



Siebel Performance Tuning Guide

Siebel Innovation Pack 2013

Version 8.1/8.2

September 2013

ORACLE®

Copyright © 2005, 2013 Oracle and/or its affiliates. All rights reserved.

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.

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.

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.

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.

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.

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

Chapter 1: What's New in This Release

Chapter 2: Siebel Architecture and Infrastructure

- About Performance and Scalability 13
- About Siebel Architecture and Infrastructure 14
- About Siebel User Request Flow 18
- Performance Tuning Terminology 19

Chapter 3: Tuning Siebel Application Object Manager

- About the Siebel Application Object Manager 21
- Siebel Application Object Manager Infrastructure 22
- Performance Factors for Siebel Application Object Manager Deployments 24
- Topology Considerations for Siebel Application Object Manager Deployments 27
- Guidelines for Siebel Application Object Manager Tuning 27
 - Tuning Siebel Application Object Manager Components for CPU and Memory Utilization 27
 - Tuning Parameters for Siebel Application Object Manager Caches 32
 - Additional Parameters Affecting Siebel Application Object Manager Performance 34
 - Memory Consumers in Siebel Application Object Manager 35
- Configuring Database Connection Pooling for Siebel Application Object Managers 36
 - About Database Connections for Siebel Application Object Manager 36
 - Database Connection Pooling Usage Guidelines 37
 - Configuring Pooling for Default Database Connections 40
 - Configuring Pooling for Specialized Database Connections 42
- Using Thread Pooling for Siebel Application Object Manager 44

Chapter 4: Tuning the Siebel Server Infrastructure

- Configuring SISNAPI Connection Pooling for Siebel Application Object Manager 47
- Tuning Server Request Broker (SRBroker) 49

Chapter 5: Tuning Siebel Web Client

About Siebel Clients	51
Performance Factors for Siebel Web Clients	52
Guidelines for Siebel Web Client Tuning	53
Providing Sufficient Web Server and Network Capacity	54
Testing Performance for Web Clients	54
Providing Sufficient Client Hardware Resources	55
Tuning System Components	55
Following Configuration Guidelines	56
Managing the Browser Cache	56
Specifying Static File Caching	57
Improving Performance Using View Layout Caching	59
Configuring the Data Block Size of HTTP Requests for the Siebel Dedicated Client	64
Managing Performance Related to the Message Bar	64
Configuring the Busy Cursor for Standard Interactivity Applications	64

Chapter 6: Tuning Siebel Communications Server

About Siebel Communications Server	67
Session Communications Infrastructure	68
Performance Factors for Session Communications	69
Topology Considerations for Session Communications	71
Guidelines for Session Communications Tuning	72
Tuning the Siebel Application Object Manager Component	73
Tuning the CommSessionMgr Component	73
Conserving Siebel Application Object Manager Server Resources Through Caching	73
Improving Performance for Communications Configurations	74
Configuring Logging for Session Communications	75
Improving Availability for Session Connections	76
Improving Screen Pop Performance	77
Reviewing Performance Impact of Activity Creation	77
Siebel Email Response Infrastructure	77
Performance Factors for Siebel Email Response	78
Topology Considerations for Siebel Email Response	79
Guidelines for Siebel Email Response Tuning	79

Chapter 7: Tuning Siebel Workflow

About Siebel Workflow	83
-----------------------	----

Monitoring Workflow Policies	84
Using the Policy Frequency Analysis View	84
Using Workflow Agent Trace Files	84
Monitoring Workflow Policies Tables	85
Tuning Workflow Policies for Performance	86
Creating Workflow Policy Groups to Manage Siebel Server Load	86
Multiple Workflow Monitor Agents and Workflow Action Agents	86
Running Workflow Agents on Multiple Siebel Servers	87
Setting Optimal Sleep Interval for Workflow Policy Groups	87
Setting Optimal Action Interval for Workflow Monitor Agent and Workflow Action Agent	88
Tuning Workflow Processes	88
Minimizing Usage of Parameter Search Specification	89
Monitoring Conditions Based on Parent and Child Business Components	89
Configuring Siebel Business Applications for Workflow Performance	90
Monitoring Memory Overhead for Workflow Processes	90
Tuning Workflow Process Manager for Performance	91
Caching Business Services	92
Caching Sessions	92

Chapter 8: Tuning Siebel Configurator

Siebel Configurator Infrastructure	93
Performance Factors for Siebel Configurator	94
Considerations for Defining Topology for Siebel Configurator	95
Running Siebel Configurator in the Siebel Application Object Manager Component	95
Running Siebel Configurator on Dedicated Servers	96
Guidelines for Siebel Configurator Tuning	98
Tuning Siebel Configurator	99
Specifying the Siebel Configurator File System Location	100
Defining Customizable Product Models and Classes	100
About Siebel Configurator Caching	101
Default Caching Behavior for Siebel Configurator	102
Cache Management for Siebel Configurator	102
Parameters for Configuring Siebel Configurator Caching	104
Determining Rough Sizing for Caching Parameters	107
Administering the Siebel Configurator Cache	108
Refreshing the Entire Siebel Configurator Cache	108
Refreshing the Siebel Configurator Cache with Product Changes	109
Updating the Siebel Configurator Cache with Product Class Changes	109

Refreshing the Siebel Configurator Cache with Product Class Changes	110
Updating the Siebel Configurator Cache with Attribute Definition Changes	110
Refreshing the Siebel Configurator Cache with Attribute Definition Changes	111

Chapter 9: Tuning Siebel EAI

About Siebel Enterprise Application Integration	113
Guidelines for Siebel EAI Tuning	113
Improving IBM WebSphere MQ Transport Performance	114
Improving HTTP Inbound Transport Performance	116
EAI Siebel Adapter Performance	117
Virtual Business Component Performance	118
Improving Workflow Process Manager Performance	119
Other Guidelines for Siebel EAI	120

Chapter 10: Tuning Siebel EIM

About Siebel EIM	121
Siebel EIM Architecture Planning Requirements	122
Database Sizing Guidelines	122
Database Layout Guidelines (Logical and Physical)	123
Siebel EIM Usage Planning	124
Defining the Siebel EIM Team	125
Mapping Data into Siebel Business Applications	125
Testing Siebel EIM Processes	126
General Guidelines for Optimizing Siebel EIM	127
Recommended Sequence for Implementing Siebel EIM Processes	128
Optimizing the .IFB File for Siebel EIM	129
Checking .IFB File Optimization for Siebel EIM	129
Separating Siebel EIM Processes by Operation	130
Troubleshooting Siebel EIM Performance	131
Optimizing SQL for Siebel EIM	132
Using the USE INDEX HINTS and USE ESSENTIAL INDEX HINTS Parameters	132
Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Example	133
Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Criteria for Passing Indexes to the Database	135
Using the SQLPROFILE Parameter	136
Additional Indexes on Siebel EIM Tables	137
Creating Proper Statistics on Siebel EIM Tables	138
Dropping Indexes in Initial Runs of Siebel EIM	139
Controlling the Size of Batches for Siebel EIM	140

Controlling the Number of Records in Siebel EIM Tables	141
Using the USING SYNONYMS Parameter with Siebel EIM	141
Using the NUM_IFTABLE_LOAD_CUTOFF Extended Parameter with Siebel EIM	141
Disabling the Docking: Transaction Logging Parameter for Siebel EIM	142
Disabling Database Triggers for Siebel EIM	142
Running Siebel EIM Tasks in Parallel	142
Database Guidelines for Optimizing Siebel EIM	143
Microsoft SQL Server and Siebel EIM	143
Oracle Database and Siebel EIM	146
IBM DB2 for z/OS and Siebel EIM	148
IBM DB2 for z/OS and Siebel EIM	150
IBM DB2 for z/OS Loading Process for Siebel EIM	151
General Recommendations for the IBM DB2 for z/OS Loading Process	151
Data Management Guidelines for Optimizing Siebel EIM	153
Run Parameter Guidelines for Optimizing Siebel EIM	153
Monitoring the Siebel Server During a Siebel EIM Task	154

Chapter 11: Tuning Siebel Remote

About Siebel Remote	155
Tuning Siebel Remote Server Components	156
Increasing Throughput for the Database Extract and Parallel Database Extract Components	156
Tuning the Transaction Router Component	157
Tuning the Mobile Web Client in a Siebel Remote Deployment	160
Optimizing Application Configuration File Parameters	160
Guidelines for Optimizing Data Synchronization Between Siebel Mobile Web Client and Siebel Remote Client	162
Choosing an Appropriate Routing Model	162

Chapter 12: Tuning Customer Configurations

General Performance Guidelines for Customer Configurations	163
Analyzing Generated SQL for Performance Issues	166
About Specifying SQL Logging and SQL Tagging for Siebel Application Object Manager Components	167
Troubleshooting Poor Performing SQL at the Database Level Using Workload Tagging	168
Specifying SQL Spooling in Siebel Developer Web Client	172
Troubleshooting Performance Using SQL Trace Files	172
Troubleshooting Performance Using SQL Query Plans	173

Performance Guidelines for Siebel Scripting	175
Using Declarative Alternatives to Siebel Scripting	176
Siebel Scripting Guidelines for Optimal Performance	177
Performance Guidelines for Data Objects Layer	179
Multilingual LOVs Query and Cache Performance	180
Managing Database Indexes in Sorting and Searching	180
Reusing Standard Columns	182
Performance Guidelines for Business Objects Layer	185
Using Cache Data Property to Improve Business Component Performance	185
Limiting the Number of Active Fields	185
Guidelines for Using Calculated Fields	186
Using Properties to Improve Picklist Performance	187
Using Primary ID Fields to Improve Performance	187
How the Check No Match Property Impacts Performance	188
Performance Guidelines for User Interface Objects Layer	189
Addressing Performance Issues Related to Grid Layout	189
Maintaining Performance When Using Applet Toggles	189

Chapter 13: Tuning the Web Server Computer for All UNIX and Linux Platforms

Tuning Microsoft Windows for Enhanced Siebel Server Performance	191
Tuning the Siebel Server for All UNIX and Linux Platforms	192
Tuning the Web Server Computer for All Applicable UNIX and Linux Platforms	193
Tuning the Siebel Web Server Extension for All UNIX and Linux Platforms	195
Tuning an Apache Web Server for Applicable UNIX and Linux Platforms	195
Tuning Siebel Business Applications for AIX	198
Tuning the IBM HTTP Server for AIX	198
Tuning the Siebel Server for AIX	198
Tuning Kernel Settings for AIX	200
Tuning Siebel Business Applications for HP-UX	202
Tuning Kernel Settings for HP-UX	202
Setting Permissions for the HP-UX Scheduler	203
Tuning Siebel Business Applications for Oracle Solaris	203
Tuning the Oracle iPlanet Web Server	203
Tuning Kernel Settings for Oracle Solaris	205
Tuning Siebel Application Object Manager Instances for Oracle Solaris	205

Chapter 14: Monitoring Siebel Application Performance with Siebel ARM

- About Siebel Application Response Measurement 207
- About Siebel ARM Parameters and Variables 208
- Enabling and Configuring Siebel ARM 211
- Guidelines for Converting Siebel ARM Files 212

Chapter 15: Analyzing Siebel ARM Data

- About Siebel ARM Files 215
- Analyzing Siebel ARM Files Using the Siebel ARM Query Tool 216
 - About the Siebel ARM Query Tool 217
 - General Commands for the Siebel ARM Query Tool 218
 - Configuring the Siebel ARM Query Tool 219
 - Configuring Input for the Siebel ARM Query Tool 220
 - Configuring Output from the Siebel ARM Query Tool 221
 - Using Selection Filters with the Siebel ARM Query Tool 224
 - Aggregating Siebel ARM Data with the Siebel ARM Query Tool 232
 - Generating Histograms with the Siebel ARM Query Tool 234
 - Using Macros with the Siebel ARM Query Tool 235
- Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool 236
 - About the Siebel ARM Analyzer Tool 237
 - Running Performance Aggregation Analysis 238
 - Running Call Graph Generation 239
 - Running User Session Trace 239
 - Running Siebel ARM Data CSV Conversion 240
 - About Siebel ARM Analyzer Output Files 241
 - About Performance Aggregation Analysis and Data 241
 - About Call Graph Generation Analysis and Data 249
 - About User Session Trace Analysis and Data 251
 - About Siebel ARM to CSV Conversion Data 253

Index

1

What's New in This Release

What's New in Siebel Performance Tuning Guide, Version 8.1/8.2

No new features have been added to this guide for this release. This guide has been updated to reflect only product name changes.

What's New in Siebel Performance Tuning Guide, Version 8.2, Rev. B

[Table 1](#) lists some of the changes in this version of the documentation to support this release of the software.

Table 1. New Product Features in Siebel Performance Tuning Guide, Version 8.2, Rev. B

Topic	Description
"Tuning Siebel Business Applications for HP-UX" on page 202	New topic. It describes how to configure and tune OS settings and Siebel Enterprise Server components to run Siebel Business Applications on HP-UX. The HP-UX operating system is supported in Siebel CRM version 8.2.2.3 and later. For more information about HP-UX support, see <i>Siebel Installation Guide for UNIX</i> . NOTE: The functionality described in this topic requires that you install Siebel CRM version 8.2.2.3 or later. For details, see the applicable <i>Siebel Maintenance Release Guide on My Oracle Support</i> .

What's New in Siebel Performance Tuning Guide, Version 8.1, Rev. A

[Table 2](#) lists some of the changes in this version of the documentation to support this release of the software.

Table 2. New Product Features in Siebel Performance Tuning Guide, Version 8.1, Rev. A

Topic	Description
"Additional Parameters Affecting Siebel Application Object Manager Performance" on page 34	Modified topic. The EnableCDA parameter for the Application Object Manager component must always be set to FALSE.
"Troubleshooting Poor Performing SQL at the Database Level Using Workload Tagging" on page 168	New topic. It describes how to use workload tagging to troubleshoot poor performing SQL at the database level.

2

Siebel Architecture and Infrastructure

This chapter provides an overview of Oracle's Siebel Business Applications architecture and infrastructure and provides introductory information about tuning the Siebel applications for performance and scalability. It contains the following topics:

- [About Performance and Scalability on page 13](#)
- [About Siebel Architecture and Infrastructure on page 14](#)
- [About Siebel User Request Flow on page 18](#)
- [Performance Tuning Terminology on page 19](#)

Related Books

Siebel Deployment Planning Guide

Siebel Installation Guide for the operating system you are using

Siebel System Administration Guide

Configuring Siebel Business Applications

About Performance and Scalability

Every implementation of Siebel Business Applications is unique. Your Siebel application architecture, infrastructure, and configurations might differ depending on your business model.

Performance and scalability are defined as follows in the context of this guide:

- **Performance.** A Siebel application's ability to function, generally measured in response time or throughput.

For example, measures of performance might include the time required to log into the Siebel application or to display a Siebel view in the Siebel Web Client, or the volume of transactions (sometimes referred to as requests) that a server component can process in a given time period.

Some typical inhibitors of performance are inadequate hardware, excessive network round trips, heavy customizations, and poor networking infrastructure.

- **Scalability.** A Siebel application's ability to continue to perform well as volumes increase.

Scalability is generally measured in hardware terms; for example, maintaining acceptable performance after adding new processors on existing computers (vertical scalability) or new Siebel Server computers (horizontal scalability) to process an increased number of users.

Some typical inhibitors of scalability are an inflexible application module structure and an inability to run parallel processes.

For more definitions of terminology related to performance and scalability, see ["Performance Tuning Terminology" on page 19](#).

About Siebel Architecture and Infrastructure

[Figure 1 on page 15](#) shows a generic representation of the architecture and infrastructure of a Siebel Business Applications deployment. Your Siebel applications might be deployed differently. For descriptions of individual entities included in this illustration, see *Siebel Deployment Planning Guide*, *Siebel System Administration Guide*, and the *Siebel Installation Guide* for the operating system you are using.

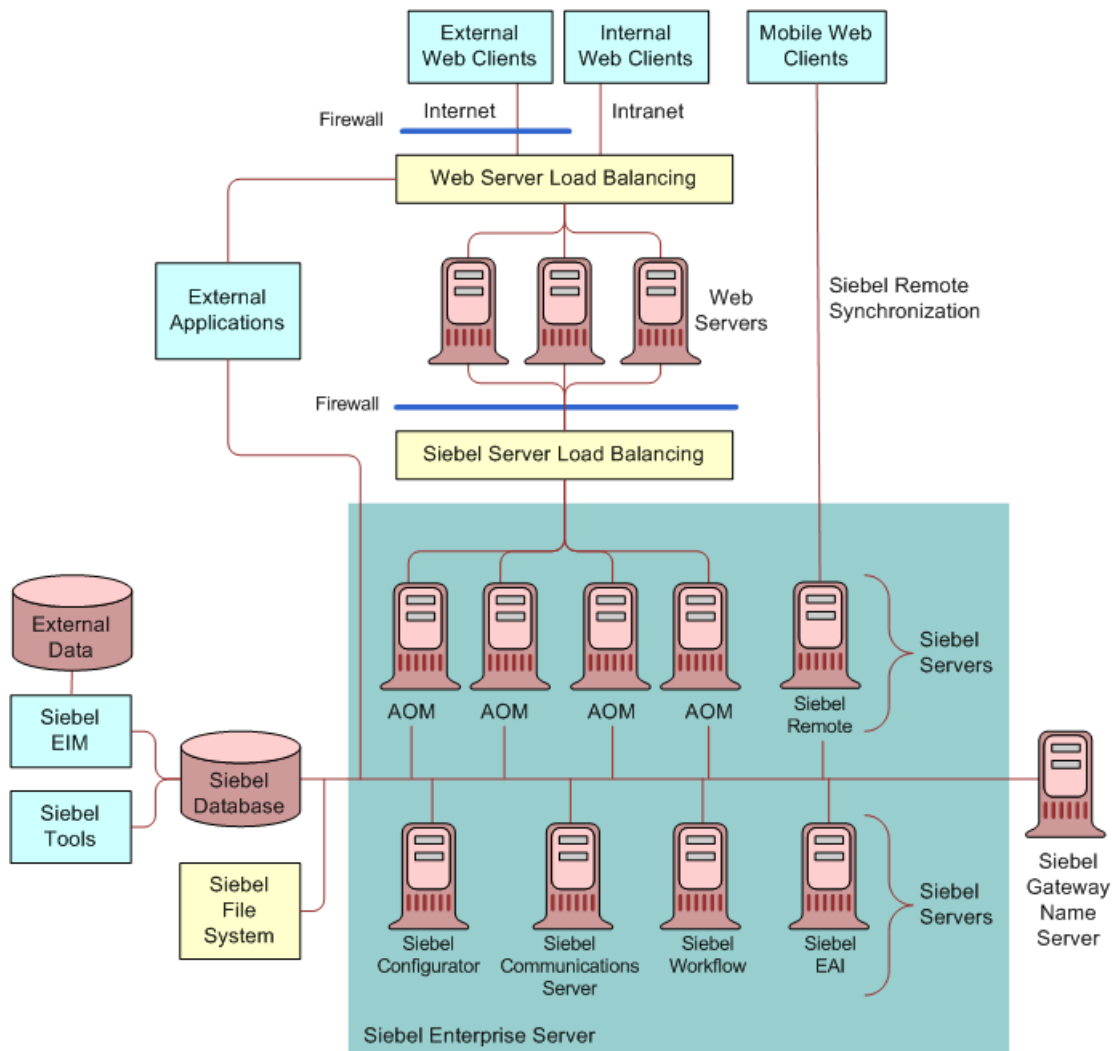


Figure 1. Generic Architecture of Siebel Business Applications

Siebel Architecture and Infrastructure Areas for Tuning

The following list provides details on tuning specific areas of the Siebel applications architecture and infrastructure.

Performance in many of these areas can be monitored and analyzed using Siebel Application Response Measurement (Siebel ARM), which is described in [Chapter 14, "Monitoring Siebel Application Performance with Siebel ARM,"](#) and [Chapter 15, "Analyzing Siebel ARM Data."](#)

- **Siebel Application Object Managers.** Siebel Application Object Managers are Siebel Server components that reside on a Siebel Server and support users accessing Siebel applications through the Siebel Web Client and a Web server, or through external applications.

Running Siebel Application Object Manager components has significant performance and scalability implications. In general, the goal for tuning a Siebel Application Object Manager is to maximize scalability with little or no performance degradation as more users use Siebel applications.

Although Siebel Application Object Manager components can be tuned for optimal performance, capacity for this and all other Siebel Server components is ultimately limited by Siebel Server computer resources such as CPU and memory. For details on tuning this area, see [Chapter 3, "Tuning Siebel Application Object Manager."](#)

- **Siebel Web Client.** The means for end users to access Siebel application features and data. Siebel Web Client uses a Web browser.

The response time experienced by the Siebel Web Client end user is subject to the configuration and tuning of Siebel Enterprise elements such as the Siebel Application Object Manager, network bandwidth and latency, Web server, Siebel Database, and the Siebel application configuration (represented in the Siebel repository file). It is also subject to local computer resources and settings, including browser settings such as those for caching. For details on tuning this area, see [Chapter 5, "Tuning Siebel Web Client."](#) See also [Chapter 12, "Tuning Customer Configurations."](#)

- **Siebel Communications Server.** Siebel Communications Server provides an infrastructure to support several kinds of communications activities for Siebel Business Applications users, including session communications (such as voice calls) and inbound and outbound communications (such as email).

Siebel Communication Server processing can affect end user response time, and might demand additional Siebel Application Object Manager resources to support user sessions. Performance and scalability is subject to third-party server configuration and capacity and Siebel Server computer resources and configuration. For details on tuning this area, see [Chapter 6, "Tuning Siebel Communications Server."](#)

- **Siebel Workflow.** Siebel Workflow is an interactive environment that automates business processes such as automating escalation of events and notification of appropriate parties; routing and assigning work; processing work; and enforcing authorization and transition rules.

Siebel Workflow processing can affect end user response time (for synchronous requests), and might demand additional Siebel Application Object Manager resources to support user sessions. Performance and scalability is subject to Siebel Server computer resources and configuration. For details on tuning this area, see [Chapter 7, "Tuning Siebel Workflow."](#)

- **Siebel Configurator.** Siebel Configurator supports order management and product configuration functions for Siebel Business Applications.

Siebel Configurator processing can affect end user response time (for configuration sessions), and might demand additional Siebel Application Object Manager resources to support user sessions. Performance and scalability is subject to Siebel Server computer resources and configuration. For details on tuning this area, see [Chapter 8, "Tuning Siebel Configurator."](#)

- **Siebel Enterprise Application Integration (Siebel EAI).** Siebel EAI provides components for integrating Siebel Business Applications with external and internal applications, and provides inbound and outbound interfaces to and from a Siebel application.

Siebel EAI processing can affect end user response time (for real-time interfaces), and might demand additional Siebel Application Object Manager resources to support user sessions. Performance and scalability is subject to Siebel Server computer resources and configuration. For details on tuning this area, see [Chapter 9, "Tuning Siebel EAI."](#)

- **Siebel Enterprise Integration Manager (Siebel EIM).** Siebel EIM is a server component in the Siebel EAI component group that transfers data between the Siebel database and other corporate data sources. For details on tuning this area, see [Chapter 10, "Tuning Siebel EIM."](#)
- **Siebel Remote.** Siebel Remote provides components that allow Siebel Mobile Web Clients (typically operating remotely, in disconnected mode on a laptop) to connect to a Siebel Server and exchange updated data and files, a process known as synchronization. For details on tuning this area, see [Chapter 11, "Tuning Siebel Remote."](#)
- **Siebel Tools.** Siebel Tools is an integrated development environment for configuring aspects of a Siebel application, including elements in the data objects, business objects, and user interface objects layers. Siebel scripting languages are also managed in the Siebel Tools environment.

Siebel Tools configurations and scripting play a critical role in the performance and scalability of a configured Siebel application. Customizations made through Siebel Tools partly determine the degree to which performance and scalability of a particular deployment differs from the original installation.

Appropriate configuration optimizes operations in the Siebel Database and does not add unnecessary overhead to supporting user sessions. (Siebel Tools itself does not play a role in the Siebel applications at run-time.) For details on tuning this area, see [Chapter 12, "Tuning Customer Configurations."](#)

- **Operating systems.** For details on tuning your Microsoft Windows or UNIX operating system, see [Chapter 13, "Tuning the Web Server Computer for All UNIX and Linux Platforms."](#)

About Siebel User Request Flow

Figure 2 on page 18 illustrates how a user request is processed within the Siebel Business Applications architecture and infrastructure (generically presented), and shows potential areas for performance tuning. For a description of each portion of this data flow, see *Siebel System Administration Guide* and other relevant documents on the *Siebel Bookshelf*.

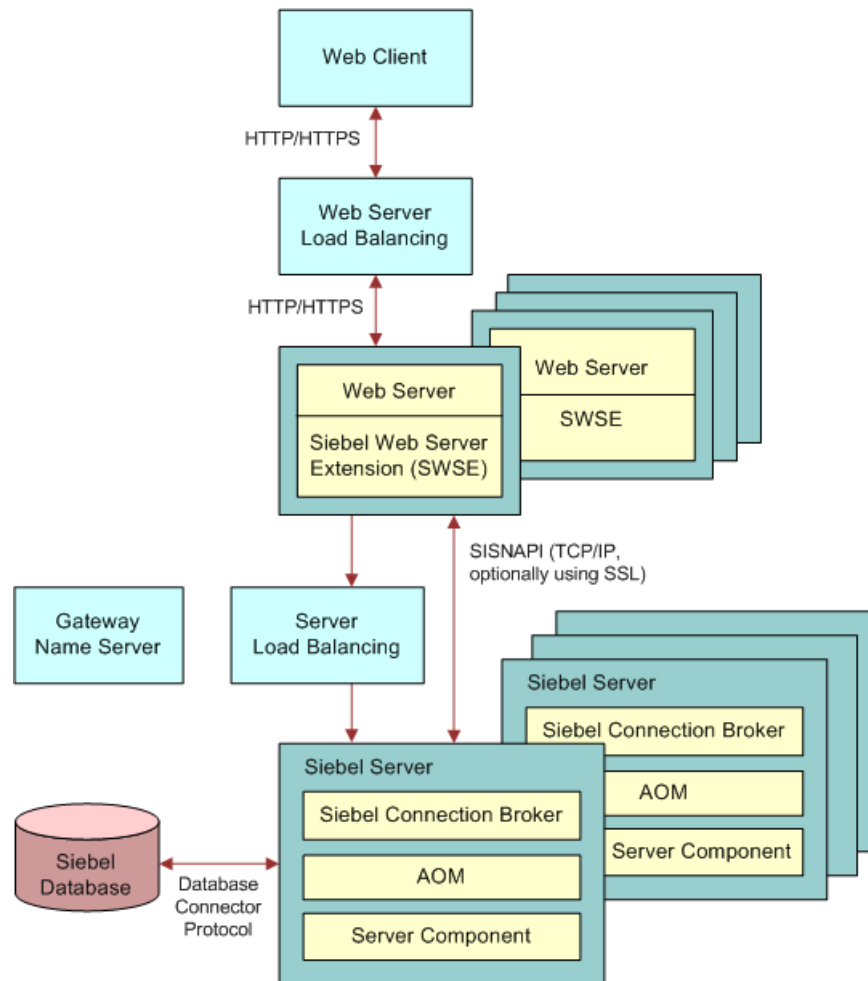


Figure 2. Generic User Request Flow in Siebel Business Applications

A typical Siebel client request flows from the user's Siebel Web Client through the system, and back again, following the general flow outlined below.

- 1 A user performs an action that initiates a request. For example, the user clicks a link in the Site Map to navigate to a particular view. The request is generated by the Web browser and Siebel Web Client framework.
- 2 The request goes through the network, using an existing or new HTTP connection. The request might go through a network router, proxy server, cache engine, or other mechanism.

- 3 If present, Web server load balancing software evaluates the request and determines the best Web server to forward the request to. It then forwards the request to a Web server.
- 4 The Web server receives the HTTP request, determines that it is a Siebel application request, and forwards the request to the Siebel Web Server Extension (SWSE) installed on the Web server.
- 5 The SWSE parses the HTTP message and generates a Siebel Internet Session Network Application Programming Interface (SISNAPI) message, based on the content of the HTTP message. SWSE also parses the incoming cookie or URL to obtain the user session ID.
 - If you are using Siebel load balancing, then SWSE forwards the request to a Siebel Server in round-robin fashion.
 - If you are using a third-party HTTP load balancer, then SWSE forwards the request to the load balancer. The load balancer uses user-configured routing rules to forward the request to a Siebel Server.

SISNAPI is a messaging format that runs on top of the TCP/IP protocol. It is used for network communication between Siebel Servers and SWSE.

- 6 On the Siebel Server, the SCBroker component receives the initial request for a session and forwards it to a Siebel Application Object Manager process. Subsequent communication for the session does not use SCBroker. For more information, see [“Siebel Application Object Manager Infrastructure” on page 22](#) and related topics.
- 7 The Siebel Application Object Manager receives and processes the SISNAPI message sent from SWSE.

If a database query is needed to retrieve the information, then the Siebel Application Object Manager formulates the SQL statement and sends the request to the Siebel Database over a database connection. The database request goes through the database connection, using a protocol format that is specific to the database connector.

- 8 The database executes the SQL statement and returns data back to the Siebel Application Object Manager. The Siebel Application Object Manager forwards the message to the Web server that originated it. If using a third-party HTTP load balancer, then the message might go through the load balancer before reaching the Web server.
- 9 The SWSE on the Web server receives the SISNAPI message, and translates it back to HTTP. It then forwards the HTTP message to the Web server. The message is now in the form of Web page content.
- 10 The Web server load balancer, if present, then forwards the Web page content through the original HTTP connection to the end user's Web browser.
- 11 The Web browser and the Siebel Web Client framework process and display the return message.

Performance Tuning Terminology

[Table 3 on page 20](#) provides definitions of specific terms related to performance and tuning Siebel Business Applications. For definitions of *performance* and *scalability*, see [“About Performance and Scalability” on page 13](#).

For more information about some of these terms and concepts (including concurrent users and think time) in the context of tuning Siebel Application Object Manager components, see [“Performance Factors for Siebel Application Object Manager Deployments” on page 24](#).

Table 3. Performance Tuning Terminology

Term	Definition
Concurrent users	The number of application users actively using and accessing the Siebel application, or a particular element, such as a Siebel Application Object Manager process, at a particular time.
Latency	Delay experienced in network transmissions as network packets traverse the network infrastructure.
Think time	<p>The wait time between user operations. For example, if a user navigates to the Account screen and reviews data for 10 seconds before performing another operation, then the think time in this case is 10 seconds.</p> <p>Average think time is a critical element in performance and scalability tuning, particularly for Siebel Application Object Manager. When think time values are correctly forecasted, then actual load levels will be close to anticipated loads.</p>
Process	An operating system (OS) process. For example, a Siebel Server component such as Siebel Application Object Manager consists of multiple OS processes, referred to as multithreaded processes.
Multithreaded process (or MT server)	A process running on a multithreaded Siebel Server component that supports multiple threads (tasks) per process. Siebel Application Object Manager components run multithreaded processes that support threads.
Task	A concept for Siebel Business Applications of a unit of work that can be done by a Siebel Server component. Siebel tasks are typically implemented as threads.
Thread	An operating system feature for performing a given unit of work. Threads are used to implement tasks for most Siebel Server components. A multithreaded process supports running multiple threads to perform work such as to support user sessions.
Response time	Amount of time the Siebel application takes to respond to a user request, as experienced by the end user. Response time is an aggregate of time incurred by all server processing and transmission latency for an operation. Response time is based on processing related to the request and to processing for other requests that might affect this user request.
Throughput	Typically expressed in transactions per second (TPS), expresses how many operations or transactions can be processed in a set amount of time.

