

# EnterpriseOne Supply Chain Planning 8.12.1 Hardware and Software Requirements Guide

---

**June 2007**

**Copyright © 2003, 2007, Oracle. All rights reserved.**

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

#### U.S. GOVERNMENT RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

#### ***Open Source Disclosure***

The Advanced Planning applications have embedded within them a variety of third party software components, each of which has its own copyright and licensing terms. The Advanced Planning applications containing third party components from the development organizations identified below are identified in the parenthetical following the name of the organization.

## Apache Software Foundation

Copyright (c) 2000 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The end-user documentation included with the redistribution, if any, must include the following acknowledgment:

“This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).”

Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.

4. The names "Apache" and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [apache@apache.org](mailto:apache@apache.org).
5. Products derived from this software may not be called "Apache", nor may "Apache" appear in their name, without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

Portions of this software are based upon public domain software originally written at the National Center for Supercomputing Applications, University of Illinois, Urbana-Champaign.

## ptmalloc

Copyright (c) 1999 Wolfram Gloger

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the name of Wolfram Gloger may not be used in any advertising or publicity relating to the software.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL WOLFRAM GLOGER BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

### **Sleepycat Software**

Copyright (c) 1990, 1993, 1994 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
4. This product includes software developed by the University of California, Berkeley and its contributors.
5. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

### **Tool Command Language (TCL)**

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, ActiveState Corporation and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

GOVERNMENT USE: If you are acquiring this software on behalf of the U.S. government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as defined in Clause 252.227-7013 (c) (1) of DFARs. Notwithstanding the foregoing, the authors grant the U.S. Government and others acting in its behalf permission to use and distribute the software in accordance with the terms specified in this license.

#### **Independent JPEG Group**

This product includes software developed by the Independent JPEG Group. Copyright (c) 1991-1998 The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

#### **Henry Spencer's Regular Expression Library (REGEX)**

This product includes software developed by Henry Spencer. Copyright (c) 1992, 1993, 1994, 1997 This software is not subject to any license of the American Telephone and Telegraph Company or of the Regents of the University of California. Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it, subject to the following restrictions:

1. The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
2. The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
3. Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
4. This notice may not be removed or altered.

**XBAE**

Copyright (c) 1991, 1992 Bell Communications Research, Inc. (Bellcore)

Copyright (c) 1995-99 Andrew Lister

All Rights Reserved.

Permission to use, copy, modify and distribute this material for any purpose and without fee is hereby granted, provided that the above copyright notices and this permission notice appear in all copies, and that the name of any author not be used in advertising or publicity pertaining to this material without the specific, prior written permission of an authorized representative of Bellcore and current maintainer.

BELLCORE AND OTHER CONTRIBUTORS MAKE NO REPRESENTATIONS AND EXTEND NO WARRANTIES, EXPRESS OR IMPLIED, WITH RESPECT TO THE INFORMATION, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR ANY PARTICULAR PURPOSE, AND THE WARRANTY AGAINST INFRINGEMENT OF PATENTS OR OTHER INTELLECTUAL PROPERTY RIGHTS. THE SOFTWARE IS PROVIDED "AS IS", AND IN NO EVENT SHALL ANY AUTHOR OR ANY OF THEIR AFFILIATES BE LIABLE FOR ANY DAMAGES, INCLUDING ANY LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES RELATING TO THE INFORMATION.

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

## Table of Contents

<b>Minimum Technical Requirements for EnterpriseOne Supply Chain Planning 8.12.1 .....</b>	<b>2</b>
Hardware Requirements .....	2
Software Requirements.....	3
Software Requirements for Thin Client Deployments .....	3
<b>EnterpriseOne Supply Chain Planning 8.12.1 Deployment Options.....</b>	<b>4</b>
EnterpriseOne Supply Chain Planning 8.12.1 Deployment Methods .....	4
Standalone Installations.....	6
Thin Client Deployments.....	7
Thin Client Deployments from Microsoft Windows Servers.....	8
Thin Client Deployments from UNIX Servers .....	10
Web Application Deployment .....	13
Deploying Applications Through Firewalls.....	15
<b>Operating System Tuning for EnterpriseOne Supply Chain Planning Applications .....</b>	<b>17</b>
Operating System Resource Limits.....	17
Extending Operating System Resources .....	20
Performance Tuning for CPU and Memory-Bound Applications.....	22
Performance Tuning for Applications with Object Databases.....	25
<b>Integrating EnterpriseOne Supply Chain Planning and EnterpriseOne Supply Chain Management Systems .....</b>	<b>27</b>
Extracting Data from EnterpriseOne Supply Chain Management Systems .....	27
Managing Large Volumes of Data.....	29
Integrating EnterpriseOne Supply Chain Planning Systems with Multiple EnterpriseOne Supply Chain Management Systems .....	29
Launching Executables on Separate Machines .....	30

# MINIMUM TECHNICAL REQUIREMENTS FOR ENTERPRISEONE SUPPLY CHAIN PLANNING 8.12.1

This section discusses:

- Hardware requirements.
- Software requirements.
- Software requirements for thin client deployments.

## Hardware Requirements

To improve the performance of EnterpriseOne Supply Chain Planning applications, you can upgrade your hardware resources from the minimum levels.

In particular, EnterpriseOne Supply Chain Planning applications require significant CPU resources for creating optimal supply chain plans and schedules. To improve the performance of a EnterpriseOne Supply Chain Planning solver, you can invest in a more powerful CPU. However, having more than one CPU per solver does not improve solve times.

Solver performance is also significantly affected by the size of a data model, the numbers of constraints and variables, and the longer the data horizon. To improve solver performance, consider dividing large data models into multiple models with separate solves.

EnterpriseOne Supply Chain Planning application performance is also affected by the number of users: the greater the number of users, the greater the CPU requirements. To improve the performance of EnterpriseOne Supply Chain Planning applications with many users, such as Production and Distribution Planning, you can invest in multiple CPUs.

For further sizing information, please contact your hardware vendors.

## Software Requirements

For information about the software required for running EnterpriseOne Supply Chain Planning applications, please see the minimum technical requirements on the Customer Connection web site:

<http://www.peoplesoft.com/corp/en/login.jsp>. The minimum technical requirements list supported operating systems, as well as databases, web servers, and browsers that have been tested with EnterpriseOne Supply Chain Planning applications.

The minimum technical requirements list the versions of the software that we have tested with each EnterpriseOne Supply Chain Planning application. We have not tested and do not support subsequent versions of the software for use with EnterpriseOne Supply Chain Planning applications.

## Software Requirements for Thin Client Deployments

The following software packages are certified for deploying EnterpriseOne Supply Chain Planning 8.12.1 applications in distributed environments:

<b>3<sup>rd</sup> Party Product</b>	<b>Certified Version(s)</b>
Microsoft Terminal Services	Windows 2003
Citrix Presentation Server	2, 3
Hummingbird Exceed	10, 2006
	<b>Note.</b> Exceed 9 requires patch level 22 or greater.
Hummingbird Exceed onDemand	5.0
Oracle Application Server	See the minimum technical requirements for EnterpriseOne Supply Chain Planning applications on the PeopleSoft Customer Connection web site.
IBM Websphere	<b>Note.</b> Oracle Application Server is not currently certified for deploying PDP Collaborative Web Client.

