when it first starts, and if the JRAD is not available, JRLY tries again whenever a new client connects to it. You can use any of the following formats for this parameter:</p> <pre>CONNECT=192.9.100.100:9100 CONNECT=//207.135.44.91:9105 CONNECT=sp-hp06:9105 CONNECT=//sp-hp06:9105</pre> <hr/> <p>Note: PeopleSoft has found that machine address formats are operating system and environment dependent. If one fails to connect to the application server, try another format.</p>
SOCKETTIMEOUT	<p><i>SOCKETTIMEOUT</i> specifies the duration (in seconds) for which the Jolt Internet Relay Windows service blocks the establishment of new socket connections to allow network activity (new connections, data to be read, closed connections) to complete. It's valid only on Windows machines.</p> <p><i>SOCKETTIMEOUT</i> also affects the Service Control Manager (SCM). When the SCM requests that the service stop, the SCM needs to wait at least the number of seconds specified by this parameter.</p>

Configuring JRAD

The JRLY connect port connects to the JRAD listener port that is specified on the application server machine. JRAD then routes the message to Jolt, either using the JSL for initial connection from a web client, or using the JSH for all subsequent connections from a web client. The return message follows the same path in reverse.

To configure JRAD:

1. Launch the PSADMIN utility.
2. Navigate to the PeopleSoft Domain Administration menu and select *Configure this domain*.
3. In the Quick Configure menu, select the number for the Jolt Relay option, to enable Jolt Internet Relay.
4. Select the JRAD Port option, and enter the appropriate port number for the JRAD Port.

Note: The JRAD (listener) port number must match the JRLY connect port that you previously configured.

Running Jolt Internet Relay

This section discusses how to:

- Use the JRJLY administration program.
- Run Jolt Relay on Windows.
- Run Jolt Relay on UNIX.

Using the JRJLY Administration Program

You use the `jrly` command located in `TUXDIR\udataobj\jolt\relay` to administer Jolt Relay on all platforms. You can use the following `jrly` command options at any time:

- `jrly -version`
Display the current version of the JRJLY binary.
- `jrly -help`
Display a summary of command-line options with brief descriptions.

Running Jolt Relay on Windows

On Windows, you set up Jolt Relay to run as a service. On other platforms you must run Jolt Relay directly.

If you want to install multiple Jolt Relay services, you must specify a string to be used as a *display suffix* that uniquely identifies each additional service you install. You subsequently use the suffix to identify each service it commands. An additional service with the suffix `MyJoltRelay`, for example, is called *Oracle Jolt Relay_MyJoltRelay*, but you refer to it using only the suffix. You can omit the suffix when installing only one of these services, which becomes the default Jolt Relay service, called *Oracle Jolt Relay*.

Note: All administrative commands in the following table except `-start` and `-stop` require that you have write access to the Windows registry. The `-start` and `-stop` commands require that you have Windows service control access. These requirements are based on Windows user restrictions.

Command	Description
<code>jrly -install [display_suffix]</code>	Install JRJLY as a Windows service.
<code>jrly -remove [display_suffix -all]</code>	Remove one instance, all instances, or the default JRJLY Windows service.

Command	Description
<code>jrly -set [-d <i>display_suffix</i>] -f<i>config_file</i></code>	Update the registry with the full path of a new configuration file for the specified JRLY service. Note: You can run multiple Jolt Relay services by specifying a different display suffix along with the name of a different configuration file for each installed service. Each configuration file must contain a unique value for the LISTEN parameter that specifies a different port. This is essential to avoid port clashes when running the services concurrently. You must run this command before the service starts.
<code>jrly -manual [<i>display_suffix</i>]</code>	Set the start/stop to manual. This command sets the specified JRLY service to be manually controlled, using either the command-line options or the Service Control Manager (SCM).
<code>jrly -auto [<i>display_suffix</i>]</code>	Set the start/stop to automatic. This command sets all the operations for a specified JRLY service to be automatically started when the OS boots and stopped when the OS shuts down.
<code>jrly -start [<i>display_suffix</i>]</code>	Start the specified JRLY service.
<code>jrly -stop [<i>display_suffix</i>]</code>	Stop the specified JRLY service.

Running Jolt Relay on UNIX

This section discusses how to start and stop Jolt Relay directly from a command line on UNIX.

To start Jolt Relay on UNIX:

1. Change directories to the Jolt Relay directory within your Tuxedo installation:

```
cd $TUXDIR/udataobj/jolt/relay
```

2. Run the following command:

```
jrly -f jrly_config &
```

Where *jrly_config* is the name of a Jolt Relay configuration file.

You can run multiple instances of Jolt Relay by using a different port for each instance. You run JRLY once for each instance, and specify a different configuration file each time. Each configuration file must contain a value for the LISTEN parameter that specifies a different port.

The & causes JRLY to run in the background.

To shut down Jolt Relay on UNIX, use the UNIX kill -9 command.

Appendix A

Securing PS_HOME and PS_CFG_HOME

Understanding PS_HOME and PS_CFG_HOME Security

With the separation of the PS_HOME and PS_CFG_HOME directories, system administrators can implement more secure PeopleSoft deployments by restricting access within each of these directory structures.

This section describes the procedures and considerations involved in configuring these additional security options.

Note: Each site can elect to implement these security measures as needed according to individual security policies.

Note: Each PeopleSoft application you have licensed may have specific instructions regarding the implementation of these security measures. Always check your application-specific documentation for any information you need to consider to ensure both a secure environment and a properly functioning application.

Understanding PS_HOME Security

Because the configuration files, by default, do not reside in PS_HOME, the PS_HOME installation can be locked down to prevent unauthorized access, by user or system process. By making the PS_HOME directory 'Read-Only', processes running in the domain cannot write to PS_HOME or any of the subdirectories therein. Likewise, any users with malicious intent are unable to delete or modify executable files in PS_HOME.

Securing PS_HOME involves making the directory read-only, yet making sure that the following system elements have sufficient access.

System Element	Description
Application server	Application server domains need read access to the executable and binary files of PS_HOME to process requests and run application logic.

System Element	Description
Process Scheduler	<p>Process Scheduler domains need read access to the executable and binary files of PS_HOME to run batch processes. Plus, keep in mind these points:</p> <ul style="list-style-type: none"> • The user creating a Process Scheduler domain on Windows needs read and write access to the Windows Registry. • The restricted OS user needs to have full privilege access to the psreports folder (and its children). Process Scheduler and the Report Distribution agent inherit the restricted user's security settings they will need to create folders in psreports. • Oracle ProcMGR (Tuxedo) should be started with the restricted OS user ID. • Configure Process Scheduler with an Admin OS user for Windows. While logged into Windows with the full privilege OS user ID, create and configure the Process Scheduler domain, so that ODBC and NVision DLLs register properly.
File server/Windows workstations	The file server and/or Windows workstations running Application Designer need read-only access to PS_HOME to facilitate three-tier connections.

Note: These instructions do not apply to the PS_HOME residing on the web server for PIA or the PS_HOME on the database server.

Note: When implementing a read-only PS_HOME, consider that environment locations to which processes write files can't be in a read-only location. Settings for "temporary" directories and "output" directories should not be located within the PS_HOME directory structure. For example, the default temporary directory locations are C:\Documents and Settings\

Note: All elements of your PeopleSoft implementation, such as Process Scheduler and SQR can operate within a secure PS_HOME configuration.

Understanding Minimum Access Required by the User Starting Domains

The bare minimum that needs write access at the time a domain boots includes:

- The domain directory: it must be possible to write content to the domain directory although most of the configuration files in this directory can be read-only.
- The domain LOGS directory: by default this is the LOGS directory beneath the domain directory. If this location is overridden in the configuration file, the relevant location must also be read-write.
- The .adm directory: this subdirectory within the domain (if present) must also be read-write. This is required by Oracle Tuxedo.

- The Archive directory: located within the domain directory, a copy of the .cfg is archived to this directory each time it is updated. This directory is also used by the Purge Cache PSADMIN option for application servers.

All other files in the PS_CFG_HOME directory tree can be made read-only to the user starting the domain.

Understanding PS_CFG_HOME Security

Some administrators may want to implement additional security and restrict access to PS_CFG_HOME. For example, in some cases you may want take further steps to limit privileges of the user starting a domain, or lock down configuration files to prevent unintended configuration changes during runtime.

Securing PS_HOME on UNIX

The UNIX operating system lends itself to a read-only configuration for PS_HOME because of the way that Inter-process Communication (IPC) resources are allocated and managed. UNIX was designed to allow multiple users concurrent access to the same physical hardware and file system while enforcing a strong privileges model.

Note: It is necessary to have access to at least two user accounts in order to setup a true and complete read-only environment on UNIX.

To illustrate the procedure, two user accounts are used.

<i>User Account</i>	<i>Description</i>
InstallAdmin	User responsible for installing PeopleTools.
DomainAdmin	User responsible for creating, configuring, and booting domains. Note: It is under this account that domain processes will run and therefore should have the most stringent permissions.

To setup read-only PS_HOME on UNIX:

1. Install PeopleTools using the InstallAdmin account.
2. Verify that PS_HOME is read-only.

After installing PeopleTools, attempt to delete both a directory and file from PS_HOME using the DomainAdmin account.

If it is not read-only to the DomainAdmin account, login as the InstallAdmin account and use the chmod command to make PS_HOME read-execute to the world.

If the DomainAdmin account is a member of the same group as the InstallAdmin account you will need to apply the read-execute restriction to the group also. For example,

```
chmod -R 755 $PS_HOME
```

3. Sign in as the DomainAdmin account, open a shell, and change directory to PS_HOME.
4. Invoke psconfig.sh to set the environment.
5. Create and configure a new domain.

This can be an application server, search server, or Process Scheduler domain.

6. Start the new domain and verify that all of the domain processes have started.

For application server domains, ensure that you can sign in through PIA and make successful page requests.

Managing a Secure PS_HOME on UNIX

When deploying a secured PS_HOME environment on UNIX, keep the items in this section in mind.

Working with User Accounts

The user account that boots the domain must be the same user who configures the domain. This is a Tuxedo requirement, not a PeopleTools requirement. This means that the user account under which the domain processes will run must have read-write access to the domain directory.

The owner of the domain processes is the user account who starts the domain. This is different from Microsoft Windows, where the domain processes are booted by the account that starts the Oracle ProcMGR service. If you use both Windows and UNIX servers to deploy PeopleSoft, keep this subtle distinction in mind between the two operating systems.

Configuring Partial PS_HOME Access

In some cases, user accounts may need to access specific parts of the PS_HOME directory tree. This is recommended through the addition of a “hybrid” user to the same group to which the “InstallAdmin” account (the user who installed PeopleTools) belongs. The InstallAdmin can then choose to allow group access to the specific parts of the PS_HOME directory tree to which the hybrid user is permitted read-write access.

For example, consider a scenario where you have installed PeopleTools at your site, but have hired a consultant to help with various implementation tasks. The InstallAdmin only wants to allow the consultant access to specific parts of the PS_HOME directory tree. The account that the consultant uses is therefore a hybrid account. It has read-write access to PS_HOME, but only to the specific subdirectories deemed necessary.

To achieve this hybrid privilege model, allow group access to those specific directories under PS_HOME to which the consultant requires write access.

Securing PS_HOME on Windows

When securing PS_HOME on Windows, you have these options:

- Multiple administrator user accounts.
- Local user accounts.

Multiple Administrator User Accounts

This method of securing PS_HOME on Windows offers the ability for many administrator user accounts to share the same PS_HOME while managing separate PS_CFG_HOMEs. This method is most appropriate for a production environment.

In this configuration, you install PeopleTools on a server machine, and share the PS_HOME installation location as read-only. Domain administrators may then map to this network drive, and invoke PSADMIN to create and start domains.

Once you have set up a secure PS_HOME, domain creation and the various prerequisites for setup are the same as before. For example, database connectivity must be available on the machine on which the domain will boot. The server where PS_HOME is located acts as a read-only file server for the domains.

To illustrate this procedure, two user accounts are used.

<i>User Account</i>	<i>Description</i>
User1	An administrator on Machine, having read-write access to PS_HOME.
User2	An administrator on Machine2, having read-only access to PS_HOME. A restricted user.

In this scenario, one administrator installs PeopleTools, and a second user, with a more limited set of privileges, creates and administers domains.

The steps for User2 on Machine2 can be applied to multiple users. Any number of users can map to a single read-only PS_HOME.

Note: This information applies to PeopleTools installations on drives assumed to be formatted as NTFS.

To install PS_HOME and restrict privileges:

1. While logged in as User1 on Machine1, run the PeopleTools installation program.
2. Choose Application Server and Batch Server installation options.
File Server is optional, depending upon whether or not it is required.
3. Ensure that you install PeopleTools such that PS_HOME is not the top most directory on the drive.

For example,

D:\PTInstalls\PT8.53

Note: This is essential if you plan on accessing PS_HOME as a mapped drive.

4. After the installation has completed, set privileges on the PS_HOME directory tree, using Windows File Sharing.

- Using Windows Explorer, select the high-level directory (as in D:\PTInstalls), right-click, and select Sharing and Security.
 - On the Properties, Sharing tab, click Share this folder.
 - Click Permissions, and for the Group of Everyone, check the Allow checkbox for Read, make sure the Allow checkbox for Change is not selected, and click Apply and/or OK.
5. Verify that the folder has been shared by making sure the ‘hand’ icon appears on the folder in Windows Explorer.

To set up access to a secure PS_HOME:

1. Sign in as User2 on Machine2.
2. Map a network drive to the shared PS_HOME.
3. Verify that you cannot add, modify, or delete any content below the mapped location.

This ensures that PS_HOME is read-only. If you can delete or change content in the mapped location, it is possible that User2 is an administrator on Machine1. User2 must not be an administrator on Machine1 for these security measures to be effective.

If User2 cannot see the shared location there may be a problem with the share or the local network. Make sure the machine can be pinged.

4. Configure Oracle ProcMGR service to allow the User2 account to access PS_HOME.

This is necessary because by default the Oracle ProcMGR service uses the Local System account.