3

Tuning Siebel Application Object Manager

This chapter describes the structure and operation of Siebel Application Object Manager components and the tuning that might be required for optimal operation. It contains the following topics:

- [About the Siebel Application Object Manager on page 21](#)
- [Siebel Application Object Manager Infrastructure on page 22](#)
- [Performance Factors for Siebel Application Object Manager Deployments on page 24](#)
- [Topology Considerations for Siebel Application Object Manager Deployments on page 27](#)
- [Guidelines for Siebel Application Object Manager Tuning on page 27](#)
- [Configuring Database Connection Pooling for Siebel Application Object Managers on page 36](#)
- [Using Thread Pooling for Siebel Application Object Manager on page 44](#)

Related Books

Siebel Deployment Planning Guide

Siebel System Administration Guide

Siebel Installation Guide for the operating system you are using

About the Siebel Application Object Manager

The term *Siebel Application Object Manager* refers to any of several Siebel Server components that support users accessing Siebel Business Applications through the Siebel Web Client and a Web server.

A different Siebel Application Object Manager component is provided for each base application among the Siebel Business Applications or Siebel Industry Applications. For example:

- Call Center Object Manager (SSCObjMgr_enu) is the Siebel Application Object Manager for Siebel Call Center in a U.S. English environment.
- Sales Object Manager (SSEObjMgr_enu) is the Siebel Application Object Manager for Siebel Sales in a U.S. English environment.
- eService Object Manager (eServiceObjMgr_enu) is the Siebel Application Object Manager for Siebel eService in a U.S. English environment.

NOTE: Separate Siebel Application Object Managers are provided for each installed language in which you can run your Siebel applications. For example, Call Center Object Manager for French is SSCObjMgr_fra.

When configured appropriately, Siebel Application Object Manager components on your Siebel Servers can use memory and CPU resources efficiently, and can communicate efficiently with the Siebel Database, the Siebel Web Server Extension (SWSE), and other components in the Siebel Enterprise.

The multiprocess, multithreaded model for Siebel Application Object Manager components provides scalability to support deployments with a wide range of concurrent Siebel Business Applications users.

The overall performance of the Siebel Application Object Manager contributes significantly to the response time as experienced by your end users.

Siebel Application Object Manager Infrastructure

A Siebel Application Object Manager component is implemented as a *multithreaded process* on the Siebel Server. At runtime, a parent process starts one or more multithreaded processes, according to the Siebel Application Object Manager configuration.

Each Siebel Application Object Manager process can host multiple user sessions (as tasks), which in turn are implemented as threads within the process. These threads might be dedicated to particular user sessions, or they might serve as a pool that can be shared by multiple user sessions. (For each process, a few threads also start that are dedicated to performing core functions for the process.)

As more users log in, additional processes can be instantiated to host these users.

- In this chapter, the term *thread* is often used interchangeably with *task*, except when you are using thread pooling. For details, see [“Using Thread Pooling for Siebel Application Object Manager” on page 44](#).
- The terms *multithreaded server* or *MT server* are alternative terms for multithreaded process (a process that supports multiple threads). For example, the names of the Siebel Application Object Manager parameters MaxMTServers and MinMTServers refer to multithreaded processes.

Siebel Application Object Manager components, which run in interactive mode, handle processing for Siebel Web Client sessions, in which the application user interface (UI) resides. The Siebel Application Object Manager task manages Siebel business objects and data objects and performs business logic for the client session.

Generally, each Siebel Application Object Manager task starts in response to a request from a Siebel Web Client running in a Web browser, and ends when the client disconnects.

Siebel Application Object Manager Communications with Other Modules

Each Siebel Application Object Manager task uses Siebel Server infrastructure capabilities to communicate with the Siebel Database, the Web server (through the SWSE), and other Siebel Enterprise Server components.

- Communication with the Siebel Database uses database connections. Database connections can also be managed and tuned for optimal performance. You can optionally configure connection pooling for database connections. For details on configuring database connection pooling, see [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#).
- Communication between the Siebel Connection Broker (SCBroker) and the Siebel Application Object Manager processes on the Siebel Internet Session Network Application Programming Interface (SISNAPI) same Siebel Server uses mechanisms internal to the operating system. SCBroker receives each SISNAPI connection request from the SWSE and forwards the connection request to a Siebel Application Object Manager multithreaded process. Once the request has been forwarded, subsequent requests for the same user session flow directly from SWSE to this Siebel Application Object Manager process.

The request is forwarded using either a least-loaded or a round-robin algorithm, according to the setting of the SCBroker parameter ConnForwardAlgorithm.

For more information about configuring SCBroker, see *Siebel Deployment Planning Guide* and *Siebel System Administration Guide*. See also the *Siebel Installation Guide* for the operating system you are using.
- Communication with the Siebel Web Server Extension uses SISNAPI, a messaging format that runs on top of the TCP/IP protocol. SISNAPI connections can be configured to use encryption and authentication based on Secure Sockets Layer (SSL). For details on tuning SISNAPI communications, see [“Configuring SISNAPI Connection Pooling for Siebel Application Object Manager” on page 47](#).
- Communication with other Siebel Enterprise Server components (including other Siebel Servers) also uses SISNAPI, going through Server Request Broker (SRBroker). For more information about tuning SRBroker, see [“Tuning Server Request Broker \(SRBroker\)” on page 49](#).

About Tuning the Siebel Application Object Manager

Tuning activities directly or indirectly applicable to Siebel Application Object Manager components might involve any or all of the following:

- Configuring your system initially using Siebel Configuration Wizards. These wizards are described in the *Siebel Installation Guide* for the operating system you are using.
- Using the Siebel Server Manager (GUI or command-line version) to tune parameters for the Enterprise Server, the Siebel Server, or the Siebel Application Object Manager component. These parameters are stored in the siebns.dat file in a directory on the Siebel Gateway Name Server.
- Selectively enabling component groups and components on each Siebel Server. Only enable the component groups and components you need.
- Tuning parameters in the eapps.cfg file on the Siebel Web Server Extension. This file is located in the bin subdirectory of the Siebel Web Server Extension installation directory, on the Web server computer. You configure the SWSE initially using configuration wizards, as described in the *Siebel Installation Guide* for the operating system you are using.

Some other chapters in this book discuss Siebel Application Object Manager tuning that relates to using other modules, such as Siebel Communications Server or Siebel Configurator.

Performance Factors for Siebel Application Object Manager Deployments

In planning to deploy Siebel Application Object Managers, or in troubleshooting performance for existing Siebel Application Object Manager deployments, you must consider several factors that determine or influence performance.

Factors that are central to the task of configuring the Siebel Application Object Manager are also called *performance drivers*. Performance drivers for Siebel Application Object Manager include concurrent users and average think time. Other important factors such as hardware resources will set limits on overall capacity or capacity per server.

Subsequent topics provide information and guidelines to help you achieve and maintain optimal performance and scalability.

These factors are critical in initially configuring your Siebel Application Object Managers, particularly when specifying values for the Siebel Application Object Manager component parameters MaxTasks, MaxMTServers, and MinMTServers, which are discussed in [“Tuning Siebel Application Object Manager Components for CPU and Memory Utilization” on page 27](#).

Concurrent Users

The number of concurrent users is the total number of user sessions supported at any one time. It also includes sessions supporting anonymous browser users. For planning and tuning purposes, you must consider concurrent users (and total users) at multiple levels:

- The entire deployment (enterprise)
- Each Siebel Server
- Each Siebel Application Object Manager component on each server
- Each multithreaded process for each Siebel Application Object Manager component

The maximum number of concurrent users per Siebel Server (assuming, for example, that a particular Siebel Server computer is dedicated to running Siebel Application Object Manager components) depends on the average think time, on your hardware resources, and on the nature of your Siebel Business Applications deployment.

In terms of configuration, the maximum number of concurrent users for the Siebel Application Object Manager is limited by the value of the MaxTasks parameter. The effective maximum is also limited by the number of multithreaded processes for this Siebel Application Object Manager and by your hardware resources.

Depending on the average think time and other factors, each multithreaded process (process within the Siebel Application Object Manager) typically supports a maximum of about 100 concurrent users. Configure enough multithreaded processes (using the MaxMTServers parameter) to support the maximum number of concurrent users required for your peak loads.

NOTE: Some complex or specialized Object Manager components support fewer concurrent users. For example, Object Managers for Siebel eCommunications (part of Siebel Industry Applications) and Siebel Configurator typically support about 25 concurrent users. For more information about the Object Manager for Siebel Configurator (Siebel Product Configuration Object Manager), see [Chapter 8, "Tuning Siebel Configurator."](#)

Think Time

Think time is the average elapsed time between operations performed by users in a Siebel application. Think time includes the time required by users to conduct customer interactions, enter data into the application, and work in other applications.

The assumed think time has a direct relationship to the number of concurrent tasks that a multithreaded process can support.

Determine the average think time based on the usage patterns typical of your user base. After the application has been configured, perform a clickstream analysis for your key processes, and try to capture the time between the user actions (operations) that are represented by the clicks. Also use the list statistics command in Siebel Server Manager to help you calculate average think time.

Consider the average time between each operation (such as clicking New) and each overall transaction (such as performing all steps for creating a new contact). Mouse clicks do not equate to operations if they do not send a request to the Siebel application infrastructure. Calculate the overall average think time based on all of these factors.

The ratio of 100 (100 tasks per process), based on a 30-second think time, is assumed in the formula for setting the MaxMTServers parameter. This formula is presented in ["Tuning Siebel Application Object Manager Components for CPU and Memory Utilization" on page 27](#).

The ratio of 100 is based on having approximately three users running operations at the exact same time (100 divided by 30 = approximately 3.3). It is generally observed that each multithreaded process can handle about three operations at the same time with minimal performance degradation.

With longer think times, one multithreaded process can support more than 100 concurrent tasks; with shorter think times, fewer tasks. For example, if the think time is 15 seconds between user operations, then about 50 tasks per process could be supported (15 times 3.3 = approximately 50, or 50 divided by 15 = approximately 3.3).

Nature of Siebel Application Deployment

Which Siebel applications and other modules you are using, how you have configured your Siebel applications, how you have deployed your applications, and other such factors also affect Siebel Application Object Manager performance and how many concurrent users you can support. Some of these factors include:

- Will you support employee applications (such as Siebel Call Center), customer applications (such as Siebel eService), partner applications (such as Siebel PRM), or some combination of these? Typically, employee applications use high interactivity and customer applications use standard interactivity.
- Will you deploy your Siebel software in a global environment using multiple languages?
- What degree and what kind of application configuration changes have you made, such as those you do using Siebel Tools? For more information, see [Chapter 12, “Tuning Customer Configurations.”](#)

The number of concurrent tasks you can support varies based on the level of customization or the use of process automation for the application the Siebel Application Object Manager supports. Recommendations in this guide generally assume that operations performed are fairly standard or typical. Depending on your deployment and the modules used, some operations initiated by a single user action can be relatively complex and demand more resources than most other operations.

- Will you use specialized functionality such as offered by Siebel Configurator (for product configuration) or Siebel CTI (computer telephony integration for call center agents)? How will you deploy such functionality? What percentage of your user base will use such functionality? These are only examples of such specialized functionality.

Hardware Resources

Hardware resources for each Siebel Server computer, particularly CPU and memory, are a factor in how many concurrent users can be supported for each Siebel Application Object Manager component. For example, a four-way computer has twice the resources of a two-way computer and can potentially support twice as many concurrent users. Key hardware resources for Siebel Application Object Manager performance include:

- **CPU.** The CPU rating and the number of CPUs per server computer.
- **Memory.** The amount of RAM, and whether it can accommodate users without excessive paging.

Disk I/O and network capacity are other important hardware factors, but they do not affect Siebel Application Object Manager tuning. They do significantly affect performance for the Siebel Database and the Siebel File System, which can severely impact the overall user response time.

The total number of computers you can devote to supporting Siebel Application Object Manager components will determine the total number of concurrent users.

Topology Considerations for Siebel Application Object Manager Deployments

Your Siebel applications can be deployed using a variety of topologies, or system layouts. Although Siebel Application Object Managers are only a part of the overall deployment, they play a direct and central role in supporting Siebel application users.

You must determine on how many computers you will run Siebel Server, and on how many of these you will run Siebel Application Object Manager components. In some cases, you might choose to run multiple components on the same Siebel Server.

NOTE: Siebel Application Object Manager components are typically the major resource consumers for your Siebel Server computers. Tuning considerations discussed in this chapter generally assume that you are not running additional components on a Siebel Application Object Manager computer that will significantly contend for available resources.

For more information about topology considerations, see *Siebel Deployment Planning Guide*.

Guidelines for Siebel Application Object Manager Tuning

Using your hardware resources optimally and configuring your system appropriately can help you to achieve your performance goals. Consider your resources and requirements carefully, and test and monitor system performance on a continual basis.

Review information presented in *Siebel System Administration Guide* and other sources. All tuning calculations must be done with some understanding of the overall system and the considerations described in [“Performance Factors for Siebel Application Object Manager Deployments” on page 24](#).

Review the following information for more details on Siebel Application Object Manager tuning:

- [“Tuning Siebel Application Object Manager Components for CPU and Memory Utilization” on page 27](#)
- [“Tuning Parameters for Siebel Application Object Manager Caches” on page 32](#)
- [“Additional Parameters Affecting Siebel Application Object Manager Performance” on page 34](#)
- [“Memory Consumers in Siebel Application Object Manager” on page 35](#)

Tuning Siebel Application Object Manager Components for CPU and Memory Utilization

This topic is part of [“Guidelines for Siebel Application Object Manager Tuning” on page 27](#). It provides background information and guidelines for tuning your Siebel Application Object Manager components, particularly for setting values for the parameters MaxTasks, MaxMTServers, and MinMTServers.

Settings for these parameters determine how well the system performs under specific user load and operations. Parameter settings provide a means of controlling the server capacity through the Siebel Server infrastructure, and directly impact the overall capacity for each server.

How you set the MaxTasks, MaxMTServers, and MinMTServers parameters is a direct function of the factors described in [“Performance Factors for Siebel Application Object Manager Deployments” on page 24](#), which determine the true capacity of the server.

The art of tuning Siebel Application Object Manager components is to come up with the right parameter settings that allow the server computers to host the largest number of users (scalability) with minimal impact on user response time (performance).

About MaxTasks, MaxMTServers, and MinMTServers

The Siebel Application Object Manager parameters MaxTasks, MaxMTServers, and MinMTServers are described below. You configure these parameters using Siebel Server Manager, which is described in detail in *Siebel System Administration Guide*.

For background information about multithreaded processes, threads, and related concepts, see [“Siebel Application Object Manager Infrastructure” on page 22](#).

- **MaxTasks (Maximum Tasks)**. Specifies the total number of tasks (threads) that can run concurrently on this Siebel Application Object Manager, for this Siebel Server. Beyond this number, no more tasks can be started to handle additional requests.
- **MaxMTServers (Maximum MT Servers)**. Specifies the maximum number of multithreaded processes that can run concurrently on this Siebel Application Object Manager. Beyond this number, no more multithreaded processes can be started to handle additional requests.
- **MinMTServers (Minimum MT Servers)**. Specifies the default minimum number of multithreaded processes that will start on this Siebel Application Object Manager when the parent process is started. The parent process can be started either explicitly (using Siebel Server Manager) or automatically (if the Siebel Server is started when the component state was last set to Running). Setting MinMTServers to 0 effectively disables the Siebel Application Object Manager component.

As more users log in, new tasks start to handle these sessions, and new multithreaded processes are started to support the additional tasks. The tasks and processes are added according to the Siebel Application Object Manager load-balancing behavior, up to the maximum number of tasks and maximum number of multithreaded processes. For details, see [“Effect of Siebel Application Object Manager Parameter Settings” on page 29](#).

NOTE: MaxTasks, MaxMTServers, and MinMTServers are generic parameters that apply to many different Siebel Server components. However, the specific behavior described in this chapter applies to Siebel Application Object Manager components. For more information, see *Siebel System Administration Guide*.

These parameters relate to one another in the following ways:

- MaxMTServers and MinMTServers are typically set to the same value. Doing this avoids any performance penalty for a user whose login causes a new multithreaded process to start. MaxMTServers must be equal to or greater than MinMTServers.

Starting all multithreaded processes up front when the parent process is started is generally acceptable. The memory overhead for running a multithreaded process itself, apart from the overhead of its threads, is minimal.
- The ratio MaxTasks/MaxMTServers (MaxTasks divided by MaxMTServers) determines the maximum number of threads (tasks) that can run concurrently on a given multithreaded process. For more information, see the discussion of think time under [“Performance Factors for Siebel Application Object Manager Deployments”](#) on page 24.

Effect of Siebel Application Object Manager Parameter Settings

The following information illustrates how a Siebel Application Object Manager behaves given particular example settings for the MaxTasks, MaxMTServers, and MinMTServers parameters. More realistic examples can be found in [“Formulas for Calculating Siebel Application Object Manager Parameter Values”](#) on page 30.

For example, if MaxTasks = 500, and MaxMTServers = 5, then MaxTasks divided by MaxMTServers = 100. This means that, at most, 100 threads (tasks) can run in a multithreaded process on this Siebel Application Object Manager.

Typically, MinMTServers would be set the same as MaxMTServers. However, in this example, assume MinMTServers = 4. In this case, four multithreaded processes start by default, which can handle a total of 400 concurrent threads.

As users start the application on the server, the number of concurrent threads rises, and the following occurs:

- As the number of concurrent threads rises, but remains below 400, these threads are distributed among the four multithreaded processes that started by default for this Siebel Application Object Manager. This is a form of load balancing internal to the Siebel Application Object Manager component.
- If the number of concurrent threads reaches 400, and a new request is received, then a fifth multithreaded process starts for this Siebel Application Object Manager. The Siebel Application Object Manager now distributes threads among five multithreaded processes for this Siebel Application Object Manager.
- If the Siebel Application Object Manager reaches 500 concurrent threads, then no more client session requests can be handled, because the existing multithreaded processes can start no more threads, and the Siebel Application Object Manager can start no more multithreaded processes. The Siebel Application Object Manager can be said to be “maxed out.”

If Siebel Application Object Manager loads fall back, then as users log out or session timeouts are enforced, then threads are freed up. In some cases, a multithreaded process whose threads have completed might also time out and stop running; this can happen only when MaxMTServers is greater than MinMTServers.

Guidelines for Configuring Siebel Application Object Manager Parameters

The following information provides formulas and guidelines for setting the MaxTasks, MaxMTServers, and MinMTServers parameters for your Siebel Application Object Manager components.

NOTE: All elements in the two formulas shown in “Formulas for Calculating Siebel Application Object Manager Parameter Values” on page 30 vary according to your deployment. The number of concurrent users a Siebel Application Object Manager can support depends on factors such as the number of processors, session timeout settings, and the average think time.

Typically, the Siebel Application Object Manager is the only component using significant resources on the Siebel Server computer. If you run multiple server components, or run non-Siebel modules, then a Siebel Application Object Manager on this computer might support fewer concurrent threads.

Follow these general steps to determine how to set these parameter values:

- Determine the total number of concurrent users, based on the average think time and other factors discussed earlier.
- Determine the number of concurrent users that must be supported on a given Siebel Server computer running Siebel Application Object Manager. In the formulas outlined below, this is the target number of users plus the number of anonymous browser users, where applicable.
- Determine how many Siebel Server computers are needed to support your concurrent users. This is typically done by Oracle's Application Expert Services or by platform vendors. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

NOTE: Because each customer application implementation is unique, customers are encouraged to engage Oracle's Application Expert Services for a detailed sizing review to best assess the hardware needs to support their deployment. Application Expert Services has experience from thousands of implementations and works closely with Oracle's Performance and Scalability team to provide sizing guidance.

- Plug your values into the formulas below, then adjust the values to meet any additional criteria. In particular:
 - If your calculated value for MaxMTServers is not an integer, then round up the value to the nearest integer.
 - After you adjust the value of MaxMTServers, if your calculated ratio for MaxTasks divided by MaxMTServers is not an integer, then round up the value of MaxTasks until this ratio is an integer.
- Test your initial parameter settings, such as to gauge the actual number of anonymous browser users required, then adjust settings further as necessary.

Formulas for Calculating Siebel Application Object Manager Parameter Values

Use the formulas below for calculating parameter values for your Siebel Application Object Manager components:

- $\text{MaxTasks} = \text{target_number_of_users plus anon_browser_users}$

- $\text{MaxMTServers} = (\text{target_number_of_users} \text{ plus } \text{anon_browser_users}) \text{ divided by } 100$
- $\text{MinMTServers} = \text{MaxMTServers}$

As necessary, after making your initial calculations, round up MaxMTServers to the nearest integer, calculate the remainder (X) of MaxTasks divided by MaxMTServers, then increment MaxTasks by adding (MaxMTServers minus X). You do this to make sure that the ratio of MaxTasks divided by MaxMTServers is an integer.

NOTE: The figure of 100 in the MaxMTServers formula represents the ratio of concurrent tasks per multithreaded process. The value of 100 is a rule of thumb only.

The following descriptions apply to the variables used in the formulas:

- *target_number_of_users* is the maximum number of concurrent user sessions your Siebel Application Object Manager will support (for users who are logged into the application).

The maximum number of concurrent users is limited by the value of the MaxTasks parameter for the Siebel Application Object Manager, by the number of multithreaded processes you are running (determined by MaxMTServers and MinMTServers), and, effectively, by your hardware resources.
- *anon_browser_users* is the number of sessions on the Siebel Application Object Manager dedicated to anonymous browser users (threads that support users who do not log in).
 - For high interactivity applications (typically, employee applications like Siebel Call Center), anonymous browser users are not supported, so this is not a factor.
 - For standard interactivity applications (typically, customer applications like Siebel eService), anonymous browser users can be approximately 25% of the target number of users.
- 100 is the approximate maximum number of concurrent threads each multithreaded process on the Siebel Application Object Manager can support. The number 100 is a rule of thumb. Use the number that is appropriate for your deployment.

NOTE: A ratio of 100 for threads per multithreaded process works for most Siebel Application Object Manager usage scenarios. However, if your deployment involves a shorter think time than 30 seconds, or a heavier than average load per thread, then each multithreaded process will support fewer concurrent threads. Conversely, a longer think time or a lighter average load will support more concurrent threads. For more information about how the ratio of threads per multithreaded process relates to think time, see [“Performance Factors for Siebel Application Object Manager Deployments” on page 24](#).

Example Settings for Siebel Application Object Manager Parameters

Along with other factors such as think time, the calculation of MaxTasks, MaxMTServers, and MinMTServers depends on your assumptions for *target_number_of_users* and *anon_browser_users*. For more information about these parameters, see [“Formulas for Calculating Siebel Application Object Manager Parameter Values” on page 30](#). Example settings follow for Siebel Call Center and Siebel eService.

Example Settings for Siebel Call Center

For Siebel Call Center, assume (for example) a think time of 30 seconds, and assume that *target_number_of_users* = 500. For this application, *anon_browser_users* is not a factor. Your parameter values would be:

MaxTasks = 500

MaxMTServers = 500 divided by 100 = 5

MinMTServers = MaxMTServers = 5

Example Settings for Siebel eService

For Siebel eService, assume (for example) a think time of 30 seconds, and assume that *target_number_of_users* = 500. Depending on your implementation, *anon_browser_users* might be about 25% of *target_number_of_users* (or 125). Your preliminary parameter values would be:

MaxTasks = (500 plus 125) = 625

MaxMTServers = (500 plus 125) divided by 100 = 6.25 = 7 (round up)

MinMTServers = MaxMTServers = 7

Adjust the value of MaxTasks. The variable X = the remainder of (625 divided by 7) = 2. Increment MaxTasks by (MaxMTServers minus X): 625 plus (7 minus 2) = 625 plus 5 = 630. Therefore, the final calculations for parameter values would be:

MaxTasks = 630

MaxMTServers = MinMTServers = 7

Tuning Parameters for Siebel Application Object Manager Caches

This topic is part of [“Guidelines for Siebel Application Object Manager Tuning” on page 27](#).

The Siebel Application Object Manager uses several caches, which affect memory usage for the Siebel Application Object Manager. Tuning Siebel Application Object Manager caches affects Siebel Application Object Manager performance and memory usage. The following are some of the major caches used by Siebel Application Object Manager that can be configured:

- SQL cursor cache
- SQL data caches

SQL Cursor Cache

The SQL cursor cache is configured using the DSMaxCachedCursors parameter. This cache can be enabled on multithreaded components (such as Siebel Application Object Manager) with database connection pooling.

The value represents the number of SQL cursors per database connection. For a Siebel Application Object Manager for which the Siebel Server computer is more likely to reach its CPU capacity before it reaches its memory capacity (for example, for Siebel Employee Relationship Management), the default value of 16 for the `DSMaxCachedCursors` parameter might be appropriate. (Such an application is sometimes referred to as *CPU-bound*.)

For a Siebel Application Object Manager for which the Siebel Server computer is more likely to reach its memory capacity before it reaches its CPU capacity (for example, for Siebel Call Center), you can set `DSMaxCachedCursors` to a lower value, even to 0. (Such an application is sometimes referred to as *memory-bound*.)

In general, the value must reflect the CPU and memory resource availability on the Siebel Server computer running a particular Siebel Application Object Manager component. The trade-off in setting this parameter is that allocating memory to caching SQL cursors means they would need to be created less often, but at a cost in memory.

The memory requirement per cursor depends on factors such as the size of the query, type of database connection, row size, and number of rows returned by the query. The utility of these cached cursors depends on the uniqueness of the queries they represent. In general, most Siebel application queries are unique and would not benefit from reusing a cached cursor.

Generally, when more users share a database connection, through connection pooling, you increase the number of cursors cached, provided that the required memory is available. For more information about database connection pooling, see [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#).

SQL Data Caches

The SQL data caches are configured using the `DSMaxCachedDatasetsPerProcess` and `DSMaxCachedDataSets` parameters. Two types of data caches are guided by these parameters:

- Global data cache, which is useful in most cases. This cache is governed by `DSMaxCachedDatasetsPerProcess`. The default value is 16.
- Per-connection data cache (which can be enabled with, or without, database connection pooling). This cache is governed by `DSMaxCachedDataSets`. The default value is 16.

For an CPU-bound Siebel Application Object Manager (for example, for Siebel Employee Relationship Management), the default values for `DSMaxCachedDatasetsPerProcess` and `DSMaxCachedDataSets` might be appropriate.

For a memory-bound Siebel Application Object Manager (for example, for Siebel Call Center), you can set `DSMaxCachedDatasetsPerProcess` and `DSMaxCachedDataSets` to a lower value, even to 0.

In general, the values must reflect the CPU and memory resource availability on the Siebel Server computer running a particular Siebel Application Object Manager component. The trade-off in setting these parameters is that allocating memory to caching SQL data sets means they would need to be created less often, but at a cost in memory. See also the discussion of the SQL cursor cache.

Additional Parameters Affecting Siebel Application Object Manager Performance

This topic is part of “[Guidelines for Siebel Application Object Manager Tuning](#)” on page 27. It provides guidelines for setting additional parameters that affect Siebel Application Object Manager performance.

- **MemProtection.** Setting the MemProtection parameter to FALSE for your Siebel Application Object Manager component might improve performance.

When this parameter is TRUE (the default), each transaction issues a large number of serialized mprotect statements, the total effect of which can degrade performance on your Siebel Server computers.

The MemProtection parameter is hidden. Click Hidden in the Component Parameters view tab to display it. Alternatively, you can set it using the command-line version of the Siebel Server Manager, as shown:

```
change param MemProtection=False for comp component_alias_name server  
siebel_server_name
```

where:

- *component_alias_name* is the alias name of the Siebel Application Object Manager component you are configuring, such as SCCObjMgr_deu for the German version of Call Center Object Manager.
- *siebel_server_name* is the name of the Siebel Server for which you are configuring the component.

In addition to improving performance, you can also improve scalability if you set the MemProtection parameter to FALSE.

NOTE: Support for setting MemProtection to FALSE is restricted to the AIX and HP-UX platforms.

- **DSMaxFetchArraySize.** This is a named subsystem parameter that controls the maximum number of records that can be returned by a business component in ForwardBackward mode. It does not restrict the number of records returned for ForwardOnly cursors. By default, DSMaxFetchArraySize has a value of 0. When this parameter is set to 0, the Siebel Application Object Manager initializes the parameter to 10,000. This means that a maximum of 10,000 records can be returned by a business component in ForwardBackward mode.
- **DSPreFetchSize and DSMaxCursorSize.** Set these parameters *only* for Siebel implementations on DB2 for z/OS and OS/390. For more information on setting these parameters, see *Implementing Siebel Business Applications on DB2 for z/OS*.

For all other databases, these parameters must be set to -1.

- **EnableCDA.** This parameter must be set to FALSE for the Siebel Application Object Manager component.

Memory Consumers in Siebel Application Object Manager

This topic is part of [“Guidelines for Siebel Application Object Manager Tuning” on page 27.](#)

In addition to the caches described earlier, this topic discusses major memory consumers in Siebel Application Object Manager components. For more information on some of these topics, see [Chapter 12, “Tuning Customer Configurations.”](#)

NOTE: Specific performance and scalability testing is required to determine the size of Siebel Application Object Manager memory required.

- **Database client libraries.** Database client libraries have their own caches, caching metadata, connections, cursors, and data. Some of these caches can be reduced in size by using Siebel database connection pooling, described in [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36.](#)
- **Scripts.** A script defined on a business component, applet, or business service is loaded into Siebel Application Object Manager memory when the script is first invoked.

For Siebel eScript, garbage collection is performed according to settings that are optimized for each release in order to use server memory and other resources appropriately.
- **Heavy configurations.** Performance is affected when an application is heavily configured, because the memory for the Siebel Application Object Manager increases proportionally in this case.
- **Navigation pattern.** Numerous scenarios that can be used to navigate in the application can make using global caches ineffective.
- **Session timeouts.** Higher session timeout values mean more active sessions on the server at a time, therefore more memory being used. Lower session timeout values can mean more frequent logins.
- **Users per Siebel Application Object Manager.** More users per Siebel Application Object Manager means more sharing of global resources between the users. While the amount of memory used *per user* on this Siebel Application Object Manager is less, more memory is used overall.
- **Number of applets on views.** More applets configured on views means more business components will be needed at a time, hence higher overall memory usage.
- **PDQ size.** The list of items in the PDQ (predefined queries) list are maintained on the server for the current business object. The higher the number of items in this list, the more memory it consumes. The size of PDQ strings also determines memory usage.

Configuring Database Connection Pooling for Siebel Application Object Managers

This topic describes database connection configuration options for Siebel Application Object Managers, particularly database connection pooling. It contains the following topics:

- [“About Database Connections for Siebel Application Object Manager” on page 36](#)
- [“Database Connection Pooling Usage Guidelines” on page 37](#)
- [“Configuring Pooling for Default Database Connections” on page 40](#)
- [“Configuring Pooling for Specialized Database Connections” on page 42](#)

NOTE: Each customer must determine whether their RDBMS has a sufficient total number of database connections for their needs. The total number of available connections is subject to limitations deriving from RDBMS and operating system platforms and other factors. Before you configure connection pooling, verify how many database connections are available for use by the Siebel Application Object Manager. RDBMS performance and usage of database connections by non-Siebel components are outside the scope of this guide.

About Database Connections for Siebel Application Object Manager

This topic is part of [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#). It provides an overview of database connections for Siebel Application Object Manager components, including nonpooled connections and pooled connections. Subsequent topics provide guidelines and instructions for configuring different types of database connection pooling.

About Nonpooled Database Connections

If you do not pool database connections, then the number of database connections corresponds to the number of Siebel Application Object Manager sessions; that is, database connections are not pooled. No special Siebel Application Object Manager configuration is required for using nonpooled database connections. When no pooling is configured, database connections are closed when the user session terminates.

- **Nonpooled default database connections.** With nonpooled database connections, during session login, a database connection is established, using the user’s database credentials. (When an external authentication system is used, such as LDAP, the user’s database credentials might not be the same as the user’s Siebel credentials.)

This database connection becomes bound to the session, and is the default database connection used for read, write, update, and delete operations.