### See Also

EnterpriseOne Supply Chain Planning 8.12.1 Deployment Options

Thin Client Deployments

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

# ENTERPRISEONE SUPPLY CHAIN PLANNING 8.12.1 DEPLOYMENT OPTIONS

This chapter discusses:

- EnterpriseOne Supply Chain Planning 8.12.1 deployment methods.
- Standalone installations.
- Thin client deployments.
- Thin client deployments from Microsoft Windows servers.
- Thin client deployments from UNIX servers.
- Web application deployments.
- Deploying applications through firewalls.

## **EnterpriseOne Supply Chain Planning 8.12.1 Deployment Methods**

EnterpriseOne Supply Chain Planning 8.12.1 includes a number of planning applications, each with its own possible deployment methods. Possible deployment methods for EnterpriseOne Supply Chain Planning 8.12.1 applications include the following:

- Standalone installations, in which EnterpriseOne Supply Chain Planning applications are installed on and accessed from a single machine.
- Thin client deployments, in which EnterpriseOne Supply Chain Planning applications are installed on one or more central servers and accessed from Microsoft Windows clients.
- Web application deployment, in which web applications are installed on central servers and accessed from web browsers.

This table lists the possible deployment strategies for each EnterpriseOne Supply Chain Planning 8.12.1 application:

EnterpriseOne Supply Chain Planning Application	Standalone Deployment	Thin Client Deployment	Web Application Deployment
Advanced Planning Agent	✓	✓	
Demand Consensus			✓
Demand Forecasting	✓	✓	
Order Promising			✓
Production and Distribution Planning	✓	✓	
PDP Collaborative Web Client			✓
Production Scheduling	✓	✓	
Production Scheduling - Process	✓	✓	
Sales & Operations Planning			✓
Strategic Network Optimization	✓	✓	
Supply Chain Business Modeler	✓	✓	

**Notes.** Although Production and Distribution Planning cannot be deployed as a web application, PDP Collaborative Web Client enables you to view and edit Production and Distribution Planning data from a web browser. However, PDP Collaborative Web Client does not enable you to run all Production and Distribution Planning functionality. For example, you cannot start Production and Distribution Planning servers or run Production and Distribution Planning solves using PDP Collaborative Web Client.

The Order Promising web application is used to configure, test and review settings for promising sales orders. To promise sales orders, you can use EnterpriseOne Sales Order Entry.

#### See Also

Standalone Installations

Thin Client Deployments

Web Application Deployment

Deploying Applications Through Firewalls

## Standalone Installations

This section discusses standalone installations of EnterpriseOne Supply Chain Planning products, and the benefits and considerations associated with standalone installations.

In a standalone installation, applications are installed on and accessed from the same machine. Although users typically install EnterpriseOne Supply Chain Planning products standalone on Microsoft Windows systems, most EnterpriseOne Supply Chain Planning 8.12.1 applications can be installed standalone on UNIX as well.

The following table lists the EnterpriseOne Supply Chain Planning 8.12.1 applications that can be installed as standalone applications, and shows the platforms where they can be installed:

<b>EnterpriseOne Supply Chain Planning Application</b>	<b>Microsoft Windows</b>	<b>UNIX</b>
Advanced Planning Agent	✓	✓
Demand Forecasting	✓	✓
Production Scheduling	✓	
Production Scheduling – Process	✓	✓
Strategic Network Optimization	✓	✓

### *Benefits of Standalone Installations*

These are the benefits of installing EnterpriseOne Supply Chain Planning applications standalone:

- You do not have to purchase, install, and administer third-party software for deploying applications as thin clients.
- If more than one EnterpriseOne Supply Chain Planning application is installed on the same machine, you can easily transfer data between applications without using file shares.

### *Considerations Associated with Standalone Installations*

Consider these issues when installing EnterpriseOne Supply Chain Planning applications in thin client deployments:

- If you install multiple EnterpriseOne Supply Chain Planning applications on a single machine, you must ensure that your hardware is sized appropriately to ensure adequate performance.
- If you install EnterpriseOne Supply Chain Planning applications standalone, users in multiple locations cannot access the same EnterpriseOne Supply Chain Planning installations.

#### **See Also**

Operating System Tuning for EnterpriseOne Supply Chain Planning Applications

Integrating EnterpriseOne Supply Chain Planning and EnterpriseOne Supply Chain Management Systems

## Thin Client Deployments

This section discusses thin client deployments of EnterpriseOne Supply Chain Planning products, and the benefits and considerations associated with thin client deployments.

In a thin client deployment, applications are installed on one or more central servers. Multiple users can then connect to and run the applications from Microsoft Windows clients.

The following table lists the EnterpriseOne Supply Chain Planning 8.12.1 applications that you can deploy as thin clients and the platforms that you can install them on:

<b>EnterpriseOne Supply Chain Planning Application</b>	<b>Microsoft Windows</b>	<b>UNIX</b>
Advanced Planning Agent	✓	✓
Demand Forecasting	✓	✓
Production and Distribution Planning	✓	✓
Production Scheduling	✓	
Production Scheduling – Process	✓	✓
Strategic Network Optimization	✓	✓
Supply Chain Business Modeler	✓	✓

### *Benefits of Thin Client Deployments*

These are the benefits of deploying EnterpriseOne Supply Chain Planning applications in thin client configurations:

- You can centrally install and administer EnterpriseOne Supply Chain Planning applications in a controlled environment.
- Users in different geographic locations can use the same EnterpriseOne Supply Chain Planning application installations.
- Users can access multiple EnterpriseOne Supply Chain Planning applications from different servers, with the appearance that they are on the same machine.
- Because the applications can be installed on multiple servers, this deployment method can prevent performance problems that could occur if multiple applications are installed on the same machine with insufficient resources.

### *Considerations Associated with Thin Client Deployments*

Consider these issues when installing EnterpriseOne Supply Chain Planning applications in thin client deployments:

- Deploying EnterpriseOne Supply Chain Planning applications in distributed environments requires extra licensing costs and administration for the third-party software that is used to deploy the applications as thin clients.

- Deploying EnterpriseOne Supply Chain Planning applications in thin client configuration involves more complex configuration than if the applications are installed on standalone machines.
- If you install EnterpriseOne Supply Chain Planning applications on multiple servers, EnterpriseOne Supply Chain Planning integration can be more complex than when all applications are installed on a single server.
- Data must then be transferred between machines through Microsoft Windows file shares, NFS, or FTP.