- a. Select Start, Programs, Administrative Tools, Services, and double-click on the Oracle ProcMGR service.
- b. On the General tab, stop the service, and set the Startup type to Manual.
- c. On the Log On tab, select the This account radio button, and enter the logon information for User2.
- d. Click OK.

Note: Do not start the service yet.

5. Set the TM_TUXIPC_MAPDRIVER user variable for User2.

This environment variable must be set to contain any mapped drives required by the domain processes, such as the drive where PS_HOME is located. Use the following format:

```
drive1:=\\machine_name1\dirpath1[;drive2:=\\machine_name2\dirpath2[...]]
```

For example,

```
N: \\10.100.200.300\PTInstalls
```

If multiple mapped drives are required, use a semicolon to separate the values, similar to the way directories are expressed in the PATH environment variable.

Note: Depending on your network, use either the DNS name or the IP address to specify the machine name.

Note: If the Oracle ProcMGR needs to run unattended, where no user is signed in, set the TM_TUXIPC_MAPDRIVER environment variable as a system environment variable instead of a user environment variable.

6. Start the Oracle ProcMGR service.

Once started, you can start PSADMIN as you normally would.

Local User Accounts

Using local user account to secure PS_HOME is a machine-bound solution that you may consider during an initial demo, development, or testing environment, where PS_HOME and PS_CFG_HOME reside on the same machine. In this method, only one machine and one domain account is required.

In this scenario, PeopleTools is installed by a user with administrative privileges. In the context of this scenario, this refers to the network domain user, a user that is a member of an existing network domain of users.

Application server domains are administered by a second user with a more limited set of privileges. This second user is created on the local system and only has access to resources on that machine.

In the following procedure, these user types are represented as:

User	Description
COMPANY\User1	User1 is an administrator on Machine1 and belongs to the network domain named <i>COMPANY</i> .
LOCAL_MACH\Guest2	Guest2 does not have administrative privileges on Machine1 but can sign on to Machine1. Guest2 is a Restricted User.

Setting up a secure PS_HOME and restricting privileges:

1. While logged in as User1 on Machine1, install PeopleTools to the server using the install program.

Choose Application Server and Batch Server installation options. File Server is optional depending upon whether or not it is required.

Ensure that you install PeopleTools such that PS_HOME is not in the top directory level on the drive. For example, an ideal location would be C:\PTInstalls\PT8.53.

2. While logged in as User1 on Machine1, create the Guest2 user account.

Create a new user with at least the following attributes:

- User name

- Password
- Select Restricted user on the Group Membership tab.
- De-select the User must change password at next logon checkbox.

See Microsoft Windows documentation for more information related to creating users on Microsoft Windows.

3. Verify that the user is a Restricted User by highlighting the user ID and clicking on the Properties tab.
You may need to exit and re-enter the User Accounts dialog to see the new user added.
4. Make the PS_HOME directory tree read-only.
 - a. Open Windows Explorer and navigate to and select the PS_HOME directory location.
 - b. Open the Properties dialog and click on the Security tab.
 - c. Deny Write access to all Local Users.

To configure access PS_HOME:

1. Setup the Oracle Proc MGR Service to allow the IPC resources to be created using the restricted user ID (Guest2).

This is necessary because by default the Oracle PRocMGR service uses 'Local System' account which provides greater access to the PS_HOME than desired.

- a. Select Control Panel, Administrative Tools, Services.
 - b. Double click on the Oracle ProcMGR and change the Startup Type to Manual.
2. Change the user and password with which the service is started to match the new local user that you created earlier (Guest2).

Note: Do not start the service, just click OK.

3. Log off and sign on to Machine1 as Guest2.
In most cases, any error messages that you see when re-signing on can be ignored as they are associated with signing on with restricted permissions.
4. Verify that you cannot delete, update, or add any content to PS_HOME.
5. Before creating a domain your database connectivity must be setup within the restrictions of the guest2 local user account that you have created.
6. Start PSADMIN, and create a new domain, and confirm that the domain boots.
7. Install PIA on a separate machine, and verify that you can signon through PIA.

Because PIA is not within the scope of the secure PS_HOME, PIA should be installed on an additional machine. This is necessary because PIA needs write access to various locations within PS_HOME.

Managing a Secure PS_HOME on Windows

When deploying a secured PS_HOME environment on UNIX, keep the items in this section in mind. Depending upon your domain configuration and usage pattern, you may need to unlock specific subdirectories of PS_HOME.

Working With Mapped Drives, UNC Paths, and TM_TUXIPC_MAPDRIVER

This page explains mapped drives (Windows share drives) and the use of UNC paths for PeopleTools application server, Process Scheduler server, and search servers. When a PeopleTools domain is started on a Windows machine, it runs under the user for whom the ProcMGR Windows service has been configured. As such, the domain processes inherit the privileges of that user and not the user logged on to the system.

By default, the ProcMGR runs as the Local System account. While the Local System account has most privileges on the local host, it can't usually access UNC paths or mapped drives.

Note: ProcMGR is the Windows service that is responsible for allocating resources to Tuxedo domains.

Accessing UNC Paths and Mapped Drives

To allow PeopleTools domain processes to access UNC paths or mapped drives, it is necessary to start the ProcMGR service using an account that has access to these resources. This is typically a Windows *domain account*. A domain account refers to an account that logs a user ID on to both a machine and the corporate network. This account has been created by the network administrator. An example of such an account would be BIGCOMPANY\TSawyer.

The ProcMGR should be configured to start as the domain account. On the Log On tab of the service configuration dialog, click This account:, and enter the credentials of the domain account.

UNC Paths

If you plan to use UNC paths to access PS_HOME you must start PSADMIN using a UNC Path. For example:

```
\\ptinstalls\pt853\APPSERV\psadmin.exe
```

With the ProcMGR service set to a domain account, you can use PSADMIN to create and configure domains as if PS_HOME were on the local file system.

Mapped Drives

Additional steps are required if you plan on using a mapped drive for your PS_HOME or PS_CFG_HOME. These additional steps are required because Windows services do not recognize mapped drives. However, Oracle Tuxedo provides a mechanism by which the ProcMGR service is permitted to access network drives, which involves defining any mapped drives using an additional

environment variable, TM_TUXIPC_MAPDRIVER. If this environment variable has not been set, the domain processes will be unable to recognize the network drives.

To make sure that the TM_TUXIPC_MAPDRIVER variable is visible to the ProcMGR service, it is necessary to set it globally, as a System environment variable. For example, set TM_TUXIPC_MAPDRIVER to:

```
N:=\\10.233.238.123\PTInstalls
```

Important! The mapped drive cannot point *directly* to PS_HOME. The mapping must point to the parent directory above PS_HOME. For example, if PS_HOME is N:\\10.233.238.123\PTInstalls\pt850, TM_TUXIPC_MAPDRIVER should point to N:\\10.233.238.123\PTInstalls. PTInstalls is the parent directory of \pt850, which would resolve to N:\pt850.

Note: When mapping network drives to the PS_HOME is located, make sure to select Reconnect at logon.

Note: The working directory for psadmin.exe, must always be mapped to a letter drive, such as C:, D:, N:, and so on. When starting psadmin.exe, you must do so from the command line (or script) by referencing the full mapped working directory, such as \\ptinstalls\pt851\APPSERV\psadmin.exe. Launching psadmin.exe outside the current working directory (as in, using Start, Run) will cause psadmin.exe to function incorrectly. This restriction is imposed by Tuxedo.

Working With Oracle ProcMGR Windows Service

When logging on with user accounts Oracle ProcMGR service should be set to *Manual* instead of *Automatic*. Failure to do so may result in your domain account becoming locked. If set to 'Automatic' the service may continually attempt to start with an expired password causing the network to lock out the domain user account due to successive failed retries.

Note: On Windows servers, it is recommended to have the Windows user logged in running PSADMIN be the same user that runs the ProcMGR service.

See PeopleTools Installation guide for Oracle.

Managing TM_TUXIPC_MAPDRIVER

The TM_TUXIPC_MAPDRIVER environment variable needs to be maintained consistent with the mapped drives upon which you access PS_HOME. If the drive mappings change, then you need to make sure the new values are specified in TM_TUXIPC_MAPDRIVER. If the drive mapping represented by the TM_TUXIPC_MAPDRIVER does not exist, the ProcMGR service will fail to start.

Resolving Initialization Timeout Issues

If the PS_HOME used by your domains is on a network drive, you may notice a delay with starting a domain. This is a result of the binaries being loaded from across the network versus from the local disk. This can cause an initialization timeout.

If you notice domain start failures, check for the following message in the Tuxedo log for the domain:

```
tmboot.16020.15792.-2: CMDTUX_CAT:1859: ERROR: Server process ID 12668
```



```
failed to initialize within 60 seconds
```

In such cases, increase the timeout values in the domain's environment variables to accommodate for the slower start time. Use the `Edit environment settings` option on the Quick Configure menu to modify environment variables.

For example,

```
TM_BOOTTIMEOUT           : [300]
TM_RESTARTSRVTIMEOUT     : [300]
```

For more information on editing domain environment variables, see [Setting Domain-Level Environment Variables](#).

Working With the Tuxedo TM_CPAU Setting

For PeopleSoft systems, the Oracle ProcMGR service (`tuxipc.exe`) is responsible for starting Tuxedo domain processes on Windows. By default, domain processes run as the same user ID that the service is running as. In a default installation, the service is configured to log on to Windows as the Local System user. Microsoft does not support assigning network privileges to the Local System user for security reasons, but the Local System user otherwise has full administrative access to the local system.

In this configuration, PeopleTools domain processes also run as the Local System user, which presents several potential issues, including:

- PeopleTools domain processes are unable to access network resources.
- PeopleTools domain processes run with more privileges than are necessary. A compromised PeopleTools process will have full access to the local system and could potentially be used to gain unauthorized access to the local system.
- All PeopleTools domain processes on the system run as the same user ID

These problems are not present on UNIX systems where domain processes are always started as the user that runs `tmadmin` (by way of `PSADMIN` for PeopleSoft installations) to boot the domain. UNIX systems therefore support multiple domains, each running under different user IDs, with only the desired local privileges, and with no undesirable restrictions to network resources.

For Windows platforms, you can use the Oracle Tuxedo `TM_CPAU` (Create Process As User) environment variable to achieve behavior similar to UNIX systems. If `TM_CPAU` is set to `YES` before `tuxipc` is started, `tuxipc` creates an Oracle Tuxedo process that belongs to the user who initiated `tmboot`. If `tuxipc` boots as a service, go to Start, Control Panel, System, Advanced, Environment Variables and set system variable `TM_CPAU=YES` on the dialog box and reboot the Oracle ProcMGR service. If the Oracle ProcMGR service (`tuxipc.exe`) is started with the `TM_CPAU=YES` environment variable set, then domain processes will run as the user ID used to run `tmadmin` (`PSADMIN`) to boot the domain.

Using the `TM_CPAU` environment variable enables a variety of configuration options, including:

- The ProcMGR service can be run as the Local System user, but domain processes can be run using a minimally privileged user. This reduces the chance of a compromised PeopleTools process being used to gain unauthorized access to the system.
- The ProcMGR service can be configured to log on to Windows using a minimally privileged user ID and PeopleTools processes can run as a user with more privileges than the Tuxedo user ID. For example, the Tuxedo user ID could have read-only access to `PS_CFG_HOME`, but the PeopleTools

user could have read-write access. The Tuxedo user ID does not actually require read-access to PS_HOME. When CreateProcessAsUser runs, access to the executable to start is evaluated using the user ID that the process will run as.

- A single Windows system can be used to host multiple PeopleTools installations that are each administered by a different user. A non-administrative user ID used to boot one domain will have no privileges to processes started with a different user ID.
- Domain processes can be identified and managed in Windows Task Manager by a non-administrative user.

For more information, see Oracle Tuxedo documentation: "File Formats, Data Descriptions, MIBs, and System Processes Reference."

Implementing PS_CFG_HOME Security

The steps in this section describe techniques for applying more stringent access to a PeopleTools environment by restricting access to PS_CFG_HOME. If you intend to secure PS_CFG_HOME, it is assumed that you have also secured PS_HOME. Securing PS_CFG_HOME enables you to prevent malicious access to content and configuration files located in PS_CFG_HOME and in domain directories.

These steps describe a security implementation where you configure a user account(s) that can create and configure domains, and user account(s) for domain administrators, who can start and stop domains.

It is possible to limit the permissions to PS_CFG_HOME, such that the domain administrator account can:

- Create files and sub-directories in PS_SERVDIR.

This is necessary for creating log files and temporary files. Tuxedo also requires read-write access to the domain directory.

- Read (but not change or delete) any existing configuration or template files in PS_SERVDIR.

These files include .cfg, .ubb, .ubx, .val files, and so on.

Note: To apply these permissions, you must do so after the domain has been configured but before it has been started.

Note: Once these permissions have been implemented, only a user account with the appropriate privileges can reconfigure the domain.

Securing PS_CFG_HOME on UNIX

You have these approaches when implementing a secure PS_CFG_HOME on UNIX:

Security Approach	Description
Use a root account, or use <i>super user do</i> (sudo), instead, to perform root-level operations with a non-root account:	This technique involves locking a user account's access to a file from the root account.

Security Approach	Description
Use two administrator user accounts:	<p>The two administrator accounts approach involves using two user accounts that are members of the same group.</p> <p>This approach permits all users in the account in the group read access to the domain configuration.</p> <p>This approach works best if the number of users in the group is kept to a minimum.</p> <p>Because there are multiple users involved, it is necessary to override the default PS_CFG_HOME environment variable such that both users will see the same domains.</p>

To set up a secure PS_CFG_HOME using sudo:

1. Create and configure the domain.

For this procedure, assume the user who does this is DomainAdmin.

2. Use `sudo` to restrict write access to the sensitive configuration files.

For example, with the `sudo` command include:

```
chmod 555 <filename>
```

In this case, only `sudo` can change the configuration files or restore write access to DomainAdmin.