In this book, such connections are called *default database connections*. These connections can alternatively be pooled, as described later in this topic.

- **Nonpooled specialized database connections.** If, during a session, specialized functionality is invoked that uses the external transaction management capabilities of the Siebel Application Object Manager, then a second database connection is opened for this specialized use.

This database connection is also bound to the session, and is used for all externally controlled transactions performed by the session. Siebel EAI components are an example of specialized code that does external transaction management.

In this book, such connections are called *specialized database connections*. These connections can alternatively be pooled, as described later in this topic.

About Pooled Database Connections

Optionally, you can configure your Siebel Application Object Manager components to support pooling for the same two types of database connections described previously for nonpooled database connections:

- **Pooled default database connections.** These database connections can be pooled to support sharing (multiplexing), persistence, or both features.
 - Shared connections support multiple user sessions at the same time, by multiplexing (sharing) database operations for multiple SQL statements over the same database connection. Using shared connections can support more users with a given number of connections.
 - Persistent connections are pooled, but are not necessarily shared. Using persistent connections can enhance performance by avoiding the cost of creating database connections. All shared connections are also persistent connections.

For details, see [“Database Connection Pooling Usage Guidelines” on page 37](#) and [“Configuring Pooling for Default Database Connections” on page 40](#).

- **Pooled specialized database connections.** These database connections are dedicated to a single session at a time, and serve a specialized purpose. Pooling such connections provides persistence, but such connections are never shared. By persistently pooling these connections, you enhance performance by avoiding the cost of creating connections.

NOTE: If you configure pooling for default database connections, but not for specialized database connections, then each specialized database connection is closed when the transaction that required it completes.

For details, see [“Database Connection Pooling Usage Guidelines” on page 37](#) and [“Configuring Pooling for Specialized Database Connections” on page 42](#).

Database Connection Pooling Usage Guidelines

This topic is part of [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#).

Observe the following guidelines to help you determine if you must use database connection pooling, or to guide your deployment of connection pooling.

For more information about configuring the Siebel Application Object Manager parameters for database connection pooling mentioned below, see [“Configuring Pooling for Default Database Connections” on page 40](#) and [“Configuring Pooling for Specialized Database Connections” on page 42](#).

When to Consider Using Database Connection Pooling

Consider implementing database connection pooling if, and *only if*, one or more of the following is true for your deployment:

- The RDBMS cannot support the number of dedicated user connections you would require if using nonpooled database connections. Pooling default database connections for shared use can reduce the number of connections you require.
 - Memory resources are scarce on the Siebel Server computer on which the Siebel Application Object Manager is running. Pooling default database connections for shared use can reduce Siebel Application Object Manager memory requirements per concurrent user.
 - Your deployment uses external authentication such as LDAP (that is, authentication other than database authentication), and creating new connections is slow on the database server. Pooling database connections can speed up login or other operations by providing persistent pooling, whether or not connections are also shared.
 - You use a Siebel Server component that requires frequent logins for special-purpose processing. Pooling database connections to provide persistent connection pooling (not sharing) can provide a significant benefit for such components.
 - For Siebel Configurator, if you are using the component Siebel Product Configuration Object Manager (alias eProdCfgObjMgr_jpn in a Japanese locale), then it is highly recommended to configure persistent connection pooling. For more information about Siebel Configurator, see [Chapter 8, “Tuning Siebel Configurator.”](#)
- NOTE:** Separate Siebel Application Object Manager components are provided for each installed and deployed language in which you can run your Siebel applications. For example, Call Center Object Manager for French is SCCObjMgr_fra.
- For some other components, such as EAI Object Manager (when run in sessionless mode), it might also be helpful to configure persistent connection pooling.

NOTE: If session caching is configured for a component (by setting the parameter ModelCacheMax), then persistent connection pooling might provide little benefit. For example, session caching is typically configured for Workflow Process Manager. For more information about session caching for Siebel Workflow, see [“Caching Sessions” on page 92](#).

Guidelines for Using Database Connection Pooling

If you decide to use database connection pooling, then observe the following guidelines:

- Start with a low ratio of MaxTasks divided by MaxSharedDbConns, such as 2:1.

NOTE: If you plan to use a ratio higher than 3:1, then it is recommended that you consult Oracle's Application Expert Services. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

- If you have short (aggressive) average think times in your user scenarios, then use a smaller ratio of MaxTasks divided by MaxSharedDbConns. Longer think times can support larger ratios.

For a 30-second think time, do not use a ratio higher than 10:1. For a 15-second think time, do not use a ratio higher than 5:1. Other factors discussed in this topic will also determine practical limits.

For more information about think time, see [“Performance Factors for Siebel Application Object Manager Deployments” on page 24](#).

- Minimize long-running queries. A database connection can process only one database operation at any particular moment. Therefore, if user sessions use all available connections in the database connection pool to execute long-running queries, then a new user session that requests a database connection is blocked until a database connection becomes available.

For example, if a long-running query is run which takes, for example, three seconds, then, for this duration, database requests from another user who requires the same database connection are queued (blocked) until the query operation completes.

A long-running query can continue to run on the RDBMS even if the user who initiated it has stopped the browser.

When using database connection pooling, it is critical to optimize database access in your environment. If long-running queries cannot be avoided, then monitor the overall database response time and use a small ratio of MaxTasks divided by MaxSharedDbConns to achieve a satisfactory response time.

Alternatively, long-running queries can be minimized or avoided by adjusting indexes to include fields that can be sorted or queried by end users, by configuring the application user interface so nonindexed fields are not exposed, or by training users to avoid operations that would perform long-running queries. For more information about how indexing can affect Siebel application performance, see [“Managing Database Indexes in Sorting and Searching” on page 180](#).

- Consider the cost of creating database connections. This cost differs based on your authentication method.

If your deployment uses database authentication, then a database connection is created for each login, for authentication purposes. Afterwards, this connection is released to the shared connection pool, if the total number of connections is less than the maximum. Or, if the pool is full, the connection is closed (terminated). Therefore, even when the pool is full and connections are available, new connections are still created temporarily for each new session login. These connections must be accounted for in determining the allocation of database connections.

With external authentication, however, you can use persistent connection pooling to reduce the cost of creating database connections. With persistent connection pooling, database connections, once created, are persistent, though such connections might or might not be shared. For pooled default database connections where connections are persistent but not shared, set $\text{MaxSharedDbConns} = \text{MaxTasks} \text{ minus } 1$.

For more information about authentication options, see *Siebel Security Guide*.

- In order to configure connection pooling for specialized database connections, you must also configure pooling for default database connections, as follows:
 - If you do not configure connection pooling for shared database connections (MaxSharedDbConns = -1 or 0), then each specialized database connection, once created, is dedicated to the user session. The value of MinTrxDBConns is ignored.
 - If you configure connection pooling for shared database connections (MaxSharedDbConns has a value greater than 0, and less than MaxTasks), then specialized database connections are not dedicated to user sessions. Such connections are handled according to the setting of MinTrxDBConns:
 - If MinTrxDBConns = -1 or 0, then, after the transaction that required it has ended, then each specialized database connection is closed (deleted).
 - If MinTrxDBConns has a value greater than 0, then, after the transaction that required it has ended, each specialized database connection can return to the connection pool.
- Siebel database connection pooling cannot be used simultaneously with MTS or the Multiplexing features of Oracle Database.

Configuring Pooling for Default Database Connections

This topic is part of [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#).

Default database connections can be used by most Siebel Application Object Manager operations.

Configuring Parameters for Pooling Default Connections

The following information describes how to enable or disable pooling for default database connections using the parameters MaxSharedDbConns (DB Multiplex - Max Number of Shared DB Connections) and MinSharedDbConns (DB Multiplex - Min Number of Shared DB Connections).

- To enable connection pooling, set MaxSharedDbConns and MinSharedDbConns to positive integer values (at least 1) that are no higher than MaxTasks minus 1. A connection will be shared by more than one user session once the number of sessions within the multithreaded process exceeds the maximum number of shared connections allowed per process.
- MaxSharedDbConns controls the maximum number of pooled database connections for each multithreaded process.
- MinSharedDbConns controls the minimum number of pooled database connections the Siebel Application Object Manager tries to keep available for each multithreaded process.

The setting of MinSharedDbConns must be equal to or less than the setting of MaxSharedDbConns. Depending on your Siebel Application Object Manager usage patterns, you can set these to the same value or set MinSharedDbConns to a lower value (if you determine this to be helpful in conserving database connection resources).

- To configure persistent and shared database connection pooling, set `MaxSharedDbConns`, using the appropriate ratio of `MaxTasks` divided by `MaxSharedDbConns`. Depending on the ratio, a greater or lesser degree of sharing will be in effect. Start with a 2:1 (or smaller) ratio for `MaxTasks` divided by `MaxSharedDbConns`. With this example ratio, two user tasks will share the same database connection.
- To configure persistent but nonshared database connection pooling, set `MaxSharedDbConns` = `MaxTasks` minus 1.
- To disable connection pooling, set `MaxSharedDbConns` and `MinSharedDbConns` to -1 (this is the default value).

`MaxSharedDbConns` and `MinSharedDbConns` are defined per Siebel Application Object Manager component, on an enterprise basis (these parameters are included in named subsystems of type `InfraDatasources`). The database connections these parameters control are not shared across multithreaded processes. The actual maximum number of database connections for each multithreaded process is determined by the ratio `MaxSharedDbConns` divided by `MaxMTServers`.

NOTE: `MaxSharedDbConns` and `MinSharedDbConns` work differently than `MinTrxDBConns`, which specifies the number of shared specialized database connections available for each multithreaded process. For details, see [“Configuring Pooling for Specialized Database Connections” on page 42](#).

Example Configuration for Pooling Default Connections

Assume, for example, the following parameter settings:

```
MaxTasks = 500
MaxMTServers = 5
MinMTServers = 5
MaxSharedDbConns = 250
MinSharedDbConns = 250
```

With these settings, the Siebel Application Object Manager component can support a maximum of 500 tasks (threads). Those 500 tasks would be spread over five multithreaded processes, each having 100 tasks. Each multithreaded process would have a maximum of 50 shared database connections, each of which would serve up to two tasks.

How Pooled Default Connections Are Assigned

When a user logs into the Siebel Application Object Manager, a database connection is established to authenticate the user, then discarded (closed) once the database or external authentication system authenticates the user. After successful authentication, the Siebel Application Object Manager's connection manager checks the connection pool for SQL statements. If this connection pool is empty, then the connection manager adds a connection.

Each time a user initializes a SQL statement, the connection manager checks the connection pool for available connections. The connection manager reserves a connection for the SQL statement in one of the following ways:

- If a connection is available to handle the SQL statement, then the connection manager assigns this connection to execute the SQL statement.

The connection manager reserves this connection until execution of the SQL statement terminates. On termination of the SQL statement, the connection manager releases the connection. At the end of a user session (due to a user logging out or session timeout), the connection manager checks the connection used by the user session. If this connection is not referenced by other user sessions and the number of connections available in the pool of database connections exceeds the value specified by `MinSharedDbConns`, then the connection manager closes this connection and releases it from the pool of database connections.

- If no connection is available in the connection pool, then the connection manager creates a new connection and assigns it to execute the SQL statement.

The connection manager continues to add connections to the connection pool until the number of connections in the connection pool equals `MaxSharedDBConns`.

Configuring Pooling for Specialized Database Connections

This topic is part of [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#).

Specialized database connections, which are not shared, are used primarily by specialized Siebel components such as Siebel EAI that need transactions to span multiple Siebel Application Object Manager operations. These connections are used for operations that use `BEGIN TRANSACTION` and `END TRANSACTION`.

Configuring Parameters for Pooling Specialized Connections

The following information describes how to enable or disable specialized connection pooling using the parameter `MinTrxDBConns` (DB Multiplex - Min Number of Dedicated DB Connections).

- `MinTrxDBConns` controls the minimum number of specialized database connections for each multithreaded process. The connections are not created until they are needed. The minimum value is thus the minimum size of the pool of specialized connections once all the connections in the pool have been created.
 - To enable specialized connection pooling, set `MinTrxDBConns` to a positive integer value (at least 1). You must also configure pooling for default database connections.
 - To disable specialized connection pooling, set `MinTrxDBConns` to -1 (this is the default value).
- There is no explicit limit to the maximum number of specialized connections. However, effectively, there cannot be more specialized connections than sessions. On average, there will be many fewer connections than sessions.

MinTrxDbConns is defined per Siebel Application Object Manager component, on an enterprise basis (this parameter is included in named subsystems of type InfraDatasources). The database connections this parameter controls are not shared across multithreaded processes. The actual minimum number of specialized database connections for each multithreaded process is what you specify as the value for MinTrxDbConns.

NOTE: MinTrxDbConns works differently than MaxSharedDbConns and MinSharedDbConns. MaxSharedDbConns and MinSharedDbConns specify the number of shared database connections available for all Siebel Application Object Manager processes while MinTrxDbConns specifies the number of specialized database connections per Siebel Application Object Manager process. For details, see [“Configuring Pooling for Default Database Connections” on page 40](#).

Example Configuration for Pooling Specialized Connections

Assume, for example, the following parameter setting, in addition to those described in [“Example Configuration for Pooling Default Connections” on page 41](#):

MinTrxDbConns = 5

With this setting, each multithreaded process would have a minimum of five specialized database connections. If all five multithreaded processes are running on this Siebel Application Object Manager, then there would be a minimum of 25 specialized connections for this Siebel Application Object Manager.

How Pooled Specialized Connections Are Assigned

When the Siebel Application Object Manager starts up, the specialized connection pool is empty. When a request is made to start a transaction, the Siebel Application Object Manager requests a database connection from the specialized connection pool. If one is available, then it is removed from the pool and given to the session for that session's exclusive use.

When the transaction completes (such as by being committed or canceled), the session returns the specialized connection to the pool. If the pool already contains more than the number of connections specified by MinTrxDbConns, then the specialized connection is closed. Otherwise, the connection is retained in the pool.

Scenario for Assigning Pooled Specialized Connections

Assume, for example, that MinTrxDbConns is set to 2. Specialized connections will be handled as follows:

- User 1 starts Transaction 1. The specialized connection pool is empty, so a new connection is created. Once Transaction 1 completes, this connection is returned to the pool.
- User 2 starts Transaction 2. If Transaction 1 is still running, then a new specialized connection is created. If Transaction 1 completed, then Transaction 2 uses the first database connection.
- When two specialized connections have been created, they will remain in the pool until the Siebel Application Object Manager terminates.

Using Thread Pooling for Siebel Application Object Manager

Optionally, you can configure your Siebel Application Object Manager components to use thread pooling. Enabling Siebel Application Object Manager thread pooling as described in this topic both pools and multiplexes (shares) multiple tasks across threads.

Using Siebel Application Object Manager thread pooling can improve performance and scalability on multiple-CPU computers that are under heavy load; for example, computers using eight or more CPUs that are running at more than 75% CPU capacity.

NOTE: Siebel Application Object Manager thread pooling is not recommended for smaller server computers or for computers that run under a relatively lower capacity.

About Thread Pooling for Siebel Application Object Manager

The pool size per multithreaded process for a Siebel Application Object Manager is determined by the combined settings of the parameters `UseThreadPool`, `ThreadAffinity`, `MinPoolThreads`, and `MaxPoolThreads`.

Siebel Application Object Manager thread pooling reduces some of system resource usage devoted to creating and closing session threads, as users log in and log out or are timed out. As when you are not using thread pooling, session threads are created as needed as session requests demand. However, instead of being closed when a session terminates, they are released to a pool, where they become available for use by a subsequent session.

NOTE: Using thread pooling introduces its own overhead, however, such as in task context-switching. For this reason, it is strongly recommended not to try to pool threads without also multiplexing them (that is, do not set `UseThreadPool = TRUE`, but `ThreadAffinity = TRUE`).

Because `ThreadAffinity = FALSE`, threads are multiplexed, as can be done with certain types of database connections or `SISNAPI` connections. At any given time, each thread can be dedicated to one or more user session (task).

Configuring Siebel Application Object Manager Thread Pooling

To use thread pooling, you set the following parameters on the Siebel Application Object Manager:

- `UseThreadPool = TRUE` (default is `FALSE`)
- `ThreadAffinity = FALSE` (default is `FALSE`)
- `MinPoolThreads = min_number_threads_in_pool` (default is 0)
where *min_number_threads_in_pool* represents the minimum number of threads in the Siebel Application Object Manager thread pool.
- `MaxPoolThreads = MinPoolThreads` (default is 0)

NOTE: You must specify a value for `MaxPoolThreads` that is equal to or greater than the value of `MinPoolThreads`.

To determine an appropriate value for MinPoolThreads and MaxPoolThreads, start slowly, monitor system performance, then introduce more multiplexing, as appropriate for your deployment. For example, start with a formula like this (based on two tasks per thread):

$$\text{MinPoolThreads} = \text{MaxPoolThreads} = (\text{MaxTasks divided by MaxMTServers}) \text{ divided by } 2$$

Subsequently, you can increase this to five tasks per thread, using this formula:

$$\text{MinPoolThreads} = \text{MaxPoolThreads} = (\text{MaxTasks divided by MaxMTServers}) \text{ divided by } 5$$

For example, if MaxTasks = 525, and MaxMTServers = 5, then this would be:

$$\text{MinPoolThreads} = \text{MaxPoolThreads} = (525 \text{ divided by } 5) \text{ divided by } 5 = 105 \text{ divided by } 5 = 21$$

Or, if MaxTasks = 725, and MaxMTServers = 5, then this would be:

$$\text{MinPoolThreads} = \text{MaxPoolThreads} = (725 \text{ divided by } 5) \text{ divided by } 5 = 145 \text{ divided by } 5 = 29$$

NOTE: Adjust values for think time, as necessary. If you cut your think time value in half, then double the number of threads in the pool.

4

Tuning the Siebel Server Infrastructure

This chapter describes the structure and operation of Siebel Application Object Manager components and the tuning that might be required for optimal operation. It contains the following topics:

- [Configuring SISNAPI Connection Pooling for Siebel Application Object Manager on page 47](#)
- [Tuning Server Request Broker \(SRBroker\) on page 49](#)

Related Books

Siebel Deployment Planning Guide

Siebel System Administration Guide

Siebel Installation Guide for the operating system you are using

Configuring SISNAPI Connection Pooling for Siebel Application Object Manager

This topic provides information about how to manage Siebel Internet Session Network Application Programming Interface (SISNAPI) connections for your Siebel Server.

SISNAPI is a messaging format that runs on top of the TCP/IP protocol, is used for network communication between Siebel Application Object Manager and Siebel Web Server Extension (SWSE), installed on the Web server. SISNAPI connections can be configured to use encryption and authentication based on Secure Sockets Layer (SSL).

Each multithreaded process for a Siebel Application Object Manager component uses a pool of SISNAPI connections managed by the SWSE. The process multiplexes (shares) many client sessions over each connection.

Each client session request opens a new connection and adds it to the pool, until the number of connections defined by the value of the parameter `SessPerSisnConn` have been created. Subsequent requests are then multiplexed over the existing pooled connections. SISNAPI connections persist until one of the following events occur:

- Web server process terminates.
- Siebel Application Object Manager component terminates.
- Value of the parameter `SISNAPI Connection Maximum Idle Time` (alias `ConnIdleTime`) is reached. For more information on this parameter, see *Siebel System Administration Guide*.
- Your firewall times out the connection.

Multiplexing traffic over a set of SISNAPI connections helps to reduce the number of open network connections.

The SISNAPI connection pool size per multithreaded process for a Siebel Application Object Manager is determined by the combined settings of MaxTasks, MaxMTServers, and SessPerSisnConn.

SessPerSisnConn determines how many sessions can be multiplexed over a single SISNAPI connection. By default, SessPerSisnConn is set to 20 for Siebel Application Object Manager components. This figure is suitable for most deployments and generally does not need to be changed. You might set this differently, depending on how you have calculated think time. For details, see [“Performance Factors for Siebel Application Object Manager Deployments” on page 24](#). For more information about configuring MaxTasks and MaxMTServers, see [“Tuning Siebel Application Object Manager Components for CPU and Memory Utilization” on page 27](#).

The number of actual SISNAPI connections per multithreaded process for the Siebel Application Object Manager is determined by the following formula:

$$(\text{MaxTasks divided by MaxMTServers}) \text{ divided by SessPerSisnConn} = \text{SISNAPI_conn_per_process}$$

where *SISNAPI_conn_per_process* represents the number of SISNAPI connections per multithreaded process.

Assume, for example, the following parameter values:

MaxTasks = 600

MaxMTServers = 5

SessPerSisnConn = 20

In this case, the formula would be:

$$(600 \text{ divided by } 5) \text{ divided by } 20 = 120 \text{ divided by } 20 = 6$$

Or, if MaxTasks = 540 and SessPerSisnConn = 18, then the formula would be:

$$(540 \text{ divided by } 5) \text{ divided by } 18 = 108 \text{ divided by } 18 = 6$$

In each case, six SISNAPI connections will be created and pooled for each Siebel Application Object Manager multithreaded process, from each SWSE. Each Web server and SWSE can potentially have its own set of six connections, so the maximum total number of connections into a Siebel Application Object Manager process is six times the number of Web servers. In the first example, 20 sessions would be multiplexed over each connection. In the second example, 18 would be multiplexed over each connection.

NOTE: In general, it is recommended that the figures used for the formula be tailored to produce an end result that is an integer. To achieve this, you might need to modify how you define MaxTasks, MaxMTServers, and SessPerSisnConn.

Some Object Manager components are not Siebel Application Object Manager components. Configuration issues for such components can differ from that applicable to Siebel Application Object Manager components. For information about the EAI Object Manager, see [Chapter 9, “Tuning Siebel EAI.”](#)

Tuning Server Request Broker (SRBroker)

The Server Request Broker (SRBroker) component routes requests between Siebel Server components, such as from a Siebel Application Object Manager to a batch component. SRBroker also handles requests between batch components. SRBroker is used whether the components run on the same computer or on different computers.

Server requests originating with a Siebel Application Object Manager component always go the SRBroker component to determine what to do with the request:

- If the destination component is running on the same Siebel Server, then SRBroker passes the request to this component. If multiple instances of the destination component are running, then SRBroker passes the request to each component instance in a round-robin fashion.
- If the destination component is not running on the same Siebel Server, then SRBroker passes the request to SRBroker running on another computer. If the destination component runs on multiple Siebel Servers, then SRBroker passes the request to each server in round-robin fashion.

The default parameter values for SRBroker work well for most deployments. If necessary, adjust the value of the MaxTasks parameter (the default value is 100). MaxTasks determines the maximum number of SRBroker threads (tasks) that can run on the Siebel Server. As necessary, set MaxTasks to a value equal to the number of batch components running on the Siebel Server, plus the number of Siebel Servers in the enterprise, plus 10 (for overhead).

MaxMTServers and MinMTServers determine the maximum and minimum number of SRBroker multithreaded processes that can run on the Siebel Server. Each multithreaded process can run a maximum of MaxTasks divided by MaxMTServers threads. Default values for MaxMTServers and MinMTServers must be set to 1. Increasing this value will not increase performance, and will not have any benefit.

CAUTION: Setting MaxTasks parameter values for SRBroker components in such a way that does not meet these guidelines might result in request failures. (HonorMaxTasks has no effect when set on the Server Request Broker (SRBroker) or Server Request Processor (SRProc) components.)

For more information about SRBroker and SRProc components, see *Siebel System Administration Guide*.

About HonorMaxTasks Parameter for Batch Components

By default, the HonorMaxTasks parameter for batch components, such as Workflow Process Manager, is set to FALSE (this setting is recommended). With this setting, if requests are routed by SRBroker to a batch component that has reached the maximum task capacity, then the requests will be queued in memory, and processed when tasks become available. Queuing such tasks minimizes the potential of request failure on the batch component due to the MaxTasks value having been reached.

You might consider setting HonorMaxTasks to TRUE for batch components in the following scenarios:

- For batch components handling asynchronous requests, consider changing HonorMaxTasks to TRUE if servers running these components have different resource levels and are therefore configured with different MaxTasks values for these components. In this case, larger servers would be forced to handle more requests. (However, if components are not running at maximum task capacity, then this effect can be hard to observe.)
- If batch components are suffering from crash or hang issues, then it might be undesirable to queue requests in component memory. If HonorMaxTasks is TRUE, then success or failure status of each request will be correctly reported. (This optional usage is a temporary measure only.)

For help resolving any component failure or hang issues, create a service request (SR) on My Oracle Support. Alternatively, you can phone Global Customer Support directly to create a service request or get a status update on your current SR. Support phone numbers are listed on My Oracle Support.

See also ["Tuning Workflow Process Manager for Performance" on page 91](#).

5

Tuning Siebel Web Client

This chapter describes configuration options that affect the performance and throughput of the Siebel Web Client, and provides guidelines for tuning the client to achieve and maintain optimal performance and scalability. It contains the following topics:

- [About Siebel Clients on page 51](#)
- [Performance Factors for Siebel Web Clients on page 52](#)
- [Guidelines for Siebel Web Client Tuning on page 53](#)

Related Books

Siebel Deployment Planning Guide

Siebel Installation Guide for the operating system you are using

Siebel System Administration Guide

Siebel Security Guide

Siebel System Requirements and Supported Platforms on Oracle Technology Network

Related Topics

For performance considerations for Siebel Application Object Manager, see [Chapter 3, "Tuning Siebel Application Object Manager."](#)

For performance considerations related to configuring Siebel Business Applications, see [Chapter 12, "Tuning Customer Configurations."](#)

For performance considerations related to Web servers and server operating systems, see [Chapter 13, "Tuning the Web Server Computer for All UNIX and Linux Platforms."](#)

About Siebel Clients

A Siebel client is a computer that operates Siebel Business Applications, accessing data and services by way of one or more servers. The Siebel clients allow users to access information managed by Siebel Business Applications. All Siebel deployments include one or more of the Siebel client types. You can deploy a mixture of clients.

The Siebel Business Applications client type covered in this book is the Siebel Web Client. This client runs in a standard third-party browser on the end user's client computer, and does not require any additional persistent software installed on the client.

Using HTTP, the browser connects to a Web server over a WAN, LAN, or VPN, or over the Internet. Through the Web server, the Siebel client connects to a Siebel Application Object Manager component on a Siebel Server, which executes Siebel application business logic and accesses data. Data is accessed from the Siebel Database and can also be accessed from other data sources using virtual business components and a variety of integration methods.

Only the user interface layer of the Siebel Business Applications architecture resides on the client computer.

For more information about the Siebel Web Client and other client types, and about supported browsers and browser settings, see the *Siebel Installation Guide* for the operating system you are using and *Siebel System Administration Guide*.

Performance Factors for Siebel Web Clients

Some performance considerations for Siebel Business Applications involve processing or tuning activities on servers only, and do not affect Siebel client performance. However, many other such factors either directly or indirectly affect Siebel client performance. This chapter highlights some of the factors most directly related to the performance of the Siebel Web Client.

The performance of the Siebel client depends on many factors, some of which are summarized below. For additional information on these topics, see applicable documents on the *Siebel Bookshelf*.

About Supporting Multiple Siebel Modules

Employee applications and customer applications have different requirements and characteristics and might use different browsers and other related technologies.

- Employee applications, such as Siebel Call Center, use high interactivity mode and run in supported Microsoft Internet Explorer browsers only.
- Customer applications, such as Siebel eService or Siebel eSales, use standard interactivity mode and can run in a wider range of browsers and browser versions.

All Siebel applications have many architectural elements in common. Multiple applications can use the same Siebel repository file (SRF). Each application uses its own Siebel Application Object Manager component. You might need to define, configure, and test multiple instances of each application to verify that your requirements are met in each usage scenario.

The performance of your Siebel applications will vary according to how you have configured the applications using Siebel Tools or custom browser scripts. See [“Following Configuration Guidelines” on page 56](#).

Client performance will also vary according to which Siebel modules you deploy. The performance of the Siebel client is highly dependent on feature functionality. Therefore, performance characteristics of Siebel modules will differ.

Some modules add special processing requirements. For example, Siebel CTI uses the Communications Session Manager (CommSessionMgr) component, and supports the communications toolbar and displaying screen pops in the client. Server and local resources support this functionality.

Supporting users who are dispersed in offices around the country or around the world introduces special deployment factors that can affect performance.

About Local Computer Resources

The resources available on each user's local computer must meet or exceed the recommendations outlined in *Siebel System Requirements and Supported Platforms* on Oracle Technology Network. Some performance enhancement measures depend directly on the available resources.

Guidelines for Siebel Web Client Tuning

Consider your hardware resources and requirements carefully prior to rolling out configuration changes, to make sure that business requirements have been met and the client performs as anticipated in the design phase.

Review guidelines presented elsewhere in this book, particularly in [Chapter 12, "Tuning Customer Configurations,"](#) and in other relevant documents on the *Siebel Bookshelf*.

Ongoing testing and monitoring of your system performance is strongly recommended as database characteristics change over time. To maintain an optimally performing system over time, you must plan for growth or other changes in your deployed application.

Activities you perform to achieve performance and scalability goals might include the following:

- Adjusting your system topology
- Configuring the Siebel application in Siebel Tools
- Configuring Siebel Server components, particularly the Siebel Application Object Manager
- Adjusting hardware resources available on local computers
- Adjusting operating system settings on server or client computers
- Adjusting Web server settings or network settings
- Adjusting Web browser settings
- Setting user preferences for Siebel applications

This topic contains the following information:

- ["Providing Sufficient Web Server and Network Capacity" on page 54](#)
- ["Testing Performance for Web Clients" on page 54](#)
- ["Providing Sufficient Client Hardware Resources" on page 55](#)
- ["Tuning System Components" on page 55](#)
- ["Following Configuration Guidelines" on page 56](#)
- ["Managing the Browser Cache" on page 56](#)
- ["Specifying Static File Caching" on page 57](#)
- ["Improving Performance Using View Layout Caching" on page 59](#)
- ["Managing Performance Related to the Message Bar" on page 64](#)
- ["Configuring the Busy Cursor for Standard Interactivity Applications" on page 64](#)

Providing Sufficient Web Server and Network Capacity

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

Make sure that your Web server is appropriately configured to meet your performance requirements. See also [“Specifying Static File Caching” on page 57](#).

Make sure that your network capacity (bandwidth) meets your performance requirements, and that your environment supports full duplex Ethernet. For details, review the general requirements for Siebel Enterprise Server installation and configuration, in the *Siebel Installation Guide* for the operating system you are using.

Several factors impact decisions regarding network bandwidth:

- **Application configuration.** Large, complex views will require bigger templates, more controls, and more data to be sent from the Web server to the client. If bandwidth is an issue, then it is important to consider user scenarios to determine the optimal size and layout per view.

For example, for views used frequently, reduce the number of fields displayed. For the high interactive client, the user can decide which columns are required in list applets. Rather than assuming a specific set, let users adjust it as necessary. Provide the minimal number of columns required in the base configuration. For more information, see [Chapter 12, “Tuning Customer Configurations.”](#)

- **View layout caching.** In high interactivity mode, administrators can determine the number of views to be cached locally. If the hardware supports a greater number of views to be cached, then adjust the value accordingly.

When a view is cached, subsequent visits will require a data update, but the Web templates need not be reloaded. This provides a substantial improvement in overall usability. For more information, see [“Improving Performance Using View Layout Caching” on page 59](#).

- **Login.** The first login is generally the most expensive operation for the high interactivity client. The client infrastructure caches the main components of the application on first login. Subsequent logins require far fewer resources. Cached objects remain on the client computer until the cache is cleared or a new version of the application configuration is available.

Testing Performance for Web Clients

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

Oracle’s Application Expert Services offers general guidance based on information known about the characteristics of the configured Siebel application. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle’s Application Expert Services.

Customer testing is advised in any case, because assumptions are based on general data. Actual experience can vary due to use-case scenarios. Select a few of the most common scenarios: those that represent the highest percentage of activity. Collect the overall bandwidth used.