#### **See Also**

Thin Client Deployments from Microsoft Windows Servers

Thin Client Deployments from UNIX Servers

Deploying Applications Through Firewalls

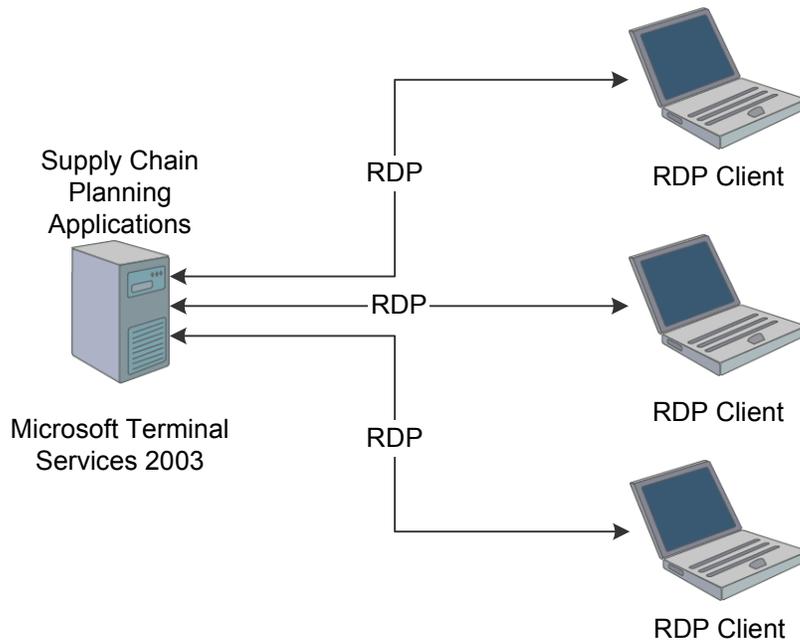
## **Thin Client Deployments from Microsoft Windows Servers**

If you install EnterpriseOne Supply Chain Planning applications on Windows servers, you can deploy the applications as thin clients using Microsoft Terminal Services Windows 2003. You can also install Citrix Presentation Server to provide additional clients and enhanced server functionality. Citrix Presentation Server is particularly important for improving application performance and color in thin client deployments with Microsoft Terminal Services Windows 2003.

Microsoft Terminal Services uses the Remote Desktop protocol (RDP) and Citrix Presentation Server use the Independent Computing Architecture (ICA) communications protocol. These communications protocols separate screen updates and user input processing from the rest of the application's logic. Using these protocols, applications appear to run locally on client machines but actually execute on servers, while screen updates, mouse movements, and keystrokes are transmitted via the Citrix ICA session.

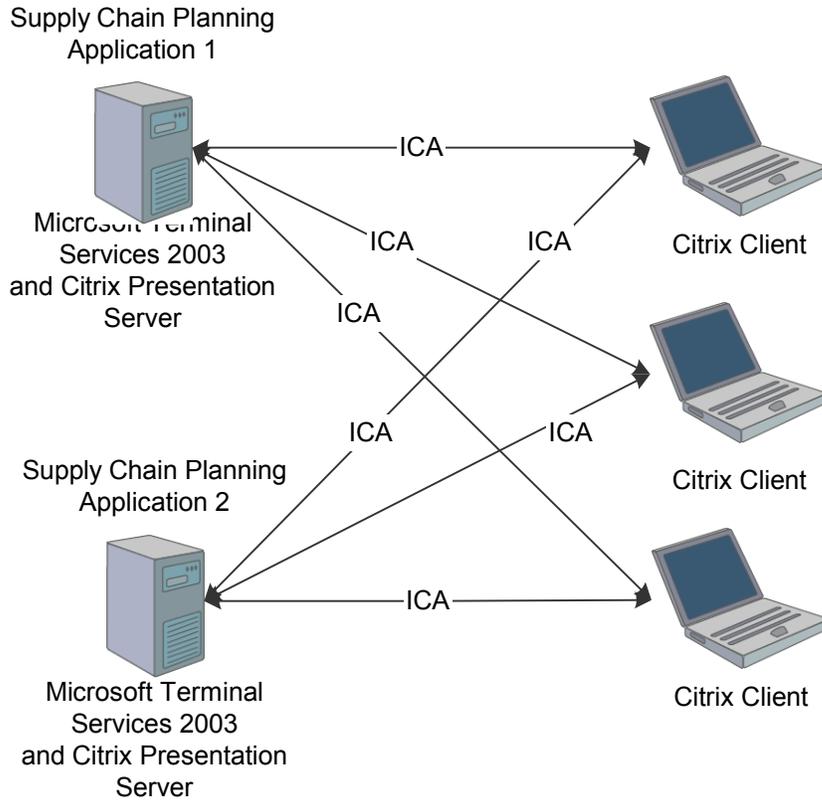
Because the RDP and ICA communications protocols are optimized for both local area networks (LAN) and wide area networks (WAN), you can deploy EnterpriseOne Supply Chain Planning applications to client machines in stable LAN or WAN environments without significant performance loss.

For example, EnterpriseOne Supply Chain Planning applications are installed on a single Microsoft Windows server running Microsoft Terminal Services 2003. Users can then access the EnterpriseOne Supply Chain Planning applications from Microsoft Terminal Services clients on Microsoft Windows machines, as illustrated in this diagram:



#### [Thin client deployment from single Microsoft Windows server](#)

You can install EnterpriseOne Supply Chain Planning applications on multiple Microsoft Windows servers in a thin client deployment. For example, in the following diagram, one EnterpriseOne Supply Chain Planning application is installed on a server running Microsoft Terminal Services 2003 and Citrix Presentation Server, and a second EnterpriseOne Supply Chain Planning application is installed on another server running Microsoft Terminal Services 2003 and Citrix Presentation Server. Users can then access both applications from Citrix clients on Microsoft Windows machines, as illustrated in this diagram:



Thin client deployment from multiple Microsoft Windows servers

**See Also**

Deploying Applications Through Firewalls

## Thin Client Deployments from UNIX Servers

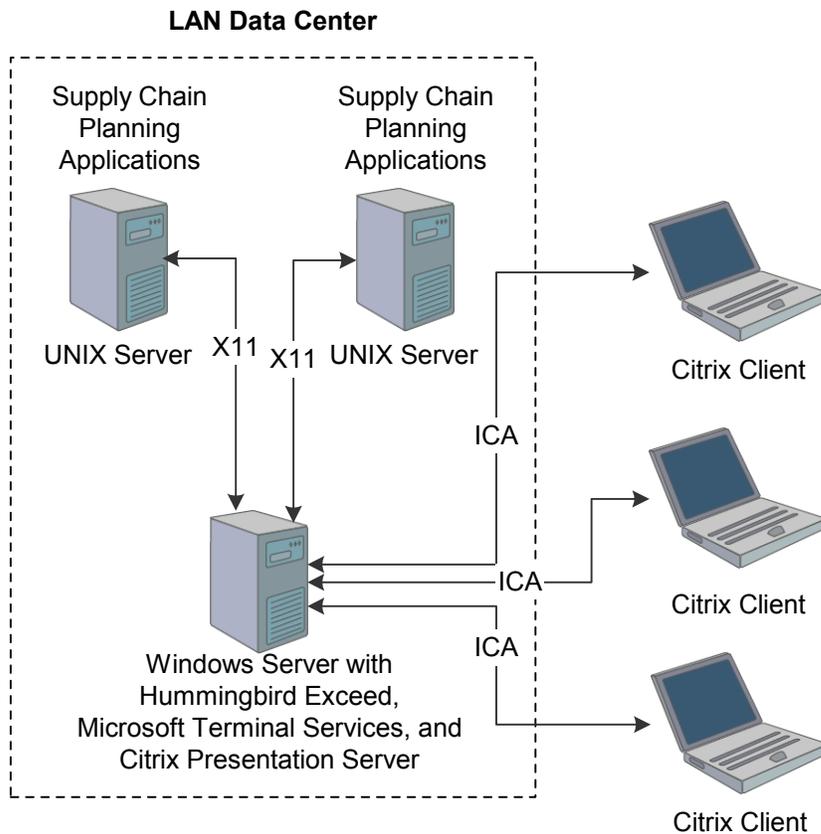
This section discusses thin client deployments of EnterpriseOne Supply Chain Planning applications from UNIX servers to Microsoft Windows clients using Hummingbird Exceed® and using Hummingbird Exceed onDemand®.