3. Log in as DomainAdmin, and verify that none of the restricted files can be changed or deleted by the DomainAdmin session.
4. Start the domain as DomainAdmin.

To set up a secure PS_CFG_HOME using two administrator accounts:

1. Create and configure a domain.

For this procedure, assume the user who does this is DomainAdmin.

2. As the DomainAdmin user, change the permissions on the domain configuration to allow write access to only those files and directories needing to be written to by the user starting the domain.
3. Signon as the DomainBootAdmin user, start PSADMIN, navigate to the Domain Administration menu, and re-configure the domain without making any changes.

This results in only the TUXCONFIG file being updated.

4. Star the domain as the DomainBootAdmin user.

Securing PS_CFG_HOME on Windows

To secure PS_CFG_HOME on Windows:

1. Ensure that your PS_CFG_HOME is created in a location that is accessible to the user account that needs to start the domain.

By default, the PS_CFG_HOME location will not be visible to the user account starting the domain because it is created in the user home of the user account creating the domain. Use one of these options to make the PS_CFG_HOME visible to the user account that needs to start the domain:

- Override the default location for PS_CFG_HOME, by manually setting the PS_CFG_HOME environment variable to a custom value.
 - Enable the restricted user to view the PS_CFG_HOME in the domain creator's user home. Set the PS_CFG_HOME as a system-level environment variable so that the restricted user will be able to see it when logged on.
2. While signed in as the domain administrator, create and configure your new application server, Process Scheduler, or search server domain.
 3. While signed on as the domain administrator, apply the necessary read-write restrictions.

It is recommended to apply read-only privileges to the entire directory path to the domain.

- a. Using Windows Explorer select the domain directory and open the Properties dialog.
 - b. Select the Security tab. If the restricted user is not displayed, click Add to append the user to the list.
 - c. Once added, highlight the restricted user ID and ensure that the following actions are checked in the Allow column: Read & Execute, List Folder Contents, Read, and Write.
4. Apply read-only privileges to the sensitive domain files that should be protected from read-write at runtime.

This typically includes the Tuxedo binary configuration (PSTUXCFG and PSBDMCFG), ASCII configuration files and templates (*.cfg, *.ubb, *.env, *.ubx, *.lst).

Using Windows Explorer, select each of these files in the domain directory and open the Properties dialog.

Note: At a minimum, it must be possible for the user starting the domain to write to the .adm and LOGS directory within PS_SERVDIR. Additionally, this user must be permitted to create new files in the domain directory.

5. Sign in as the restricted user, boot the domain, and verify that it is not possible to delete or modify any of the restricted domain files.

Note: Sign in using the same user account as the one entered in the Oracle ProcMGR service.

Note: If subsequent configuration changes are required it is necessary to configure the domain while signed in as the administrator account that created and configured the domain.

Appendix B

WebLogic Managed Server Architecture

PeopleSoft Internet Architecture Servlets and Applications

PIA is packaged as a JEE Enterprise Archive and is comprised of a collection of JEE web applications, commonly referred to as webapps or servlets. For the most part, in the context of PeopleSoft, the term *servlets* is used. The PeopleSoft servlets are:

PORTAL	PeopleSoft Interaction Hub
PSIGW	Integration gateway
PSEMHUB	PeopleSoft Environment Management Framework
PSINTERLINKS	PeopleSoft Business Interlinks

Note: PeopleSoft Business Interlinks is a deprecated product. These options exist for upgrade compatibility and transition.

In addition, these servlets are added when you install PIA on a WebLogic server machine. These elements are not part of the PeopleSoft Enterprise Archive, but instead are defined individually.

HttpProxyServlet	Reverse Proxy Server – Proxy to a single content server per URL. Each URL can provide unique content.
HttpClusterServlet	Reverse Proxy Server – Proxy to multiple WebLogic servers. All content servers provide access to the same content for load balancing.
Console	Administrative console for WebLogic Server.

WebLogic Domain Types

This section provides an overview of Weblogic domain types and discusses:

- Single-server domain.
- Multi-server domain.
- Distributed managed server.
- Common default settings.
- Single-server and multi-server/distributed server analogy.

- Domain topology.

Understanding WebLogic Domain Types

During PIA setup, you can choose between two different WebLogic domain configurations: a single-server domain and a multi-server domain. A multi-server domain can be expanded across multiple machines using the *distributed managed server* option. So, a distributed managed server implementation is a variation of the multi-server domain. Each of these domain configurations has a specific purpose but is fully customizable beyond that purpose.

This section discusses:

- single-server domains.
- multi-server domains.
- distributed managed servers.

Single-Server Domains

In a single-server configuration, the WebLogic domain's administration console and the J2EE components of PIA are all provided on a single instance of WebLogic Server. This configuration is intended for single-user or very small scale, non-critical production environments. It can be used as a starting point for you to familiarize yourself with WebLogic Server.

In a single-server domain, the resources used to administer WebLogic Server and your PeopleSoft application are not isolated from one another, therefore you don't administer each element individually. While each element must complete for the same resources, the low resource requirements of this configuration make it ideal for small scale and non-production usage.

The single-server domain in a PeopleSoft implementation consists of only one server: PIA.

Single-Server Deployment

Some of the servlets deployed in a single-server domain configuration must be accessed using a modified URL:

`http://server:port/servlet_name/...`

The single-server domain configuration deploys servlets as follows:

Application	Deployed to Server	Servlet Name in URL
PORTAL	PIA	(not needed)
PSIGW	PIA	PSIGW
PSEMHUB	PIA	PSEMHUB
PSINTERLINKS	PIA	PSINTERLINKS
Console	PIA	console

Application	Deployed to Server	Servlet Name in URL
HttpProxyServlet	Defined but not deployed.	(not needed)
HttpClusterServlet	Defined but not deployed.	(not needed)

Single-Server Domain Specific Settings

To configure the single-server domain specific settings, launch the Administration Console.

In the console, expand the Environment tree and select Servers. Click on the PIA server. Select the Configuration tab, and the General sub-tab. The default web application for the PIA server is PORTAL. The single-server domain specific default settings for the PIA server are as follows:

Setting	Default Value
Listen address	* (all local IPs).
Listen Port	80 (set during PIA setup).
SSL Listen Port Enabled	Enabled with demonstration self-signed digital certificates.
SSL Listen Port	443 (set during PIA setup).

Note: To configure SSL, you must also define SSL certificates.

See [Implementing WebLogic SSL Keys and Certificates](#).

Example: Single-Server Domain

The following illustrates sample contents of a typical single-server domain:

Server	Process Type	WebLogic and PIA Elements
Single server machine	WebLogic administration	Weblogic Administration Server
	Server instance	PIA
	Servlet or application	Administration Console PORTAL PSEMHUB PSIGW

Multi-Server Domains

The multi-server domain configuration is intended for production environments. This configuration takes advantage of WebLogic's administration server and managed server architecture. In a multi-server

configuration, multiple instances of WebLogic server are used, each contributing a specific function. The WebLogic console is provided on the domain's administration server, WebLogicAdmin, and the J2EE components of PIA are provided on individual or shared WebLogic managed servers.

A production application warrants process and resource pool isolation for greater stability and optionally tighter security controls, which this configuration provides. In a multi-server configuration, the resources used for WebLogic domain administration and monitoring are isolated from similar resources used to support the PIA application. A server process named *WebLogicAdmin* performs nothing but WebLogic administration, which includes domain administration and monitoring. Continuing that separation, the individual PIA servlets are (usually) isolated from each other. The PIA servlets are targeted and deployed across a portion of the six remaining server definitions, all of which are classified as *managed servers*, which are delivered in the multi-server configuration.

A multi-server domain in a PeopleSoft implementation creates the following servers:

WebLogicAdmin	Administration server for WebLogic domain administration.
PIA	Server for the PeopleSoft Interaction Hub and integration gateway.
PIA1	Server for the PeopleSoft Interaction Hub and integration gateway.
PIA2	Server for the PeopleSoft Interaction Hub and integration gateway.
PSEMHUB	Server for the PeopleSoft Environment Management Framework application.
RPS	Server for WebLogic reverse proxy server applications.

Multi-Server Servlet Deployment

Some of the servlets deployed in a multi-server domain configuration must be accessed using a modified URL:

`http://server:port/servlet_name/...`

The multi-server domain configuration deploys servlets as follows:

Application	Deployed to Server, Cluster (members)	Servlets Name in URL
PORTAL	PIA, PeopleSoftCluster (PIA1, PIA2)	(not needed)
PSIGW	PIA, PeopleSoftCluster (PIA1, PIA2)	PSIGW
PSEMHUB	PSEMHUB	PSEMHUB
PSINTERLINKS	PIA, PeopleSoftCluster (PIA1, PIA2)	PSINTERLINKS
Console	WebLogicAdmin	console

Application	Deployed to Server, Cluster (members)	Servlets Name in URL
HttpProxyServlet	RPS	(not needed)
HttpClusterServlet	Defined but not deployed.	(not needed)

Multi-Server Domain Specific Default Settings

To configure the multi-server domain specific settings, launch the Administration Console. In the Domain Structure tree, expand Environment, select Servers and click on the appropriate servlet, server, or application. Then select the Configuration tab, and the General sub-tab.

The domain specific default settings for the WebLogicAdmin server are as follows:

WebLogicAdmin Setting	Default Value
Listen Address	* (all local IPs)
Listen Port	9999
SSL Listen Port Enabled	Disabled

The domain specific default settings for the PIA server are as follows:

PIA Setting	Default Value
Listen Address	* (all local IPs)
Listen Port	80 (set during PIA setup)
SSL Listen Port Enabled	Enabled with demonstration self-signed digital certificates.
HTTPS Listen port	443 (set during PIA setup)

The domain specific default settings for the PIA1 server are as follows:

PIA1 Setting	Default Value
Listen Address	Locally determined hostname.
Listen Port	80 (set during PIA setup)
SSL Listen Port Enabled	Enabled with demonstration self-signed digital certificates.
SSL Listen Port	443 (set during PIA setup)

The domain specific default settings for the PIA2 server are as follows:

PIA2 Setting	Default Value
Listen Address	127.0.0.1
Listen Port	80 (set during PIA setup)
SSL Listen Port Enabled	Enabled with demonstration self-signed digital certificates.
SSL Listen Port	443 (set during PIA setup)

The default web application for the PSEMHUB server is PSEMHUB. The domain specific default settings for the PSEMHUB server are as follows:

PSEMHUB Setting	Default Value
Listen Address	* (all local IPs)
Listen Port	8001
SSL Listen Port Enabled	Disabled

In the console, navigate to Environments, RPS, Configuration, General to configure the RPS server. The default web application for the RPS server is HttpProxyServlet. The domain specific default settings for the RPS server are as follows:

RPS Setting	Default Value
Listen Address	* (all local IPs)
Listen Port	8080 (set during PIA setup)
SSL Listen Port Enabled	Enabled with demonstration self-signed digital certificates.
SSL Listen Port	8443 (set during PIA setup)

Note: To configure SSL, you must also define SSL certificates.

See [Implementing WebLogic SSL Keys and Certificates](#).

Example: Multi-Server Domain

The following illustrates the elements running within a typical multi-server domain:

Server	Process Type	WebLogic and PIA Elements
Single server machine	WebLogic administration	WebLogic Administration Server (WebLogicAdmin) WebLogic Node Manager

Server	Process Type	WebLogic and PIA Elements
	Managed Server	PIA 1, PIA 2 for PeopleSoft Interaction Hub
	Managed Server	PIA for PSIGW
	Managed Server	RPS for httpproxyservlet to proxy content for PIA 1 and PIA 2

Distributed Managed Servers

The *distributed managed server* configuration, although listed alongside the single-server and multi-server domain types, is not a *true* domain type. It's an optional extension for an existing multi-server configuration. As with the multi-server domain type, this configuration takes advantage of WebLogic's managed server architecture. The distributed managed server configuration is intended for production environments encompassing multiple machines.

A distributed managed server configuration enables you to spread a *logical* WebLogic domain configuration *physically* across multiple machines in a heterogeneous network. One server machine might act as the domain administration server, while the other server machines act as distributed managed servers, running various managed servers, such as PIA1, PIA2, and so on. Typically, you would also take advantage of the WebLogic Node Manager, which enables a system administrator to control the managed servers remotely.

See your WebLogic Node Manager documentation.

In a distributed, multi-server configuration you can have multiple machines, multiple collections of resources and program files, multiple web server processes, and a replicated domain configuration file. In this model, if any of system resources or server instances becomes unavailable, work shifts to the next instance of that resource. In the multi-server configuration, an increase in PeopleSoft Interaction Hub usage, for example, can be accommodated by configuring an additional WebLogic server instance or machine to also serve the PeopleSoft Interaction Hub. Using distributed managed servers provides flexibility in shaping your hardware allocation to meet your system demands.

Benefits of the Distributed Managed Server

A distributed managed server configuration provides the same benefits as a multi-server configuration with the added benefit of hardware isolation. This option requires a multi-server installation to be performed to some other location, which will contain the configuration for this distributed managed server.

The fundamental benefits of a multi-server configuration include:

- **Dedicated service providers:** Web servers can be dedicated to providing PeopleSoft Interaction Hub and are insulated from other portions of PIA such as PeopleSoft Integration Gateway or PeopleBooks.
- **Redundant service providers:** Multiple web servers can be used to serve different aspects of PIA, providing load balancing and failover support.
- **Distributed resources:** Multiple web server machines can be used, each capable of serving different or redundant aspects of PIA.

- Centralized and replicated configuration: Master domain configuration is centralized and distributed server information is replicated locally.

The server configuration settings for a distributed managed server are maintained by that domain's administration server are stored locally on that administration server. Configuration settings are replicated to a managed server during its startup, but are only maintained as a read-only backup copy for that individual managed server in the event that the administration server isn't available the next time this particular managed server needs to be started.