Make sure you are testing with warm views (already visited and cached) if this is how the application will be used most of the time, assuming that users log in and use the application for 4 to 8 hours at a time before logging off and starting a new session.

When you estimate bandwidth required for several users sharing a low-bandwidth connection, consider use-cases carefully and plan accordingly. Rather than planning for worst-case network-performance scenarios (such as all users simultaneously pressing the Enter key or visiting a new view), it is likely that very few users are actually using the network at the same time. For more information about performance monitoring, see [Chapter 14, “Monitoring Siebel Application Performance with Siebel ARM.”](#)

Providing Sufficient Client Hardware Resources

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53.](#)

For best client performance for high interactivity applications, provide sufficient or generous hardware resources to your end users (typically, employees). Requirements can vary according to your deployment.

The more memory that is available on your client computers, the greater the number of views that can be cached. For more information, see:

- [“Managing the Browser Cache” on page 56](#)
- [“Specifying Static File Caching” on page 57](#)
- [“Improving Performance Using View Layout Caching” on page 59](#)

The speed of the processors (CPU) on your client computers will affect how quickly the Siebel application user interface is rendered.

For best performance for the high interactivity client, which is used by employee applications like Siebel Call Center, it is generally recommended to include the latest supported version of Microsoft Internet Explorer in your testing. More recent versions often include fixes and performance enhancements.

For best performance for the standard interactivity client, which is used by customer applications like Siebel eService, you must determine the minimum capabilities of customer environments, such as browser to support, processor speed, or expected Internet connection speed. Customer applications must support a wide range of customer environments. Accordingly, you generally minimize the complexity of such applications.

For Siebel client hardware and other platform requirements and recommendations, see *Siebel System Requirements and Supported Platforms* on Oracle Technology Network. For information about browser settings for Siebel Business Applications, see *Siebel System Administration Guide*.

Tuning System Components

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53.](#)

Overall end user performance is affected by the processing on the client as well as by everything from the Web server to the Siebel Database and back. Explore all applicable areas for opportunities to improve overall performance.

Most performance tuning involving Siebel Server components focuses on the Siebel Application Object Manager. For more information, see [Chapter 3, “Tuning Siebel Application Object Manager.”](#)

You can use Siebel ARM to monitor transactions through the Siebel infrastructure. Note areas which require substantial time and resources, and investigate them further for tuning opportunities.

For example, a custom configuration might have resulted in an unintentionally complex SQL statement for which the database instance has not been optimized. Small configuration adjustments in Siebel Tools, or database tuning, can improve both client performance and application scalability on Siebel Servers. For more information about Siebel ARM, see [Chapter 14, “Monitoring Siebel Application Performance with Siebel ARM.”](#)

Following Configuration Guidelines

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53.](#)

For best performance by the Siebel client, carefully assess all customer configuration initiatives. All configuration changes must be justifiable in terms of the cost of configuration itself, and in terms of possible impact on performance.

Some application administration tasks might also affect application performance, and must also be carefully assessed.

Follow guidelines presented in [Chapter 12, “Tuning Customer Configurations,”](#) or in other books on the *Siebel Bookshelf*.

Managing the Browser Cache

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53.](#)

Some types of Siebel application elements are stored in the browser cache, to improve performance when users log in or access Siebel views.

NOTE: When measuring performance, take into account view layout caching or other types of caching. For example, performance is better when a Siebel view layout is retrieved from a cache than it is when the view layout is not cached and must be retrieved from the system. For more information, see [“Improving Performance Using View Layout Caching” on page 59.](#)

Cache usage varies according to what browser is being used, what applications are running, and application settings. For example, high interactivity applications use the browser cache more than standard interactivity applications.

For high interactivity applications, it is generally recommended that users do not clear their browser cache, including when the browser is closed. The following settings for Microsoft Internet Explorer are recommended:

- Choose Tools, then Internet Options. Click the Advanced tab. In the Security options, uncheck the setting Empty Temporary Internet Files Folder when browser is closed.

NOTE: If you do not use this setting, then persistent view layout caching and preloading, described in [“Improving Performance Using View Layout Caching” on page 59,](#) will not work.

- Choose Tools, then Internet Options. Click the Advanced tab. In the Security options, uncheck the setting Do not save encrypted pages to disk.

NOTE: If you do not use this setting, then persistent view layout caching and preloading, described in [“Improving Performance Using View Layout Caching” on page 59](#), will not work for encrypted views (views encrypted using SSL).

- Choose Tools, then Internet Options. In the General tab, click Settings. For the option Check for newer versions of stored pages, use the setting Automatically.
- Browser caching is also subject to the size of the temporary Internet files folder. This setting is located in Tools, then Internet Options. In the General tab, click Settings, then specify the amount of disk space to use for this folder.

NOTE: Setting the size of the temporary Internet files folder to 0 disables persistent view layout caching and preloading, which are described in [“Improving Performance Using View Layout Caching” on page 59](#).

For more information about browser settings for Siebel Business Applications, see *Siebel System Administration Guide*.

Caching in the browser is also subject to Web server settings controlling static file caching. For details, see [“Specifying Static File Caching” on page 57](#).

Specifying Static File Caching

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

Browser caching behavior is also subject to Web server settings for static file caching. Appropriate settings allow files that are rarely updated, such as image files, JavaScript files, or style sheet files, to be cached on the browser. Caching static files reduces network utilization and enhances Siebel Web Client response time.

Caching for Siebel Web template files is described in the persistent view caching content in [“Improving Performance Using View Layout Caching” on page 59](#).

Because some static files can in fact be updated periodically, there is some risk that outdated versions of static files might be served from the cache. Therefore, some appropriate content expiration time must be specified. In general, setting an expiration time of 7 days might be appropriate.

If static files are rarely updated, then you can specify a larger number, for less frequent expiration. If static files are updated more often, then you can specify a smaller number, for more frequent expiration.

Instructions follow for specifying static file caching on Microsoft Internet Information Services (IIS), Apache Web Servers, and Oracle iPlanet Web Server. You must restart your Web server for the settings to take effect. For details, see your third-party Web server vendor documentation.

For more information about supported Web servers and versions, see *Siebel System Requirements and Supported Platforms* on Oracle Technology Network. For more information about tuning Web servers, see [Chapter 13, “Tuning the Web Server Computer for All UNIX and Linux Platforms.”](#)

Static File Caching for Microsoft IIS

For Microsoft IIS, follow the procedure below to specify static file caching and content expiration.

To specify static file caching on Microsoft IIS

- 1 On the Web server computer, choose Start, then Settings, then Control Panel, and then Administrative Tools.
- 2 Run Internet Service Manager.
- 3 In Internet Service Manager, right-click Default Web Site.
- 4 In Default Web Site Properties, click the HTTP Headers tab.
- 5 Check the Enable Content Expiration check box.
- 6 Select Expire After, and specify the value of 7 (to expire static files after 7 days), or another value appropriate for your deployment.

Static File Caching for an Apache Web Server

For supported Apache Web Servers (Oracle HTTP Server, IBM HTTP Server, and HP Apache Web Server), follow the procedure below to specify static file caching and content expiration.

To specify static file caching on an Apache Web Server

- 1 On the Web server computer, open the file httpd.conf for editing. This file is located in the Web server installation directory.
- 2 Verify that the following line is included and not commented out:

```
LoadModule expires_module modules/mod_expires.so
```

- 3 Add the following lines, if not already present, to the file (below the line shown in [Step 2 on page 58](#)). Or, instead of 7 days, specify another value appropriate for your deployment.

```
#####
ExpiresActive On

<IfModule mod_expires.c>

ExpiresByType image/gif                "access plus 7 days"
ExpiresByType image/jpeg               "access plus 7 days"
ExpiresByType application/x-javascript "access plus 7 days"
ExpiresByType text/css                  "access plus 7 days"

</IfModule>

#####
```

- 4 Save the file.

Static File Caching for Oracle iPlanet Web Server

For Oracle iPlanet Web Server, follow the appropriate procedure to specify static file caching and content expiration. For example, for Oracle iPlanet Web Server, follow the steps below.

To specify static file caching on Oracle iPlanet Web Server

- 1 From a browser, connect to the Web server administration page (for example, `http://web_server_name/8080`).
- 2 Select the server, and click Manage.
- 3 Click the link for Class Manager, in the upper right area.
- 4 Among the horizontal tabs at the top, click Content Mgmt.
- 5 Click the link for Cache Control Directives, in the left tab area.
- 6 Under Cache Control Response Directives, select Maximum Age (sec), and input 604800 (seconds) for a valid cache of 7 days.
- 7 Click Apply to apply the change.

Improving Performance Using View Layout Caching

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

View layout caching in the browser (also referred to as *layout caching* or *view caching*) improves the performance of accessing views in a high interactivity application. It speeds up the rendering of views in a Siebel application session by caching the following on the browser:

- Static HTML (from the templates) used for interpreting the view.
- Dynamic HTML generated on the client for rendering controls.

Appropriate caching settings can optimize client performance and network utilization for Siebel client sessions. Caching behavior is subject to considerations described under [“Managing the Browser Cache” on page 56](#).

Two kinds of view layout caching are used for the Siebel Web Client. These types of caching work together and must be configured together.

- **Caching in browser memory.** For details, see [“View Layout Caching in Memory” on page 60](#).
- **Persistent caching (in the browser cache directory on the local disk).** For details, see [“Persistent View Layout Caching” on page 61](#).

NOTE: Whether views can be cached depends on the underlying requirements described in [“Determining If Views Are Available for Layout Caching” on page 63](#).

View Layout Caching in Memory

View layout caching in memory creates multiple HTML frames on a browser to store the layout for a view. The number of these frames represents the view cache size. When a view is displayed, the HTML frame containing the layout for that view will be sized to occupy all (100%) of the available browser space, while the other frames will be hidden (that is, sized to occupy 0% of the space). For information on setting the view cache size, see [“Setting the View Cache Size” on page 61](#).

View layout caching uses the following logic:

- If a user navigates to a view whose layout is already available in the browser memory cache, then the HTML frame containing that view will be made visible and the currently visible frame will be hidden.
- If a user navigates to a view whose layout is not in the browser memory cache, then one of the available HTML frames will be used to load the layout of the view into memory. The view layout will be loaded from the persistent cache, if possible. The view layout in memory will be cached subject to the View Cache Size setting.
- If the view layout is not currently stored in the persistent cache, then it is loaded from the server. It will also be stored in the persistent cache. The view layout in memory will be cached subject to the View Cache Size setting.

For more information, see [“Persistent View Layout Caching” on page 61](#).

NOTE: The high interactivity framework separates the retrieval of the Siebel application user interface from the server and the retrieval of database records. Database records to be displayed in views are always retrieved from the server.

The memory cache contains the layouts of views that the user has visited and that are available for view caching. When the view cache is full and another view is visited, the first view visited is removed from the cache. The memory cache contents are thus managed on an LRU or least recently used basis.

The HTML frames are loaded into memory when a user navigates to the view. To cache a startup view (one that is cacheable), the user must visit the view twice (that is, visit the view a second time, after visiting another view).

NOTE: Views specifically created as home page views, such as Home Page View (WCC) for Siebel Call Center, are standard interactivity views and are not cacheable.

The view caching framework is designed so that if the frames containing cached views are deleted (for example, by performing a browser refresh, which removes any previously cached views), the framework begins reloading the layout cache, starting with the next cacheable view the user visits.

At startup, view layouts for recently visited views can be preloaded into the memory cache from the persistent browser cache on disk. This behavior is specified using the parameter ViewPreloadSize. For more information, see [“Persistent View Layout Caching” on page 61](#) and [“Preloading Cached Views into Memory” on page 62](#).

Setting the View Cache Size

For browser memory caching, the size of the view layout cache is controlled by the View Cache Size user preference setting for each user, as described below.

NOTE: Setting View Cache Size to 1 turns off view caching. This has the same effect as setting the EnableViewCache parameter to FALSE, as described in [“Disabling View Layout Caching”](#) on page 62.

To set the size of the view layout cache

- 1 From the application-level menu, choose View, then User Preferences.
- 2 From the Show drop-down list, choose Behavior.
- 3 In the View Cache Size field, select a value from the drop-down list, or type in a value.

The default value for View Cache Size is 10. This value specifies that 10 HTML frames are cached in memory to represent Siebel view layouts. One of these frames is displayed at any one time.

- Using a figure that is too low might not provide enough caching if your users access many views and client computers have sufficient memory.
- Using a figure that is too high might impair performance by using more memory than is available on the computer.

Persistent View Layout Caching

Persistent layout caching stores the layout of certain views in a local client's browser cache on disk. The stored layout is then reused for subsequent visits to this view, in the same session or in a subsequent session. (For subsequent visits in the same session, the view layout is accessed from the browser memory cache, if available.)

Persistent view layout caching helps improve performance by reducing the number of pages that have to be generated from the server from session to session.

The parameter WebTemplateVersion determines whether the Siebel Web Engine will use a view layout stored in the browser's cache or build a new view layout.

- For Siebel Web Clients, set this parameter for the Siebel Application Object Manager component using Server Manager.
- For Siebel Mobile Web Clients, this parameter is located in the [InfraUIFramework] section of the application configuration file, such as uagent.cfg for Siebel Call Center.

When you modify Web templates for Siebel views, set the WebTemplateVersion parameter to the value to 1. Subsequently, each time you change any of the Web templates, increment the value of the parameter by 1. Doing so forces loading view layouts from the Web templates on the server.

When a view is requested, the Siebel Web Engine includes in the URL a checksum value that encapsulates the value of the WebTemplateVersion parameter.

- If the parameter value and the value encapsulated in the URL match, then it is assumed that the view layout for this view has not been updated. If it is available, then the view layout stored in the persistent cache can be used.

- If no match is found, then a new view layout is loaded from the server. The Web template on the server is presumably more current than the view layout stored in the browser's persistent cache.

Preloading Cached Views into Memory

For recently visited views, view layouts that are cached in the persistent cache on the browser can be preloaded into browser memory when the user logs in. The number of views that can be preloaded depends on the content of the persistent cache and is limited by the setting of the View Cache Size setting for each user.

For better performance at login time, it can be helpful to further limit the number of view layouts that are preloaded into memory during startup. To do this, use the parameter ViewPreloadSize.

- For Siebel Web Clients, set this parameter for the Siebel Application Object Manager component using Server Manager.
- For Siebel Mobile Web Clients, this parameter is located in the [InfraUIFramework] section of the application configuration file, such as uagent.cfg for Siebel Call Center.

NOTE: ViewPreloadSize only affects a user session when it is set to a positive integer value lower than the View Cache Size value. If the parameter is not set, then the default behavior is to preload the number of view layouts corresponding to the View Cache Size value, minus one. (One of the frames specified using View Cache Size is reserved for the application startup view, where applicable.)

If ViewPreloadSize is set to 0, then no view layouts are preloaded into memory. In scenarios where users frequently log into the Siebel application, such as to access a single view, then log out again, login performance might be more important than precaching multiple views. In this case, you can choose to set this parameter to 0.

Disabling View Layout Caching

You can disable browser memory caching of view layouts for your application users by changing the parameter EnableViewCache to FALSE.

- For Siebel Web Clients, set this parameter for the Siebel Application Object Manager component using Server Manager.
- For Siebel Mobile Web Clients, this parameter is located in the [InfraUIFramework] section of the application configuration file, such as uagent.cfg for Siebel Call Center.

NOTE: In general, setting EnableViewCache to TRUE is recommended. If some users do not need view layout caching, then they can set View Cache Size to 1, as described in [“Setting the View Cache Size” on page 61](#).

Setting EnableViewCache to FALSE disables browser memory view layout caching only. It does not disable persistent view layout caching.

Determining How the Current View Layout Was Loaded

If you are running an application and want to determine how the current view was retrieved, then go to the view, press SHIFT, and choose Help, then About View. The Cache Mode identified for the current view indicates how the application retrieved the view layout. Possible values include:

- **Not Cached.** The view layout was not cached (and cannot be cached).
- **Memory.** The view layout was retrieved from the browser memory cache.
- **Server.** The view layout was retrieved from the Siebel Server and the Web server. If the view is cacheable, and you visit another view and then return to this one, then the Cache Mode value changes to Memory.
- **Disk.** The view layout was retrieved from the browser disk cache (persistent caching). If the view is cacheable, and you visit another view and then return to this one, then the Cache Mode value changes to Memory.

The longer you go without clearing the cache, the more likely that a rarely visited view will be retrieved from the persistent cache on the browser, rather than from the server.

Determining If Views Are Available for Layout Caching

Not all Siebel views are available for layout caching. Views that contain applets that have dynamic layouts or controls that are data-dependent cannot be cached. Only applets that support high interactivity are available for view layout caching.

Layout caching is a feature of the C++ class that implements an applet. The ability to be cached is determined by a property of each applet's class object definition. Using Siebel Tools, check the value of the High Interactivity Enabled property of a class object definition to determine whether applets for this class support layout caching. For a view to be available for caching, the class objects for *all* of the applets in the view must have High Interactivity Enabled values of 2 or 4 (available for caching).

For detailed information about settings for the High Interactivity Enabled property for a class, see *Siebel Object Types Reference*.

View layout caching is also disabled for a view in the following cases:

- If personalization rules are defined for any of the applets
- If any of the applets are dynamic toggle applets
- If any of the applets are hierarchical list applets or explorer (tree) applets
- If HTML frames are used within the view template (for example, for explorer views)

Configuring the Data Block Size of HTTP Requests for the Siebel Dedicated Client

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

Siebel Business Applications use HTTP requests and responses to exchange information between your browser and the Siebel Web Server. You can change the maximum length of the HTTP request data that is passed. By default, the maximum limit for the HTTP request is 524288 bytes. However, you can change this limit by configuring the `DataBlockSize` parameter in the [siebel] section of the `siebel.cfg` file.

NOTE: The `DataBlockSize` parameter by default is set to 512 kilobytes, which is typically sufficient. However, if you choose to change this parameter, then the `siebel.exe` memory usage will reflect whatever value you set.

Managing Performance Related to the Message Bar

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

Employee applications such as Siebel Call Center include a message bar feature. The message bar requires network resources and local resources on the client computer to continually update the displayed text.

- If your deployment does not require it, then turn off the message bar feature to save processing resources.
- If some of your users require the message bar, then you can specify that users will be able to turn it off by choosing Tools, then User Preferences, and then Message Broadcasting.

For more information about message broadcasting using the message bar, see *Siebel Applications Administration Guide*.

Configuring the Busy Cursor for Standard Interactivity Applications

This topic is part of [“Guidelines for Siebel Web Client Tuning” on page 53](#).

When the parameter `EnableSIBusyCursor` is set to TRUE (default), the Siebel application cannot accept new requests until it has finished serving the current request. An example of a new request is a user's attempt to drill down on a record.

The default setting of this parameter helps prevent the occurrence of JavaScript errors as the user cannot click on another link while a page is loading. Setting `EnableSIBusyCursor` to FALSE can improve network bandwidth usage for Siebel applications that use standard interactivity mode. It disables the appearance of the hourglass icon and allows the user to click on other links before the current request has been served.

To set a value for the EnableSIBusyCursor parameter:

- For Siebel Web Clients, set this parameter for the Siebel Application Object Manager component using Server Manager.
- For Siebel Mobile Web Clients, this parameter is located in the [InfraUIFramework] section of the application configuration file, such as uagent.cfg for Siebel Call Center.

6

Tuning Siebel Communications Server

This chapter describes some issues that affect the performance and throughput of selected functionality for Siebel Communications Server and related modules, and provides guidelines for tuning these modules to achieve and maintain optimal performance and scalability. Functionality described in this chapter includes session communications (typically, Siebel CTI) and Siebel Email Response. Other communications-related modules are not described.

This chapter contains the following topics:

- [About Siebel Communications Server on page 67](#)
- [Session Communications Infrastructure on page 68](#)
- [Performance Factors for Session Communications on page 69](#)
- [Topology Considerations for Session Communications on page 71](#)
- [Guidelines for Session Communications Tuning on page 72](#)
- [Siebel Email Response Infrastructure on page 77](#)
- [Performance Factors for Siebel Email Response on page 78](#)
- [Topology Considerations for Siebel Email Response on page 79](#)
- [Guidelines for Siebel Email Response Tuning on page 79](#)

Related Books

Siebel CTI Administration Guide

Siebel Email Administration Guide

Siebel System Administration Guide

Siebel Deployment Planning Guide

About Siebel Communications Server

Siebel Communications Server provides an infrastructure to support several kinds of communications activities for Siebel Business Applications users. For session communications performance tuning information, see [“Session Communications Infrastructure” on page 68](#) and subsequent topics. For Siebel Email Response performance tuning information, see [“Siebel Email Response Infrastructure” on page 77](#) and subsequent topics.

- **Session communications.** Supports interactive (session) communications for contact center agents who use the multichannel communications toolbar to:
 - Make or receive voice calls using computer telephony integration supported by CTI middleware or other third-party products.
 - Receive inbound email messages (for Siebel Email Response).

- **Inbound communications.** Supports integrating with third-party email servers and processing inbound email (when using Siebel Email Response).
 - **Outbound communications.** Supports integrating to a variety of third-party communications systems, such as email servers or wireless messaging providers, to send outbound communications.
 - Supports agents sending email replies using Siebel Email Response.
 - Supports the Send Email and Send Fax commands for Siebel Business Applications users. (Send Page is also available, but uses the Page Manager server component.)
 - Supports users sending outbound communications content (email, fax, or page) using communication requests. Requests can be created and submitted either programmatically or manually through a user interface described in *Siebel CTI Administration Guide*.
- Many Siebel modules invoke business service methods through workflows to send outbound communications.

Session Communications Infrastructure

Session communications refers to using Communications Server components to enable contact center agents or other users to handle interactive communications work items. For example, Siebel CTI supports this capability, enabling agents to handle voice calls using the communications toolbar.

It is important to understand the infrastructure that supports session communications in order to prevent or address performance issues in this area.

Session communications performance is addressed in this topic and in:

- [“Performance Factors for Session Communications” on page 69](#)
- [“Topology Considerations for Session Communications” on page 71](#)
- [“Guidelines for Session Communications Tuning” on page 72](#)

Key Siebel Server Components

Session communications are supported in the Siebel Server environment primarily by the following components:

- **Communications Session Manager (CommSessionMgr).** This server component manages interactive communications work items such as voice calls.
- **Siebel Application Object Manager.** This server component manages application sessions for end users who use the Siebel Web Client, including users who handle communications work items (agents). Interactive communication requests from agents typically go through Siebel Application Object Manager. For more information about Siebel Application Object Manager, see [Chapter 3, “Tuning Siebel Application Object Manager.”](#)

- **Server Request Broker (SRBroker)**. This server component handles communications between the Siebel Application Object Manager and certain other Siebel Server components, including CommSessionMgr.

For example, when a Siebel CTI agent makes a call through the communications toolbar, the request goes from Siebel Application Object Manager to CommSessionMgr by way of SRBroker.

SRBroker is used whether CommSessionMgr runs on the same computer as the Siebel Application Object Manager, or on a different computer. For more information about such scenarios, see [“Topology Considerations for Session Communications” on page 71](#). For more information about SRBroker, see [“Tuning Server Request Broker \(SRBroker\)” on page 49](#).

Other Siebel Server Components

You might also be using the following components in your Siebel Server environment and communications infrastructure:

- **Communications Configuration Manager (CommConfigMgr)**. Optionally used to cache communications configuration data.
- **Communications Inbound Receiver (CommInboundRcvr)**. For details, see [“Siebel Email Response Infrastructure” on page 77](#).
- **Communications Inbound Processor (CommInboundProcessor)**. For details, see [“Siebel Email Response Infrastructure” on page 77](#).
- **Communications Outbound Manager (CommOutboundMgr)**. Sends outbound email or other types of messages.

Third-Party Product Modules

You might be using third-party product modules such as CTI middleware, driver, and configuration; routing products; predictive dialers; interactive voice response modules; email servers; fax servers; and so on. For information about supported email servers, see *Siebel System Requirements and Supported Platforms* on Oracle Technology Network.

NOTE: To use Siebel CTI, you must obtain a CTI middleware package and work with your vendor to integrate this module. Oracle Contact Center Anywhere can be used for this purpose. For more information, see the Virtual CTI content in *Siebel CTI Administration Guide*.

Performance Factors for Session Communications

This topic describes factors that drive or affect performance for session communications deployments.

Depending on your deployment, your agents can be handling phone calls (Siebel CTI), email messages (Siebel Email Response), work items of other communications channels, or a combination of these.

- **Inbound calls processed per hour.** The number of inbound calls (or other types of work items) processed per hour (or some other time period) by your communications infrastructure.
- **Outbound calls processed per hour.** The number of outbound calls processed per hour (or some other time period) by your communications infrastructure. (For outbound predictive dialer calls, only the calls that are answered and processed by Communications Server are relevant here.)
- **Number of user communications actions per minute (load).** The average number of communications-related user actions per minute, and the average think time between such user actions. Communications-related actions typically refers to actions performed using the communications toolbar.

Longer think times mean less load on the Siebel Database and Siebel Server. Think time is an important factor in overall system load. Estimation of think time must approximate actual user usage. For more information about think time and Siebel Application Object Manager tuning, see [Chapter 3, “Tuning Siebel Application Object Manager.”](#)

- **Number of concurrent communications users (agents).** The number of concurrent users of session communications features (typically, contact center agents). This figure will be some percentage of the total number of concurrent users on the Siebel Application Object Manager.

You also need to understand how agents work with these features, the average number of inbound and outbound work items per agent, and how these factors relate to your organization’s service goals. Some agents receive a large number of work items from ACD queues or initiate a large number of work items. Supervisors or other users can be defined as agents but might receive only escalated work items, for example. For more information about concurrent users and Siebel Application Object Manager tuning, see [Chapter 3, “Tuning Siebel Application Object Manager.”](#)

- **Volume of customer data.** The total volume of customer data.

Data volume affects how quickly data can be retrieved for various purposes, such as to perform lookups for screen pops, route work items, or populate the customer dashboard. In many cases, data volume directly affects response times seen by agents. The volume of data must be realistic and the database needs to be tuned to reflect real-world conditions.

These and many other factors (such as the average call time, average time between calls for an agent, and so on) will affect system performance as experienced by contact center agents. An agent will be concerned with general response time, screen pop response time, and other perceived measures of performance.

Third-Party Product Considerations

Review information presented in applicable third-party documentation for any requirements that affect your deployment. For example:

- Some CTI middleware software might place limitations on the number of agents that can be served at a single contact center site.
- Integration with ACD queues, predictive dialers, or other modules can affect your configurations, affect network traffic, or have other impacts.

- The capacity of your telephony link (between the ACD switch and the CTI middleware) can affect performance.

Topology Considerations for Session Communications

Generally, Siebel Communications Server components for session communications, such as CommSessionMgr, must be run on the same Siebel Server computers as those running Siebel Application Object Managers. In some cases, however, you must run CommSessionMgr on a different computer than the Siebel Application Object Managers. These options are described in detail below.

CTI middleware generally runs on servers located at each contact center facility.

Running CommSessionMgr on Siebel Application Object Manager Computers

Generally, Siebel Communications Server components for session communications must be run on the same Siebel Server computers as those running Siebel Application Object Managers. Such a topology allows the Siebel Application Object Manager load-balancing mechanism to indirectly balance Communications Server load. CommSessionMgr loads are fairly light and do not, in themselves, present a reason to run this component on dedicated computers.

Set the Enable Communication parameter to TRUE for all Siebel Application Object Managers to which your agents will connect. If you are using Siebel Server load balancing, then all Siebel Application Object Managers to which requests are distributed must be configured the same way.

Running CommSessionMgr on Dedicated Computers

Sometimes you *must* run CommSessionMgr on a different computer than the Siebel Application Object Manager components.

CommSessionMgr must run on the same computer where the communications driver for your CTI middleware is running. If your driver requires a particular operating system platform, then you must install Siebel Server and run CommSessionMgr on a computer of this platform. (Communications drivers are required to be able to run on one of the supported Siebel Server platforms, as described in *Siebel System Requirements and Supported Platforms* on Oracle Technology Network.)

If your Siebel Application Object Manager components (Call Center Object Manager) run on computers using a different platform, then you set several parameters in the communications configuration, including CommSessionMgr and RequestServer, in order to designate the computer where CommSessionMgr is running. All communications session requests from a Siebel Application Object Manager supporting users for this communications configuration will be routed to the CommSessionMgr component on the dedicated computer. For related information, see [“Tuning the CommSessionMgr Component” on page 73](#). For more information about these parameters, see *Siebel CTI Administration Guide*.

Guidelines for Session Communications Tuning

Using your hardware resources optimally, and configuring your system appropriately, can help you to achieve your performance goals. Consider your resources and requirements carefully, and test and monitor system performance on a continual basis.

Review information presented in *Siebel CTI Administration Guide*, *Siebel System Administration Guide*, relevant third-party documentation, and other sources.

Activities you can perform to achieve performance and scalability goals include, but are not limited to, the following:

- Adjusting your system topology. For more information, see [“Topology Considerations for Session Communications” on page 71](#).
- Configuring the Siebel Application Object Manager component.
- Configuring CommSessionMgr and related components.
- Modifying communications configurations, communications driver settings, and so on. Many of the activities described in the topics that follow are of this nature.

To maintain an optimally performing system over time, you must plan for changes in the volume of incoming communications, number of users, and so on. Verify that your CTI middleware can support an anticipated increase in the volume of incoming communications and in the number of users. Then additional hardware might be required to run more Siebel Application Object Manager components and CommSessionMgr components to support the increase in volume of communications and in number of users.

This topic contains the following information:

- [“Tuning the Siebel Application Object Manager Component” on page 73](#)
- [“Tuning the CommSessionMgr Component” on page 73](#)
- [“Conserving Siebel Application Object Manager Server Resources Through Caching” on page 73](#)
- [“Improving Performance for Communications Configurations” on page 74](#)
- [“Configuring Logging for Session Communications” on page 75](#)
- [“Improving Availability for Session Connections” on page 76](#)
- [“Improving Screen Pop Performance” on page 77](#)
- [“Reviewing Performance Impact of Activity Creation” on page 77](#)

Tuning the Siebel Application Object Manager Component

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

CommSessionMgr and CommConfigMgr components use a small percentage of the resources of the Siebel Server on which it runs. Siebel Application Object Manager performance has the greatest effect on overall system performance, even when CommSessionMgr or CommConfigMgr components are present.

Siebel Application Object Manager memory requirements for agent sessions depend on many factors. Siebel Application Object Manager memory usage for an agent using session communications is greater than for other users (those who are not defined as agents in a communications configuration).

Siebel Application Object Manager tuning also depends on your communications configuration caching methods. See also [“Conserving Siebel Application Object Manager Server Resources Through Caching” on page 73](#). For more information about Siebel Application Object Manager tuning, see Chapter 3, [“Tuning Siebel Application Object Manager.”](#)

Tuning the CommSessionMgr Component

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

For the CommSessionMgr component, the MaxTasks parameter determines the maximum number of communications events that can be processed at one time.

Generally, the default values are appropriate for the MaxTasks, MinMTServers, and MaxMTServers parameters, particularly if CommSessionMgr runs on each Siebel Application Object Manager computer.

If you use a dedicated Siebel Server computer to run the CommSessionMgr component, then it might be appropriate to set these parameters to higher values to optimize usage of server resources such as CPU and memory. See also [“Topology Considerations for Session Communications” on page 71](#).

Conserving Siebel Application Object Manager Server Resources Through Caching

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

You can use two caching mechanisms to make communications configuration data load faster for each agent session and to reduce demand on server resources on the Siebel Application Object Manager.

These caching mechanisms can be used together or separately. For more information, see *Siebel CTI Administration Guide*.