### *Deployments using Hummingbird Exceed*

With Hummingbird Exceed, you can create startup scripts that run EnterpriseOne Supply Chain Planning applications on UNIX machines and display the X-windows output on the Microsoft Windows machines where Exceed is installed. However, because Exceed uses the X11 communication protocol, which is not optimized for a WAN environment, the UNIX and Microsoft Windows machines must be contained in a LAN environment.

To access EnterpriseOne Supply Chain Planning applications on UNIX machines in a WAN environment, you can install Hummingbird Exceed on a Microsoft Windows server running Microsoft Terminal Services or Citrix Presentation Server. Users can then connect to the Microsoft Windows server from Windows clients and run EnterpriseOne Supply Chain Planning application startup scripts, with the display redirected from the Windows server to the client machines. Because Microsoft Terminal Services uses the RDP communications protocol and Citrix Presentation Server uses the ICA communications protocol, users can access the EnterpriseOne Supply Chain Planning application startup scripts from Microsoft Windows clients in a WAN environment.

For example, EnterpriseOne Supply Chain Planning applications are load balanced across two UNIX servers. A Microsoft Windows server running Hummingbird Exceed connects to the UNIX servers, and then redirects the EnterpriseOne Supply Chain Planning application displays to Citrix clients. Hummingbird Exceed is installed on and accessed from the Microsoft Windows server; it does not have to be installed on all client machines, as shown in this diagram:



[Deploying EnterpriseOne Supply Chain Planning applications from UNIX application servers](#)

In this configuration, because the X11 protocol is not optimized for WAN environments, the UNIX server and the Microsoft Windows Citrix server with Hummingbird Exceed are contained in a data center or LAN environment. However, because the server communicates with the Citrix clients using the ICA protocol, the client machines can be in a WAN environment.

It can be useful to use this configuration if EnterpriseOne Supply Chain Planning applications are installed on both Microsoft Windows and UNIX servers. Users can then access EnterpriseOne Supply Chain

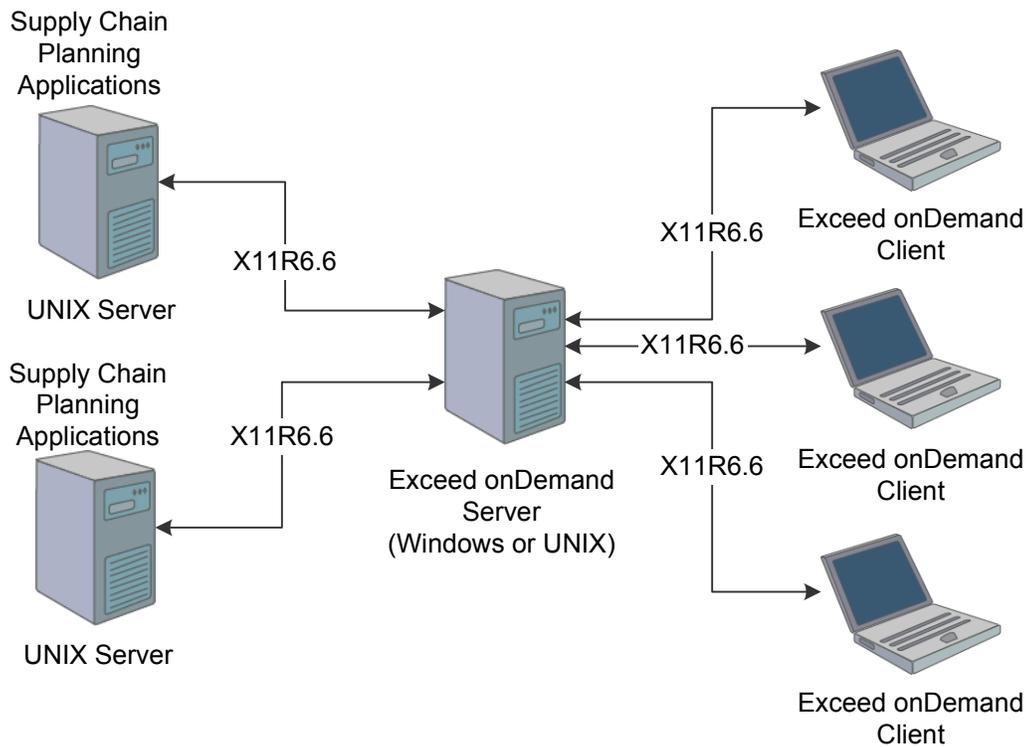
Planning applications from Microsoft Windows servers and Exceed start-up scripts for running EnterpriseOne Supply Chain Planning applications from UNIX servers from the same Microsoft Windows server.

**See Also**

Deploying Applications Through Firewalls

### *Deployments Using Exceed onDemand*

Users can also access EnterpriseOne Supply Chain Planning applications installed on UNIX machines from Microsoft Windows desktops using Exceed onDemand. As shown in this diagram, users can access UNIX applications from an Exceed onDemand server using Exceed onDemand clients that are installed on Microsoft Windows desktops:



#### [Deploying EnterpriseOne Supply Chain Planning applications from UNIX application servers using Exceed onDemand](#)

Because the Exceed onDemand server communicates with the UNIX servers using the X11R6.6 communication protocol, these machines can be installed in either a LAN or WAN environment without significantly reduced performance. Similarly, because the Exceed onDemand server and clients communicate using the X11R6.6 protocol, the Exceed onDemand server and clients can be located in a LAN or WAN environment. As a result, you do not have to purchase and administer additional server software such as Microsoft Terminal Services or Citrix Presentation Server.

## See Also

Deploying Applications Through Firewalls

## Web Application Deployment

This section discusses EnterpriseOne Supply Chain Planning web applications, and the benefits and considerations associated with web application deployments.

The following EnterpriseOne Supply Chain Planning 8.12.1 applications are web applications:

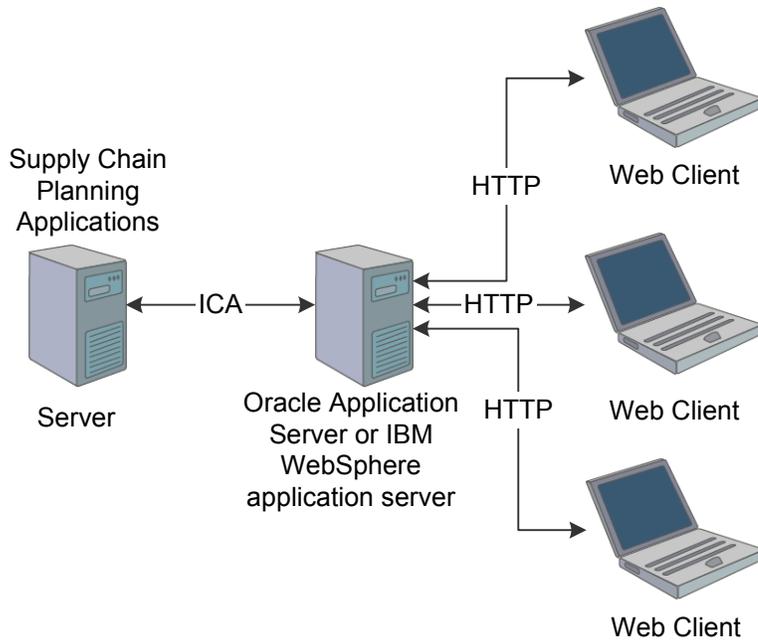
- Demand Consensus.
- Sales and Operations Planning.
- PDP Collaborative Web Client.
- Order Promising.