Note: Only one managed server can be run per distributed managed server domain directory. If you intend to run multiple distributed managed servers on a single machine, perform the PIA install and create unique distributed managed server domain directories, one for each distributed managed server that you intend to run on that machine.

Example: Distributed Managed Server

The following illustrates the multiple servers, their roles, and the various elements running on the multiple server machines within a sample distributed managed server configuration:

Server	Role	WebLogic and PIA Elements
Server 1	Domain Administration Server	WebLogic Administration Server (WebLogicAdmin)
Server 2	Distributed Managed Server	PIA1 (for PeopleSoft Interaction Hub) PIA2 (for PeopleSoft Interaction Hub) WebLogic Node Manager
Server 3	Distributed Managed Server	PIA3 (for PeopleSoft Interaction Hub) PIA4 (for PeopleSoft Interaction Hub) WebLogic Node Manager
Server 4	Distributed Managed Server	PIA (for Integration Gateway) WebLogic Node Manager
Server 5	Distributed Managed Server	RPS (for httpproxyservlet WebLogic Node Manager

Common Default Settings

Single-server and multi-server domain configurations have many settings in common.

Domain Defaults

Many of these common settings can be configured in the WebLogic Server Console, but some are configured in other environments. Default values are listed when available.

Setting	Default Value	Where To Configure In the Administration Console
SSL functionality	Enabled with demonstration self-signed digital certificates.	Environment, Servers <i>server name</i> , Configuration, Keystores tab and SSL tab. Command line: pskeymanager
Server logs	<i>Weblogic domain</i> \logs\server name_*.log	Environment, Servers, <i>server name</i> , Logging.
HTTP access log	Disabled	Environment, Servers, <i>server name</i> , Logging, HTTP.
HTTP keep-Alive	30 seconds	Environment, Servers, <i>server name</i> , Protocols, HTTP.
HTTPS keep-Alive	60 seconds	Environment, Servers, <i>server name</i> , Protocols, HTTP.
Low memory settings	various	Environment, Servers, <i>server name</i> , Configuration, Tuning.
Stuck Thread Max Time	600	Environment, Servers, <i>server name</i> , Configuration, Tuning.
Stuck Thread Timer Interval	60	Environment, Servers, <i>server name</i> , Configuration, Tuning.
Enable Administration Port	Disabled	Click the <i>domain name</i> (such as peoplesoft) in the Domain Structure section, then select Configuration, General.
PORTAL HTTP session monitoring	On (applies only to servers running PORTAL)	Deployments, Applications, peoplesoft, PORTAL, Monitoring.
System administrator user ID	system (set during PIA setup)	Security Realms, myrealm, Users and Groups, Users.
System administrator password	password (set during PIA setup)	Security Realms, myrealm, Users and Groups, Users.
System operator user ID	operator	Security Realms, myrealm, Users and Groups, Users.
System operator password	password	Security Realms, myrealm, Users and Groups, Users.

Setting	Default Value	Where To Configure In the Administration Console
System monitor user ID	monitor	Security Realms, myrealm, Users and Groups, Users.
System monitor password	password	Security Realms, myrealm, Users and Groups, Users.

Script and Environment Defaults

Modify these settings by editing a setEnv script or applying command line parameter overrides to WebLogic control scripts.

The following settings specify the names and structure of various directories on the web server machine.

Setting	Default Value	Description/Override
PS_HOME	(none)	PeopleSoft home directory (set during PIA setup).
BEA_HOME	(none)	High-level install directory (set during PIA setup). Where Tuxedo and WebLogic may be installed.
WL_HOME	(none)	WebLogic home directory (set during PIA setup).
DOMAIN_NAME	peoplesoft	Name of this WebLogic domain (set during PIA setup).
JAVA_HOME	(Depends on the operating system platform.)	Location of Java. Set during PIA setup or with a call to WebLogic's CommEnv script.

Note: You configure Java VM options including JVM memory size using the `JAVA_OPTIONS_OSplatform` parameter, during PIA setup.

The following are miscellaneous settings.

Setting	Default Value	Description/Override
HOSTNAME	<i>Local hostname</i>	Set during PIA setup.
PRODUCTION_MODE	TRUE	Enable WebLogic production mode (set during PIA setup).

Setting	Default Value	Description/Override
DISCOVERY_MODE	FALSE	Disable auto detection of unregistered applications. Script: setEnv
WLS_USER	Operator	Use to stop WebLogic with stop scripts and run it as a Windows service.
WLS_PW	Password	Use to stop WebLogic with stop scripts and run it as a Windows service.
ADMINSERVER_PROTOCOL	HTTP	Protocol used for managed server to connect to administration server (not used in single-server domain).
ADMINSERVER_HOSTNAME	Single-server: <i>local hostname</i> . Multi-server: <i>local hostname</i> . Distributed server: (none — set manually).	Administration server's hostname that managed servers attempt to connect to by default when started. Set during PIA setup (except distributed server).
ADMINSERVER_PORT	Single-server: <i>HTTP port of PIA server</i> . Multi-server: 9999. Distributed server: (none — set manually).	Administration server's Listen port that managed servers attempt to connect to by default when started. Set during PIA setup (except distributed server).
ADMINSERVER_SERVERNAME	Single-server: PIA. Multi-server: WebLogicAdmin. Distributed server: WebLogicAdmin.	WebLogic server instance name of this domain's administration server, used for stopping and starting the server.
WL_VERSION	<i>Detected major WebLogic version</i> .	WebLogic major version, such as 10.
WL_SERVICEPACK	<i>Detected minor WebLogic version</i> .	WebLogic service pack level.
WL_PATCH	<i>Detected WebLogic patch version</i> .	WebLogic patch level.
BACKGROUND_PROCESS	TRUE	Run WebLogic server as a background process. On UNIX you can force foreground execution using the start script's -foreground option.

The following are debugging output settings.

Setting	Default Value	Description/Override
SET_CAPTURE_STDOUT_STDERR	FALSE	(Windows only) Capture standard output and standard error of a WebLogic server running as a foreground process. You can also set this with the start script's –capture option.
ENABLE_JDPA_DEBUG	FALSE	(PeopleSoft development only) Enable JDPA debug support. You can also set this with the start script's –debug option.
ENABLE_VERBOSE_GC	FALSE	Enable verbose output of Java's garbage collector. You can also set this with the start script's –verbose:gc option.
ENABLE_VERBOSE_SSL	FALSE	Enable SSL debug support. Produces verbose SSL output. You can also set this with the start script's –verbose:ssl option.
ENABLE_VERBOSE_WL	FALSE	Enable verbose output for the core WebLogic server (not verbose output of PIA). You can also set this with the start script's –verbose:wl option.
MAX_FILE_DESCRIPTOR	4096	The number of open file descriptors set for any Weblogic server process.

The following are HTTP forward proxy support settings.

Setting	Default Value	Description/Override
ENABLE_HTTP_PROXY	FALSE	Enable the use of the forward http proxy.
HTTP_PROXY_HTTPHOST	(none)	IP address or hostname of the forward HTTP proxy server for HTTP requests.
HTTP_PROXY_HTTPPORT	(none)	HTTP Port number of the forward HTTP proxy server for HTTP requests.
HTTP_PROXY_HTTPSHOST	(none)	IP address or hostname of the forward HTTP proxy server for HTTPS requests.
HTTP_PROXY_HTTPSPORT	(none)	HTTP Port number of the forward HTTP proxy server for HTTPS requests.
HTTP_PROXY_NONPROXY_HOSTS	<i>localhost,local hostname, anddomainname.</i>	Host names and domain names of content servers that will not be proxied.

WebLogic Domain Directory Structure and Files

This section discusses:

- WebLogic domain directory structure.
- WebLogic domain file listing by type.
- J2EE application files.

WebLogic Domain Directory Structure

In a PeopleSoft implementation, the WebLogic domains are installed into your PS_HOME directory. The major components of the directory structure layout of a PIA install on WebLogic Server are:

Directory	Description
<i>PIA_HOME</i> \webserv	Parent, high-level WebLogic domain directory.
<i>PIA_HOME</i> \webserv\peoplesoft	Default WebLogic domain directory.
<i>PIA_HOME</i> \webserv\peoplesoft\applications\peoplesoft	PeopleSoft application directory.
<i>PIA_HOME</i> \webserv\peoplesoft\bin	WebLogic domain's bin directory, containing numerous administration scripts.
<i>PIA_HOME</i> \webserv\peoplesoft\config	WebLogic domain's configuration directory.
<i>PIA_HOME</i> \webserv\peoplesoft\servers\PIA\logs	WebLogic domain's log directory.
<i>PIA_HOME</i> \webserv\peoplesoft\keystore	WebLogic key store location for storing keys for configuring SSL.

WebLogic Domain File Listing by Type

This section lists the WebLogic domain files installed by the PeopleSoft installation, organized by file type. Where necessary, each table includes columns that indicate whether a given file is used in a single-server, multi-server, or distributed server configuration.

This listing does not include Java classes or PIA configuration files. On UNIX an equivalent Bourne shell script is provided where a Windows script is listed.

The following table lists WebLogic server administration scripts. All the life cycle scripts are stored in *PIA_HOME*\webserv*domain*\bin.

Script	Single-Server	Multi-Server	Distributed Server	Description
setEnv.cmd	X	X	X	Use this script to set required environment variables for the WebLogic server, for example: CLASSPATH, PATH, UNIX Library Path, and JVM options.
startPIA.cmd	X			Use this script to start the WebLogic domain's administration server (the PIA server) in a single-server configuration. On Windows this starts WebLogic as a foreground process. On UNIX this starts WebLogic as a background process. Run the script with – help for usage.
startWebLogicAdmin.cmd		X		Use this script to start the WebLogic domain's administration server (the WebLogicAdmin server) in a multi-server configuration. On Windows this starts WebLogic as a foreground process. On UNIX this starts WebLogic as a background process. Run the script with – help for usage.
startManagedWebLogic.cmd		X	X	Use this script to start a WebLogic managed server. All WebLogic servers in a WebLogic domain other than the administration server are WebLogic managed servers. Run the script with – help for usage.

Script	Single-Server	Multi-Server	Distributed Server	Description
stopPIA.cmd	X			Use this script to stop the WebLogic PIA server. Run the script with – help for usage.
stopWebLogic.cmd		X	X	Use this script to stop WebLogic servers. Run the script with – help for usage.
InstallNTservicePIA.cmd	X			(Windows only) Use this script to install the WebLogic PIA server as a Windows service. The service name is <i>WebLogic_domain-PIA</i> . Run the script with – help for usage.
InstallNTservice.cmd		X	X	(Windows only) Use this script to install a WebLogic server as a Windows service. The service name is <i>WebLogic_domain-server_name</i> . Run the script with – help for usage.
uninstallNTServicePIA.cmd	X			(Windows only) Use this script to uninstall the WebLogic PIA server Windows service. Run the script with – help for usage.
uninstallNTService.cmd		X	X	(Windows only) Use this script to uninstall a WebLogic server Windows service. Run the script with – help for usage.

Script	Single-Server	Multi-Server	Distributed Server	Description
pskeymanager.cmd	X	X	X	Use this script to manage the JKS keystore used by WebLogic Server, which is in <i>WebLogic_domain\keystore\pskey</i> . SSL certificates for WebLogic Server are stored in this keystore. PeopleSoft Integration Gateway can also share this keystore. Run the script with – help for usage.
startWebLogicBuilder.cmd	X	X	X	Use this script to start WebLogic Builder, which is used to change local application deployment descriptors.
createThreadDump.cmd	X	X	X	Use this script to create a JVM Thread dump. Run the script with – help for usage.

The following table lists WebLogic server configuration files stored in *PS_HOME/webserv/domain/config*.

File	Single-Server	Multi-Server	Distributed Server	Description
config.xml	X	X		This file stores the WebLogic domain configuration, including information about server names, ports, IP addresses, webapps, and SSL. Edit these settings using the WebLogic administration console: http://webserver:port/console .
msi-config.xml		X	X	This is a version of config.xml that's copied for use with a distributed managed server configuration. It's automatically replicated from the original config.xml after a managed server successfully starts.

File	Single-Server	Multi-Server	Distributed Server	Description
boot.properties	X	X	X	This file contains the WebLogic system ID and password used for administering the WebLogic domain.
fileRealm.properties	X	X	X	This file is used by WebLogic's internal LDAP server for system administration.
DefaultAuthenticatorInitIdift	X	X	X	This file is used by WebLogic's internal LDAP server for system administration.
DefaultRoleMapperInitIdift	X	X	X	This file is used by WebLogic's internal LDAP server for system administration.
SerializedSystemIni.dat	X	X	X	This file is used by WebLogic's internal LDAP server for system administration.

The following table lists PeopleSoft J2EE application scripts, which are all used with Integration Broker, and can be used with every WebLogic server configuration.

Script	Description
BatchProjectExecutor.bat	Use this script for Integration Broker batch EIP testing.
HashKeyGenerator.bat	Use this script to generate a hash key used for Integration Gateway playback.
MessageExport.bat	Use this Integration Broker script for extracting transaction data from request and response data.
StartSendMaster.bat	This is a Integration Broker test utility.

The following table lists miscellaneous files, which can be used with every WebLogic server configuration.

File	Description
Businterlink.txt	This file is used by PeopleSoft's Business Interlinks servlet for loading PeopleSoft libraries when needed.
piaInstallLog.xml	This is the PIA install log.
PSCipher.bat	Use this script for encrypting Integration Broker passwords.

Related Links

PeopleSoft Integration Broker Administration

J2EE Application Files

In addition to WebLogic domain configuration files, application descriptors are installed with the PeopleSoft J2EE enterprise application. The following table lists these descriptor files. The path shown for each file is relative to *PIA_HOME*\webserv\WebLogic_domain\applications\.