- **CommConfigCache parameter.** Setting the CommConfigCache parameter on the Siebel Application Object Manager to TRUE caches communications configuration data when the first agent logs in. Configuration data is cached until the Siebel Application Object Manager is restarted. For agents associated with the same communications configuration, each agent session uses the same cached data. See also [“Tuning the Siebel Application Object Manager Component” on page 73](#).
 - Performance is improved for subsequent agent logins, because the configuration data is loaded from the cache rather than from the database.
 - Siebel Application Object Manager scalability is also improved because configuration data is shared in Siebel Application Object Manager memory across agent sessions, therefore conserving server resources even as the number of agent sessions increases.
- **CommConfigMgr server component and CommConfigManager parameter.** The CommConfigMgr server component caches communications configuration data when the first agent logs in. Setting the CommConfigManager parameter on the Siebel Application Object Manager to TRUE enables this server component.
 - Performance is improved for subsequent agent logins, because the configuration data is loaded from the cache rather than from the database.
 - Using the CommConfigMgr component to cache data speeds up the login process and reduces memory usage per agent session because the component uses configuration data that was already cached on the Siebel Application Object Manager component.
 - Although it is not required to use the CommConfigMgr component in conjunction with the CommConfigCache parameter for Siebel Application Object Manager, if you use them together, communications configuration data gets cached at the enterprise level rather than only for the Siebel Application Object Manager. Overall performance can be enhanced compared to using each of these mechanisms separately.

Improving Performance for Communications Configurations

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

When you deploy session communications, you create communications configurations, define employees as agents, and associate each agent with one or more configurations. How you do these things affects performance and scalability.

In a deployment supporting a large number of agents across multiple physical sites, you must determine criteria for grouping your agents within configurations.

For example, some dialing filters you define, using the parameter DialingFilter.Rule*N*, might be appropriate for agents at a specific place, such as within the same country or area code. Other dialing filters might be suitable for a different set of agents.

In addition, some switch, teleset, or CTI middleware settings are reflected in your communications configuration, and can differ between physical locations.

It might be helpful to define a communications configuration to apply to users at a single location only. In addition to simplifying the process of defining communications configurations, telesets, or other elements, this approach can help you reduce demand on server resources such as Siebel Application Object Manager memory or CPU.

If call transfers or similar functions are to be supported between contact centers, then additional configuration issues apply. For more information about defining communications configurations and agents, see *Siebel CTI Administration Guide*.

Configuring Logging for Session Communications

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

Logging data can be analyzed as part of performance monitoring or tuning, as described in [Chapter 14, “Monitoring Siebel Application Performance with Siebel ARM.”](#)

Higher levels of logging provide more data to help you resolve system errors or performance issues; this is appropriate for system testing. For production systems, however, logging levels must be reduced to improve performance.

Log-related parameters applicable to session communications are summarized below. The Siebel Application Object Manager component logs activity related to the user’s client session, including usage of the communications toolbar, screen pops, and so on. The CommSessionMgr logs activity related to this component, such as commands and events for the communications driver.

The logging for Siebel Application Object Manager and CommSessionMgr are written to separate files for each user. Typically (though not necessarily), these logging mechanisms both write into the same set of files. This makes it easier to monitor or troubleshoot issues related to session communications for a particular user session. For details on these logging parameters, see *Siebel CTI Administration Guide*.

Siebel Application Object Manager Logging Parameters

Siebel Application Object Manager parameters that log session communications activity include:

- **CommLogFile.** Specifies the name of the log file (default value is SComm.log). A separate log file is created for each agent session, in the form SComm_username.log.
- **CommLogDebug.** Specifies whether log files must contain extra detail. Setting to FALSE provides better performance.
- **CommMaxLogKB.** Specifies the maximum size of log files.
- **CommReleaseLogHandle.** Specifies that the log file handle must be released periodically. The default setting of TRUE provides better performance.

CommSessionMgr Logging Parameters

CommSessionMgr parameters that log session communications activity include:

- **LogFile.** Specifies the name of the log file (default value is SComm.log). A separate log file is created for each agent session, in the form SComm_username.log.

- **LogDebug.** Specifies whether log files must contain extra detail. Setting to FALSE provides better performance.
- **MaxLogKB.** Specifies the maximum size of log files.
- **ReleaseLogHandle.** Specifies that the log file handle must be released periodically. The default setting of TRUE provides better performance.

Improving Availability for Session Connections

This topic is part of [“Guidelines for Session Communications Tuning” on page 72.](#)

When agents log into the Siebel application after experiencing browser failure or a dropped connection, session communications can sometimes remain unavailable.

Session communications availability can be considered a performance issue. In addition to affecting agent productivity, loss of availability of session communications wastes server resources that could support other functions.

You can improve session communications availability using the following mechanisms:

- **Push Keep Alive driver.** Using the Push Keep Alive communications driver pushes empty messages (heartbeat messages) to agents at regular intervals. In this manner, it helps keep the communications push channel alive. This feature can help in environments where enforced timeouts sometimes cause communications session connections to be dropped.

For example, many customers deploy some kind of network appliance to load-balance Web servers. By default, such network appliances can time out connections to browsers, causing communication interruptions for agents. The Push Keep Alive driver generates periodic traffic so connections do not time out due to inactivity.

To use the Push Keep Alive driver, you create a driver profile, and specify a heartbeat interval (such as 180 seconds) using the PushKeepAliveTimer driver parameter. Then you add this profile to your communications configurations.

- **ChannelCleanupTimer parameter (communications configuration).** The ChannelCleanupTimer parameter for communications configurations reduces reconnection delays related to session timeouts. This parameter allows the system to identify when a connection is no longer functioning; for example, due to dropped connections or browser failure.

NOTE: If you are using the Push Keep Alive driver, then you *must* also use the ChannelCleanupTimer parameter.

- **CommMaxMsgQ and CommReqTimeout parameters (AOM).** In addition to setting general application timeouts, setting the Siebel Application Object Manager parameters CommMaxMsgQ and CommReqTimeout can also help you manage agent connections effectively.
- **Backup Communications Session Manager (CommSessionMgr) component.** A backup CommSessionMgr component can be specified using communications configuration parameters. The backup CommSessionMgr component runs on another Siebel Server computer and can be accessed without agent interruption in case the primary CommSessionMgr component fails and does not restart.

For more information about using these features, see *Siebel CTI Administration Guide*.

Improving Screen Pop Performance

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

Screen pop response time as experienced by contact center agents is an important indicator of acceptable performance. A screen pop is the display of a view and, optionally, specific records, in response to a communications event. Such events are typically received from CTI middleware; for example, an incoming call is ringing, or the agent has answered the call.

Screen pop behavior is determined by call-handling logic that applies to a particular call based on data attached to the call. Behavior for individual agents is also affected by user settings in the Communications section of the User Preferences screen.

Screen pop performance is affected by the relative complexity of your communications configuration elements, such as event handlers and event responses, and by scripts or business services that can be invoked. Query specifications, database performance, and network capacity and latency also affect screen pop performance. For related information, see [“Improving Performance for Communications Configurations” on page 74](#). For more information about Siebel Web Client response time, see [Chapter 5, “Tuning Siebel Web Client.”](#)

Reviewing Performance Impact of Activity Creation

This topic is part of [“Guidelines for Session Communications Tuning” on page 72](#).

By default, for each communications work item, an activity record is created in the S_EVT_ACT table and related tables.

As you plan your deployment, you must consider how or whether such records are created, review the indexing and layout of applicable database tables, and review the performance impact of generating activity records.

Siebel Email Response Infrastructure

Siebel Email Response uses Communications Server components to enable contact center agents to read and respond to inbound email messages.

It is important to understand the infrastructure that supports Siebel Email Response communications in order to prevent or address performance issues in this area.

Siebel Email Response performance is addressed in this topic and in:

- [“Performance Factors for Siebel Email Response” on page 78](#)
- [“Topology Considerations for Siebel Email Response” on page 79](#)
- [“Guidelines for Siebel Email Response Tuning” on page 79](#)

Key Server Components

Siebel Email Response is supported in the Siebel Server environment primarily by the following components:

- **Communications Inbound Receiver (CommInboundRcvr)**. Receives and queues inbound work items, and queues them for processing by Communications Inbound Processor.
 - For nonreal-time work items, such as email messages for most deployments of Siebel Email Response, Communications Inbound Receiver queues work items it has received for further processing by Communications Inbound Processor.
 - For real-time work items, such as phone calls for Siebel CTI or email messages for some deployments of Siebel Email Response, Communications Inbound Receiver processes work items it has received. Communications Inbound Processor is not used.
- **Communications Inbound Processor (CommInboundProcessor)**. Processes inbound work items that were queued by Communications Inbound Receiver.
- **Communications Outbound Manager (CommOutboundMgr)**. Sends outbound email or other types of messages.
- **Siebel File System Manager (FSMSrvr)**. Writes to and reads from the Siebel File System. This component stores inbound messages prior to processing and stores attachments to inbound and outbound email messages.

Other Siebel Components or Modules

In addition to Siebel Email Response, you might be using Siebel Assignment Manager and Siebel Workflow for routing email messages to agents.

Third-Party Email Server

Siebel Email Response works in conjunction with your third-party email server. Review information presented in documentation for your email server for any requirements that affect your deployment. For information about supported email servers, see *Siebel System Requirements and Supported Platforms* on Oracle Technology Network.

Performance Factors for Siebel Email Response

This topic describes factors that drive or affect performance for Siebel Email Response deployments.

- **Inbound email messages processed per hour**. The number of inbound email messages processed per hour (or some other time period) by your communications infrastructure.

Requirements for processing outbound messages are relatively minor and are tied to inbound message volume. However, other usage of the CommOutboundMgr component or of the email system must also be considered. For example, the Send Email command can be configured to send email through CommOutboundMgr.

- **Volume of customer data.** The total volume of customer data, including templates or categories, literature items, and so on. Template format (HTML or plain text) is a related factor.

Other factors include the size and complexity of inbound email messages and outbound replies.

Also relevant are user settings in the Outbound Communications section of the User Preferences screen, such as whether a reply contains the original message (Include Original Message in Reply setting), or whether HTML or plain text is an agent's default message format (Default Message Format setting).

NOTE: Siebel Email Response coverage in this book focuses on inbound and outbound email processing. In a multichannel environment, session communications performance issues also apply.

Topology Considerations for Siebel Email Response

Processing inbound email messages makes more demands on server resources, particularly CPU usage levels, than processing outbound messages.

Processing of inbound messages associated with a single response group must be handled on a single computer.

If inbound message volume warrants it and if multiple server computers are available to run CommInboundRcvr, CommInboundProcessor, and other components, then consider running CommInboundRcvr and CommInboundProcessor on separate computers (or computers) from other Communications Server components. Topology options for these component are different for real-time and nonreal-time processing. For more information about CommInboundRcvr and CommInboundProcessor, see *Siebel CTI Administration Guide* and *Siebel Email Administration Guide*.

Combining processing of messages for multiple email accounts in a single response group can make processing of inbound messages more efficient. However, if message volume is expected to grow, then limiting the number of email accounts processed by each response group will give you more flexibility to distribute processing across multiple servers, and thereby avoid processing bottlenecks.

Guidelines for Siebel Email Response Tuning

Using your hardware resources optimally, and configuring your system appropriately, can help you to achieve your performance goals. Consider your resources and requirements carefully, and test and monitor system performance on a continual basis.

Review information presented in *Siebel Email Administration Guide*, *Siebel CTI Administration Guide*, relevant third-party documentation, and other sources.

Configuring CommInboundRcvr Threads

Each CommInboundRcvr task runs multiple threads to process inbound email. To determine the number of threads, set the parameters MinThreads and MaxThreads. If extra CPU capacity exists on a given server computer, then you can run more threads for each applicable CommInboundRcvr task.

Managing Email Processing Directories

By default, CommInboundRcvr temporarily writes the content of inbound email messages into subdirectories of the Siebel Server installation directory, until the messages can be processed by the applicable response group and workflow process.

You can use parameters for the Internet SMTP/POP3 Server or Internet SMTP/IMAP Server communications driver to specify alternative directory locations for incoming email, processed email, sent email, and email messages representing certain other processing statuses. You can also set certain driver parameters to specify whether to save or delete processed email messages, for example.

- You must consider the resource requirements for temporary email processing directories when you set up your system.
- Do not delete messages from incoming or queued email directories. Email messages written to processed or sent directories can subsequently be deleted or saved, according to your needs.
- Because of the frequency by which CommInboundRcvr processing writes to temporary email processing directories, the disk must be defragmented regularly.

For more information about email processing directories, see *Siebel CTI Administration Guide* and *Siebel Email Administration Guide*.

Reviewing Performance Impact of Activity Creation

For each email work item, an activity record is created in the S_EVT_ACT table and related tables.

Attachments to such activity records, for inbound and outbound messages, are stored in the Siebel File System.

As you plan your deployment, you must consider how such records are created and managed, review the indexing and layout of applicable database tables, and review the performance impact of generating activity records.

In addition, you must consider the resource requirements for the Siebel File System for storing activity attachments.

The FSMSrvr server component generally runs on the same Siebel Server computers where you are running CommInboundRcvr and CommOutboundMgr.

NOTE: Because of the frequency by which Siebel Email Response processing writes to the Siebel File System, the disk must be defragmented regularly.

For more information about activity attachments stored for inbound email, see *Siebel CTI Administration Guide* and *Siebel Email Administration Guide*.

Configuring Logging for Siebel Email Response

Logging data can be analyzed as part of performance monitoring or tuning, as described in [Chapter 14, “Monitoring Siebel Application Performance with Siebel ARM.”](#)

Higher levels of logging provide more data to help you resolve system errors or performance issues; this is appropriate for system testing. For production systems, however, logging levels must be reduced to improve performance.

For the Internet SMTP/POP3 Server or Internet SMTP/IMAP Server communications driver, an applicable parameter is LogDebug. For details, see *Siebel CTI Administration Guide*.

Applicable event log levels for Siebel Email Response include those for task execution, workflow step execution, workflow process execution, and workflow performance.

7

Tuning Siebel Workflow

This chapter provides guidelines for tuning workflow processes and policies to achieve and maintain optimal performance and scalability. It contains the following topics:

- [About Siebel Workflow on page 83](#)
- [Monitoring Workflow Policies on page 84](#)
- [Tuning Workflow Policies for Performance on page 86](#)
- [Tuning Workflow Processes on page 88](#)
- [Tuning Workflow Process Manager for Performance on page 91](#)

Related Books

Siebel Business Process Framework: Workflow Guide

Configuring Siebel Business Applications

Siebel System Administration Guide

About Siebel Workflow

Siebel Workflow is an interactive software tool that automates business processes.

Workflow processes are designed and administered using the Business Process Designer, a graphical user interface provided through Siebel Tools. Designing, planning, creating, and testing individual workflow processes using the Business Process Designer are described in detail in *Siebel Business Process Framework: Workflow Guide*.

Workflow Policies and Workflow Processes are two components of Siebel Workflow that are designed and created when automating a business process. These components are defined as follows:

- **Workflow Processes.** The representation of a business process. A workflow process comprises one or more steps that indicate when a business process starts and ends and includes information about individual activities within the business process.
- **Workflow Policies.** A systematic expression of a business rule. A workflow policy contains one or more policy conditions and one or more policy actions. If all the policy conditions for a workflow policy are true, then the policy action occurs when all the policy conditions are met. A workflow policy is contained by one workflow policy group and is related to one workflow policy object. A workflow policy contains additional properties that govern its behavior.

Monitoring Workflow Policies

You need to monitor Workflow Policies regularly to check that all events are handled correctly and that the Siebel Server uses its resources optimally. Purging your log files periodically prevents them from becoming too large. Workflow Policies use the General Events event for logging. To see informational messages, set the log level to 3. To see debugging information, set the log level to 4.

You can monitor Workflow Policies using the following views, log files, and tables:

- **Policy Frequency Analysis view.** For details, see [“Using the Policy Frequency Analysis View” on page 84.](#)
- **Workflow Agent trace files.** For details, see [“Using Workflow Agent Trace Files” on page 84.](#)
- **Workflow Policies tables.** For details, see [“Monitoring Workflow Policies Tables” on page 85.](#)

Using the Policy Frequency Analysis View

This topic is part of [“Monitoring Workflow Policies” on page 84.](#)

The Policy Frequency Analysis view provides a list of all executed policies. The Policy Frequency Analysis view is available to analyze how frequently policies are executed over time.

This view displays a log of all the policies executed, as evidenced by a Workflow Monitor Agent process. The policy maker can monitor Workflow Agent process activity to determine if the current policies are adequate, if new policies need to be created, or if policies need to be refined.

The Policy Frequency Analysis view lets you view Policy Log data in a graphical format. The log information is generated by Siebel Server components for Workflow Policies. You access the Policy Frequency Analysis view from the Siebel client by navigating to Policy Frequency Analysis view in the Administration - Business Process screen.

The Policy Frequency Analysis view contains the following fields:

- **Policy.** The name of the policy that was executed.
- **Workflow Object.** The name of the assigned workflow policy object.
- **Object Identifier.** The ID of the workflow policy object for which the policy was executed.
- **Object Values.** Identifying information for the row that executed the policy.
- **Event.** The date and time of the policy execution event.

Using Workflow Agent Trace Files

This topic is part of [“Monitoring Workflow Policies” on page 84.](#)

Workflow agent trace files include the following:

- **Workflow Monitor Agent trace file.** Workflow Monitor Agent provides detailed information about its processing in its trace file.

- **Workflow Action Agent trace file.** Workflow Action Agent provides detailed information about its processing in its trace file.

Setting tracing on the Workflow Action Agent task is required only when the parameter Use Action Agent for Workflow Monitor Agent is set to TRUE. In this case, Workflow Action Agent must be started manually. (It also must be started manually when you use email consolidation.)

Use Action Agent is FALSE by default: Workflow Action Agent is started automatically by Workflow Monitor Agent.

- **Email Manager and Page Manager trace files.**

- Run Email Manager and Page Manager components with Trace Flag set to 1 for detailed reporting on email activity.
- Query the S_APSRVR_REQ table for status information on email and page requests that were logged by Workflow Action Agent.

Monitoring Workflow Policies Tables

This topic is part of [“Monitoring Workflow Policies” on page 84](#).

Workflow Policies use three database tables for processing and tracking requests:

- S_ESCL_REQ
- S_ESCL_STATE
- S_ESCL_ACTN_REQ

Monitor these tables to verify that policies are being processed correctly.

When a trigger fires against a Workflow Policy condition, a record is inserted in the escalation request table, S_ESCL_REQ. Records in this table identify rows in the database that could trigger a Workflow Policy to take action. After the workflow Monitor Agent processes a request, it removes the row from this table.

The S_ESCL_STATE time-based table identifies all the rows that have been executed (all conditions are true) and are waiting for the time duration element to expire.

The S_ESCL_ACTN_REQ table identifies all the rows that are awaiting action execution. These rows have violated the policy; and the time duration element, if any, has expired.

If one of these tables (S_ESCL_REQ, S_ESCL_STATE, and S_ESCL_ACTN_REQ) becomes very large, then this could indicate that the number of policies being monitored is too large, and new Workflow Policies processes need to be created to share the load and improve performance.

If rows are being monitored, but are not being removed from a table after the time interval is met, then this could indicate that a policy was deactivated without removing the database triggers. The triggers are continuing to send data that is not being acted on by a Workflow Policies process. These tables will become very large if you do not restart Generate Triggers.

If you expire or delete any active Workflow Policies, then confirm that no outstanding records remain in the S_ESCL_REQ, S_ESCL_STATE, or S_ESCL_ACTN_REQ tables.

Maintain the S_ESCL_REQ, S_ESCL_ACTN_REQ, and S_ESCL_STATE tables by adjusting parameters related to storage, access, and caching. See database documentation for additional information on properly adjusting such parameters. Also, make sure the database administrator (DBA) is aware of these key tables.

Tuning Workflow Policies for Performance

Workflow Policies can be tuned to optimize your resources while also meeting the policy's timing requirements by grouping similar policies and assigning these policy groups to Siebel Servers that can handle the workload. Performance tuning can be handled in several interrelated ways. The following topics provide more information:

- [“Creating Workflow Policy Groups to Manage Siebel Server Load” on page 86](#)
- [“Multiple Workflow Monitor Agents and Workflow Action Agents” on page 86](#)
- [“Running Workflow Agents on Multiple Siebel Servers” on page 87](#)
- [“Setting Optimal Sleep Interval for Workflow Policy Groups” on page 87](#)
- [“Setting Optimal Action Interval for Workflow Monitor Agent and Workflow Action Agent” on page 88](#)

Creating Workflow Policy Groups to Manage Siebel Server Load

This topic is part of [“Tuning Workflow Policies for Performance” on page 86](#).

Workflow policy groups allow you to group policies with similar polling intervals. This distributes the load to allow efficient processing. For example, if you have very critical policies that must be responded to within minutes of the policy trigger event and you have other policies that need a response within a day, then you can assign them to different workflow policy groups.

The advantage of selective grouping is that a Workflow Agent's polling resources are focused on a smaller number of policies, which helps make monitoring and action execution more effective.

Multiple Workflow Monitor Agents and Workflow Action Agents

This topic is part of [“Tuning Workflow Policies for Performance” on page 86](#).

Each Workflow Agent combination monitors the policies within its assigned workflow policy group. If you are a high-volume call center or you have a large number of policies that need very short polling intervals, then you might want to create multiple groups with Workflow Agent processes to run in parallel. A single Workflow Agent process that is monitoring and handling a large number of events can become slow to respond and not meet the time interval commitments set by the policy.

Running multiple Workflow Monitor Agent and Workflow Action Agents in parallel:

- Focuses a component's polling resources on a smaller number of workflow policies.
- Allows faster throughput by shortening the time between when the workflow policy event is triggered and when the component notices the event.

Running Workflow Agents on Multiple Siebel Servers

This topic is part of [“Tuning Workflow Policies for Performance” on page 86](#).

You can run Workflow Agent processes on different Siebel Servers to ease the workload on each Siebel Server. You can then adjust the polling interval for each group so that polling for noncritical policies does not prevent efficient processing of critical policies.

By distributing workflow policy processes across Siebel Servers:

- High-maintenance policies can be grouped on a Siebel Server with sufficient resources to handle the workflow CPU requirements.
- Low-maintenance policies can be run on a Siebel Server that shares resources with other Siebel processes.

Setting Optimal Sleep Interval for Workflow Policy Groups

This topic is part of [“Tuning Workflow Policies for Performance” on page 86](#).

By creating groups with similar polling intervals, you can assign the workflow policy group to a Workflow Agent process with a polling rate that matches the workflow policy group. Different polling intervals can be assigned to each workflow policy group using the Sleep Time parameter. For more information about Workflow Policies server administration, see *Siebel Business Process Framework: Workflow Guide*.

After Workflow Agents process all requests, the agent processes sleep for the interval specified by this argument before processing begins again. Set the sleep intervals as large as is possible, but at an interval that still meets your business requirements.

NOTE: Setting sleep intervals at values that are too small can put undue stress on the entire infrastructure. Make sure the sleep interval is as large as possible within the context of the business process.

Adjust the sleep interval for each Workflow Agent process to meet the requirements of each workflow policy group.

For example, workflow policy group A contains accounts that require a response to a Severity 1 service request within 10 minutes. Workflow policy group B contains policies that require a customer follow-up call within 14 days.

Workflow policy group A is very time-critical, so you could set the sleep interval to 60 seconds so that the assigned Workflow Policies instance polls frequently. Workflow policy group B is not as time-critical, so you could set the sleep interval to 48 hours and the Workflow Policy instance can still meet its commitments.

Another example where optimal configuration of the Sleep Time parameter can be required is in the case of multiple users who might need to update the same record. If you have, for example, a workflow policy that monitors service requests and you have multiple users who retrieve and modify open service request records, then set the sleep time parameter so that users will have enough time to update the text fields.

If the sleep interval is not set high enough, then you can encounter an error message stating “The selected record has been modified by another user since it was retrieved. Please continue.” In this case, you will lose your changes as the new field values for this record are displayed.

NOTE: If you find that Workflow Policies runs significantly slower during a certain time period, then investigate what other processes might be contending for CPU resources on the Siebel Server. You might discover that the Siebel Server has certain time periods with high activity that interfere with the ability of the Workflow Policies process to monitor or act. Arrange the Workflow Policies processes on the Siebel Servers so that the polling periods are compatible with the resources available.

Setting Optimal Action Interval for Workflow Monitor Agent and Workflow Action Agent

This topic is part of [“Tuning Workflow Policies for Performance” on page 86](#).

For each Workflow Monitor Agent or Workflow Action Agent component, you can set the Action Interval parameter, which determines when actions for a given policy are re-executed on a given base table row. This setting limits the number of times actions are executed if a row keeps going in and out of a matching condition.

You set the Action Interval parameter for Workflow Monitor Agent (rather than Workflow Action Agent) if you have set the parameter Use Action Agent to TRUE for Workflow Monitor Agent. Use Action Agent is FALSE by default.

For example, if a service request severity is set to critical and triggers a policy, then you do not want to re-execute the policy action if it is changed and has been reset to critical during this interval.

Tuning Workflow Processes

In order to improve performance when running workflow processes, you can follow the guidelines explained in the following topics:

- [“Minimizing Usage of Parameter Search Specification” on page 89](#)
- [“Monitoring Conditions Based on Parent and Child Business Components” on page 89](#)
- [“Configuring Siebel Business Applications for Workflow Performance” on page 90](#)
- [“Monitoring Memory Overhead for Workflow Processes” on page 90](#)

NOTE: This performance tuning information is provided as general guidelines for tuning and optimizing performance of workflow processes. Every implementation of Siebel Business Applications is unique, so every use of workflow processes is also unique.

Minimizing Usage of Parameter Search Specification

This topic is part of [“Tuning Workflow Processes” on page 88](#).

Although the server component parameter Search Specification (alias SearchSpec) is a feature of Siebel Workflow, it is recommended that you minimize your use of this parameter with workflow processes that are frequently invoked.

Minimizing SearchSpec use, especially for frequently invoked processes, improves Workflow engine performance during runtime because the engine does not have to construct the SearchSpec string.

It is important, however, that you do not completely avoid using SearchSpec. Not using this parameter can indicate actions taking place on the current row in some cases, and on all rows in other cases. For specific guidelines, note the following:

- For Siebel operations, minimize usage of SearchSpec.
- For batch process requests, use SearchSpec on the business object to limit the number of rows processed.

Indexing Fields in SearchSpec

If you determine that SearchSpec does need to be used, then make sure that all the fields being used are properly indexed. Proper indexing of the fields helps Siebel Workflow and the underlying database to efficiently build queries.

Monitoring Conditions Based on Parent and Child Business Components

This topic is part of [“Tuning Workflow Processes” on page 88](#).

When a condition is being evaluated at a decision step or any other step using a combination of parent and child business components, it is recommended that you closely benchmark the expression or the condition. In some cases, this will require spooling the SQL. For more information, see [“Analyzing Generated SQL for Performance Issues” on page 166](#).

NOTE: The query plan of the SQL might show an extended and poorly performing query. In such cases, it is better to break the conditions up into multiple decision steps and evaluate the conditions separately.

Configuring Siebel Business Applications for Workflow Performance

This topic is part of [“Tuning Workflow Processes” on page 88](#).

In some cases, you might need to perform a comparison between different objects.

Assume, for example, a service request is assigned to a candidate depending on the industry of the account associated with it. In this case, it is necessary to perform a query against Account to fetch the appropriate industries, or to check an industry against all the industries with which the account is associated.

If the workflow process in this example is going to be evaluated frequently, then consider exposing Account Industry on Service Request by the appropriate configuration in order to enhance workflow performance.

Monitoring Memory Overhead for Workflow Processes

This topic is part of [“Tuning Workflow Processes” on page 88](#).

Overhead and performance and scalability characteristics varies depending on whether you are running workflows locally in the Siebel Application Object Manager or in Workflow Process Manager (WfProcMgr), and also on where you run WfProcMgr. The performance and scalability characteristics also depend on whether you are using asynchronous mode for workflow process requests. For more information, see *Siebel Deployment Planning Guide* and *Siebel Business Process Framework: Workflow Guide*.

Running Workflows Locally in Siebel Application Object Manager

A workflow instance (that is, one run of a workflow definition) can run within a Siebel Application Object Manager. In this case, the workflow runs locally, within the current thread that the logged-in user is using. This means that if N users are connected and they all need to run a workflow definition, then the definition would run in that user thread.

In this mode, Workflow adds a fixed overhead (100 to 200 KB) to the user session memory (sometimes referred to as the model) plus memory taken up by other objects (such as business components) contained in the tasks within that workflow.

In general, this option provides the best performance, but is suitable only where scalability is not an important factor.

Running Workflows in Workflow Process Manager

The workflow itself runs within a separate component, which uses a fixed set of resources (parameters MaxMTServers, MaxTasks) to schedule the workflow. The Workflow Process Manager (component alias WfProcMgr) is a multithreaded process that runs multiple workflows and is more scalable because it uses a pool of threads and models.

Generally, the mode of the workflow used depends on what the application is trying to achieve. It is generally recommended that you try to schedule a workflow task in the WfProcMgr, especially if the results of a run are not immediately needed.

You can optionally run WfProcMgr on the same Siebel Server (colocating) as the Siebel Application Object Manager where the workflow is invoked, or run it on dedicated Siebel Server computers. Compared to running workflows locally, running workflows in WfProcMgr might reduce performance, but improve scalability. Running WfProcMgr on dedicated Siebel Servers typically provides the best scalability, while colocating WfProcMgr and Siebel Application Object Manager might provide better performance.

About Asynchronous Mode for Workflow Process Requests

For all Workflow Processes deployment options described previously, workflow process requests can be handled synchronously or using asynchronous mode. Using asynchronous mode comes with the advantages and disadvantages listed in [Table 4 on page 91](#).

Table 4. Advantages and Disadvantages of Using Asynchronous Mode

Advantages	Disadvantages
<ul style="list-style-type: none"> ■ All user threads are not loaded. ■ More scalable as long as: <ul style="list-style-type: none"> ■ There are maximum N simultaneously connected users. ■ There are maximum X simultaneous running workflows. ■ If X is smaller than N, then a WfProcMgr with X tasks can handle a much larger pool (N) of users. 	<ul style="list-style-type: none"> ■ On error, you must look at the log files because there is no automatic notification. ■ The SRBroker could have a timeout or retry feature. ■ Slightly more latency. Additional cost (minimal) of one request per response.

Tuning Workflow Process Manager for Performance

This topic provides general approaches to tune and optimize performance of Workflow Process Manager.

NOTE: Every implementation of Siebel Business Applications is unique, and so every use of workflow processes is also unique. It is recommended that you test, continually monitor, and tune your workflow processes to achieve optimal throughput.

You can follow the guidelines explained in the following information:

- [“Caching Business Services” on page 92](#)
- [“Caching Sessions” on page 92](#)

NOTE: Consider the information provided in this topic as general background information. No attempt is made to detail the many variables that affect tuning at specific sites. This content is not a substitute for specific tuning recommendations made by Global Customer Support or other Oracle service organizations.

Caching Business Services

This topic is part of [“Tuning Workflow Process Manager for Performance” on page 91](#).

Business services invoked through Workflow Process Manager must have the Cache property set to TRUE. This feature makes it possible for the Workflow engine to not reload and reparse the business service, and therefore enhances the performance of workflows that invoke business services.

NOTE: Predefined Siebel business services that have the Cache property set to FALSE must *not* be reset to TRUE.

Caching Sessions

This topic is part of [“Tuning Workflow Process Manager for Performance” on page 91](#).

The parameter OM - Model Cache Maximum (alias ModelCacheMax) for Workflow Process Manager determines the size of the cache for model objects (also known as cached sessions). Cached sessions maintain database connections and session data for locale, user preferences, and access control.

NOTE: Session caching applies only to noninteractive Object Manager-based server components like Workflow Process Manager. It does not apply to Siebel Application Object Manager or EAI Object Manager components.

This feature maintains and reuses existing sessions rather than creating a new session each time one is requested. Using this feature can improve login performance for Workflow Process Manager.

Each model in the cache creates two database connections for the life of the model (one connection for insert, update, and delete operations; the other connection for read-only operations).