**Notes.** PDP Collaborative Web Client enables you to view and edit Production and Distribution Planning data from a web browser. PDP Collaborative Web Client does not enable you to run all Production and Distribution Planning functionality. For example, you cannot start Production and Distribution Planning servers or run Production and Distribution Planning solves using PDP Collaborative Web Client.

Oracle Application Server is not currently certified for deploying PDP Collaborative Web Client.

To configure, test and review settings for promising sales orders, you use the Order Promising web application. To promise sales orders, you use EnterpriseOne Sales Order Entry.

To deploy a EnterpriseOne Supply Chain Planning web application, you can install the application on a Microsoft Windows or UNIX machine and then publish the application's Web Archive (WAR) file to an Oracle Application Server or IBM WebSphere application server. Users can then connect to the application from web browsers, as shown in this diagram:



### [Deploying EnterpriseOne Supply Chain Planning web applications](#)

You can also deploy EnterpriseOne Supply Chain Planning applications from multiple servers through an Oracle Application Server or IBM WebSphere application server. Because the applications can be installed on multiple servers, this deployment method can prevent performance problems that could occur if multiple applications are installed on the same machine with insufficient resources.

## ***Benefits of Web Application Deployments***

EnterpriseOne Supply Chain Planning web applications offer the following benefits:

- You can centrally install and administer EnterpriseOne Supply Chain Planning applications in a controlled environment.
- Users in different geographic locations can use the same EnterpriseOne Supply Chain Planning application installations.
- Users can access multiple EnterpriseOne Supply Chain Planning applications from different servers, with the appearance that they are on the same machine.
- You can leverage the use of the Oracle Application Server or IBM WebSphere Portal.

## ***Considerations Associated with Web Application Deployments***

When you install EnterpriseOne Supply Chain Planning web applications, you must purchase, configure, and administer an IBM WebSphere Portal Server or Oracle Applications Server and its components.

### **See Also**

Deploying Applications Through Firewalls

## Deploying Applications Through Firewalls

When deploying applications as thin clients or deploying web applications, you must ensure that users can access the applications securely through firewalls. A firewall is usually a combination of hardware and software that restricts access to an internal network.

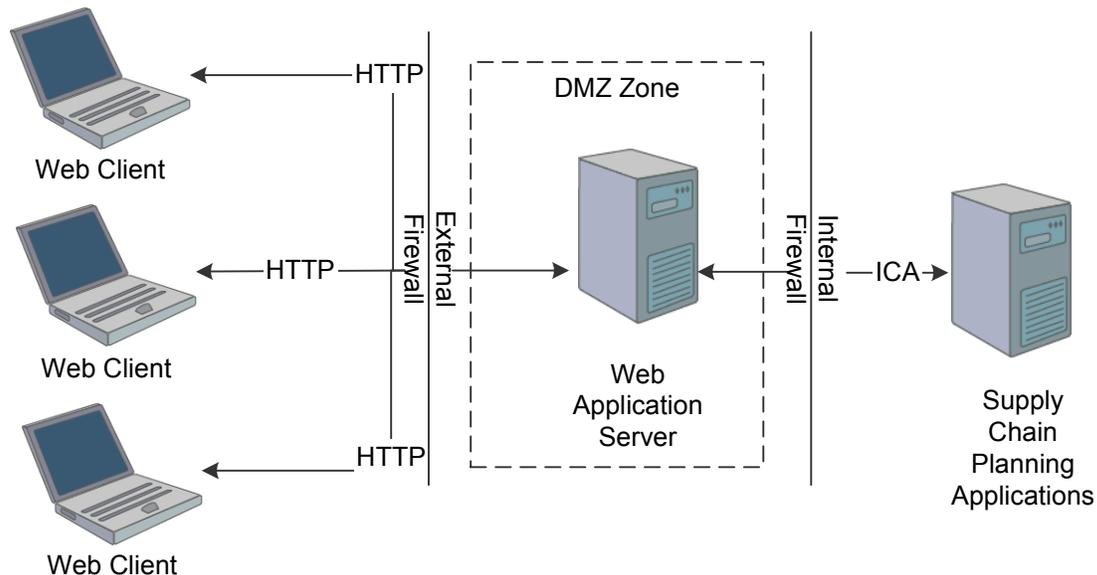
To enable users to access EnterpriseOne Supply Chain Planning applications behind a network firewall, you must ensure that the correct communication ports are open and available in your firewall. For example, to deploy EnterpriseOne Supply Chain Planning applications in thin clients using Hummingbird Exceed, you could ensure that port 6000 is open and is not being used for other network communications.

Opening ports in a firewall can raise security concerns. However, you can configure user authentication to restrict network access to authorized users. For example, you could restrict access to EnterpriseOne Supply Chain Planning applications to users from particular IP addresses or users with particular network accounts.

You can also increase security by encrypting data traveling across the network. In particular, if you are deploying EnterpriseOne Supply Chain Planning applications as thin clients, Citrix Presentation Server offers data encryption through the network gateway. If you are deploying web-based applications, you can use Secure Socket Layer (SSL) encryption and ensure that the SSL port on the firewall is open.

For further firewall protection, you can implement a demilitarized zone (DMZ). A DMZ is an area between your internal network and an external network that is protected by two firewalls. When deploying EnterpriseOne Supply Chain Planning applications, you can install the web application server in the DMZ, while protecting EnterpriseOne Supply Chain Planning applications behind a second firewall as further protection from unauthorized access.

For example, in this diagram, a web application server is located in a DMZ behind an external firewall:



### Deploying EnterpriseOne Supply Chain Planning Web Application using a DMZ

The appropriate port in the external firewall is opened to allow the web application server to communicate with web clients. To provide further protection, a second firewall is positioned between the web server and

the EnterpriseOne Supply Chain Planning applications, with the appropriate port opened in the internal firewall. The web application can talk to the EnterpriseOne Supply Chain Planning applications beyond the internal firewall through the web server.

**See Also**

Thin Client Deployments

Web Application Deployment

# OPERATING SYSTEM TUNING FOR ENTERPRISEONE SUPPLY CHAIN PLANNING APPLICATIONS

As with all applications, EnterpriseOne Supply Chain Planning applications are dependent on underlying operating systems and can be affected by kernel and operating system tuning parameters. This section describes how to tune your system to best support EnterpriseOne Supply Chain Planning 8.12.1 applications.

This section discusses:

- Operating system resource limits.
- Extending operating system resources.
- Performance tuning for CPU and memory-bound applications.
- Performance tuning for applications with object databases.