File	Description
peoplesoft\META-INF\MANIFEST.MF	Use this script to set required environment variables for the WebLogic server, for example: CLASSPATH, PATH, UNIX Library Path, and JVM options.
peoplesoft\META-INF\application.xml	This file contains a list of the webapps that comprise the PeopleSoft J2EE enterprise application.
peoplesoft\PORTAL\WEB-INF\web.xml	This file is the web application descriptor for the PORTAL webapp. It lists all of the servlets deployed as part of that application.
peoplesoft\PORTAL\WEB-INF\weblogic.xml	This file is the PORTAL web application extension descriptor. It specifies, among other things, the HTTP session cookie name, optional cookie domain, and context path of this application.
peoplesoft\PSIGW\WEB-INF\web.xml	This file is the web application descriptor for the PeopleSoft Integration Gateway (PSIGW) webapp. It lists all of the servlets deployed as part of that application.
peoplesoft\PSIGW\WEB-INF\weblogic.xml	This file is the PSIGW web application extension descriptor. It specifies the context path of this application.
peoplesoft\PSEMHUB\WEB-INF\web.xml	This file is the web application descriptor for the PeopleSoft Environment Framework (PSEMHUB) webapp. It lists all of the servlets deployed as part of that application.
peoplesoft\PSEMHUB\WEB-INF\weblogic.xml	This file is the PSEMHUB web application extension descriptor. It specifies the context path of this application.
peoplesoft\PSINTERLINKS\WEB-INF\web.xml	This file is the web application descriptor for the PeopleSoft Business Interlinks (PSINTERLINKS) webapp. It lists all of the servlets deployed as part of that application.
peoplesoft\PSINTERLINKS\WEB-INF\weblogic.xml	This file is the PSINTERLINKS web application extension descriptor. It specifies the context path of this application.
HttpProxtServlet\WEB-INF\web.xml	This file is the web application descriptor for the WebLogic Server Reverse Proxy Server (RPS) webapp that's used to proxy content from a single WebLogic server. It lists all of the servlets deployed as part of that application.
HttpProxyServlet\WEB-INF\weblogic.xml	This file is the single-server RPS web application extension descriptor. It specifies the context path of this application.

File	Description
HttpClusterServlet\WEB-INF\web.xml	This file is the web application descriptor for the WebLogic Server Reverse Proxy Server (RPS) webapp that's used to proxy content from a cluster of WebLogic servers. It lists all of the servlets deployed as part of that application.
HttpClusterServlet\WEB-INF\weblogic.xml	This file is the multi-server RPS web application extension descriptor. It specifies the context path of this application.

PIA Install and Reinstall Options

The PeopleSoft Internet Architecture (PIA) installer enables you to create a new WebLogic server domain or update a valid existing WebLogic domain. A valid domain is a domain built by the PIA installer in the *PS_HOME* directory that you specify.

Depending on which option you select, you're prompted for additional information relevant to that selection. When creating a new domain, you're prompted to select from three configuration types: Single-server, multi-server and distributed managed server. If you select to update an existing domain, you're prompted to indicate which domain you would like to update and what type of update you would like to perform. These options are described in detail in your PeopleTools installation documentation.

For more information, see *PeopleTools 8.53 Installation* for your database platform: "Setting Up the PeopleSoft Pure Internet Architecture in GUI Mode."

Administering a WebLogic Server Life Cycle

This section provides an overview of the WebLogic server life cycle and discusses how to:

- Start and stop single-server processes.
- Start and stop multi-server processes.
- Start and stop a distributed managed server.

Related Links

[Starting WebLogic](#)

[Stopping WebLogic](#)

Understanding the WebLogic Server Life Cycle

You control a WebLogic server's life cycle primarily using a collection of scripts provided in that server's WebLogic domain directory. Each instance of a WebLogic server runs in an isolated Java Runtime Environment (JRE), regardless of whether you're testing with a single-server configuration or implementing a multi-server configuration for production. All scripts must be launched from the WebLogic domain directory; and provide usage syntax if run with `-help`.

Starting and Stopping Single-Server Processes

In a single-server configuration, there's only one server to administer: PIA. You can control the life cycle of the PIA server using scripts or in the WebLogic console.

Scripts

For all platforms:

startPIA Use this script to start the WebLogic server locally.

stopPIA Use this script to connect to a locally running WebLogic server and issue a shutdown command through WebLogic APIs.

Note: When you shut down the server, a warning is displayed since the shutdown command uses a non-SSL http connection to connect to the WebLogic Server. This shutdown command can be changed to use the SSL connection by editing the stopPIA.sh script. To use the SSL connection the shutdown command will be the following.

For Windows only:

installNTservicePIA Use this script to register the PIA WebLogic server as a Windows service that runs as a background process. The service is named as *WebLogicDomainName-PIA*, for example: peoplesoft-PIA.

uninstallNTservicePIA Use this script to deregister the PIA Windows service.

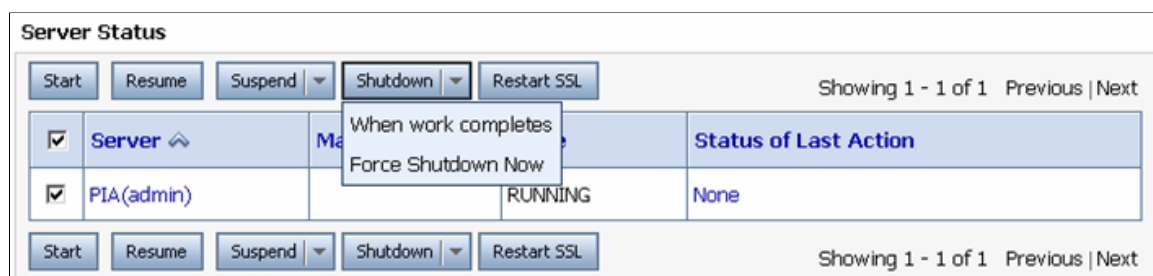
WebLogic Console

A WebLogic server can also be shut down from the Administration Console. Sign in to the console at <http://webserver:port/console> and perform either of the following.

Note: Before you perform any action in the WebLogic console, you have to click the Lock and Edit button and then the Activate Changes button after the changes are done.

Image: Shutting down from the console

This example illustrates the fields and controls on the Shutting down from the console. You can find definitions for the fields and controls later on this page.



In the navigation tree on the left, expand your domain, click Environment, Servers. Click PIA and select the Control tab. Select the check box for the server that you would like to shutdown, and click Shutdown. You have these options:

When work completes	This option enables transactions in progress to complete before shutting down the server. To terminate all HTTP sessions immediately, you can first select Ignore Sessions During Shutdown.
Force Shutdown Now	Immediately terminate all HTTP sessions and transactions in progress, and shut down the server.

Starting and Stopping Multi-Server Processes

In a Multi-server configuration, as the title implies there are multiple servers to administer. Controlling the life cycle of these servers can be done via scripts, the WebLogic console and the WebLogic Node Manager.

Scripts

For all platforms:

startWebLogicAdmin	Use this script to start the WebLogicAdmin server.
startManagedWebLogic	Use this script to start a WebLogic managed server. All of the servers defined in a multi-server domain, except the WebLogicAdmin server, are controlled as managed servers. For example, to start PIA1 as a managed server run <code>startManagedWebLogic PIA1</code> .
stopWebLogic	Use this script to connect to a locally running WebLogic server and issue a shutdown command using WebLogic APIs. A remote distributed managed server can be shut down using a local administration server.

For Windows only:

installNTservice.cmd	Use this script to register a WebLogic server as a Windows service that runs as a background process. The service is named as <i>WebLogicDomainName-ServerName</i> . For example, to define the PIA1 managed server as a Windows service, run <code>installNTservice PIA1</code> . To define the WebLogicAdmin server as a Windows service, simply run <code>installNTservice</code> .
UninstallNTservice.cmd	Use this script to deregister a WebLogic server that's defined as a Windows service.

Consider the following when using scripts with managed servers:

Operating System	Consideration
All	<p>When starting a WebLogic managed server it will attempt to connect to its administration server. A managed server's administration server is specified either as a command line parameter when starting the managed server, or using the three administration server environment variables in setEnv, specifically ADMINSERVER_PROTOCOL, ADMINSERVER_PORT, and ADMINSERVER_HOSTNAME. The first time a managed server starts, it <i>must</i> connect to its administration server. If on subsequent startups the administration server is not available, the managed server starts up in Managed Server Independence (MSI) mode by using its locally replicated msi-config.xml. A managed server running in MSI mode can't be administered from a console, so this situation should only be encountered when it is imperative that the managed server be started even though the administration server is not running. Once the administration server is back online, running managed servers that were not previously known by the administration server to be running may be rediscovered using WebLogic's command line utility <code>java weblogic.Admin DISCOVERMANAGEDSERVER</code>, or you can just restart the managed server.</p> <p>To use WebLogic's java command line utility classes run setEnv to set up your environment, then run <code>java weblogic.Admin</code> for usage.</p>
Windows	<p>When running a WebLogic managed server as a Windows service, the managed server's administration server <i>must</i> be running. When installing a managed server as a Windows service, the managed server service can be configured to be dependent on its local administration server. To configure a managed server service to be dependent on its local admin server service use the <code>-depends</code> option of <code>installNTservice.cmd</code> when defining the Windows service for the managed server. In addition, when the administration server is also a Windows service, you must define it using the following command:</p> <pre>installNTservice.cmd -delay <i>interval</i></pre> <p>Where <i>interval</i> is a period in milliseconds, for example <code>6000</code>. This allows the administration server sufficient time to start before the managed server starts.</p>

WebLogic Server Console

A WebLogic server can also be shut down from its administration console in the same way as single server environments.

WebLogic Node Manager

The WebLogic Node Manager provides the ability to start a WebLogic managed server from the WebLogic Server Console. In addition, the console provides a way to automatically restart a failed server. As with all WebLogic servers, the WebLogic Node Manager runs isolated in its own JRE, and on Windows it can also run as a Windows service. The WebLogic Node Manager binds to a unique IP address and port at startup and accepts lifecycle commands from a WebLogic administration server.

Multiple WebLogic domains running on a single machine can have its managed servers administered by a shared WebLogic Node Manager, as long as each WebLogic domain uses the same version of WebLogic.

The following table lists the WebLogic Node Manager files that are provided with WebLogic server, not the PIA install. These files are located in *WLS_HOME*\wlserver 10.3\server\bin\, not your WebLogic domain directory created within the PeopleSoft directory structure.

File	Description
startNodeManager.cmd	Use this script to start the WebLogic Node Manager as a foreground process.
installNodeMgrSvc.cmd	Use this script to define the WebLogic Node Manager as a Windows service that runs as a background process. The service is called WebLogic Platform NodeManager.
uninstallNodeMgrSvc.cmd	Use this script to uninstall the WebLogic Node Manager as a Windows service.
nodemanager.properties	This is the WebLogic Node Manager configuration file.

WLS_HOME\wlserver 10.3\common\nodemanager\NodeManagerLogs\ is the default logs directory for WebLogic Node Manager.

Additional configuration files are located in ...common\nodemanager.

To configure Node Manager:

1. Start the Admin Server.
2. Signon to the Administration Console.
3. Select Environment, Machine.

Create one local machine (as in MachineLocal) for all of the managed servers running in the local server and one remote machine (as in MachineRemoteX) for each of the managed servers running in remote servers.

4. Configure machines.

Click each of the machines created. Under the Configuration tab, select the Node Manger tab.

Choose *Plain* from the Type drop down list.

For Listen Address, enter the IP address of the server on which Node Manager is running.

Enter the Node Manger port number into the Listen Port text box. The default port number is 5556. This can be changed by updating the `NodeManager.properties` file.

5. Add managed servers.

Under the Servers tab, add the managed servers into the machine. For example, you may have two managed servers running locally, with *PIA1* and *RPS* assigned to *MachineLocal*.

6. Start the Node Manager.

Start the Node Manger from your local WebLogic install directory (BEA Home for the WebLogic Sever on which your PeopleSoft domain is referring to). Node Manager can be started by `startNodeManager.cmd/.sh` script under `BEA_Home\wlserver_10.3\server\bin` directory.

After the Node Manger has been started for the first time, a file called `nodemanager.properties` appears under the `BEA_Home\wlserver_10.3\common\nodemanager` directory.

7. Modify the `nodemanager.properties` file.

Stop the Node Manger and open the `nodemanager.properties` file in a text editor, and make the following changes and save the file.

`SecureListener=false`

`ListenAddress=` the IP of the box where the Node Manager is running

Note: In the same file you can see the default port number is set as `ListenPort=5556`. Changing the port setting here and restarting the Node Manager will set the Node Manager to be listening on the newly configured port number. The port setting here must match with the one set in the machine configuration.

8. Restart the Node Manager.

9. Once the Node Manger is up and running, it should be reachable from the machine configured to listen to the local Node Manager (as in, *MachineLocal*). This can be confirmed by the Node Manager Status tab under the Monitoring tab.

10. In the remote box hosting the managed server, follow the previous steps to configure the Node Manager.

11. From the bin directory under your `PS_Domain`, run `setEnv.cmd/sh` to set up the environment.

12. Start WLST using the command `java weblogic.WLST`.

13. Connect the remote server to the Admin Server running on the local box.

Following is the syntax of the `connect` command:

```
connect('username','password','t3://Local Machine's IP:port on which the admin
server is running on Local machine')
```

14. Enroll the remote domain and Node Manger into the Admin Server running on the local box.

The command is `nmEnroll` with the following syntax:

```
nmEnroll('full path to the distributed PeopleSoft Domain ', 'BEA_Home\wlserver_10.3\common\nodemanager')
```

For example:

```
nmEnroll('D:\PT850-809R2\websevr\Dis4103Installed603', 'D:\Wls103Installed603\wlserver_10.3\common\nodemanager')
```

15. Once the domain and Node Manager have been enrolled, check the Admin Console to confirm that the Node Manager can be reached from the server that was configured to listen to the remote Node Manager.

Once a Node Manager is reached from a "machine," the Admin Console will be able to start and stop the managed server assigned to that "machine."

The Servers tab shows the list of servers. The Control tab grants the control of those servers to the Admin Console.

Starting and Stopping a Distributed Managed Server

In a multi-server configuration, a distributed managed server is simply a managed server that isn't started from the same physical location as its domain's administration server. You can control the life cycle of these servers using scripts, the WebLogic Server Console and the WebLogic Node Manager.