The default value is 10. A value of 0 disables this parameter. The maximum value is 100. In general, you set ModelCacheMax to a value approximately equal to the number of concurrent sessions the Workflow Process Manager component is expected to support.

NOTE: When component sessions use multiple user IDs, session caching provides less benefit relative to its cost. The benefit is greatest for component sessions using the same user ID.

See also *Siebel System Administration Guide*.

8

Tuning Siebel Configurator

This chapter describes some issues that affect the performance and throughput of server-based deployments of Siebel Configurator, and provides guidelines for tuning this module to achieve and maintain optimal performance and scalability. It contains the following topics:

- [Siebel Configurator Infrastructure on page 93](#)
- [Performance Factors for Siebel Configurator on page 94](#)
- [Considerations for Defining Topology for Siebel Configurator on page 95](#)
- [Guidelines for Siebel Configurator Tuning on page 98](#)
- [About Siebel Configurator Caching on page 101](#)

Related Books

Siebel Product Administration Guide

Siebel System Administration Guide

Siebel Deployment Planning Guide

Siebel Pricing Administration Guide

Siebel Order Management Guide

Siebel eSales Administration Guide

Siebel Advisor Administration Guide

Siebel Product Administration Guide

Siebel Configurator Infrastructure

Siebel Configurator is one of the Siebel Order Management modules. These modules work together to support various phases in conducting commerce, including online selling. Siebel Configurator provides product configuration and solution-computing capabilities, and can be deployed as a server-based or browser-based module. It uses several infrastructure elements to manage configuration sessions. Siebel Configurator is supported in the Siebel Server environment by the following components:

- **Siebel Application Object Manager.** Siebel Configurator functions can be performed within the Siebel Application Object Manager, such as Call Center Object Manager (alias SCCObjMgr_enu in a U.S. English environment) for Siebel Call Center.

- **Siebel Product Configuration Object Manager.** An optional component, suitable for some Siebel Configurator deployments, that processes configuration requests for user sessions submitted from a Siebel Application Object Manager component. This component has the alias `eProdCfgObjMgr_locale`, such as `eProdCfgObjMgr_jpn` in a Japanese locale. Typically, this component is run on a separate Siebel Server computer than the one running the Siebel Application Object Manager. Multiple instances of this component can be run on the separate Siebel Server computer where it is possible to distribute requests across the various instances. For more information, see [“Considerations for Defining Topology for Siebel Configurator” on page 95](#).

NOTE: The three-letter extension to the alias of the Siebel Product Configuration Object Manager component (such as `jpn` in the example `eProdCfgObjMgr_jpn`) must correspond to the value for the Locale Code parameter (alias `LocaleCode`) associated with the invoking Siebel Application Object Manager. For more information about this requirement, see *Siebel Deployment Planning Guide*. For more information about the Locale Code parameter, see *Siebel Global Deployment Guide*.

For more information about elements of the internal architecture of Siebel Configurator, including Instance Broker (Complex Object Instance Service business service) and Object Broker (Cfg Object Broker business service), see *Siebel Product Administration Guide*.

Performance Factors for Siebel Configurator

In planning Siebel Configurator server-based deployments, or in troubleshooting performance for existing deployments, you must consider several key factors that determine or influence performance.

Subsequent topics provide information and guidelines to help you achieve and maintain optimal performance and scalability.

Performance contexts to consider include response times for:

- **Loading customizable products.** This is the time elapsed from the moment a user clicks Customize in a quote or order until the user interface for the customizable product has been loaded and displayed to the user.
- **Responding to user selections.** This is the time elapsed from the moment a selection is made by the user until Siebel Configurator returns a response such as an update to the customizable product or a conflict message.

The factors below, particularly customizable product size and complexity, are relevant in both of these contexts.

Some of the key performance factors for server-based deployments of Siebel Configurator include:

- **Number of concurrent configuration users.** The number of concurrent users who access customizable product models. This figure will be some percentage of the total number of concurrent users on the Siebel Application Object Manager.

More specifically, you would be concerned with the total number of configuration sessions per hour, and the average length of those sessions.

- **Size and complexity of product models.** The total size and complexity of each customizable product model, particularly where multiple hierarchical levels, many constraints, and a complex user interface are defined.

A major potential performance factor is custom scripting attached to update events on applicable business components, such as Quote, Quote Item, Quote Item Attribute, Order, Order Item, and Order Item Attribute.

- **Number of product models.** The number of customizable product models accessed by users. It is assumed that each user accesses no more than one customizable product model at one time. A given group of concurrent users can access multiple models, however, each of which must be separately cached.

Considerations for Defining Topology for Siebel Configurator

This topic describes considerations for defining the topology for Siebel Configurator server-based deployments. There are two major topology approaches to deploying Siebel Configurator:

- Running Siebel Configurator in the Siebel Application Object Manager component. For more information, see [“Running Siebel Configurator in the Siebel Application Object Manager Component” on page 95](#).
- Running Siebel Configurator on one or more dedicated Siebel Servers. (Such servers are sometimes referred to as remote servers, because they are remote to the computer on which Siebel Application Object Manager is running. In general, this topic uses the term dedicated servers.) For more information, see [“Running Siebel Configurator on Dedicated Servers” on page 96](#).

The optimal deployment approach for Siebel Configurator, and the optimal number of server computers you require for this module, depends on factors such as those described in [“Performance Factors for Siebel Configurator” on page 94](#).

Running Siebel Configurator in the Siebel Application Object Manager Component

This topic is part of [“Considerations for Defining Topology for Siebel Configurator” on page 95](#).

You can run Siebel Configurator in the Siebel Application Object Manager component, such as for Siebel Call Center.

If a small number of concurrent users require configuration sessions, or there are a small number of customizable product models, then this deployment option can yield reasonable performance and make the most effective use of your hardware resources.

With this option, you set all parameters for managing Siebel Configurator caching on each applicable Siebel Application Object Manager. For details, see [“About Siebel Configurator Caching” on page 101](#).

Running Siebel Configurator on Dedicated Servers

This topic is part of [“Considerations for Defining Topology for Siebel Configurator” on page 95](#).

You can run Siebel Configurator on one or more dedicated Siebel Server computers using a server component other than the Siebel Application Object Manager. This component is Siebel Product Configuration Object Manager (alias eProdCfgObjMgr_jpn in a Japanese locale).

Possible variations on this general topology option include:

- Running one eProdCfgObjMgr component with one Siebel Application Object Manager component
- Running multiple eProdCfgObjMgr components with one Siebel Application Object Manager component
- Running one eProdCfgObjMgr component with multiple Siebel Application Object Manager components
- Running multiple eProdCfgObjMgr components with multiple Siebel Application Object Manager components

If a large number of concurrent users require configuration sessions, or there are a large number of customizable product models, then this deployment option (using one or more dedicated servers) can yield the best performance and make the most effective use of your hardware resources.

With this option, you set some parameters for managing Siebel Configurator caching on each applicable Siebel Application Object Manager, and some on each applicable dedicated Siebel Configurator server. For details, see [“About Siebel Configurator Caching” on page 101](#).

Configuring Siebel Application Object Manager for Dedicated Siebel Configurator Deployments

When you designate one or more dedicated server computers to run the Siebel Product Configuration Object Manager (alias eProdCfgObjMgr_jpn in a Japanese locale) component, then you must configure any Siebel Application Object Manager components from which users will initiate configuration sessions to route configuration requests to these computers.

The Siebel Application Object Manager forwards each configuration session request to the dedicated Siebel Configurator server with the fewest concurrent users.

Table 5 on page 97 lists server parameters for managing dedicated Siebel Configurator deployments. Using Server Manager, set these parameters on each Siebel Application Object Manager (do not set them on the dedicated Siebel Configurator server computer).

Table 5. Server Parameters for Dedicated Siebel Configurator Server Deployment

Parameter Name	Display Name	Data Type	Default Value	Description
eProdCfgRemote	Product Configurator-Use remote service	Boolean	FALSE	Set this parameter to TRUE if you are running the eProdCfgObjMgr component on one or more dedicated servers. Set this parameter to FALSE for Siebel Configurator deployments using Siebel Application Object Manager only.
eProdCfgServer	Product Configurator-Remote Server Name	Text	Not applicable	When you have not enabled explicit product mapping for products to a Siebel Configurator server, set this parameter to the names of the dedicated computers on which you are running eProdCfgObjMgr. Otherwise, set the value of this parameter to NULL.
eProdCfgTimeOut	Product Configurator-Time out of connection	Integer	20	Sets the length of time, in milliseconds, that the Siebel Application Object Manager tries to connect to a dedicated Siebel Server running eProdCfgObjMgr. After the timeout has been reached, an error is returned to the user.

Table 5. Server Parameters for Dedicated Siebel Configurator Server Deployment

Parameter Name	Display Name	Data Type	Default Value	Description
eProdCfgKeepAliveTime	Product Configurator - Keep Alive Time of Idle Session	Integer	900	<p>Setting in seconds to determine the maximum interval of inactivity during a configuration session.</p> <p>If the interval of inactivity reaches this value, then the user session is stopped and the worker returns to the pool.</p> <p>If this parameter is not set, then an infinite interval is assumed.</p> <p>Set this parameter on the Siebel Application Object Manager only. It does not apply on the remote Configurator server.</p> <p>NOTE: On the remote Configurator server (eProdCfgObjMgr component), set the parameter ConnIdleTime to a value like eProdCfgKeepAliveTime plus 1 second.</p>

Guidelines for Siebel Configurator Tuning

Using your hardware resources optimally, and configuring your system appropriately, can help you to achieve your performance goals. Consider your resources and requirements carefully, and test and monitor system performance on a continual basis.

Review information presented in *Siebel Product Administration Guide*, *Siebel System Administration Guide*, and other sources.

Activities you can perform to achieve performance and scalability goals include:

- Adjusting your system topology. For more information, see [“Considerations for Defining Topology for Siebel Configurator” on page 95](#).
- Configuring Siebel Server server components for Siebel Configurator.
- Designing and deploying your customizable product models. For more information, see [“Defining Customizable Product Models and Classes” on page 100](#).

This topic applies to deployments using Siebel Web Client. It contains the following information:

- [“Tuning Siebel Configurator” on page 99](#)
- [“Specifying the Siebel Configurator File System Location” on page 100](#)
- [“Defining Customizable Product Models and Classes” on page 100](#)

Tuning Siebel Configurator

This topic is part of [“Guidelines for Siebel Configurator Tuning” on page 98](#).

How you configure your Siebel Server components for Siebel Configurator server deployments, for appropriate tuning, depends in part upon which deployment method you use, as described in [“Considerations for Defining Topology for Siebel Configurator” on page 95](#).

- If you deploy Siebel Configurator on the Siebel Application Object Manager, then your Siebel Configurator tuning calculations must be made in combination with your Siebel Application Object Manager tuning calculations.
- If you deploy Siebel Configurator using the Product Configurator Object Manager (eProdCfgObjMgr) server component on a dedicated Siebel Server computer, then your Siebel Configurator tuning calculations will be only indirectly related to your Siebel Application Object Manager tuning calculations and will be determined primarily by configuration-related concurrent users and request loads.

In particular, note that, for a dedicated Siebel Configurator server, the MaxTasks parameter is generally set much lower than for a Siebel Application Object Manager. By default, the ratio of MaxTasks to MaxMTServers is 20:1 for eProdCfgObjMgr.

In addition, depending on request load, MaxTasks is generally set lower for a Siebel Application Object Manager running Siebel Configurator than for a Siebel Application Object Manager that is *not* running Siebel Configurator.

You can follow this general procedure to determine how to set these parameters:

- Determine what percentage of users for your Siebel application are also users of Siebel Configurator. For example, for every 100 users, 60 work with Quotes.
- Calculate what percentage of time these users spend using Siebel Configurator. For example, out of the 60 users mentioned previously, only 30 are concurrently using Siebel Configurator.
- Maintain the default ratio of 20:1 for MaxTasks divided by MaxMTServers.

If you deploy Siebel Configurator using eProdCfgObjMgr on a dedicated Siebel Server computer and database connection (login and log out) is slow, then it is recommended that you do the following:

- Enable database connection pooling

To enable connection pooling, set the parameters MaxSharedDbConns and MinSharedDbConns to positive integer values (at least 1) that are no higher than MaxTasks minus 1.

This pools all user connections without sharing and avoids the creation and deletion of a new database connection for each eProdCfgObjMgr session.

■ Use third-party user authentication

Using third-party user authentication, such as LDAP, rather than database authentication avoids creating an additional database connection for authentication. For more information about authentication options, see *Siebel Security Guide*.

For more information about database connection pooling, see [“Configuring Database Connection Pooling for Siebel Application Object Managers” on page 36](#).

Specifying the Siebel Configurator File System Location

This topic is part of [“Guidelines for Siebel Configurator Tuning” on page 98](#).

Siebel Configurator uses a file system directory to cache all configuration related object definitions. The server parameter, Product Configurator - FS location (alias eProdCfgCacheFS), specifies the location. Specify a value for this parameter to reference a server directory path which has write permission. For example, \\MyServer\Si bFS\Si ebConfi g.

NOTE: The value for the directory that you specify must be network-accessible.

It is recommended that you do not specify a top-level directory. For example, if the directory Si bFS is a top-level directory, then specify a subdirectory such as Si ebConfi g.

If you do not specify a value for eProdCfgCacheFS, then Siebel Configurator attempts to use the Siebel File System. If the Siebel File System uses the File System Manager (alias FSMSrvr) component, then Siebel Configurator does not cache object definitions to the file system. For more information about the Siebel File System, see *Siebel System Administration Guide*.

Defining Customizable Product Models and Classes

This topic is part of [“Guidelines for Siebel Configurator Tuning” on page 98](#). It describes some guidelines about creating customizable products and classes in a manner that will optimize performance:

- To maintain good performance, do not make your customizable products or classes any larger or more complex than absolutely necessary.
- Complexity is a function of the number of hierarchical levels and constraints built into the customizable product models and of the structure of the class.
- For defining class relationships, use specific classes as much as possible. For example, avoid defining class relationships without specifying classes, or use a subclass rather than a parent class if it is so defined.
- Minimize the complexity of user interface elements you associate with your customizable product models.
- Generally, using interactive or automatic pricing updates for customizable products is recommended. If performance is adversely affected, then consider switching to manual pricing updates.

- When creating rules, using the Set Preference template allows you to create soft constraints that guide the Siebel Configurator engine in producing solutions, but which the engine can ignore if needed to avoid conflicts or performance problems.
- By default, when you add a customizable product to a quote, for example, default products and selections will be included, and Siebel Configurator can be invoked to create this default instance. If the customizable product default selections are large and complex, and if users are required to immediately customize the product, then turning off the Default Instance Creation feature will enhance performance with no loss of functionality.

For more information on these issues, see *Siebel Product Administration Guide*.

About Siebel Configurator Caching

Siebel Configurator supports several types of caching of customizable product information, to optimize response time for configuration-session users. Caching options include:

- Caching in memory

Siebel Configurator caches versions of customizable products, product classes, and attribute definition objects in memory. When the size limit for this cache is reached, the versions of the objects that were least recently used are discarded. For more information, see [“Default Caching Behavior for Siebel Configurator” on page 102](#).

- Caching in the Siebel Configurator File System

This directory caches versions of the customizable products, product classes, and attribute definition objects that were loaded into memory. This is default behavior. For more information, see [“Default Caching Behavior for Siebel Configurator” on page 102](#).

- In addition to the caching options, you can also specify which server or component caches versions of customizable products, product classes, and attribute definition objects.

The specified cache can be updated at regular intervals. Using these options can improve response times to requests for a specific customizable product. For more information, see [“Cache Management for Siebel Configurator” on page 102](#).

NOTE: The memory resources for your Siebel Configurator server computer must be sufficient to support your caching requirements.

In addition to the topics referred to previously, this topic contains the following information:

- [“Parameters for Configuring Siebel Configurator Caching” on page 104](#)
- [“Determining Rough Sizing for Caching Parameters” on page 107](#)
- [“Administering the Siebel Configurator Cache” on page 108](#)

Default Caching Behavior for Siebel Configurator

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

The default caching behavior for Siebel Configurator is as follows:

- When a user starts a configuration session, Siebel Configurator looks for new cache update requests that affect objects in the cache. If there are new cache update requests that affect objects currently in the cache, then Siebel Configurator updates or removes these objects.
- Siebel Configurator looks to see if the requested customizable product is cached in memory.
- If the customizable product is not already cached in memory, then Siebel Configurator looks in the Configurator File System.

NOTE: The location of the Configurator File System is specified by the value of the Product Configurator - FS location parameter (alias eProdCfgCacheFS). If no value is specified for eProdCfgCacheFS, then Siebel Configurator looks in the Siebel File System. For more information about the Configurator File System, see [“Specifying the Siebel Configurator File System Location” on page 100](#).

- If the customizable product is not in the Configurator File System, then it is loaded from the Siebel database. The product is added to the memory cache and to the Configurator File System.
- Thereafter, when a configuration session starts, the customizable product is loaded from the memory cache or the Configurator File System.
- Before loading the customizable product from the Configurator File System, the system checks the Siebel database to make sure each item in the product is the current version.
- If the cached product has changed in the database, then the current version of the item is loaded from the database. This makes sure that the most recent version of a customizable product and its contents are loaded.
- When the product administrator releases a new version of a customizable product, the changes are written to the Siebel database and a cache update request is posted for the modified customizable product. The memory cache and the Configurator File System are not updated with the changes until the next configuration session is requested for the customizable product.

NOTE: It is recommended that you avoid the use of start or end dates in rules for customizable products. The arrival of a date does not cause the customizable product to be refreshed in the cache.

Cache Management for Siebel Configurator

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

When a user starts a configuration session, Siebel Configurator loads the requested customizable product into memory. You can specify a cache (server or component) to serve requests for frequently requested customizable products to improve response times. This requires that you select the setting *Explicit Product Mappings Only* on the specified cache. The specified cache loads the customizable products that are mapped to it into memory before any user requests are received.

You can also specify a time interval so that the specified cache updates the customizable products it holds at regular intervals. This reduces the possibility that a user request requires the retrieval of data from the database and the loading of a revised customizable product. To specify the time interval, you set values for the following parameters on the eProdCfgObjMgr component:

- Server Session Loop Sleep Time (alias ServerSessionLoopSleepTime)
- Product Configurator - Cache Engine Objects (alias eProdCfgCacheEngineObjects)

For more information on these parameters, see [Table 6 on page 104](#).

Requests for other customizable products that are not mapped to a specific cache are served by a cache that has the setting *Explicit Product Mappings Only* disabled.

The following procedure describes how to configure product caching by mapping a product to a cache that has the setting *Explicit Product Mappings Only* enabled.

To configure product caching

- 1 Navigate to the Administration - Product screen, then the Cache Administration view.

The Cache Administration view appears.

- 2 In the Cache applet, select a cache.

NOTE: Only one cache can be active at a time.

- 3 In the Cache Type field, select a value as described in the following list:

- **Server.** Select Server as the cache type to route configuration requests to the Siebel Servers associated with the cache.
- **Component.** Select Component as the cache type to route configuration requests to the components associated with the cache. These components can span multiple Siebel Servers depending on where components are active.

NOTE: If you select Component as the cache type, then you must set the same value for the component parameter Enable internal load balancing (alias EnableVirtualHosts) on both the Siebel Application Object Manager and on the eProdCfgObjMgr component. For example, if EnableVirtualHosts is set to TRUE on the Siebel Application Object Manager component, then it must also be set to TRUE on the eProdCfgObjMgr component.

- 4 In the Components applet, specify a Siebel Server name or a component name to associate with the cache that you selected in [Step 2 on page 103](#).

The value that you specify depends on the value that you selected for Cache Type in [Step 2 on page 103](#). For example, if you set Cache Type equal to Server, then you enter the name of a Siebel Server. If you set Cache Type equal to Component, then you enter the name of a component.

- 5 If you want a server cache or component cache to only serve products that are mapped to that server cache or component cache, then select *Explicit Product Mappings Only*.
- 6 In the Product applet, select the product(s) that you want to associate with the component that you selected in [Step 4 on page 103](#).

- 7 In the Cache applet, click Validate.

The application validates that the Siebel Server or component names that you select are valid for the cache type that you specified.

- 8 If the configuration that you created validates correctly, then click Release to enable the cache that you selected cache instances of the products that are mapped to it.

Parameters for Configuring Siebel Configurator Caching

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

Siebel Configurator caching is enabled by default (eProdCfgSnapshotFlg is set to TRUE). Other parameters must be sized following guidelines such as those described in [“Determining Rough Sizing for Caching Parameters” on page 107](#).

[Table 6 on page 104](#) lists the server parameters for configuring Siebel Configurator caching. Set these parameters on the Siebel Application Object Manager component for a Siebel Application Object Manager deployment of Siebel Configurator. For a dedicated Siebel Configurator server deployment, set these parameters on the Siebel Application Object Manager *and* on the eProdCfgObjMgr component. For information on how to configure server parameters, see *Siebel System Administration Guide*.

Table 6. Server Parameters for Configuring Siebel Configurator Cache Behavior

Parameter Alias	Parameter Name	Data Type	Default Value	Description
eProdCfgCacheFS	Product Configurator - FS location	String	Not applicable	Specifies the location of the Configurator File System. If no value is specified for eProdCfgCacheFS, then Configurator looks in the Siebel File System. For more information about the Configurator File System, see “Specifying the Siebel Configurator File System Location” on page 100 .

Table 6. Server Parameters for Configuring Siebel Configurator Cache Behavior

Parameter Alias	Parameter Name	Data Type	Default Value	Description
eProdCfgAttrSnapshotFlg	Product Configurator - Collect and Use the snapshots of the ISS_ATTR_DEF Ob	Boolean	TRUE	Set to TRUE to enable caching for attribute definitions. This caches attribute definitions in memory. It is strongly recommended that you do not change this parameter.
eProdCfgNumOfCachedAttrs	Product Configurator - Number of Attribute Definitions Cached in Memory	Integer	100	Sets the number of attribute definitions kept in memory at any given time during configuration.
eProdCfgClassSnapshotFlg	Product Configurator - Collect and Use the snapshots of ISS_CLASS_DEF Ob	Boolean	TRUE	Set to TRUE to enable caching for product class definitions. It is strongly recommended that you do not change this parameter.
eProdCfgNumOfCachedClasses	Product Configurator - Number of Class Definitions Cached in Memory	Integer	100	Sets the number of class definitions kept in memory at any given time during configuration.
eProdCfgProdSnapshotFlg	Product Configurator - Collect and Use the snapshots of ISS_PROD_DEF Ob	Boolean	TRUE	Set to TRUE to enable caching for product definitions. This caches product definitions in memory. It is strongly recommended that you do not change this parameter.

Table 6. Server Parameters for Configuring Siebel Configurator Cache Behavior

Parameter Alias	Parameter Name	Data Type	Default Value	Description
eProdCfgNumOfCachedProducts	Product Configurator - Number of Product Definitions Cached in Memory	Integer	1000	Sets the number of product definitions kept in memory at any given time during configuration.
eProdCfgSnapshotFlg	Product Configurator- Collect and use snapshots of the Cfg objects	Boolean	TRUE	Set to TRUE to turn on Siebel Configurator caching. It is strongly recommended that you do not change this parameter.
eProdCfgNumbOfCachedCatalogs	Product Configurator- Number of cached catalogs	Integer	10	Sets the maximum number of Model Manager catalogs that can be cached in memory. NOTE: This parameter provides functionality provided by the parameter eProdCfgNumbOfCachedFactories in some previous releases.
eProdCfgNumbOfCachedWorkers	Product Configurator- Number of workers cached in memory	Integer	50	Sets the maximum number of workers that can be cached in memory. This number applies to all Model Manager catalogs.
eProdCfgCacheEngineObjects	Product Configurator - Cache Engine Objects	Boolean	TRUE	Set to TRUE to enable content cache and pre-caching.

Table 6. Server Parameters for Configuring Siebel Configurator Cache Behavior

Parameter Alias	Parameter Name	Data Type	Default Value	Description
ServerSessionLoopSleepTime	Server Session Loop Sleep Time	Integer	300	<p>Specify an interval time (in seconds) to refresh cached products that are mapped to a Configurator server cache or component cache using the explicit product mapping setting.</p> <p>NOTE: Product Configurator - Cache Engine Objects must be set to TRUE.</p>

Determining Rough Sizing for Caching Parameters

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

To help you determine how to set the Siebel Configurator caching parameters, a general suggestion is to measure the incremental memory required for a customizable product.

Requirements for Model Manager and Worker caching are more relevant than those for object caching. Object caching has a small requirement, and applies to multiple users. Model Manager caching applies to multiple users (using the same customizable product). Worker caching also applies to multiple users.

You can try this on a Siebel Developer Web Client (a Mobile Web Client using a dedicated database connection) by checking the memory used by the siebel.exe process before and after you click Customize for a customizable product included in a quote or order, and again after you have further configured the customizable product (to reach maximum likely memory usage).

For example, X might be the before-loading memory size, Y might be the after-loading size, and Z might be the memory size after additional product configuration.

Of the incremental memory observed, consider the following breakdown:

- The size of a Model Manager for a customizable product is about 75% of the incremental memory required to instantiate the product (that is, 75% of Y minus X).
- The size of a Worker for a customizable product varies during runtime, generally increasing as user selections are made. This size can be approximated by subtracting the Model Manager size from the difference of Z less X.

Administering the Siebel Configurator Cache

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

Siebel administrators or product administrators can refresh or update the Siebel Configurator cache in several ways. The following topics describe procedures to refresh or update the Siebel Configurator cache with changes for customizable products, product classes, and attribute definitions.

Related Topics

[“Refreshing the Entire Siebel Configurator Cache” on page 108](#)

[“Refreshing the Siebel Configurator Cache with Product Changes” on page 109](#)

[“Updating the Siebel Configurator Cache with Product Class Changes” on page 109](#)

[“Refreshing the Siebel Configurator Cache with Product Class Changes” on page 110](#)

[“Updating the Siebel Configurator Cache with Attribute Definition Changes” on page 110](#)

[“Refreshing the Siebel Configurator Cache with Attribute Definition Changes” on page 111](#)

Refreshing the Entire Siebel Configurator Cache

This topic is part of [“About Siebel Configurator Caching” on page 101](#). It describes how you can refresh the Siebel Configurator cache with changes made to customizable products, product classes, and attribute definitions in one operation.

Consider performing the task described here as part of standard maintenance for your Siebel Configurator deployment or, for example, if you intend to migrate data from a development to a production environment.

Other related topics describe how you can update the Siebel Configurator cache with changes you made to customizable products, product classes, and attribute definitions in separate operations. For more information, see [“Administering the Siebel Configurator Cache” on page 108](#).

To refresh the entire Siebel Configurator cache

- 1 Navigate to the Administration - Product screen, then the Cache Administration view.
- 2 Select the record for the cache that you want to refresh.
- 3 Click the Menu button in the Cache list, then choose Refresh Product Cache.

Refreshing the Siebel Configurator Cache with Product Changes

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

While editing a product record, a product administrator can select Refresh Product Cache to refresh the Siebel Configurator cache with changes made to the selected product. The next time a user requests the customizable product, the user receives a freshly instantiated version reflecting the product change and the cache is refreshed with this version. For example, you could change the product description and then refresh the cache. For more information about refresh options for the Siebel Configurator cache, see [“Administering the Siebel Configurator Cache” on page 108](#).

To refresh the cache with product changes

- 1 Navigate to the Administration - Product screen.
- 2 Select the record for a customizable product that has been changed or that is to be refreshed.
- 3 Click the Menu button in the Products list, then choose Refresh Product Cache.

Updating the Siebel Configurator Cache with Product Class Changes

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

While editing a product class record, a product administrator can select Update Cache to remove the version access keys of a cached product class from the Siebel Configurator cache for all versions of the selected product class. This forces the application to consult the database the next time it requires a version access key. A product administrator must select Update Cache if he or she modifies a nonversioned cached property such as, for example, a product name. If the product administrator does not select Update Cache, a Siebel Application Object Manager that has a version already cached uses the old version. For more information about administration options for the Siebel Configurator cache, see [“Administering the Siebel Configurator Cache” on page 108](#).

To update the cache with class changes

- 1 Navigate to the Administration - Product screen, then the Product Classes view.
The Product Classes list applet appears.
- 2 Select a product class and modify it or its attribute definitions as needed.
- 3 From the menu in the Product Classes list, choose Update Cache.

Refreshing the Siebel Configurator Cache with Product Class Changes

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

While editing a product class record, a product administrator can select Refresh Cache to refresh the customizable products in the Siebel Configurator cache with changes made to the product class record. The next time a user requests the customizable product, the user receives a freshly instantiated version reflecting the product change and the cache is refreshed with this version. This new instance reflects the changes you made to the product class. For more information about refresh options for the Siebel Configurator cache, see [“Administering the Siebel Configurator Cache” on page 108](#).

To refresh the cache with class changes

- 1 Navigate to the Administration - Product screen, then the Product Classes view.
The Product Classes list applet appears.
- 2 Select a product class and modify it or its attribute definitions as needed.
- 3 From the menu in the Product Classes list, choose Refresh Cache.

Updating the Siebel Configurator Cache with Attribute Definition Changes

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

While editing an attribute definition record, a product administrator can select Update Cache to remove the version access keys of an attribute definition record from the Siebel Configurator cache for all versions of the selected attribute definition. This forces the application to consult the database the next time it requires a version access key. A product administrator must select Update Cache if he or she modifies a nonversioned cached property such as, for example, an attribute definition name. If the product administrator does not select Update Cache, then a Siebel Application Object Manager that has a version already cached uses the old version. For more information about administration options for the Siebel Configurator cache, see [“Administering the Siebel Configurator Cache” on page 108](#).

To refresh the cache with attribute definition changes

- 1 Navigate to the Administration - Product screen, then the Attribute Definitions view. The Attribute Definitions list applet appears.
- 2 Select an attribute definition and modify it as needed.
- 3 From the menu in the Attribute Definitions list, choose Update Cache.

Refreshing the Siebel Configurator Cache with Attribute Definition Changes

This topic is part of [“About Siebel Configurator Caching” on page 101](#).

While editing an attribute definition record, a product administrator can select Refresh Cache to refresh customizable products in the Siebel Configurator cache with the changes made to the attribute definition record.

The next time a user requests the customizable product, the user receives a freshly instantiated version reflecting the attribute definition change and the cache is refreshed with this version. This new instance reflects the changes you made to the attribute definition. For more information about refresh options for the Siebel Configurator cache, see [“Administering the Siebel Configurator Cache” on page 108](#). For more information, see [“Administering the Siebel Configurator Cache” on page 108](#).

To refresh the cache with attribute definition changes

- 1 Navigate to the Administration - Product screen, then the Attribute Definitions view. The Attribute Definitions list applet appears.
- 2 Select an attribute definition and modify it as needed.
- 3 From the menu in the Attribute Definitions list, choose Refresh Cache.

9

Tuning Siebel EAI

This chapter discusses tuning for Siebel Enterprise Application Integration (Siebel EAI) that might be required for optimal performance. It contains the following topics:

- [About Siebel Enterprise Application Integration on page 113](#)
- [Guidelines for Siebel EAI Tuning on page 113](#)

Related Books

Overview: Siebel Enterprise Application Integration

Integration Platform Technologies: Siebel Enterprise Application Integration

Transports and Interfaces: Siebel Enterprise Application Integration

Business Processes and Rules: Siebel Enterprise Application Integration

XML Reference: Siebel Enterprise Application Integration

About Siebel Enterprise Application Integration

Siebel EAI provides components for integrating Siebel Business Applications with external applications and technologies within your company. Siebel EAI works with technologies, standards, or applications that include XML, HTTP, Java, and various third-party middleware products and application integration solutions.

Siebel EAI provides bidirectional real-time and batch solutions for integrating Siebel Business Applications with other applications. Siebel EAI is designed as a set of interfaces that interact with each other and with other Siebel components.