## Operating System Resource Limits

Some UNIX kernels have built-in limits on how many system resources can be accessed by individual applications and users. On a large multi-user system, these limits can be useful for sharing resources fairly. However, this mechanism is not useful for EnterpriseOne Supply Chain Planning applications, which tend to require large amounts of memory and CPU time.

This section lists the kernel parameter values that we use when testing our applications. You can adjust the parameter values to be appropriate for your environment. You should consult your vendors or system administrators before changing these parameters.

## HP-UX Kernel Parameters

The following table lists recommended kernel parameter values for HP-UX:

Kernel Parameter	Recommended Value	Description
maxdsiz	0XF0000000 4 GB for 32-bit processes	This parameter specifies the maximum process data segment size. By default, each 32bit process can access only 2GB.
maxdsiz_64bit	Maximum of estimated memory requirement for your model running with a 64-bit application such as Production and Distribution Planning or Strategic Network Optimization solvers	This parameter specifies the maximum process data segment size for 64-bit processes.
maxfiles	200 or larger	This parameter specifies the soft limit for open files.
maxfiles_lim	1024	This parameter specifies the hard limit for open files.
maxssiz	128000000	This parameter specifies the maximum process storage segment size.
maxtsiz	0X4000000	This parameter specifies the maximum process text segment size.
maxuprc	200	This parameter specifies the maximum number of processes per user.
Nfile	2500 or larger	This parameter specifies the system-wide open-files limit.
nproc	1024	This parameter specifies the maximum number of processes system-wide.
Npty	60 or larger	This parameter specifies the maximum number of pseudo-teletypes (ptys) allowed system-wide.

You might need to further configure the shared memory if you have installed database software. For more information, see your database application manuals.

## IBM AIX® Kernel Parameters

The following table lists recommended kernel parameter values for AIX.

Kernel Parameter	Recommended Value	Description
Cpu	-1	This parameter specifies the per process CPU time limit in seconds. The value is a soft limit. A value of -1 indicates unlimited.

Kernel Parameter	Recommended Value	Description
Data	-1	This parameter specifies the data segment size in 512 blocks. The value is a soft limit. A value of -1 indicates unlimited.
Rss	-1	This parameter specifies the largest amount of physical memory for a process. The value is a soft limit. A value of -1 indicates unlimited.
stack	65536	This parameter specifies the largest process stack segment. The value is a soft limit. A value of -1 indicates unlimited.

You can set these parameters in the `/etc/security/limits` file in AIX, or set these parameters as the default values for all users. Only the root user can change the `/etc/security/limits` file. You do not have to reboot the system after changing these parameters.

## Sun Solaris™ Kernel Parameters

This table shows recommended Solaris kernel parameters for EnterpriseOne Supply Chain Planning applications with embedded object databases:

Kernel Parameter	Recommended Value	Description
shmsys:shminfo_shmmax	1073741824	This parameter specifies the maximum size, in bytes, of a single shared memory segment.
shmsys:shminfo_shmseg	120	This parameter specifies the maximum number of shared memory segments that can be attached (i.e. used) by a single process.
shmsys:shminfo_shmmni	512	This parameter specifies the number of shared memory segments that can be on the system.
msgsys:msginfo_msgtql	4096	This parameter specifies the number of system message headers.
msgsys:msginfo_msgmnb	65535	This parameter specifies the maximum number of bytes on queue.

These parameters must be set in the `/etc/system` file, which can only be changed by the root user. The system must be rebooted after these parameters are changed.

### See Also

Performance Tuning for Applications with Object Databases

## Extending Operating System Resources

The standard memory layout of both Unix and Microsoft Windows limits the amount of memory that applications can access. As EnterpriseOne Supply Chain Planning applications often require large amounts of memory, this can limit the size of a solvable model.

For applications that support 64-bit memory models, the operating system limits are of no practical importance and will not limit performance. Currently, EnterpriseOne Strategic Network Optimization and the EnterpriseOne Production and Distribution Planning solver run as 64-bit applications on UNIX systems.

For 32-bit applications, operating system vendors have provided extensions to allow access to larger amounts of memory.

This table lists and describes the four types of memory or quadrants of an operating system (OS) memory map:

Quadrant	Description
DataSegment or Heap	All variable data is stored in this memory quadrant. EnterpriseOne Supply Chain Planning data models are loaded into Heap.
Text Segment	This memory quadrant stores actual application code.
Stack	This memory quadrant is used for short-term storage. One key use is to track subroutines in a program.
Shared memory	Shared memory is used by applications to share data. This memory quadrant is heavily used by databases, including databases embedded in EnterpriseOne Supply Chain Planning products.

The UNIX memory model for 32 bits, which translates to 4 GB of total addressable memory, allocates one GB to each type of memory. This is not an efficient use of memory as both the stack and text segment tend to be very small and, in the case of shared memory, might not be required by the application.

EnterpriseOne Supply Chain Planning applications can often require a heap size greater than 1 GB. Each supported operating system has a different mechanism to achieve this, but the general method is to combine the various quadrants together in some manner.

### *Extending Operating System Resources on HP/UX*

By default, EnterpriseOne Supply Chain Planning applications allow access to 2 GB of heap memory. In order to increase the heap further, the third or fourth quadrants must be allocated as heap.

To enable the third or fourth quadrants as heap memory on HP/UX, you can use the `chattr` command.

Use this command to enable the third quadrant as heap:

```
chattr +qp3 enable executable_name
```

The application will then have up to 2.8 GB of data storage.

Use this command to enable the fourth quadrant as heap memory:

```
chatr +qp4 enable executable_name
```

The application will have up to 3.8 GB of data storage. This command automatically enables the +qp3.

In some cases, the application might still not be able to address the full 3.8 GB of memory. This could occur for the following reasons:

- The kernel parameter *maxssize* is set to a very large value.
- The value of *maxssize* is subtracted from the total memory available for heap.

For information about tuning this kernel parameter, see Operating System Resource Limits.

The details are out of scope for this document, however, by default most EnterpriseOne Supply Chain Planning applications reserve a stack of 8 times the *maxssize* kernel parameter.

This can be fixed by setting an environment variable:

```
_M_ARENA_OPTS = 1:4096
```

These extensions should not be used for applications that use the object database. Using the third and fourth quadrant for heap makes shared memory unavailable. If only the third quadrant is used for heap, available shared memory is reduced from 2 GB to 1 GB.

Shared memory can be increased to 2.8 GB, but this reduces the available heap to 1 GB. It is unlikely that this type of configuration is necessary. This configuration is again enabled through the chatr command:

```
chatr -M executable
```

## *Extending Operating System Resources on Solaris*

Solaris automatically allows 32-bit applications to access 4 GB of heap memory. No tuning is required.

## *Extending Operating System Resources on AIX*

By default, EnterpriseOne Supply Chain Planning applications allow access to 2 GB of heap memory. The heap can be extended to 3.25 GB on a per-executable basis using the AIX very large memory model. The following two methods are available:

- Use this command to modify the executable header:

```
/usr/ccs/bin/ldedit -b maxdata=0xD0000000/dsa executable
```

- Use this command to set an environment variable:

```
LDR_CNTRL=MAXDATA=0xD0000000@DSA a.out
```

Do not use the very large memory model for applications with object databases.