See [Starting and Stopping Multi-Server Processes](#).

Tuning Performance and Monitoring Resources

Monitoring the performance of a WebLogic instance is primarily performed in the Administration Console. This section discusses how to:

- Manage JVM heap size and execute thread usage.
- Monitor HTTP session count for PeopleSoft Interaction Hub.

Related Links

[Using WebLogic Server Administration Console to Monitor PeopleSoft Sessions](#)

Managing JVM Heap Size

This section discusses how to:

- Monitor JVM heap.
- Change JVM heap size.

Monitoring JVM Heap

The JVM heap size is the amount of memory that a particular JRE (Java Runtime Environment) gives to the JVM (Java Virtual Machine) that it creates. The `java.exe` command on Windows, `java` on UNIX and `beasvc.exe` when running WebLogic as a Windows service is the JRE, and the JVM exists within the

JRE's memory space. The primary sources for monitoring the amount of memory that is in use within a JVM are the Administration Console and the WebLogic logs.

To monitor the amount of the JVM heap size available and in use:

1. Sign on to the Administration Console by entering the following URL in a browser:

`http://webserver:9999/console`

Where *webserver* is the hostname of the WebLogic server.

2. Traverse the following in the navigation tree on the left:
 - a. Expand your WebLogic domain (for example, peoplesoft).
 - b. Expand Servers.
3. Click the server you intend to monitor (for example, PIA).
4. Select the Monitoring tab, and the Performance sub-tab.

Changing the JVM Heap Size

If you need to adjust any of the Java options, most commonly the JVM heap size, you must manually edit that WebLogic domain's local setEnv script. The parameters, -Xms and -Xmx, control the JVM memory minimum and maximum heap size respectively.

Following are examples of the JVM heap size as specified in setEnv using the `JAVA_OPTIONS_OSplatform` environment variable. You only need to set the variables that correspond to the operating system where the WebLogic server is running.

See [Adjusting the JVM Heap Size](#).

Note: If you do adjust any of the Java options, most commonly the JVM heap size, you must restart WebLogic for these changes to take effect.

If you're running WebLogic Server as a Windows service you must rerun the `installNTservice` script to propagate this change into the Windows registry.

The WebLogic Node Manager does not use the Java options set in setEnv, but instead uses Java options set from the WebLogic console.

To modify the Java options for a WebLogic instance started via the WebLogic Node Manager:

1. Sign on to the WebLogic Server Console by entering the following URL in a browser:

`http://webserver:9999/console`

Where *webserver* is the hostname of the WebLogic server.

2. Expand your WebLogic domain (for example, peoplesoft) and click Environment, then Servers.
3. Select the managed server you intend to modify.
4. Select the Configuration tab, and the Server Start sub-tab.
5. Update the Arguments field.

6. Click Save.

Related Links

[Adjusting the JVM Heap Size](#)

Monitoring HTTP Session Count for PeopleSoft Interaction Hub

In addition to memory and thread usage, it's also possible to monitor the number of established HTTP sessions used in conjunction with the PeopleSoft PORTAL application. This number, although not necessarily directly related to current performance, is an indicator of the following performance factors:

- JVM memory used to store HTTP session data.
- Current number of logged on clients.
- Peak number of logged on clients.
- Idle time of logged on clients.

To monitor the total number of HTTP sessions:

1. Sign on to the Administration Console.
2. In the Domain Structure section, click Deployments and select the check box next to peoplesoft.
3. Select the Monitoring tab.

See the following My Oracle Support document for more information: *Clustering and High Availability for PeopleTools [ID 747378.1]*.

Changing Configuration Settings

This section provides an overview of the WebLogic server configuration files, and discusses how to:

- Change the WebLogicAdmin server's listen ports.
- Change application and server deployment targets.

Understanding the WebLogic Server Configuration Files

WebLogic server configuration settings are stored in a collection of files, primarily these include: script, config.xml, and the web.xml and weblogic.xml for each webapp.

Configuration File	Description
setEnv script	SetEnv contains statically and dynamically defined environment variables. It's called from all of the WebLogic administration scripts to assist in building the Java command line. You modify this file using a text editor.

Configuration File	Description
config.xml	Config.xml contains server runtime settings, such as the HTTP port. You modify this file using the Administration Console.
web.xml weblogic.xml	Located in the WEB-INF directory of each servlet, providing web application descriptors and settings relevant to their application.

Changing the WebLogicAdmin Server's Listen Ports

In the multi-server configuration, several parameters are set based on the environment detected and delivered defaults. One such parameter is the HTTP port of the WebLogicAdmin server. By default the WebLogicAdmin server's HTTP listen port is 9999.

To change this value:

1. Start the WebLogicAdmin server using the startWebLogicAdmin script.
2. Sign on to the WebLogic Server Console by entering the following URL in a browser:

`http://webserver:9999/console`

Where *webserver* is the hostname of the WebLogic server.

3. Navigate to Servers, WebLogicAdmin, Configuration, General.
4. Modify the value of the Listen Port field.
5. Click Apply.
6. Restart the WebLogic server.

If you can't initially start the server due to a port conflict, you can manually edit the value of the ListenPort parameter in that domain's config.xml file. Creating a backup of config.xml is recommended before manually changing this file.

After changing the ListenPort value in your domain's config.xml, either directly or using the console, you must also update your setEnv script. Update the ADMINSERVER_PORT environment variable to reflect the new HTTP port. This setting is used by the stopWebLogic and startManagedWebLogic scripts as the default administration server HTTP port.

Changing Application and Server Deployment Targets

With WebLogic, J2EE applications are targeted to any combination of WebLogic servers and WebLogic clusters. A WebLogic cluster is a logical grouping of servers, generally all providing the same application, though that's not a requirement. To change the servers or clusters that to which an application is targeted and deployed, sign on to the Administration Console and update the application's target assignments. You can view and modify application and server target assignments on the Deployments, Applications tabs, and on the Targets tab for each server.

Following is an example of how to change the target assignments of the PeopleSoft Integration Gateway (PSIGW) web application so it's the only application targeted to the PIA server, and is the sole application on that instance.

To change the target assignments of the PeopleSoft Integration Gateway web application:

1. Sign on to the WebLogic Server Console.
2. In the Domain Structure section:
 - a. Expand peoplesoft.
 - b. Expand Deployments.
 - c. Expand Applications.
 - d. Expand peoplesoft.
3. Select PSIGW.
4. Select the Targets tab.
5. In the Clusters section, clear the peoplesoftCluster check box.
6. Click Apply.
7. In the navigation tree, select PORTAL.
8. Select the Targets tab.
9. In the Independent Servers grid, clear the check box for targeting the PORTAL to this server.
10. Click Apply.
11. Repeat steps 7 to 10 for the PSEMHUB and remaining servlets.

To deploy an application to a cluster, target the server to the cluster and target the application to the cluster.

PeopleSoft Timeout Settings

Web Server Timeouts

You specify web server timeouts using the Web Profile Configuration component (WEB_PROFILE). To access these settings in PIA, select PeopleTools, Web Profile, Web Profile Configuration, then select the appropriate page.

The following table provides basic information about the web server timeout settings, which are more completely documented in the *PeopleTools Portal Technology* PeopleBook.

Page Element	Page Name	Description	Default
Inactivity Warning	Security	<p>Specify how long the portal should wait before warning users that their browser session is about to expire. They can continue with their current session by clicking the OK button in the message.</p> <p>If a user doesn't respond, the session ends and the expired connection page appears.</p> <p>Suppress this warning by setting this value to be greater than the sessionTimeout value.</p>	1080 seconds (18 minutes)
Inactivity Logout	Security	<p>Specify the inactivity timeout interval of the PeopleSoft application for which the user is currently authenticated. When the interval passes with no user activity, the user's browser displays the page specified by the Expire Page - Page field on the Web Profile Configuration - Look and Feel page.</p> <hr/> <p>Note: Depending on the application implementation, authenticated users might also experience an HTTP session inactivity timeout.</p> <hr/>	1200 seconds (20 minutes)

Page Element	Page Name	Description	Default
Authenticated Users - HTTP Session Inactivity	Security	<p>Specify the HTTP session inactivity timeout interval that applies to authenticated users.</p> <p>When the interval passes with no user activity, the web server discards all session information, including cached page states. The next time the user submits a request, the web server creates a new HTTP session.</p> <p>If not set, the HTTP interval for an authenticated user is the same value as the inactivity logout.</p>	0 seconds for all profile types.
Public Users - HTTP Session Inactivity	Security	<p>Specify in seconds the inactivity timeout interval that applies to public users.</p> <p>When the interval passes with no user activity, the web server discards all session information, including cached page states. The next time the user submits a request, the web server creates a new HTTP session.</p> <p>Unlike authenticated users, public users are not signed out of their PeopleSoft application when this interval expires.</p> <p>However, PIA releases their application states from memory. If users click a link, they regain access to the application at the search dialog. This setting prevents an overload of web server resources for inactive public users.</p>	<p>DEV, KIOSK profile: 1200 seconds (20 minutes).</p> <p>TEST, PROD profile: not set.</p>

Page Element	Page Name	Description	Default
Disconnect Timeout	Security	<p>Specify the amount of time to wait before disconnecting the Jolt connection.</p> <p>A value of 0 seconds (the default) means no limit. This means that the client connection must be retained throughout the session. If the connection becomes invalid (due to one of the other timeouts) the session will be expired.</p> <hr/> <p>Note: If you specify 0 seconds, the Jolt client attempts to connect the Jolt Server Handler (JSH) in RETAINED mode. If any positive value is specified, the Jolt client attempts to connect the JSH in RECONNECT mode.</p> <hr/>	0 seconds
Send Timeout	Security	<p>Specify the maximum time permitted between the sending of the Jolt Request by the client servlet and its full receipt on the application server.</p> <hr/> <p>Note: You might need to increase this value where a large amount of data is being sent to the application server, or the network is slow.</p> <hr/>	50 seconds

Page Element	Page Name	Description	Default
Receive Timeout	Security	<p>Specify how long the client servlet should wait after issuing a Jolt Request for a response from the application server.</p> <p>This value should be considerably larger than the Send Timeout. Make sure that this value is also greater than your application server online service timeouts, such as the Service Timeout setting for PSAPPSRV that appear in the PSAPPSRV.CFG configuration file on the application server.</p> <hr/> <p>Note: Ideally, this timeout should also be greater than the Tuxedo SANITY_SCAN setting (BLOCKTIME * SCANUNIT).</p> <hr/>	1300 seconds

Related Links

"Configuring Portal Security" (PeopleTools 8.53: Portal Technology)

Session-Timeout

You specify the web server *session-timeout* using the Inactivity Logout property and HTTP Session Inactivity property in the web profile.

Web Server Default System Timeout

PeopleSoft Interaction Hub normally depends on a content reference timeout setting to determine how long to wait for a pagelet to load before it considers the pagelet to be unavailable. However, if the remote server is unavailable, the content reference timeout setting is ignored. If the portal can't establish a connection to the remote host, it uses the default system timeout.

The default system timeout defaults to 20 seconds. If you expect the remote server to be slow or down for longer than 20 seconds, you should specify a longer default system timeout, by configuring your web server to set the defaultConnectTimeout JVM environment variable to an appropriate value using one of the following procedures.

For example,

```
SET JAVA_OPTIONS_WIN32=-server -Xmsnnm -Xmxnnm -XX:MaxPermSize=nnm
-Dsun.net.client.defaultConnectTimeout=default_timeout
```

Where *default_timeout* is the number of milliseconds that the portal should wait to establish the connection to the host.

See Your web server documentation for instructions on modifying this JVM environment variable.

Application Server Timeouts

All configurable settings for the application server require modification in PSADMIN:

Name	In This File	Description	Default
JOLT Listener/Client CleanupTimeout	psappsrv.cfg	<p>Specify the inactivity interval permitted for the server-side JoltSession.</p> <p>Specifying too low a value can cause unnecessary instantiation of resources for clients who surpass this inactivity interval. However, specifying too high a value can keep unnecessary server-side resources allocated.</p> <hr/> <p>Note: This setting doesn't affect the user experience, but it has an impact on server-side performance.</p>	10 minutes
JOLT Listener/Init Timeout	psappsrv.cfg	<p>Specify the amount of time that's allowed for the JSL process to start.</p> <hr/> <p>Note: It's not necessary to adjust this setting.</p>	5 minutes
Workstation Listener/Client Cleanup Timeout	psappsrv.cfg	<p>Specify the inactivity interval permitted for the server-side Workstation Listener Session.</p> <p>Specifying too low a value can cause unnecessary instantiation of resources for clients who surpass this inactivity interval. However, specifying too high a value can keep unnecessary server-side resources allocated.</p> <hr/> <p>Note: This value is required only for three-tier connections.</p>	60 minutes
Workstation Listener/init Timeout	psappsrv.cfg	<p>Specify the amount of time that's allowed for the WSL process to start.</p> <hr/> <p>Note: It's not necessary to adjust this setting. This value is required only for three-tier connections.</p>	5 minutes

Name	In This File	Description	Default
Spawn Threshold	psappsrv.cfg	<p>Applies only if spawning is enabled.</p> <p>Specify the rates at which PSAPPSRV processes spawn and decay.</p> <p>The spawn rate is determined by the last two numbers, and the decay rate is determined by the first two numbers.</p> <p>Using the default value as an example, for the spawn rate of 1,1 an extra PSAPPSRV process is spawned if there is at least 1 outstanding service request on the application server request queue for 1 second or more. This spawning will continue until the PSAPPSRV Max Instances value is reached.</p> <p>For the decay rate of 1,600 a server process is decayed if less than 1 service request is in the application server request queue for 600 seconds (ten minutes) or more.</p> <hr/> <p>Note: This parameter applies only if, for PSAPPSRV, the value of <i>Max Instances</i> is greater than that of <i>Min Instances</i>.</p> <hr/>	1,600:1,1