Guidelines for Siebel EAI Tuning

This topic describes guidelines for maintaining acceptable performance using Siebel EAI.

General guidelines are followed by recommendations specific to Siebel EAI features such as IBM WebSphere MQ (formerly MQSeries) Transport adapter, HTTP Inbound Transport adapter, EAI Siebel Adapter, virtual business components, and Workflow Process Manager used with Siebel EAI.

Follow these general guidelines to improve overall performance for data integration and throughput of Siebel Business Applications:

- Try to minimize round trips between systems. For example, if an integration needs to request three pieces of data, then do not send a request for one piece of data, wait for the response, and then send the next request. If you need multiple pieces of data, then gather the data in a single request.

- Try to keep processing in a single session wherever possible, to avoid having to make calls between server components.
- Within a session, try to minimize the nesting of calls between components such as workflow, scripting, and the EAI Siebel Adapter. For example, use a workflow process to sequence the calling of business services and keep scripting code in self-contained steps. Workflow subprocesses can be used to package together commonly called sequences of services.
- Use alternatives to scripting, where possible. If you use scripting, then use it minimally and economically and apply documented guidelines. For more information, see ["Performance Guidelines for Siebel Scripting" on page 175](#).
- Configure business components, business services, caching, and other application functionality that supports integration processing to obtain optimal performance. For more information, see other topics in this chapter and see [Chapter 12, "Tuning Customer Configurations."](#)
- Perform capacity planning for all servers that support integration processing. For sizing reviews, consult Oracle's Application Expert Services. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.
- Try to represent the incoming external data in the same code page and encoding that the Siebel application uses internally (UCS-2). This eliminates the need to use the Transcode business service in your workflow process, thus improving performance.

Related Topics

- ["Improving IBM WebSphere MQ Transport Performance" on page 114](#)
- ["Improving HTTP Inbound Transport Performance" on page 116](#)
- ["EAI Siebel Adapter Performance" on page 117](#)
- ["Virtual Business Component Performance" on page 118](#)
- ["Improving Workflow Process Manager Performance" on page 119](#)
- ["Other Guidelines for Siebel EAI" on page 120](#)

Improving IBM WebSphere MQ Transport Performance

This topic is part of ["Guidelines for Siebel EAI Tuning" on page 113](#).

The performance of an IBM WebSphere MQ queue is highly dependent on the disk performance of the queue manager computer and the layout of the queue's files on the disk. Test your queue with stand-alone utilities so that you have an upper boundary for the performance that can be expected in an active application.

To achieve higher throughput, consider the following options:

- **Run multiple MQ Receiver tasks.** Run multiple MQ Receiver tasks in parallel on the same computer or across several computers. The optimal number of MQ Receiver tasks depends on the transaction type.

NOTE: This guide refers to *MQ Receiver*, where the actual Siebel Server component you are using can be either MQSeries Server Receiver (alias MqSeriesSrvRcvr) or MQSeries AMI Receiver (alias MqSeriesAMIRcvr).

The default number of MQ Receiver tasks is 1. You can set this to 10 or more, depending on the nature of your transactions and on available server capacity.

Adding MQ Receivers is generally most helpful for handling CPU-bound transactions, where the dequeuing rate is low and MQ contention is not experienced.

Sometimes contention is still experienced after adding MQ Receiver tasks, such as when multiple MQ Receivers connect to the same MQ queue manager or queue. See the next item for more information.

- **Run multiple MQ queue managers.** If you experience diminishing returns from adding MQ Receiver tasks, then you might benefit from running additional MQ queue managers. Doing so can help to reduce contention of MQ resources stored in physical folders on disk.
- **Turn off persistent queuing if it is not necessary.** Performance issues for non CPU-bound transactions or for persistent queuing are often related to MQ contention, which is not helped by adding receivers. If you do not require persistent queuing, then turn it off.

Persistent queuing is significantly slower than normal queuing for WebSphere MQ. If you do not use this feature, however, messages will be lost if the queue manager goes down.
- **Set Maximum Number of Channels parameter.** Set the Maximum Number of Channels parameter in the WebSphere MQ queue manager to be greater than or equal to the maximum number of simultaneous clients you have running.

In addition, there are specific actions you can take to improve WebSphere MQ Transport performance for outbound and inbound transports, as detailed below.

Inbound Messages

For inbound WebSphere MQ messages, run multiple MQ Receivers in parallel to increase throughput.

Outbound Messages (Send, SendReceive)

Caching of WebSphere MQ Transport business services can improve outbound performance by eliminating the need to connect to the queue for each message. Caching is disabled by default because it is not usable in every situation. Follow these tips to enable caching:

- Cache in client sessions only. Do not use caching if your transport will be called within the Workflow Process Manager (WfProcMgr) component. The threading model of this component is not compatible with the WebSphere MQ APIs.
- To enable caching for a business service, set the Cache property to TRUE in Siebel Tools, then recompile the SRF file.

- If you need to call the WebSphere MQ Transport in Workflow Process Manager and in a client session, then make a separate copy of the service (one cached and one uncached) for each situation.
- Caching occurs on a per-queue basis and only one connection is kept open at a time. If a single session is going to talk to multiple queues, then consider making a copy of the transport for each outbound queue.

NOTE: See your IBM WebSphere MQ documentation for performance and sizing guidelines.

Performance Events

You can get detailed performance tracing of the WebSphere MQ Transport by setting the EAITransportPerf event to level 5.

You can set this event level for multiple Siebel Server components that play a role in Siebel EAI functionality, including Workflow Process Manager (WfProcMgr), EAI Object Manager (EAIObjMgr), MQ Receiver, or other components. For example, you can use `srvrmgr` to set the event level for MQ Receiver:

```
change evtloglvl EAITransportPerf=5 for comp MqSeriesSrvRcvr
```

Improving HTTP Inbound Transport Performance

This topic is part of [“Guidelines for Siebel EAI Tuning” on page 113](#).

The HTTP Inbound Transport supports two modes, session mode and sessionless mode:

- In session mode, the session stays active until a logout call
- In sessionless mode, logging in and logging out occur automatically for each request

Use session mode whenever possible, because the time required to log into the application is usually significantly longer than the time required to process an average request.

You can also use the `SessPerSisnConn` component parameter to control the number of sessions sharing the same physical Siebel Internet Session Network Application Programming Interface (SISNAPI) connection between the Web server and the EAI Object Manager.

Setting this parameter to 1 provides a dedicated physical connection for each Siebel session. The default value is 20, to allow up to 20 sessions to share the same SISNAPI connection. For the EAI Object Manager, it is recommended that you set `SessPerSisnConn` to 1. If setting `SessPerSisnConn` to 1 results in an excessive number of sessions, then consider increasing the value of `SessPerSisnConn` or provide additional hardware resources.

You can change this parameter using `srvrmgr` at the Enterprise or Server level. For example, to set the parameter at the Enterprise level for the EAI Object Manager, you enter the following command:

```
change param SessPerSisnConn=1 for compdef eaiobjmgr_enu
```

For more information about configuring `SessPerSisnConn`, see [“Configuring SISNAPI Connection Pooling for Siebel Application Object Manager” on page 47](#).

EAI Siebel Adapter Performance

This topic is part of [“Guidelines for Siebel EAI Tuning” on page 113](#).

Use the techniques described here to improve the EAI Siebel Adapter performance and throughput.

Reviewing Scripting

Avoid scripting events on business components used by the EAI Siebel Adapter. Perform any scripting task either before or after the EAI Siebel Adapter call, rather than within it. For general scripting guidelines, see also [“Performance Guidelines for Siebel Scripting” on page 175](#).

Disabling Logging

Disable logging for performance-critical processes that are functioning correctly to gain about 10% faster performance. You can disable logging for the EAI Object Manager (or other applicable server components, such as MQ Receiver) by setting the BypassHandler parameter to TRUE.

Minimizing Integration Object Size

The size of an integration object and its underlying business components can have an impact on the performance of the EAI Siebel Adapter. To minimize this impact, you can:

- Consider copying business objects and business components and modifying them to remove any elements (such as scripts, joins, multi-value fields, user properties, and so on) that you do not require in the Siebel EAI context. Base your integration objects on these relatively streamlined object definitions. Verify that user keys on your integration objects make effective use of indexes when queries are performed.
- Inactivate unnecessary integration components and integration component fields in your integration objects. Activate only the components and fields needed for message processing, according to your business needs.
- Inactivate unnecessary fields for each underlying business component. For fields that are unneeded, if Force Active is set to TRUE, then set it to FALSE. Setting Force Active to FALSE prevents the EAI Siebel Adapter from processing these fields. If you do not inactivate these fields, then the adapter processes them even when they are not actually included in the integration object.

For more information, see [“Limiting the Number of Active Fields” on page 185](#).

Analyzing SQL Produced by EAI Siebel Adapter

Requests to the EAI Siebel Adapter eventually generate SQL to be executed against the Siebel Database. By setting the event SQL to level 4 in the component running in the EAI Siebel Adapter, you can get a trace of the SQL statements being executed, along with timings for each statement, in milliseconds.

You can get timings for each EAI Siebel Adapter operation by setting the event EAISiebAdptPerf to 4 or 5. Do this to correlate the EAI Siebel Adapter calls with their associated SQL.

After you have this information, look through the logs to find any SQL statements taking significantly longer than average. To improve the performance of such statements, look at the business component (perhaps eliminating unnecessary joins and fields) or at the physical database schema (perhaps adding indexes).

NOTE: The overall timing across operations (equivalent to the TotalTimeForProcess event) cannot be determined by adding the individual logged values associated with the EAI SiebAdptPerf event, because the EAI Siebel Adapter requires some additional overhead. Overhead is greater when EAI SiebAdptPerf is set to a high value. Set this event to a lower value for a production system for optimal performance.

Running EAI Siebel Adapter in Parallel

A common technique to improve throughput is to run multiple instances of the EAI Siebel Adapter in parallel.

For the MQ Receiver, you do this by running multiple receiver tasks. For more information, see [“Improving IBM WebSphere MQ Transport Performance” on page 114](#).

For the EAI Object Manager, you do this by setting the MaxTasks, MaxMTServers, and MinMTServers parameters, in order to run more threads (tasks) on more multithreaded processes for the EAI Object Manager component. Also start multiple simultaneous HTTP sessions. There is little interaction between each instance of the EAI Siebel Adapter.

If the Siebel Database Server is large enough, then almost linear scalability of the EAI Siebel Adapter is possible until either the limits of the CPU or the memory limits of the Siebel Server are reached.

CAUTION: If two sessions attempt to simultaneously update or insert the same record, then one will succeed and one will produce an error. Therefore, when running the EAI Siebel Adapters in parallel, you need to prevent the simultaneous update of the same record in multiple sessions. You can prevent this by either partitioning your data or retrying the EAI Siebel Adapter operation where the error occurs.

Caching Business Objects

The EAI Siebel Adapter caches business objects by default. The default cache size is five objects. Using caching, subsequent runs on the adapter are significantly faster because the business objects do not need to be re-created for each run.

Use the BusObjCacheSize parameter on the EAI Siebel Adapter to change the size of the cache, if required. However, the five-object cache size is enough for most purposes. Making this number too large creates an unnecessarily large memory.

Virtual Business Component Performance

This topic is part of [“Guidelines for Siebel EAI Tuning” on page 113](#).

Because users must wait for the virtual business component (VBC) response to display the GUI component for the integration on their screens, this type of integration is especially sensitive to latency.

To improve virtual business component performance when your integration has multiple requests, put the requests for a given system in a single batch.

Improving Workflow Process Manager Performance

This topic is part of [“Guidelines for Siebel EAI Tuning” on page 113](#). It discusses performance issues for the Workflow Process Manager component. For more information about Siebel Workflow performance, see [Chapter 7, “Tuning Siebel Workflow.”](#) Also see *Siebel Business Process Framework: Workflow Guide*.

Workflow Process Manager is a task-based server component. A new thread is created for each request. However, sessions for Object Manager components (such as EAI Object Manager or Siebel Application Object Managers that can invoke workflow processes) are cached and reused for subsequent requests. When sizing a system, consider the maximum number of workflow tasks you expect to have active at a given time. This determines the maximum number of Object Manager sessions Siebel applications create. In general, creating smaller workflow processes is recommended. If you cannot avoid creating a large workflow process, then divide the workflow process into subprocesses.

CPU and Memory Consumption

The exact CPU and memory consumption of each task depends on the actions performed in your workflow processes. To estimate CPU and memory consumption in your production environment, run a single task, measure its resource consumption, and make an estimation based on your maximum concurrent sessions. Take session caching into account when making these measurements.

If you need a large number of sessions, then you might want to run Workflow Process Manager on multiple Siebel Server computers. You can then use Siebel Server load balancing to load-balance requests across the Siebel Servers. If you plan to run a significant number of tasks per server (such as 100 or more), then you might also want to run multiple multithreaded processes.

If you are going to run several different types of workflows, then run each type in a separate process. This makes it easier to monitor the overall CPU and memory usage of each process type.

The number of multithreaded processes and the number of tasks per process are controlled through the parameters MaxMTServers (Maximum MT Servers), MinMTServers (Minimum MT Servers), and MaxTasks (Maximum Tasks).

NOTE: These parameters are per Siebel Server. For example, MaxMTServers refers to how many multithreaded processes to run on each Siebel Server computer. For details, see *Siebel System Administration Guide*.

Performance Events

You can get performance tracing of workflows by setting the event WfPerf for the component in which your workflow is running. Setting the event to level 4 gives timing for the execution of the overall process. Setting the event to level 5 provides timing for each step as well.

You can set this event level for any Siebel Server component that invokes a workflow process as part of Siebel EAI functionality. For example, to set this event level for the MQ Receiver using `srvrmgr`, enter the following:

```
change evtloglvl WfPerf=5 for comp MqSeriesSrvRcvr
```

These events can be useful not just for measuring workflow performance but also for measuring the performance of business services executed within these workflows.

Other Guidelines for Siebel EAI

This topic is part of [“Guidelines for Siebel EAI Tuning” on page 113](#).

Review the following issues for applicability to your deployment, for optimizing Siebel EAI performance:

- **Check disks on the computer.** Do a preliminary test on the queue manager you are using to see how many sends and corresponding receivers it can support per second (use multiple drivers). Queue vendors such as IBM WebSphere MQ provide test programs you can use to drive these and determine how much the queue itself can scale. The speed of the disks on the computer is important.
- **Optimize messages.** In the messages, reference only the columns you require.
- **Create smaller business components.** Messages might use only a small portion of the actual business components.

Create copies of the business components you are using. In the copies, keep active all fields used by the optimized integration object or otherwise used for correctly processing of messages (like the visibility fields or status fields). Deactivate all other fields. Also deactivate the join definitions and multi-value links (MVLs) that are not needed for processing of the messages.

The original business components are often large and complex and contain elements you will not need for your integration purposes. Use the smaller business components and business objects and links created when creating the optimized integration object.

Business components can have fields with Force Active set to TRUE. Check this property for fields in the business components, using Siebel Tools. If the fields are not needed, then set Force Active to FALSE.
- **Set user property All Mode Sort to FALSE.** Set the user property All Mode Sort to FALSE for optimized business components (if not already set). Do this only for the smaller business components created for use with Siebel EAI, because this user property changes the order in which rows are retrieved, which might not be appropriate or normal clients. For more information about All Mode Sort, see *Siebel Developer's Reference*.
- **Optimize database queries.** Review queries generated by the receiver process and verify that they are optimized.
- **Turn off logging.** Turn off server-side logging that you do not require.

10 Tuning Siebel EIM

This chapter describes recommended guidelines for improving the performance of Siebel EIM. It contains the following topics:

- [About Siebel EIM on page 121](#)
- [Siebel EIM Architecture Planning Requirements on page 122](#)
- [Siebel EIM Usage Planning on page 124](#)
- [General Guidelines for Optimizing Siebel EIM on page 127](#)
- [Recommended Sequence for Implementing Siebel EIM Processes on page 128](#)
- [Troubleshooting Siebel EIM Performance on page 131](#)
- [Database Guidelines for Optimizing Siebel EIM on page 143](#)
- [Data Management Guidelines for Optimizing Siebel EIM on page 153](#)
- [Run Parameter Guidelines for Optimizing Siebel EIM on page 153](#)
- [Monitoring the Siebel Server During a Siebel EIM Task on page 154](#)

About Siebel EIM

Siebel Enterprise Integration Manager (Siebel EIM) is a server component in the Siebel EAI component group that transfers data between the Siebel database and other corporate data sources. This exchange of information is accomplished through intermediary tables called EIM tables. (In earlier releases, EIM tables were known as interface tables.) The EIM tables act as a staging area between the Siebel application database and other data sources.

Siebel EIM is your primary method of loading mass quantities of data into the Siebel database. Use Siebel EIM to perform bulk imports, updates, merges, and deletes of data.

In the Siebel application database, there are application tables (known as base tables), which Siebel Business Applications use. For data to come from other corporate data sources (external databases) into Siebel application tables, the data must go through EIM tables. So the data exchanges between the Siebel database and external databases occurs in two phases:

- 1 Load data into EIM tables.
- 2 Run Siebel EIM to import the data from the EIM tables into the Siebel base tables.

NOTE: While the first phase of this data-exchange process involves the intermediary tables that are called EIM tables, only the second phase involves the functionality of Siebel EIM.

When data is entered through the Siebel user interface, the application references properties set at the business component object type. However, when data is entered into Siebel base tables through Siebel EIM, EIM references properties set at the table object type.

NOTE: You must use Siebel EIM to perform bulk imports, exports, merges, and deletes, because it is not supported to use native SQL to load data directly into Siebel base tables (the tables targeted to receive the data). Additionally, be aware that Siebel EIM translates empty strings into NULL.

Siebel EIM Architecture Planning Requirements

You must consider the size and complexity of the implementation before executing any single item with the Siebel application. Aspects that have a direct impact on how the production application will perform might not be your highest priority when you initially begin your Siebel implementation. However, the decisions made during the initial phases of an implementation have a far reaching impact, not only on performance and scalability but also on the overall maintenance of the Siebel application.

It is strongly recommended to have a Siebel certified principal consultant or architecture specialist from Oracle's Application Expert Services involved in designing the most effective logical and physical architecture for your organization. This includes capacity planning and system sizing, physical database layout, and other key architecture items. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services. For more information, see the following information:

- ["Database Sizing Guidelines" on page 122](#)
- ["Database Layout Guidelines \(Logical and Physical\)" on page 123](#)

Database Sizing Guidelines

This topic is part of ["Siebel EIM Architecture Planning Requirements" on page 122](#).

One of the most important factors to determine about the database is its overall size. During the planning phase, you need to allocate space for system storage, rollback segments and containers, temporary storage space, log files, and other system files required by the relational database management system (RDBMS), as well as space for the Siebel application data and indexes. If you allocate too little space for the system, then performance will be affected and, in extreme cases, the system itself can be halted. If you allocate too much space, then it can cause inefficiency.

The space needed by the database depends on the total number and types of supported users. It is recommended that you consult your vendor RDBMS technical documentation for more information on these requirements.

The space required for Siebel data and indexes depends on the functionality being implemented and the amount and nature of data supporting this functionality.

The process for making accurate database size calculations is a complex one involving many variables. Use the following guidelines:

- Determine the total number, and types, of users of Siebel Business Applications (for example, 500 sales representatives and 75 sales managers).
- Determine the functionality that you will implement and the entities required to support them. Typically, the largest entities are as follows:
 - Accounts
 - Activities
 - Contacts
 - Forecasts
 - Opportunities
 - Service Requests
- Estimate the average number of entities per user (for example, 100 accounts per sales representative) and calculate an estimated total number of records per entity for the total user base.
- Using standard sizing procedures for the specific database, and *Siebel Data Model Reference* on My Oracle Support, calculate the average record size per entity and multiply by the total number of records. Typically, these entities span multiple physical tables, all of which must be included in the row size calculation. This determines the estimated data sizes for the largest entities.
- You must add additional space for the storage of other Siebel application data. A rough guideline for this additional amount would be one-half the storage required for these key entities.
- Indexes typically require approximately the same amount of space as data.
- Be sure to allow for a margin of error in the total size calculation.
- Be sure to factor growth rates into the total size calculation.

Database Layout Guidelines (Logical and Physical)

This topic is part of [“Siebel EIM Architecture Planning Requirements” on page 122](#).

The overall performance of Siebel Business Applications largely depends on the input/output (I/O) performance of the database server. To achieve optimal I/O performance, it is critical that the tables and indexes in the database be arranged across available disk devices in a manner that evenly distributes the I/O load.

The mechanism for distributing database objects varies by RDBMS, depending on the manner in which storage space is allocated. Most databases have the ability to assign a given object to be created on a specific disk. These objects, and guidelines for some of them, are provided in the following list.

A redundant array of independent disks, or RAID, can provide large amounts of I/O throughput and capacity, while appearing to the operating system and RDBMS as a single large disk (or multiple disks, as desired, for manageability). The use of RAID can greatly simplify the database layout process by providing an abstraction layer above the physical disks while ensuring high performance.

Regardless of the implemented RDBMS and the chosen disk arrangement, be sure that you properly distribute the following types of database objects:

- Database log or archive files.
- Temporary workspace used by the database.
- Tables and Indexes: In most implementations, the tables and corresponding indexes in the following list tend to be some of the more heavily used and must be separated across devices. In general, the indexes listed below must be placed on different physical devices from the tables on which they are created.

- | | |
|------------------|---------------|
| ■ S_ACCNT_POSTN | ■ S_PARTY_REL |
| ■ S_OPTY | ■ S_PARTY |
| ■ S_ADDR_ORG | ■ S_SRV_REQ |
| ■ S_OPTY_POSTN | ■ S_EVT_ACT |
| ■ S_CONTACT | ■ S_OPTY |
| ■ S_POSTN_CON | ■ S_ORG_EXT |
| ■ S_DOCK_TXN_LOG | |

NOTE: If you plan on making extensive use of Siebel EIM, then put the key EIM tables (based on the unique business requirements) and their corresponding indexes on different devices from the Siebel base tables and indexes, because all of them are accessed simultaneously during Siebel EIM operations.

Siebel EIM Usage Planning

This topic provides several general guidelines for effective and efficient implementations of Siebel EIM, regardless of the size of the overall Siebel implementation. You must take a strategic perspective when implementing Siebel EIM to make sure your deployment is successful.

For more information, see the following information:

- [“Defining the Siebel EIM Team” on page 125](#)
- [“Mapping Data into Siebel Business Applications” on page 125](#)
- [“Testing Siebel EIM Processes” on page 126](#)

Defining the Siebel EIM Team

This topic is part of [“Siebel EIM Usage Planning” on page 124](#).

Based on customer experience, it is recommended that a team of individuals be assigned to manage and maintain the Siebel EIM processes required for your organization. Consider using individuals with the following skill sets:

- For small to medium-sized Siebel Business Applications implementations:
 - A database administrator with a detailed understanding of not only the RDBMS used by your organization, but also the Siebel Data Model. This individual would be responsible for identifying the actual data to be loaded into the EIM tables and making sure that the physical layout of the database provides optimal performance. This person would also be responsible for the task of mapping the data into the Siebel base tables. For more information on performing this task, see *Siebel Enterprise Integration Manager Administration Guide*.
 - A system administrator with a strong background in the systems used by your organization. This individual would be responsible for developing scripts unique to your organization to automate the loading of data into the EIM tables, and to execute Siebel EIM in order to process the data into the Siebel base tables.

NOTE: Your organization might have one individual with both these skill sets and so you might rather dedicate only a single individual to these tasks. If this is the case, then consider having a backup person, so that when this primary individual is unavailable, the backup person is capable of performing what needs to be done to keep the Siebel implementation operational.

- For larger to very large-sized Siebel implementations:
 - A database administrator with a detailed understanding of not only the RDBMS used by your organization, but also the Siebel Data Model. This individual would be responsible for identifying the actual data to be loaded into the EIM tables and to make sure that the physical layout of the database provides optimal performance. This team member would also be responsible for the crucial task of mapping the data into the Siebel base tables. For more information on performing this task, see *Siebel Enterprise Integration Manager Administration Guide*.
 - A system administrator with a strong background in the systems (both the database server and application server) used by your organization. This individual would be responsible for developing scripts unique to your organization to automate the loading of data into the EIM tables, and to execute Siebel EIM in order to process the data into the Siebel base tables.
 - A business analyst with a strong understanding of the Siebel Data Model and its intended usage in the Siebel implementation. This individual would act as a liaison between the business and technical members of the Siebel EIM team.

Mapping Data into Siebel Business Applications

This topic is part of [“Siebel EIM Usage Planning” on page 124](#).

Siebel EIM uses EIM table mappings to map columns from EIM tables to Siebel base tables. Predefined Siebel EIM mappings are fixed and cannot be remapped.

NOTE: Siebel EIM uses only EIM table mappings to determine table relationships. Siebel EIM does not use configuration logic in the Siebel repository to determine table relationships.

Using Siebel Tools, you can view:

- Mappings of EIM tables to Siebel base tables
- Mappings of EIM table columns to Siebel base table columns
- Mappings of Siebel base tables to EIM tables
- Mappings of Siebel base table columns to EIM table columns

Some base tables might not be mapped to a corresponding EIM table. In such cases, use Siebel Visual Basic (VB) to load data into these base tables and inform Global Customer Support regarding the missing mapping. For information on using Siebel VB, see *Siebel VB Language Reference*.

If you have licensed Database Extensibility and created extensions, then you can use the Column Mapping screen to specify mappings to the new fields. Database extensibility and Siebel EIM support mappings between columns in extension tables and EIM tables only if these columns share the same base table. To map EIM table extensions to base table extensions, you must specify which column the extended field will point to in the base table. For more information on Database Extensibility, see *Configuring Siebel Business Applications*.

CAUTION: Manually mapping new extension columns to columns in EIM tables presents risks of errors during Siebel EIM execution. Whether or not you have licensed Database Extensibility, it is strongly recommended for you to request an EIM Data Mapping and Design review or other assistance from Oracle's Application Expert Services to help you perform the necessary tasks. This review can be used to make sure that the EIM mappings are correct and will accomplish intended goals. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

To map data into a Siebel application

- 1 Determine which Siebel base table columns need to be populated for the Siebel implementation, along with the external data that will be loaded into these base tables.
- 2 Determine which EIM table and columns will be used to import from the source to the destination.
- 3 Analyze this external data to determine which attributes need to be stored and the relationship this data has to other entities.

Testing Siebel EIM Processes

This topic is part of ["Siebel EIM Usage Planning" on page 124](#).

Fully and completely testing Siebel EIM processes must not be overlooked. Testing is more than simply mapping the data and then running a Siebel EIM process using the default Siebel EIM configuration file. Complete testing requires you to run a large number of identical Siebel EIM jobs with similar data. Testing in this way allows you to not only find areas that you might have overlooked, but also provides some insight into optimal sizing of the Siebel EIM batches and exposure to scenarios that can occur in a production environment.

Before using Siebel EIM, a database administrator must populate the EIM tables with data to be processed by Siebel EIM. Then, you can invoke Siebel EIM to process this data, with Siebel EIM making multiple passes through the tables to complete the specified process.

Siebel EIM reads a special configuration file that specifies the Siebel EIM process to perform (import, merge, delete, or export) and the appropriate parameters. The Siebel EIM configuration file (the default file is `default.ifb`) is an ASCII text file of extension type `.IFB` that resides in the `admin` subdirectory under the Siebel Server directory. Before running a Siebel EIM process, you must edit the contents of the Siebel EIM configuration file to define the processes that Siebel EIM will perform.

The Siebel EIM log file can contain information at different levels of detail depending on the values of three flags: the Error flag, the SQL flag, and the Trace flag. For more information on these flags, see *Siebel Enterprise Integration Manager Administration Guide*. Some of the recommended settings are described in the following list:

- As a starting point, it is recommended to set the Error flag = 1, the SQL flag = 1, and the Trace flag = 1. This setting will show errors and unused foreign keys. Setting the Trace flag = 1 will provide a summary (after each batch) of the elapsed time after Siebel EIM updates primary child relationships in the Siebel database tables as necessary and runs optional miscellaneous SQL statements.
- Set the Error flag = 1, the SQL flag = 8, and the Trace flag = 3. These settings will produce a log file with SQL statements that include how long each statement took, which is useful for optimizing SQL performance.
- Set the Error flag = 0, the SQL flag = 0, and the Trace flag = 1. These settings will produce a log file showing how long each Siebel EIM step took, which is useful when figuring out the optimal batch size as well as monitoring for deterioration of performance in a particular step.

General Guidelines for Optimizing Siebel EIM

The following guidelines are recommended for improving Siebel EIM performance:

- Verify that all indexes exist for the tables involved. Keep in mind, however, that for large loads you must drop most of the indexes from the target tables to increase the speed of the process, rebuilding those indexes afterward when the process is finished.
- Limit tables and columns to be processed using ONLY BASE TABLES/COLUMNS configuration parameters to minimize Siebel EIM processing.
- Consider disabling the Docking: Transaction Logging system preference during the Siebel EIM run. Switching off transaction logging improves performance; however, this benefit must be balanced with the need for mobile users to reextract afterward.
- Altering batch sizes to find the optimal batch size for a given business component typically helps resolve performance issues. The batch size is dependent upon the quantity of data and which type of Siebel EIM process you are running.

NOTE: Although the limit of rows you can process is directly related to the capabilities of your database server, executing batches greater than 100,000 rows is strongly discouraged.

- For Siebel EIM delete processes that use the DELETE EXACT parameter, use a batch size of 20,000 rows or less.
- Try using batch ranges (BATCH = X-Y). This allows you to run with smaller batch sizes and avoid the startup overhead on each batch. The maximum number of batches that you can run in a Siebel EIM process is 1,000.
- Perform regular table maintenance on EIM tables. Frequent insert or delete operations on EIM tables can cause fragmentation. Consult your database administrator to detect and correct fragmentation in the EIM tables.
- Delete batches from EIM tables on completion. Leaving old batches in the EIM table wastes space and could adversely affect performance.
- Run independent Siebel EIM jobs in parallel.
- Set the USING SYNONYMS parameter to FALSE in the .IFB file to indicate that account synonyms do not need to be checked.
- If no other strategy appears to be successful, then use the SQLPROFILE parameter to identify slow-running steps and queries. For more information, see [“Using the SQLPROFILE Parameter” on page 136](#).

Recommended Sequence for Implementing Siebel EIM Processes

The following sequence is recommended for implementing Siebel EIM processes:

- 1 Customize and test the .IFB file to meet the business requirements.
- 2 Tune the .IFB parameters.
- 3 Separate the Siebel EIM processes.
- 4 Set the database parameters, making sure the basic requirements are met, including the hardware, the settings, and no or minimal fragmentation.

Before you start optimizing Siebel EIM processes, make sure there are no network problems or server performance problems that can affect the results. Oracle's Application Expert Services recommends using at least 100 MB network segments and network-interface cards (NICs) to connect the Siebel Server and the Siebel Database, and also recommends using a network switch or similar technology, rather than a hub, to maximize throughput. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

This topic contains the following information:

- [“Optimizing the .IFB File for Siebel EIM” on page 129](#)
- [“Checking .IFB File Optimization for Siebel EIM” on page 129](#)
- [“Separating Siebel EIM Processes by Operation” on page 130](#)

Optimizing the .IFB File for Siebel EIM

This topic is part of [“Recommended Sequence for Implementing Siebel EIM Processes”](#) on page 128.

When you have finished coding and testing the .IFB file to meet your business requirements, the next step is to optimize the .IFB file. The selected parameters in each section of the .IFB file determine the focus of each Siebel EIM task. The following recommendations are provided for each section of the .IFB file:

- **ONLY BASE TABLES** or **IGNORE BASE TABLES**. These parameters specify and restrict the selected base tables for the Siebel EIM process. A single EIM table (sometimes referred to as an interface table) is mapped to multiple user or base tables. For example, the EIM table EIM_ACCOUNT is mapped to S_PARTY, S_ORG_EXT, and S_ADDR_ORG, as well as other tables. The default configuration is to process all base tables for each EIM table.