## See Also

Performance Tuning for Applications with Object Databases

## *Extending Operating System Resources on Windows 2003*

By default, Supply Chain Planning applications on Windows 2003 can access 2 GB of heap memory. It is possible to extend this limit to 3 GB using the 4GT RAM tuning on windows. Currently the use of this extension has only been certified for EnterpriseOne Order Promising, EnterpriseOne Demand Management and EnterpriseOne Supply Chain Business Modeler.

For other applications the use of the RAM tuning extensions may not provide the expected memory increase and the application may become unstable once memory utilization increases past 2GB. EnterpriseOne Production and Distribution Planning and EnterpriseOne Strategic Network Optimization are not supported using the RAM tuning extension for this reason.

To enable 4GT RAM tuning on Windows 2003 please refer the appropriate Microsoft Document or the Microsoft Developer Network (MSDN).

## **Performance Tuning for CPU and Memory-Bound Applications**

EnterpriseOne Supply Chain Planning applications such as Strategic Network Optimization, Production and Distribution Planning, and Production Scheduling can be considered CPU and memory-bound applications. CPU and memory-bound application require sufficient physical memory and a fast CPU.

EnterpriseOne Supply Chain Planning solves are floating point intensive operations. For this reason, CPUs with better floating point performance produce faster solves. The choice of CPU depends on the choice of vendor.

While tuning options for these applications are limited, performance can be affected by:

- Page sizes.
- Processor affinity.

## *Page Sizes*

Any paging during solves, where memory is moved between disk and RAM, significantly affects solve times.

The page size is the minimum unit of memory that an application can allocate. A page is a block of memory that is addressed by the virtual memory system. Typically, a page is 4096 bytes. Each time an application tries to access memory, a lookup in a table is required to translate the virtual memory address to a physical address in RAM. This lookup table is cached in the CPU memory and is referenced by page address.

Applications that access very large amounts of memory, such as a gigabyte or more, have too many pages to fit into the CPU buffer. The virtual memory system has to use CPU cycles to complete the lookup in the much slower system memory. A large number of cache misses impacts application performance.

Operating systems that support large page sizes can allow individual applications to use a page size greater than 4096 bytes. A larger page size requires fewer entries in the lookup table and reduces cache misses. For example, if an application requires 32 MB of memory, a page size of 2 MB requires 16 entries in the lookup table while a page size of 16 MB requires only 2 entries in the lookup table.

The downside of using a large page size is that applications with a small memory footprint require more virtual memory. For example, if an application requires 2 MB of memory and the page size is set to 1 MB, the application uses 2 MB of memory (two pages). If the page size is set to 16 MB, the application uses 16 MB of memory (one page).

Tuning the page size for performance typically requires benchmarking the application with different page sizes.

On HP/UX, the recommended page size for solver executables, including Strategic Network Optimization and the Production and Distribution Planning, is 16 MB. You can use the following command on HP/UX to change an executable's page size:

```
chattr executable +pd 16M
```

The maximum value is currently 256 MB. You can also set this value to L, which specifies that the kernel should select the largest available page size.

Page size tuning for Solaris is only available on Solaris 9 and has not yet been certified for use in EnterpriseOne Supply Chain Planning applications. For more information, please contact your Sun support representative.

Page size tuning on AIX has not been certified for use with EnterpriseOne Supply Chain Planning applications. For more information, please contact your IBM support representative.

Page size tuning is not available on Microsoft Windows systems.

## *Processor Affinity*

Multi-CPU operating systems load balance applications on all available CPUs. For typical application workloads, load balancing results in an efficient use of hardware resources. For long running CPU-bound applications, however, switching between CPUs can negatively impact performance.

To improve solve times for applications like Strategic Network Optimization and Production and Distribution Planning, you can bind the application to one CPU. This process is called processor affinity.

**Note.** Binding processes to processors is not supported for EnterpriseOne Supply Chain Planning applications on systems running Dynamic Logical Partitioning (DLPAR).

This section describes how to bind a process to a processor on HP/UX, Solaris, AIX, and Windows 2003.

### Binding Processes to Processors on HP/UX

On HP/UX, to bind a process to a processor when starting the process, use the following command:

```
mpsched -c processor_id process_name
```

If the process is already running, use this command to assign the process to a processor:

```
mpsched -c processor_id -p process_id
```

### Binding Processes to Processors on Solaris

On Solaris, you can use the `psrinfo` command to print a list of processors that are online. You can use this command to bind a process to a particular processor:

```
pbind -b processor_id process_id
```

For example, if the process id is 6543, this command binds it to CPU 0.

```
$ pbind -b 0 6543
```

Use this command to create a new processor set consisting of CPU 1 and 9 and print the ID of the new processor set:

```
$ psrset -c 1 9
```

Use this command to bind process 6543 to this newly created processor set:

```
$ psrset -b processor_set_id 6543
```

To unbind a process, use the `-u` flag.

### Binding Processes to Processors on AIX

On AIX, the kernel threads of a process can be bound or unbound to a processor using the following command:

```
bindprocessor process processorNum | -q | -u process
```

If you include the `-q` flag, the system displays the processors that are available. If you include the `-u` flag, the system unbinds the threads of the specified process.

**Notes.** The `bindprocessor` command is meant for multiprocessor systems. Although it also works on uniprocessor systems, binding has no effect on such systems.

You need root authority to bind or unbind threads in processes you do not own.

Use this command to see which processors are available (processorNum values):

```
bindprocessor -q
```

For a system with four processors, the output could be:

```
The available processors are: 0 1 2 3
```

Use this command to bind the threads in process 19254 to processor 1:

```
bindprocessor 19254 1
```

### **Binding Processes to Processors on Windows 2003**

To bind a process to a processor on a multi-CPU Windows machine:

1. Start the Task Manager.
2. On the Processes tab, right-click the solver process and select Set Affinity from the shortcut menu.
3. Select the processor that you want to run the solver process.

To set a higher priority for a process in a Windows system:

1. Start the Task Manager.
2. On the Processes tab, right-click the solver process and select **Set Priority** from the shortcut menu.
3. Select the appropriate higher priority for your process.

To improve performance for background applications:

1. From the Start menu, select **Control Panel, System**.
2. On the Advanced Tab, click **Performance Options**.
3. Select the services that you want to optimize performance for: Applications or Background services.

**Note.** To set processor affinity from a command line, you must use a third-party product.

## **Performance Tuning for Applications with Object Databases**

The following Supply Chain Planning applications include object databases:

- Supply Chain Business Modeler.

- Demand Forecasting and Demand Consensus.
- Sales and Operations Planning.

In addition to requiring large amounts of memory and CPU and being affected by system IO and shared memory availability, you must monitor and manage the sizes of the object databases in these applications.

**See Also**

EnterpriseOne Demand Consensus 8.12.1 Design Studio Implementation Guide, Database Management

EnterpriseOne Sales and Operations Planning 8.12.1 Installation for Windows, Managing the Sales and Operations Planning Database

EnterpriseOne Sales and Operations Planning 8.12.1 Installation for UNIX, Managing the Sales and Operations Planning Database

EnterpriseOne Supply Chain Business Modeler 8.12.1 Implementation Guide, Managing Supply Chain Business Modeler Databases

# INTEGRATING ENTERPRISEONE SUPPLY CHAIN PLANNING AND ENTERPRISEONE SUPPLY CHAIN MANAGEMENT SYSTEMS

EnterpriseOne Supply Chain Planning applications use data from EnterpriseOne Supply Chain Management systems to create optimal supply chain plans and schedules.