Name	In This File	Description	Default
Service Timeout	psappsrv.cfg	<p>Each server process has its own instance of this setting in its section of the psappsrv.cfg file.</p> <p>Specify the maximum interval for services to run in a given process. If a service has not completed within the specified interval, Tuxedo terminates the server processing and restarts the server process.</p> <p>For each server process, specify the longest time that any service is expected to take.</p> <hr/> <p>Note: A value of 0 produces an indefinite timeout for any service.</p>	<p>PSAPPSRV: 300 seconds (5 minutes)</p> <p>PSSAMSRV: 300 seconds</p> <p>PSQCKSRV: 300 seconds</p> <p>PSQRYSRV: 1200 seconds (20 minutes)</p> <p>PSBRKHND_dflt: 1200 seconds</p> <p>PSSUBHND_dflt: 1200 seconds</p> <p>PSPUBHND_dflt: 1200 seconds</p>
Restart Period (PSBRKDSP_dflt, PSSUBDSP_dflt, PSPUBDSP_dflt)	psappsrv.cfg	Specify how long each dispatcher should wait before redispersing a message if the associated handler has not started processing it.	120 seconds
TM_RESTARTSRV TIMEOUT	psappsrv.ubx (which is the template for psappsrv.env)	<p>Specify the time period that a domain server process (for example, PSAPPSRV, PSWATCHSRV, PSSAMSRV) is permitted to remain in REStarting mode before it is killed by Tuxedo. This setting resolves processes hanging during restart.</p> <hr/> <p>Note: To modify this setting, you must change the value in the .UBX template file, then recreate your domain.</p>	60 seconds (one minute)

Process Scheduler Timeouts

All configurable settings for PeopleSoft Process Scheduler require modification through domain configuration within PSADMIN:

Name	In This File	Description	Default
Process Scheduler/ Reconnection Interval	psprcs.cfg	Specify the interval between attempts to reconnect to the database when the connection is lost.	300 seconds (5 minutes)
Process Scheduler/ Authentication Timeout	psprcs.cfg	Specify how long PeopleSoft Security has to authenticate a process that's released by PeopleSoft Process Scheduler. The timer starts when Process Scheduler initiates the request.	5 minutes
RemoteCall/RCCBL Timeout	psprcs.cfg	Specify the maximum interval for a remote call from an Application Engine program to run before it's terminated. This is similar to a general Tuxedo service timeout.	300 seconds (5 minutes)

For Spawn Threshold, see the application server timeout settings.

Search Server Timeouts

The following are the configurable timeout settings for the search server.

Name	File	Description	Default
Domain Settings / Spawn Threshold	pssrchrsv.cfg	<p>Applies only if spawning is enabled.</p> <p>This is the rate at which PSSRCHSRV processes will spawn and decay. The spawn ratio is determined by the last two digits. The decay ratio is determined by the first two digits. Using the default value as an example, we see that an extra PSSRCHSRV process will be spawned if there is at least 1 outstanding service request on the request queue for one second or more. This spawning will continue until Max Instances is reached. For the decay rate of 1, 600, if less than 1 service request is on the request queue for ten minutes (600 seconds), a server process is decayed. Note: This value is only relevant if $\text{PSSRCHSRV} / \text{Max instances} > \text{PSSRCHSRV} / \text{Min Instances}$.</p>	1,600:1,1
PSSRCHSRV / Service Timeout	pssrchrsv.cfg	This parameter indicates the duration in seconds to run a Search service within a Search domain.	300 secs
TM_RESTARTSRVTIMEOUT	pssrchrsv.ubx (and then UBBGENned into pssrchrsv.env)	The time period that a domain server process PSSRCHSRV, is allowed to remain in Restarting mode before it is killed by the BBL. This resolves processes hanging during restart. This setting is defaulted in the <code>\$PS_CFG_HOME/appserv/Search/*</code> . UBX files. If this value needs to be changed, you must change the value in the UBX template file and then recreate your domain.	60 secs

PIA Timeouts

A number of additional timeouts may be set through PIA. These settings reflect changes at the database level that may pertain to different groups of users.

Note: The timeout settings in PIA are optional and are not required to run PIA. However, an understanding of how these settings can contribute to a user's session duration is important in the context of other timeout values that appear in this topic.

Name	Navigation Path	Description	Default
Authentication Token expiration time	PeopleTools, Security, Security Objects, Single Signon	<p>Specify the interval during which the system can trust a single signon token (PS_TOKEN) from the same or another content provider.</p> <hr/> <p>Note: As long as users remain signed in, the expiration of PS_TOKEN does not affect them. This setting is relevant only for the GetCertificate request during single signon.</p>	720 minutes (12 hours)
Permission List - Time-out Minutes	PeopleTools, Security, Permissions & Roles, Permission Lists	<p>Specify an interval during which a given permission list applies. The interval starts for a user to which the permission list is assigned when that user signs in. When the timeout period elapses, the user's online session is terminated.</p> <p>If a user belongs to multiple permission lists, the largest timeout value from among those permission lists is applied to the user's session during signon. The permission list timeout is effective only if its value is less than the web server session-timeout. This means that all of the permission list timeouts for a given user must be less than the web server session-timeout to be effective. However, the Inactivity Warning timeout still applies.</p> <hr/> <p>Note: A value of 0 produces an indefinite timeout.</p>	0 minutes

Related Links

[Web Server Timeouts](#)

Appendix D

Ensuring Session Stickiness

Working With Session Cookie Names

Session stickiness, generally, refers to ensuring that the requests for a particular session get routed to the same physical machine that serviced the first request for that session. This helps to reduce the situation where sessions are lost as a result of requests for a session being routed to different servers.

For any of the following scenarios, all PeopleSoft web servers need to have the same session-cookie name:

- The Load-Balancer uses the PeopleSoft web server's session-cookie as the stickiness criterion.
- Any of the any the supported proxy servers is used to forward requests to the PeopleSoft web servers.

Ensure that the session-cookie name is the same for all web server instances to which the Load Balancer or proxy dispatches requests. The cookie name needs to be configured in the appropriate configuration section of the Load-Balancer or proxy server.

Using PeopleSoft Interaction Hub and Content Providers

Typically, PeopleTools recommends that the portal and content sites run in separate PeopleSoft Internet Architecture domains. If these sites are running in the same PeopleSoft Internet Architecture domain (as in having multiple sites within the same domain, with one site for the portal and the other sites for content), then there is no need to load-balance requests separately to the content sites, because the requests will be routed to the same JVM's.

When the portal and content sites run in the same PeopleSoft domain as multiple sites, they share the same session because the web application is the same for all sites. PeopleTools does not recommend configuring portal and content sites in the same PeopleSoft domain for performance and reliability reasons.

So assuming that portal and content sites are running in separate PeopleSoft domains, they will have separate ports and, usually, separate hosts. In this case, if portal and content sites are using the same physical load balancer, the routing rules will have to be set up for these sites separately.

When using PeopleSoft Interaction Hub and Content Providers:

- The session-cookie name of PeopleSoft Interaction Hub and the Remote Content Provider must be *absolutely* distinct. For example,

PeopleSoft Interaction Hub's session-cookie name: PTENPRTL-8080-PORTAL-PSJSESSIONID

Content Provider's session-cookie name: PTFIN-9090-PORTAL-PSJSESSIONID

- Neither of the systems' session-cookie name must be part of the other's session-cookie name. The following example illustrates a *bad choice* of a session-cookie name:

PeopleSoft Interaction Hub's session-cookie name: PORTAL-PSJSESSIONID

Content Provider's session-cookie name: PTFIN-9090-PORTAL-PSJSESSIONID

- If there are multiple PeopleSoft Interaction Hub or Content Provider web servers with Load-Balancer or Proxy, see carefully choose a session-cookie name.

For example, consider an environment having a Load Balancer in front of two PeopleSoft Interaction Hub web servers, with the session-cookie's name used as Load Balancer's stickiness criterion. Suppose the name of the session-cookie is PTENPRTL-8080-PORTAL-PSJSESSIONID. In this case, both PeopleSoft Interaction Hub web servers must have the same session-cookie name: PTENPRTL-8080-PORTAL-PSJSESSIONID.

If this environment has an HCM Load balancer with four HCM web servers, with the session-cookie name as stickiness criterion, each of the HCM web servers must have the same cookie, like HTR-9090-PORTAL-PSJSESSIONID and this name must be distinct from the name of the session-cookie used for PeopleSoft Interaction Hub.

Working With Absolute URLs

PeopleSoft Internet Architecture constructs absolute URLs in certain scenarios (such as the Homepage URL to which the user is re-directed after successful login). Certain Load Balancers and proxy servers change the HOST header of the request to the target web server's name before forwarding the request. This causes PeopleSoft to construct the URL with the host-name of the web-server, rather than the name of the Load Balancer or proxy, which means that subsequent requests would be sent directly to the web server, instead of routing requests by way of the Load Balancer or proxy. In order to ensure that URLs can be constructed with the Load Balancer or proxy's host name, the Virtual Addressing page of the Web Profile needs to be set with the Load Balancer or proxy information.

To modify virtual addressing information:

1. Select PeopleTools, Web Profile, Web Profile Configuration, and open the web profile configured for the site.
2. Click on the Virtual Addressing tab and populate the Default Addressing section with the Load-Balancer or Proxy's host name, scheme and port.

Example: Suppose the system has been set up to be accessed through a Load Balancer which has a host-name of mycompany.com and accepts http requests at port 8080. Suppose the PeopleSoft site's name is ps. The URL for accessing the site would be as follows:

`http://mycompany.com:8080/ps/signon.html`

In this case, the Default Addressing section will need to be set with Protocol: http, Port: 8080, and Host: mycompany.com.

3. Save the Web Profile.
4. Re-start the web server.

Setting The Cookie Domain

Browsers enforce the rules stated by RFC 2965 for sending cookies to servers. If the cookie's domain is not set when the browser receives the cookie, that cookie is associated with the server that sent the cookie. If the cookie's domain has been set but does not match the name or the domain suffix of the server, then the browser does not send the cookie. Therefore, not having the correct cookie domain can lead to failure in maintaining a PeopleSoft session.

The cookie domain value is configured in web server-specific configuration by the PeopleSoft Internet Architecture setup if the Authentication Token Domain value has been entered. If this field was left blank during PeopleSoft Internet Architecture setup, either re-install PeopleSoft Internet Architecture or configure the cookie domain manually in the web server-specific configuration.

Working with Load Balancer Timeouts

If session-stickiness is configured with a timeout on the Load Balancer, that timeout value must be higher than the PeopleSoft Inactivity Logout timeout configured in the current Web Profile.

Keep in mind the session cookie name format. The session-cookie name must not start with a digit (0-9) and must not contain dots.

Example of bad name for session-cookie: 11.12.13.14-80-PORTAL-JSESSIONID

Example of good name for session-cookie: FINTSTWEB-80-PORTAL-JSESSIONID

Configuring The Session Cookie Name for WebSphere

To configure session cookie name:

1. Open a browser and access the WebSphere Administrative Console.
2. Go to Enterprise Applications, and select the PeopleSoft application and select the Manage Modules link.
3. Click on Portal.war.
4. Click on Session Management under Additional properties.
5. Make sure that the check box Override Session Management is enabled.
6. Click the Enable cookies link.
7. Set the cookie name.

The recommended format is:

<hostname>-<port>-PORTAL-PSJSESSIONID

Example: PSTST-8080-PORTAL-PSJSESSIONID

The default name of session-cookie is JSESSIONID.

8. Re-start the WebSphere Server to make this change effective.

Configuring The Cookie Domain for WebSphere

To configure the cookie domain:

1. Open a browser and access the WebSphere Administrative Console.
2. Go to Enterprise Applications, and select the PeopleSoft application and select Manage Modules link.
3. Click on Portal.war.
4. Click on Session Management under Additional properties.
5. Make sure that the check box Override Session Management is enabled.
6. Click Enable cookies link
7. Enter the value in the Cookie domain field.
8. Re-start the WebSphere Server to make this change effective.

Configuring the Session Cookie Name for WebLogic

The session cookie name is configured in weblogic.xml of the portal servlet (web application). This file can be found in the following directory: <PIA_HOME>/webserv/<DOMAIN-NAME>/applications/peoplesoft/PORTAL/WEB-INF

Locate the tag CookieName in weblogic.xml, as in the following example:

```
<session-param>
  <param-name>CookieName</param-name>
  <param-value>PSTST-8080-PORTAL-PSJSESSIONID</param-value>
</session-param>
```

The name of session cookie is assigned in the tag <param value>. In this example, the session-cookie name is PSTST-8080-PORTAL-PSJSESSIONID.

Re-start the WebLogic Server to make this change effective.

Configuring Session Cookie Domain for WebLogic

Update the CookieDomain manually in weblogic.xml file, as in the following example:

```
<session-param>
  <param-name>
    CookieDomain
  </param-name>
  <param-value>
    .company.com
  </param-value>
```

In this example, the value .company.com is the domain that will be assigned to the session-cookie.

Re-start the WebLogic Server to make this change effective.

Troubleshooting Server Issues

Uploading Files Using Non-Latin Characters

When uploading an attachment file with a filename that uses non-Latin characters, such as Hebrew or Japanese, the system uses the exact filename or UTF-8 % sequenced numbers. In previous releases, the system renamed the file and converted to strings of numbers.

When uploading non-Latin characters, the following error can occur: "Attachment Filename is too long (maximum is 64 characters)."

Solution For UNIX

To resolve this issue on UNIX servers:

1. Make sure that the corresponding language locale, such as the UTF-8 character set, is installed at the operating system level.

Use "locale -a" to get the list of installed locales.

2. Use the LANG environment variable to set for the session that runs the application server to a valid locale for the character set.

For example, en_US.utf8, iw_IL.utf8, ja_JP.utf8, and so on. See your "locale -a" output.

3. Use PSADMIN to set the Character Set parameter to the appropriate value, as in Character Set=utf8.

The character set values specified in the shell locale (LANG) and PSAPPSRV.CFG (Character Set) must agree, and, in this case, PSAPPSRV.CFG requires utf8.