This section discusses issues to consider when setting up the following processes:

- Extracting data from EnterpriseOne Supply Chain Management systems.
- Managing large volumes of data.
- Integrating EnterpriseOne Supply Chain Planning systems with multiple EnterpriseOne Supply Chain Management systems.
- Launching executables on separate machines.

**Note.** In this section, “EnterpriseOne Supply Chain Management systems” refers to Oracle Enterprise, EnterpriseOne, and World EnterpriseOne Supply Chain Management systems.

## Extracting Data from EnterpriseOne Supply Chain Management Systems

EnterpriseOne Supply Chain Business Modeler and EnterpriseOne Supply Chain Planning systems read and write XML data in UTF-8 format. When extracting data from EnterpriseOne Supply Chain Management systems for use in EnterpriseOne Supply Chain Planning systems, you must ensure that the data is encoded in UTF-8 format. By default, EnterpriseOne encodes data extracts in UTF-16 format so you must change the flat file encoding to UTF-8.

When you extract data from a EnterpriseOne Supply Chain Management system, you must ensure that the transactional data presents a consistent, synchronized view of the enterprise. For example, if a work order is partially completed according to the data that you extract, the appropriate inventory adjustments must also

be made to reflect the consumption on the work order. If the transactional data does not present a consistent view of the enterprise, the data sent to EnterpriseOne Supply Chain Planning applications could be unreliable. For example, if you ship an item from a location before the work order for producing the item has been entered in the EnterpriseOne Supply Chain Management system, the data extracts would indicate a negative beginning inventory value for the item.

Because EnterpriseOne Supply Chain Management system databases can be updated at the same time as data extracts are running, you should perform these tasks to ensure the integrity and consistency of the data:

- Extract data from EnterpriseOne Supply Chain Management systems at a time of day when there are few transactions. Choosing a time to extract data can be especially difficult if the enterprise has locations in multiple time zones. For example, if you have a location in the United States and a location in Singapore, you must choose a time to extract data when there are few transactions at either location.
- Limit the amount of time that elapses between the first and last data extract from your EnterpriseOne Supply Chain Management system. It is especially important to limit the amount of time that elapses when you extract data that changes frequently; the longer the time between extracts, the greater the chance that transactions could occur and extracts could become inconsistent. For this reason, you could extract all data that does not change frequently, such as customer addresses and product names, before you extract all dynamic data, such as inventory levels and sales orders.

When extracting data from a EnterpriseOne Supply Chain Management system, you should also try to limit the size of the data files. Extracting and transferring large data files can require a significant amount of time. In addition, some technology for transforming data does not work with large data files. For example, you cannot use XSLT to transform extremely large data files.

To reduce the amount of data that you extract and transfer between systems, you can do the following:

- Only extract data from your EnterpriseOne Supply Chain Management system if it is required to support your business processes. For example, if you are using Oracle Production Scheduling to create production schedules, you do not need to extract sales order data. Similarly, if you do not extract data for expired operations, you will not waste time transferring unnecessary data between systems.
- Run multiple extracts from your EnterpriseOne Supply Chain Management system to create relatively small data files, and then use incremental imports to add the data into Supply Chain Business Modeler.
- Only extract transactional data that is consistent with your enterprise data extracts. Supply Chain Business Modeler does not import transactional data for items that are not defined in the Base package in a model. If you extract transactional data for items that are not defined in the Base package, you will spend time transferring data that cannot be used.
- Extract static data less frequently than dynamic data. For example, you could extract static data on weekends and dynamic data every day. Alternatively, you could extract static and dynamic data every day, except for customer data. Customer data is usually quite stable and very large.

If you are integrating supply chain systems using Advanced Planning Agent, you must ensure that the data does not include restricted characters. For example, if quotation marks are used as delimiters in flat files, data values cannot include quotation marks. This issue is not a concern for data in XML format, which is used in integration systems that use Supply Chain Business Modeler. In addition, Advanced Planning Agent 8.12.1 is not certified to handle data with international characters.

## Managing Large Volumes of Data

When transferring large data files between systems, it can be faster to compress the data file on the source system, transfer the compressed data files, and uncompress the data on the destination system than to send the data uncompressed.

If you are using Advanced Planning Agent to integrate supply chain systems, remember that Advanced Planning Agent loads all data into memory. Depending on the data volume, you might need to break Advanced Planning Agent data flows into smaller flows. For large data volume flows, you should remove unnecessary temporary files, unnecessary folders, and use selection filters on a data node to reduce memory requirements.

## Integrating EnterpriseOne Supply Chain Planning Systems with Multiple EnterpriseOne Supply Chain Management Systems

Global enterprises can have multiple locations with different EnterpriseOne Supply Chain Management systems. When using data from multiple supply chain systems to create optimal plans in EnterpriseOne Supply Chain Planning applications, you must ensure that the data presents a reliable, accurate view of the enterprise. This means that you must accommodate differences such as:

- Time zone. To accommodate time zone differences, you might have to convert data from different time zones to a standard time zone for your enterprise.
- Date formats. Storing dates in standard ISO format can be useful so that you do not have to transform dates from location-specific formats.
- Units of measure. Unit of measure conversion rates can vary by location. For example, locations at different altitudes can require different quantities because of varying weight and density ratios.
- Item codes. For example, different EnterpriseOne Supply Chain Management systems might use the same short item number to refer to different items. To solve this problem, you could select an alternate item number to use or change item codes from one system to be unique by adding a character to the front of each item number from one system.
- Currencies.
- Business processes. If you create plans for locations with separate EnterpriseOne Supply Chain Management systems, you must be careful to send the resulting plans to the correct EnterpriseOne Supply Chain Management system for implementation. For example, if you create net production requirements in Production and Distribution Planning for an entire enterprise, you must be careful to direct the plans to the appropriate EnterpriseOne Supply Chain Management systems. This process can be especially difficult if your plans span locations. For example, a transportation plan could have a source location in one EnterpriseOne Supply Chain Management system and a destination location in another system.

## Launching Executables on Separate Machines

This section discusses issues that you must consider when integrating systems that are installed on separate machines.

EnterpriseOne supply chain systems each have their own batch interface. To automate integration systems, you must create scripts that switch between command shells and run component scripts.

EnterpriseOne Supply Chain Management and EnterpriseOne Supply Chain Planning 8.12.1 systems can be installed on separate servers or on different logical partitions on the same server. If applications and connectors are installed on different machines, you must remotely execute the supply chain system and connector executables that are involved in the integration. You must also address user name and password security issues.

If applications are installed on different Microsoft Windows machines, you must install software for remotely executing commands or opening connections to servers. Software for remote execution does not automatically ship with Microsoft Windows. You must also address security concerns associated with opening connections to the servers.

If executables are installed on machines with different operating systems, you must address differences in file systems and naming conventions on UNIX and Microsoft Windows. Maximum file sizes can differ from one server to another, especially between Microsoft Windows and UNIX servers. Additional software might be required because UNIX and Microsoft Windows do not support shared directories across platforms by default.

### **See Also**

Deploying Applications Through Firewalls