Note: If the shell locale is set to something other than the default of C, Tuxedo will issue warnings when the application server boots. In such a case, add a symlink from \$TUXDIR/prod/locales/<your_shell_locale> to \$TUXDIR/prod/locales/C to resolve that issue.

Solution For Windows

On a Windows-based application server, make sure the server operating system has the corresponding language installed.

To resolve this issue on Windows servers:

1. Make sure that the corresponding language locale, such as CP932 (for Japanese) character set, is installed at the operating system level.

Use the Regional and Language Options utility in the Control Panel to get the list of installed locales, and use "chcp" in the command prompt window to see the current default code page.

2. Use Control Panel/Regional and Language Options to set the system locale for the session that runs the application server with a valid language.

For example, on the Advanced tab for the Language for non-Unicode programs option, select *Japanese* for the Japanese system locale of CP932.

3. Use PSADMIN to set the Character Set parameter so that it reflects the appropriate code set.

Using the Japanese example: the default code page CP932 (equivalent to SJIS) and the character set value in PSAPPSRV.CFG must agree. So in this case, PSAPPSRV.CFG requires “sjis” to be set as the character set, as in Character Set=sjis.

WebSphere: Port Set In the Host Header of a Request Returned Incorrectly

This issue applies only to configurations involving the IBM WebSphere web server and either of the following reverse proxy servers:

- Microsoft IIS
- Oracle Sun Java System Web Server

Scenario

This issue appears in the following situation (or similar). Assume there is PeopleSoft Interaction Hub deployed on WebSphere 1, say Site A, which accesses a pagelet in a PeopleSoft Interaction Hub deployed on WebSphere 2, say Site B. In front of Site A, there is reverse proxy server (Oracle Sun Java System Web Server or IIS) configured.

In such a scenario, when a HTTP request is made to the PeopleSoft Interaction Hub on Site A through an RPS port, Site A returns the HTTP request to fetch the remote pagelet from Site B. The HTTP response from Site B contains the RPS port number instead of the HTTP port of Site A. This is due to the default behavior of WebSphere, which reads the port number from the Request URL instead of from the HTTP Header (HOST field).

See <http://www-01.ibm.com/support/docview.wss?uid=swg1PK55330>.

Solution

This default behavior can be altered by modifying the web container custom properties in WebSphere on Site B.

To modify the custom properties:

1. Open the Administrative Console and select Servers, Server Types, WebSphere application servers, server1, Web Container Settings, web container.
2. Set these custom properties:
 - trusted = false

- `trusthostheaderport = true`
 - `com.ibm.ws.webcontainer.extractHostHeaderPort = true`
3. Restart the web server.

Replicating an Installed Environment

Understanding Environment Replication

Environment Replication involves taking a working, well-tested server elements, and copying the PeopleTools binary and configuration files to a new location to create a new environment that required only making minor modifications to the new configurations.

The server elements that can be replicated appear in the following:

<i>Server Element</i>	<i>Description</i>
PeopleSoft Internet Architecture	A web server environment consists of one or more site directories under the <i>PIA_HOME</i> \webserv directory within a single <i>PIA_HOME</i> location. Each directory contains configurations that point to a single application server domain through a JSL port designation. Although a site may point to more than one redundant application server machine name for failover or load balancing purposes, each machine name is given a unique JSL port number. Additional <i>PIA_HOME</i> directory locations on the same web server machine are considered separate environments.
application server domains	A single application server environment consists of one or more domain directories under the \appserv directory within a single <i>PS_CFG_HOME</i> location. Each domain contains configuration settings that point to a single database. Multiple domains can be configured to point to the same database for failover or load balancing. Each domain has its own server processes and must be configured to have unique WSL and JSL port numbers.
Process Scheduler domains	A single Process Scheduler Environment consists of one or more domain directories under the \appserv\prcs directory within a single <i>PS_CFG_HOME</i> location. Each database that needs to schedule processes must have its own Process Scheduler server configured separately. The configuration files are kept in their own directory location under the \prcs directory.

Important! When you configure and run an Environment Management agent, hub, or viewer, files associated with that Environment Management Framework are created and updated with information that refers to the absolute directory structure of your PeopleSoft system. When you replicate your installed environment, the information in those files is no longer valid for the new location. You must reconfigure the Environment Management Framework to recognize the environment changes.

Replicating a Web Server Environment

Use the following steps to replicate the web server environment:

1. Copy the web server *PIA_HOME*\webserv directory structure to a new location, whether it is a new machine, or same machine with a new high level *PIA_HOME* directory name.
2. Review the following configuration files to make sure they reflect the new location:
 - piainstall.xml: PS_HOME and host name
 - psconfig.sh: PS_HOME and path information
 - peopletools.properties: install location
3. Run *PS_HOME*\setup\PsmPPIAInstall\setup.exe or equivalent from the copied/cloned location.
4. Select Redeploy PeopleSoft Application.
5. Choose a site name.
6. Enter the application server name and port information.

Note: Be sure to enter a different JSL port value when replicating on the same machine.

Replicating PS_CFG_HOME

The application server and Process Scheduler domains exist within PS_CFG_HOME directory structure. You replicate both of these server elements by using the Replicate Config Home option in PSADMIN. The Replicate Config Home option enables you to clone a PS_CFG_HOME to a second location using PSADMIN internal functions to ensure that the environment is correct and cloned domains are readily bootable. The destination PS_CFG_HOME can reside on the same or a separate host. PSADMIN accounts for and modifies any host-specific settings and omits files that do not require cloning, such as log files.

This feature enables you to create a PS_CFG_HOME containing all the required domains that is reliable, tuned, and configured correctly for your site and deploy it to multiple locations across the enterprise. Replicating PS_CFG_HOME can be useful in a variety of scenarios, such as creating:

- backup environments.
- testing environments.
- scaling environments.
- failover environments.

Note: When replicating a source PS_CFG_HOME, PSADMIN overwrites the destination PS_CFG_HOME and any domains currently existing within it. If the destination PS_CFG_HOME contains domains, PSADMIN alerts you before proceeding.

Note: PS_CFG_HOME/data is also copied from the source PS_CFG_HOME to the destination PS_CFG_HOME.

Note: File cache is also replicated in the destination PS_CFG_HOME. Consider that replicating a PS_CFG_HOME with 30 domains with file cache will take longer to copy over a network to a remote host than replicating a PS_CFG_HOME with two domains and no cache.

Important! Using the PSADMIN replication feature is not designed to support replication from one operating system to another. For example, you *cannot* configure a domain on UNIX and then replicate that domain on a Windows server. You can replicate between environments on the same operating system platform only.

To replicate a PS_CFG_HOME:

1. Set PS_CFG_HOME to reflect the *destination* PS_CFG_HOME location.
2. Launch PSADMIN.
3. Select 5) Replicate Config Home from the PeopleSoft Server Administration menu.

```
-----
PeopleSoft Server Administration
-----
```

```

1) Application Server
2) Process Scheduler
3) Search Server
4) Service Setup
5) Replicate Config Home
q) Quit
```

```
Command to execute (1-5, q): 5
```

4. When prompted, enter the location of the destination PS_CFG_HOME and press ENTER.
5. When the process completes, verify the contents of the destination PS_CFG_HOME.
6. Use PSADMIN to make any adjustments to the replicated domain environment information as needed.

For example, consider the following items for the new host:

- database connectivity software location.
- port availability
- cache location
- file system settings
- (if replicating to the same machine) JSL and WSL port conflicts

Reconfiguring Replicated Environment Management Components

This section discusses how to:

- Reconfigure an environment management agent.
- Reconfigure the environment management hub.
- Reconfigure the environment management viewer.
- Update the GUID value of cloned databases.

Reconfiguring an Environment Management Agent

Configuring PSEMAgent in a replicated PS_HOME involves deleting the agent local cache and modifying the PS_HOME location specified within the PSEMAgent configuration files.

Note: *New_PS_HOME* refers to the directory where PeopleTools is located in your *replicated* environment. If the same PS_HOME location is used by your source and destination environment, you do not need to make any changes.

To reconfigure the agent:

1. Delete all subdirectories below *New_PS_HOME\PSEMAgent\envmetadata\data*

Note: Don't delete the *matchers.xml* file in this location.

2. Delete the following directories:

- *New_PS_HOME\PSEMAgent\envmetadata\PersistentStorage*
- *New_PS_HOME\PSEMAgent\envmetadata\scratchpad*
- *New_PS_HOME\PSEMAgent\envmetadata\transactions*

3. Modify *StartAgent.bat* (Windows), or *StartAgent.sh* (UNIX).

Ensure that references to the drive and path of the PSEMAgent directory are correct for the *New_PS_HOME* location.

4. Modify *StopAgent.bat* (Windows) or *StopAgent.sh* (UNIX).

Ensure that references to the drive and path of the PSEMAgent directory are correct for the *New_PS_HOME* location.

5. Verify the settings in the agent configuration file: *New_PS_HOME\PSEMAgent\envmetadata\config\configuration.properties*.

Properties that specify path locations must be valid for the replicated agent in its new location, but the same hub is still addressed by all agents.

- *hubURL*

- windowsdrivestocrawl
- unixdrivestocrawl

Reconfiguring the Environment Management Hub

The Environment Management hub is replicated as part of replicating the web server environment. After redeploying PeopleSoft Internet Architecture, clear the cache by performing the instructions below:

Note: *New_PIA_HOME* is the directory where PeopleSoft Internet Architecture is located in your replicated environment.

To reconfigure the hub:

1. Delete all subdirectories below *New_PIA_HOME*\webserv\peoplesoft\applications\peoplesoft\PSEMHUB\envmetadata\data\

Note: Don't delete the data.txt file in this location.

2. Delete the following directories:

- *New_PIA_HOME*\webserv\peoplesoft\applications\peoplesoft\PSEMHUB\envmetadata\PersistentStorage
- *New_PIA_HOME*\webserv\peoplesoft\applications\peoplesoft\PSEMHUB\envmetadata\scratchpad
- *New_PIA_HOME*\webserv\peoplesoft\applications\peoplesoft\PSEMHUB\envmetadata\transactions

Reconfiguring the Environment Management Viewer

In the replicated environment, local cache files need to be cleared for the *New_PS_HOME*\PSEMViewer directory.

Note: *New_PS_HOME* is the directory where PeopleTools is located in your replicated environment.

To reconfigure the PSEMViewer:

1. Delete all subdirectories below *New_PS_HOME*\PSEMViewer\envmetadata\data\

Note: Don't delete the files in this location, just the subdirectories.

2. Delete the following directories:

- *New_PS_HOME*\PSEMViewer\envmetadata\PersistentStorage
- *New_PS_HOME*\PSEMViewer\envmetadata\scratchpad

Updating the GUID Value of Cloned Databases

When copying databases, it is extremely important to delete the GUID value in the new (copied) database. If not deleted, the hub will assume that the old and new environments are the same, leading to confusing environment records. To resolve this, set the value of the GUID field in the PSOPTIONS table to <space> in the cloned database using the SQL tool at your site. The next time an application server connects to the database, the system generates a new, unique GUID for the cloned database.

Working with PSADMUTIL

Understanding PSADMUTIL

The PSADMUTIL utility runs outside of PSADMIN for the purposes of:

- preloading file or database cache.
- purging database cache.

Preload cache projects as well as stored file and database cache repositories can be very large. Because of this, it is not advantageous to have caching operations subject to Tuxedo timeout limits. PeopleTools uses PSADMUTIL so that the operations it handles run independently of PSADMIN and Tuxedo.

When running the preload cache or purge cache options from PSADMIN, the PSADMIN utility calls PSADMUTIL to perform these operations.

The PSADMUTIL utility can also be invoked outside of PSADMIN from the command line.

Using PSADMUTIL from the Command Line

The psadmutil.exe is located in the PS_HOME\bin\server directory for the application server installation. The command line interface for psadmutil.exe uses the following syntax:

```
psadmutil -CT dbtype [-CS server name] -CD database_name -CO user_ID -CP user_⇒  
password  
-CI connect_ID -CW connect_password [-Preload [ProjectName]] [-Purge]
```

The following table describes the accepted command line parameters.

Parameter	Argument	Example
-CT	Specify the database type. Valid values are DB2ODBC, DB2UNIX, INFORMIX, MICROSOFT, ORACLE, and SYBASE. <hr/> Note: Notice the spelling of MICROSOFT. DB2ODBC is the database type for DB2 z/OS. <hr/>	-CT ORACLE
-CS	(Optional) Specify the name of the database server for the database to which you're connecting. Note. This parameter is required only if you specify INFORMIX or SYBASE as the database type.	-CS pt-sun05
-CD	Specify the name of the database to connect to, as you would when signing in to PeopleSoft.	CD- FSDMO
-CO	Specify the PeopleSoft user ID you're using to sign in.	-CO VP1
-CP	Specify the user password for the PeopleSoft user ID you specified.	-CP VP1
-CI	Specify the connect ID used to connect to the database server.	-CI people
-CW	Specify the password for the Connect ID you specified.	-CW people
-Preload [preload cache project]	Invokes the preload cache operation. If this parameter is used, you must also specify the name of the preload project. <hr/> Note: If you do not specify a cache project name, the system reads the value from the server configuration file (PSAPPSRV.CFG or PSPRCS.CFG file), based on the PS_SERVER_CFG file. <hr/>	-Preload PLC_HR_PAYROLL
-Purge	Purges the cache for the specified database. <hr/> Note: This option only applies if database caching is implemented. <hr/>	-Purge

Note: If any of these parameters are not specified, the system reads them from the PSAPPSRV.CFG or PSPRCS.CFG file. For example, for security reasons, the database connection related parameters do not need to be specified on the command line, and the system can get these values from the server configuration files.

Note: Before running psadmutil.exe from the command line, make sure PS_SERVER_CFG and PS_SERVDIR are set correctly. When PSADMIN calls psadmutil.exe, these environment variables are already set by PSADMIN to the valid values.