NOTE: Oracle's Application Expert Services strongly recommends that you always include these parameters in every section of the .IFB file, and list only those tables and columns that are relevant for a particular Siebel EIM task. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

- **ONLY BASE COLUMNS** or **IGNORE BASE COLUMNS**. These parameters specify and restrict the selected base columns for the Siebel EIM process. The default is to process all base columns for each base table. It is likely that you are not using every column in a base table, and these parameters will make sure that Siebel EIM is only processing the desired columns in the table.

You will see an additional performance increase if you exclude those columns that are defined as foreign keys (FKs) and are not used by the Siebel configuration, because Siebel EIM does not need to perform the interim processing (using SQL statements) to resolve the values for these FKs. Set the Siebel EIM task parameter Error Flags to a value of 1 to see which FKs are failing to be resolved by Siebel EIM (you might have missed excluding that FK with this parameter).

NOTE: Do not use the IGNORE BASE COLUMNS parameter for merge processes or export processes. Use this parameter only for import processes and delete processes.

Checking .IFB File Optimization for Siebel EIM

This topic is part of [“Recommended Sequence for Implementing Siebel EIM Processes”](#) on page 128.

One method to find out if the .IFB file is optimized is to check the status of the records being processed in the EIM tables. This indicates if there are tables or columns that are being processed unnecessarily. The following query can be used to check the status of records in an EIM table:

```
select count(*), IF_ROW_STAT from EIM Table
where IF_ROW_BATCH_NUM = ?
group by IF_ROW_STAT;
```

If many rows have a status of PARTIALLY IMPORTED, then it is likely that further tuning can be done by excluding base tables and columns that are not necessary. For example, two tests were run to import 5000 accounts from the EIM_ACCOUNT table. The first test included all of the base tables while the second test only focused on the four necessary tables by including the following line in the .IFB file:

```
ONLY BASE TABLES = S_ORG_EXT, S_ADDR_ORG, S_ACCNT_POSTN, S_ORG_TYPE
```

The first test took 89 minutes to import (excluding the Updating Primaries step), while the second test only took 2 minutes to import (excluding the Updating Primaries step).

Separating Siebel EIM Processes by Operation

This topic is part of [“Recommended Sequence for Implementing Siebel EIM Processes” on page 128](#).

Wherever possible, divide the Siebel EIM batches into insert-only transactions and update-only transactions. For example, assume that you are loading 50,000 records into an EIM table as part of a weekly process. 10,000 records represent new data and 40,000 records represent updates to existing data.

By default, Siebel EIM can determine which records are to be added and which records are to be updated in the base tables, however, Siebel EIM will need to perform additional processing (through SQL statements) to make these determinations. If you were able to divide the 50,000 records into different batch numbers based on the type of transaction, then you could avoid this additional processing.

In addition, the columns being processed as part of the update activity might be less than those for the insert activity (resulting in an additional performance increase). To illustrate this, the .IFBs in the preceding example can be coded with the following sections:

■ .IFB for mixed transactions:

```
[Weekly Accounts]
TYPE = IMPORT
BATCH = 1-10
TABLE = EIM_ACCOUNT
ONLY BASE TABLES = S_ORG_EXT
IGNORE BASE COLUMNS = S_ORG_EXT. ?
```

■ .IFB for separate insert or update transactions:

```
[Weekly Accounts - New]
TYPE = IMPORT
BATCH = 1-2
TABLE = EIM_ACCOUNT
ONLY BASE TABLES = S_ORG_EXT
```

```

IGNORE BASE COLUMNS = S_ORG_EXT. ?
INSERT ROWS = TRUE
UPDATE ROWS = FALSE
[Weekly Accounts - Existing]
TYPE = IMPORT
BATCH = 3-10
TABLE = EIM_ACCOUNT
ONLY BASE TABLES = S_ORG_EXT
ONLY BASE COLUMNS = S_ORG_EXT. NAME, S_ORG_EXT. LOC, S_ORG_EXT. ?
INSERT ROWS = FALSE
UPDATE ROWS = TRUE

```

Troubleshooting Siebel EIM Performance

Before troubleshooting Siebel EIM performance, verify that there are no performance bottlenecks on the Siebel Server computer or network. This topic contains the following information:

- ["Optimizing SQL for Siebel EIM" on page 132](#)
- ["Using the USE INDEX HINTS and USE ESSENTIAL INDEX HINTS Parameters" on page 132](#)
- ["Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Example" on page 133](#)
- ["Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Criteria for Passing Indexes to the Database" on page 135](#)
- ["Using the SQLPROFILE Parameter" on page 136](#)
- ["Additional Indexes on Siebel EIM Tables" on page 137](#)
- ["Creating Proper Statistics on Siebel EIM Tables" on page 138](#)
- ["Dropping Indexes in Initial Runs of Siebel EIM" on page 139](#)
- ["Controlling the Size of Batches for Siebel EIM" on page 140](#)
- ["Controlling the Number of Records in Siebel EIM Tables" on page 141](#)
- ["Using the USING SYNONYMS Parameter with Siebel EIM" on page 141](#)
- ["Using the NUM_IPTABLE_LOAD_CUTOFF Extended Parameter with Siebel EIM" on page 141](#)
- ["Disabling the Docking: Transaction Logging Parameter for Siebel EIM" on page 142](#)
- ["Disabling Database Triggers for Siebel EIM" on page 142](#)
- ["Running Siebel EIM Tasks in Parallel" on page 142](#)

Optimizing SQL for Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

During this process, you need to be able to run several similar batches. If you do not have enough data with which to experiment, then you might need to back up and restore the database between runs, so that you can continue processing the same batch.

First, run a Siebel EIM job with the following flag settings: Error flag = 1, SQL flag = 8, and Trace flag = 3. This will produce a log file that contains SQL statements and shows how long each statement took. Identify SQL statements that are taking too long (on a run of 5000 rows in a batch, look for statements that took longer than one minute). These are the statements that you want to concentrate on, and it is recommended that you consult an experienced database administrator at this point. The process of optimizing the SQL for Siebel EIM involves the following:

- Use the respective database vendor’s utility or a third-party utility to analyze the long-running SQL statements.
- Based on the review of the data access paths, review the SQL statements for proper index usage. There might be cases where an index is not used at all or the most efficient index is not being chosen. This can require a thorough analysis.
- Based on this analysis, use a systematic approach to tuning these long-running statements. Perform one change at a time and then measure the results of the change by comparing them to the initial benchmarks. For example, you might find that dropping a particular index to improve the performance of one long-running statement might negatively impact the performance of other SQL statements.

Base the decision on whether to drop the index on the impact to the overall process as opposed to the individual long-running SQL statement. For this reason, it is important that one change be implemented at a time in order to measure the impact of the change.

- After repetitively going through and optimizing each long-running SQL statement, the focus can be shifted to other tuning measures, such as increasing the number of records processed in the EIM table at a time and the running of parallel Siebel EIM tasks.

Using the USE INDEX HINTS and USE ESSENTIAL INDEX HINTS Parameters

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Perform testing with the .IFB file parameters USE INDEX HINTS and USE ESSENTIAL INDEX HINTS, trying both settings (TRUE and FALSE). The default value for USE INDEX HINTS is FALSE. The default value for USE ESSENTIAL INDEX HINTS is TRUE.

NOTE: If your configuration file has more than one process section, then you must specify USE INDEX HINTS within each one.

If these parameters are set to FALSE, then Siebel EIM does not generate hints during processing. By setting the value to FALSE, you can realize performance gains if the TRUE setting means that hints are being generated that direct the database optimizer to use less than optimal indexes. Siebel EIM processing must be tested with both the TRUE and FALSE settings to determine which one provides better performance for each of the respective Siebel EIM jobs.

NOTE: The `USE INDEX HINTS` parameter is applicable only for Oracle Database. The `USE ESSENTIAL INDEX HINTS` parameter is applicable only for Oracle Database and Microsoft SQL Server.

These two parameters work for different queries, so you need to enable both to get all of the index hints on Oracle Database.

Related Topics

- [“Using `USE INDEX HINTS` and `USE ESSENTIAL INDEX HINTS`: Example” on page 133](#)
- [“Using `USE INDEX HINTS` and `USE ESSENTIAL INDEX HINTS`: Criteria for Passing Indexes to the Database” on page 135](#)

Using `USE INDEX HINTS` and `USE ESSENTIAL INDEX HINTS`: Example

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

The example in [Table 7 on page 133](#) illustrates the results achieved for an SQL statement with index hints and without index hints. This example was performed on the Microsoft SQL Server platform.

Table 7. Example Results for SQL Statement With and Without Index Hints

SQL User Name	CPU	Reads	Writes	Duration	Connection ID	SPID
SADMIN	549625	38844200	141321	626235	516980	9

```

UPDATE dbo.S_ASSET5_FN_I F
SET T_APPLDCVRG__RI D =
(SELECT MIN(BT. ROW_ID)
FROM dbo.S_APPLD_CVRG BT (INDEX = S_APPLD_CVRG_U2)
WHERE (BT. COVERAGE_CD = I T. CVRG_COVERAGE_CD AND
BT. TYPE = I T. CVRG_TYPE AND
BT. ASSET_ID = I T. T_APPLDCVRG_ASSETI AND
(BT. ASSET_CON_ID = I T. T_APPLDCVRG_ASSETC OR
(BT. ASSET_CON_ID IS NULL AND I T. T_APPLDCVRG_ASSETC IS NULL))) AND
(BT. I NSI TEM_ID = I T. T_APPLDCVRG_I NSI TE OR

```

```

(BT.INSITEM_ID IS NULL AND IT.T_APPLDCVRG_INSITE IS NULL)))
FROM dbo.S_ASSET5_FN_IF IT
WHERE (CVRG_COVERAGE_CD IS NOT NULL AND
CVRG_TYPE IS NOT NULL AND
T_APPLDCVRG_ASSETI IS NOT NULL AND
IF_ROW_BATCH_NUM = 10710001 AND
IF_ROW_STAT_NUM = 0 AND
T_APPLDCVRG__STA = 0)
SET STATISTICS PROFILE ON
GO
SET STATISTICS IO ON
GO
select
(SELECT MIN(BT.ROW_ID)
FROM dbo.S_APPLD_CVRG BT (INDEX = S_APPLD_CVRG_U2)
WHERE (BT.COVERAGE_CD = IT.CVRG_COVERAGE_CD AND
BT.TYPE = IT.CVRG_TYPE AND
BT.ASSET_ID = IT.T_APPLDCVRG_ASSETI AND
(BT.ASSET_CON_ID = IT.T_APPLDCVRG_ASSETC OR
(BT.ASSET_CON_ID IS NULL AND IT.T_APPLDCVRG_ASSETC IS NULL)) AND
(BT.INSITEM_ID = IT.T_APPLDCVRG_INSITE OR
(BT.INSITEM_ID IS NULL AND IT.T_APPLDCVRG_INSITE IS NULL))))
FROM dbo.S_ASSET5_FN_IF IT
WHERE (CVRG_COVERAGE_CD IS NOT NULL AND
CVRG_TYPE IS NOT NULL AND
T_APPLDCVRG_ASSETI IS NOT NULL AND
IF_ROW_BATCH_NUM = 10710001 AND
IF_ROW_STAT_NUM = 0 AND
T_APPLDCVRG__STA = 0)

```

With Hints

Table 'S_APPLD_CVRG'. Scan count 1, **logical reads 394774**, physical reads 0, read-ahead reads 280810.

Table 'S_ASSET5_FN_IF'. Scan count 1, logical reads 366, physical reads 0, read-ahead reads 0.

Without Hints

Table 'S_APPLD_CVRG'. Scan count 1268, **logical reads 10203**, physical reads 697, read-ahead reads 0.

Table 'S_ASSET5_FN_IF'. Scan count 1, logical reads 366, physical reads 0, read-ahead reads 0.

Related Topics

- [“Using the USE INDEX HINTS and USE ESSENTIAL INDEX HINTS Parameters” on page 132](#)
- [“Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Criteria for Passing Indexes to the Database” on page 135](#)

Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Criteria for Passing Indexes to the Database

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#). It explains how Siebel EIM determines which indexes to include on the hint clause passed to the database when using the USE INDEX HINTS and USE ESSENTIAL INDEX HINTS parameters.

See also:

- [“Using the USE INDEX HINTS and USE ESSENTIAL INDEX HINTS Parameters” on page 132](#)
- [“Using USE INDEX HINTS and USE ESSENTIAL INDEX HINTS: Example” on page 133](#)

When determining which indexes to pass on to the database as index hints, Siebel EIM takes the following steps:

- 1 Before generating a query, Siebel EIM makes a list of columns for which it has determined that an index is needed.
- 2 Siebel EIM then checks all of the indexes in the repository to find the index with the most matching columns.

Siebel EIM uses the following selection criteria in choosing indexes:

- Unique indexes have priority over nonunique indexes.
- Required columns have priority over nonrequired columns.

If a new index is created and it is declared in the repository, then there is a chance that Siebel EIM will choose it and pass it to the database on a hint.

Using the SQLPROFILE Parameter

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

The inclusion of this parameter greatly simplifies the task of identifying the most time-intensive SQL statements. By inserting the following statement in the header section of the .IFB file, the most time-intensive SQL statements will be placed in the file:

```
SQLPROFILE = c:\temp\ei msql .sql
```

Below is an example of the file eimsql.sql.

Start of the file: list of most time-intensive queries:

```
EIM: Integration Manager v6.0.1.2 [2943] ENU SQL profile dump (pid 430).
```

```
*****
```

```
Top 34 SQL statements (of 170) by total time:
```

```
Batch Step Pass Total Rows Per Row What
```

```
-----
```

```
106 10 401 1334.48 5000 0.27 update implicit primaries to child
```

```
106 9 114 242.56 5000 0.05 copy
```

...list of queries continues

Statistics by step and by pass:

```
*****
```

```
Statements per step by total time:
```

```
Step Stmt Total Min Max Avg %
```

```
-----
```

```
10 15 2627.27 0.00 1334.48 175.15 83.73
```

```
9 11 329.52 0.00 242.56 29.96 10.50
```

...list of statistics continues...

SQL statements:

```
*****
```

```
batch 106, step 10, pass 401: "update implicit primaries to child":
```

```
(total time 22:14m (1334s), 5000 rows affected, time/row 0.27s)
```

```
UPDATE siebel.S_CONTACT BT
```

```
SET PR_BL_PER_ADDR_ID =
```



```

(SELECT VALUE(MIN(ROW_ID), 'No Match Row Id')
FROM siebel.S_ADDR_PER CT
WHERE (CT.PER_ID = BT.ROW_ID)),
LAST_UPD = ?,
LAST_UPD_BY = ?,
MODIFICATION_NUM = MODIFICATION_NUM + 1
WHERE (ROW_ID IN (
SELECT T_ADDR_PER_PER_ID C1
FROM siebel.EIM_CONTACT
WHERE(
T_ADDR_PER_PER_ID IS NOT NULL AND
IF_ROW_BATCH_NUM = 106 AND
T_ADDR_PER__STA = 0 AND
T_ADDR_PER__EXS = 'N' AND
T_ADDR_PER__UNQ = 'Y' AND
T_ADDR_PER__RID IS NOT NULL)
GROUP BY T_ADDR_PER_PER_ID)
AND
(PR_BL_PER_ADDR_ID IS NULL OR PR_BL_PER_ADDR_ID = 'No Match Row Id'))
*****

```

...list of SQL statements continues...

Additional Indexes on Siebel EIM Tables

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

An examination of the data access path will assist you in determining whether additional indexes are necessary to improve the performance of the long-running SQL. In particular, look for table scans and large index range scans. In the following example, after evaluating the inner loop of the nested select, it was recommended to add an index on all T2 columns:

Inner loop:

```

(SELECT MIN(ROW_ID)
FROM siebel.EIM_ACCOUNT T2

```

```
WHERE (T2.T_ADDR_ORG__EXS = 'Y' AND
T2.T_ADDR_ORG__RID = T1.T_ADDR_ORG__RID AND
T2.IF_ROW_BATCH_NUM = 105 AND
T2.IF_ROW_STAT_NUM = 0 AND
T2.T_ADDR_ORG__STA = 0))
```

The index was created to consist of T2 columns used in the WHERE clause with ROW_ID at the end of the index. This influenced the database optimizer to choose this index for index-only access. Since the query wants the minimum (ROW_ID), the very first qualifying page in the index will also contain the lowest value.

NOTE: Having the ROW_ID column as the leading index column would also be a good strategy. Since the ROW_ID is unique, the index is likely to be more selective.

Adding Indexes to Improve Performance of S_ORG_EXT

The S_ORG_EXT table has indexes on many columns, but not all columns. If you have a large number of records (such as several million accounts) in S_ORG_EXT, then you can get a performance improvement in deleting and merging by adding an index to one or more of the following:

- PR_BL_OU_ID
- PR_PAY_OU_ID
- PR_PRTNR_TYPE_ID
- PR_SHIP_OU_ID

Before implementing any additional indexes, first discuss this with qualified support personnel.

Creating Proper Statistics on Siebel EIM Tables

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Use of the .IFB file parameter UPDATE STATISTICS is only applicable to IBM DB2 for z/OS. This parameter can control whether Siebel EIM dynamically updates the statistics of EIM tables. The default setting is TRUE. This parameter can be used to create a set of statistics on the EIM tables that you can save and then reapply to subsequent runs. After you have determined this optimal set of statistics, you can turn off the UPDATE STATISTICS parameter in the .IFB file (UPDATE STATISTICS is FALSE) thereby saving time during the Siebel EIM runs.

To determine the optimal set of statistics, you need to run several test batches and RUNSTATS commands with different options to see what produces the best results.

Before and after each test, execute the db2look utility in mimic mode to save the statistics from the database system catalogs. For example, if you are testing Siebel EIM runs using EIM_CONTACT1 in database SIEBELDB, then the following command generates UPDATE STATISTICS commands in the file EIM_CONTACT1_mim.sql:

```
db2look -m -a -d SIEBELDB -t EIM_CONTACT1 -o
EIM_CONTACT1_mim.sql
```

The file EIM_CONTACT1_mim.sql contains SQL UPDATE statements to update database system catalog tables with the saved statistics.

You can experiment with running test Siebel EIM batches after inserting the RUNSTATS commands provided in [“IBM DB2 for z/OS Options” on page 139](#). After you find the set of statistics that works best, you can apply that particular mim.sql file to the database.

NOTE: Do not forget to save statistics with db2look between runs.

IBM DB2 for z/OS Options

The syntax for IBM DB2 for z/OS commands provides more options, as follows:

- shrlevel change
- allow write access
- allow read access

The clauses allow read access and shrlevel change provide the greatest concurrency.

Dropping Indexes in Initial Runs of Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Typically, the Siebel EIM initial load is a very database-intensive process. Each row that is inserted into the base table requires modifications on the data page and the index pages of all the affected indexes. However, most of these indexes are never used during a Siebel EIM run. Index maintenance is very time-consuming for most database managers and must be avoided as much as possible.

Therefore, the goal is to determine any indexes that are unnecessary for Siebel EIM and that can be dropped for the durations of the Siebel EIM run. You can create these indexes later in batch mode by using parallel execution strategies available for the respective database platform. Using this approach can save a significant amount of time.

NOTE: Under normal operations, using parallel execution strategies is *not* recommended.

- **Target Table Indexing Strategy.** For a target base table (such as S_ORG_EXT) you only need to use the Primary Index (Px, for example P1), and the Unique Indexes (Ux, for example U1), and then drop the remaining indexes for the duration of the Siebel EIM import. Past experience has determined that the Fx and Mx indexes can be dropped after an extensive SQL analysis of sample Siebel EIM runs.

- **Nontarget Table Indexing Strategy.** For child tables (such as S_ADDR_ORG) you only need to use the Primary Index (Px), the Unique Indexes (Ux), and the Foreign Key Indexes (needed for setting primary foreign keys in the parent table). Past experience has determined that the Fx and Mx indexes can be dropped after an extensive SQL analysis of sample Siebel EIM runs.

NOTE: Always perform testing when dropping indexes (or adding indexes) to make sure that expected results are achieved.

Controlling the Size of Batches for Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

After tuning the long-running SQL statements, further tests can be run to determine the optimal batch size for each entity to be processed. The correct batch size varies and is influenced by the amount of buffer cache available. Optimal batch ranges have been observed to range anywhere between 500 and 15,000 rows. Run several tests with different batch sizes to determine the size that provides the best rate of Siebel EIM transactions per second. Using the setting Trace Flag = 1 while running Siebel EIM helps in this task because you are then able to see how long each step takes and how many rows were processed by the Siebel EIM process.

NOTE: Also monitor this throughput rate when determining degradation in parallel runs of Siebel EIM.

Recommended Number of Rows for a Single Batch

For an initial load, you can use 30,000 rows for a large batch. For ongoing loads, you can use 20,000 rows for a large batch. You must not exceed 100,000 rows in a large batch.

Furthermore, for Microsoft SQL Server and Oracle Database environments, limit the number of records in the EIM tables to those that are being processed. For example, if you have determined that the optimal batch size for your implementation is 19,000 rows per batch and you are going to be running eight parallel Siebel EIM processes, then you must have 152,000 rows in the EIM table. Under no circumstances can you have more than 250,000 rows in any single EIM table because this reduces performance.

The restrictions mentioned in the preceding example do not apply to IBM DB2 for z/OS environments. As long as an index is being used to access the EIM tables, the numbers of rows in the EIM tables does not matter in DB2 environments.

NOTE: The number of rows you can load in a single batch can vary depending on your physical computer setup and on which table is being loaded. To reduce demands on resources and improve performance, generally try to vary batch sizes to determine the optimal size for each entity to be processed. In some cases, a smaller batch size can improve performance. But for simpler tables such as S_ASSET, you might find that loads perform better at higher batch sizes than for more complex tables such as S_CONTACT.

Controlling the Number of Records in Siebel EIM Tables

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Determine the number of records that can reside at one time in an EIM table while still maintaining an acceptable throughput rate during Siebel EIM processing. One observed effect of increasing the number of records in an EIM table is reduced performance of Siebel EIM jobs. This is often caused by object fragmentation or full table scans and large index range scans.

NOTE: In an IBM DB2 for z/OS environment, EIM table size is not an important factor that impacts performance, because it is easy to correct table scans and nonmatching index scans. So a large number of records in an EIM table is not likely to reduce performance in a DB2 environment.

After addressing any object fragmentation and after the long-running SQL statements have been tuned, it is likely that you can increase the number of records that can reside in the EIM tables during Siebel EIM processing. When loading millions of records, this can result in a significant time savings because it reduces the number of times that the EIM table needs to be staged with a new data set.

When performing large data loads (millions of records) it is recommended that you perform initial load tests with fewer records in the EIM table. For example, while identifying and tuning the long-running SQL, you start with approximately 50,000 records. After tuning efforts are complete, you run additional tests while gradually increasing the number of records. For example you can incrementally increase the number of records to 100,000, then 200,000, and so on until you have determined the optimal number of records to load.

Using the USING SYNONYMS Parameter with Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

The USING SYNONYMS parameter controls the queries of account synonyms during import processing. This parameter is also related to the S_ORG_SYN table. When set to FALSE, this parameter saves processing time because queries that look up synonyms are not used. The default setting is TRUE. Set this parameter to FALSE only when account synonyms are not needed.

Using the NUM_IPTABLE_LOAD_CUTOFF Extended Parameter with Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Setting this extended parameter to a positive value will reduce the amount of time taken by Siebel EIM to load repository information. This is because when you set this parameter to a positive value, only information for the required EIM tables is loaded. For more information on this parameter, see *Siebel Enterprise Integration Manager Administration Guide*.

NOTE: While this parameter is especially important for merge processes, it can also be used for any of the other types of processes.

Here is an example of using this parameter while running on Microsoft Windows from the server command-line mode:

```
run task for comp eim server siebserver with config=account2.iffb,
ExtendedParams="NUM_I FTABLE_LOAD_CUTOFF=1", traceflags=1
```

Disabling the Docking: Transaction Logging Parameter for Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Typically, a disabled Docking: Transaction Logging setting is only used during initial data loads. The value of the system preference, Docking: Transaction Logging, is set from the System Preferences view within the Siebel application. This setting indicates whether or not the Siebel application logs transactions for the purpose of routing data to Siebel Mobile Web Clients.

The default for this parameter is TRUE. If there are no Siebel Mobile Web Clients, then you can set this system preference to FALSE. If you have Siebel Mobile Web Clients, then this parameter must be set to TRUE in order to route transactions to the Siebel Mobile Web Clients. However, during initial data loads, you can set this parameter to FALSE to reduce transaction activity to the Siebel docking tables. After the initial loads are complete, set the parameter back to TRUE.

NOTE: For incremental data loads, Docking: Transaction Logging must remain set to TRUE if there are mobile clients. If this setting is changed for incremental data loads, then you will need to perform a reextract of all of the mobile clients.

Disabling Database Triggers for Siebel EIM

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Disabling database triggers, by removing them through the Server Administration screens, can also help improve the throughput rate. This can be done by running the Generate Triggers server task with both the REMOVE and EXEC parameters set to TRUE. Be aware that components such as Workflow Policies and Assignment Manager will not function for the new or updated data. Also, remember to reapply the triggers after completing the Siebel EIM load.

Running Siebel EIM Tasks in Parallel

This topic is part of [“Troubleshooting Siebel EIM Performance” on page 131](#).

Running Siebel EIM tasks in parallel is the last strategy you want adopt in order to increase the Siebel EIM throughput rate. In other words, do not try this until all long-running SQL statements have been tuned, the optimal batch size has been determined, the optimal number of records to be processed at a time in the EIM table has been determined, and the database has been appropriately tuned. Before running tasks in parallel, check the value of the Maximum Tasks parameter. This parameter specifies the maximum number of running tasks that can be run at a time for a service. For more information about this parameter, see *Siebel System Administration Guide*.

NOTE: UPDATE STATISTICS must be set to FALSE in the .IFB file when running parallel Siebel EIM tasks on IBM DB2 for z/OS. Failure to do so can cause Siebel EIM tasks and executing RUNSTATS to take a longer time to complete. Also, when running parallel Siebel EIM tasks, deadlock and timeout will occur if UPDATE STATISTICS is set to TRUE in the .IFB file.

Database Guidelines for Optimizing Siebel EIM

This topic describes Siebel EIM tuning tips for the database platforms supported by Siebel Business Applications. It contains the following information:

- [“Microsoft SQL Server and Siebel EIM” on page 143](#)
- [“Oracle Database and Siebel EIM” on page 146](#)
- [“IBM DB2 for z/OS and Siebel EIM” on page 148](#)
- [“IBM DB2 for z/OS and Siebel EIM” on page 150](#)
- [“IBM DB2 for z/OS Loading Process for Siebel EIM” on page 151](#)
- [“General Recommendations for the IBM DB2 for z/OS Loading Process” on page 151](#)

Microsoft SQL Server and Siebel EIM

This topic is part of [“Database Guidelines for Optimizing Siebel EIM” on page 143](#).

The information that follows describes Siebel EIM tuning tips for the Microsoft SQL Server database platform.

Fixing Table Fragmentation

Table and index fragmentation occurs on tables that have many insert, update, and delete activities. Because the table is being modified, pages begin to fill, causing page splits on clustered indexes. As pages split, the new pages can use disk space that is not contiguous, hurting performance because contiguous pages are a form of sequential input/output (I/O), which is faster than nonsequential I/O.

Before running Siebel EIM, it is important to defragment the tables by executing the DBCC DBREINDEX command on the table's clustered index. This applies especially to those indexes that will be used during Siebel EIM processing, which packs each data page with the fill factor amount of data (configured using the FILLFACTOR option) and reorders the information on contiguous data pages. You can also drop and recreate the index (without using the SORTED_DATA option). However, using the DBCC DBREINDEX command is recommended because it is faster than dropping and recreating the index, as shown in the following example:

```
DBCC SHOWCONTIG scanning '**S_GROUPIF' table...
```

```
Table: '**S_GROUPIF' (731969784); index ID: 1, database ID: 7
```

```
TABLE level scan performed.
```

```

Pages Scanned.....: 739
Extents Scanned.....: 93
Extent Swi tches.....: 92
Avg. Pages per Extent.....: 7.9
Scan Densi ty [Best Count:Actual Count].....: 100.00% [93:93]
Logi cal Scan Fragmentati on .....: 0.00%
Extent Scan Fragmentati on .....: 1.08%
Avg. Bytes Free per Page.....: 74.8
Avg. Page Densi ty (full).....: 99.08%

DBCC executi on completed. If DBCC printed error messages, contact the system
admini strator.

```

To determine whether you need to rebuild the index because of excessive index page splits, look at the Scan Density value displayed by DBCC SHOWCONTIG. The Scan Density value must be at or near 100%. If it is significantly below 100%, then rebuild the index.

Purging an EIM Table

When purging data from the EIM table, use the TRUNCATE TABLE statement. This is a fast, nonlogged method of deleting all rows in a table. DELETE physically removes one row at a time and records each deleted row in the transaction log. TRUNCATE TABLE only logs the deallocation of whole data pages and immediately frees all the space occupied by that table's data and indexes. The distribution pages for all indexes are also freed.

Parallel Data Load for EIM tables Using bcp

Microsoft SQL Server allows data to be bulk copied into a single EIM table from multiple clients in parallel, using the bcp utility or BULK INSERT statement. Use the bcp utility or BULK INSERT statement when the following conditions are true:

- SQL Server is running on a computer with more than one processor.
- The data to be bulk copied into the EIM table can be partitioned into separate data files.

These recommendations can improve the performance of data load operations. Perform the following tasks, in the order in which they are presented, to bulk copy data into SQL Server in parallel:

- 1 Set the database option truncate log on checkpoint to TRUE using sp_dboption.(*)

- 2 Set the database option `select into/bulkcopy` to `TRUE` using `sp_dboption`.

In a logged bulk copy all row insertions are logged, which can generate many log records in a large bulk copy operation. These log records can be used to both roll forward and roll back the logged bulk copy operation.

In a nonlogged bulk copy, only the allocations of new pages to hold the bulk copied rows are logged. This significantly reduces the amount of logging that is needed and speeds the bulk copy operation. Once you do a nonlogged operation, immediately back up so transaction logging can be restarted.

- 3 Make sure that the table does not have any indexes, or if the table has an index, make sure it is empty when the bulk copy starts.
- 4 Make sure you are not replicating the target table.
- 5 Make sure the `TABLOCK` hint is specified using `bcf_control` with `eOption` set to `BCPHINTS`.

NOTE: Using ordered data and the `ORDER` hint will not affect performance because the clustered index is not present in the EIM table during the data load.

- 6 After data has been bulk copied into a single EIM table from multiple clients, any clustered index on the table must be recreated using `DBCC DBREINDEX`.

TempDB

This is the database that Microsoft SQL Server uses for temporary space needed during execution of various queries. Set the initial size of the `TEMPDB` to a minimum of 100 MB, and configure it for auto-growth, which allows SQL Server to expand the temporary database as needed to accommodate user activity.

Configuration Parameters

Additional parameters have a direct impact on SQL Server performance and must be set according to the following guidelines:

- **SPIN COUNTER.** This parameter specifies the maximum number of attempts that Microsoft SQL Server will make to obtain a given resource. The default settings are adequate in most configurations.
- **MAX ASYNC I/O.** This parameter configures the number of asynchronous inputs/outputs (I/Os) that can be issued. The default is 32, which allows a maximum of 32 outstanding reads and 32 outstanding writes per file. Servers with nonspecialized disk subsystems do not benefit from increasing this value. Servers with high-performance disk subsystems, such as intelligent disk controllers with RAM caching and RAID disk sets, can gain some performance benefit by increasing this value because they have the ability to accept multiple asynchronous I/O requests.
- **MAX DEGREE OF PARALLELISM.** This option is used to configure Microsoft SQL Server's use of parallel query plan generation. Set this option to 1 to disable parallel query plan generation. This setting is mandatory to avoid generating an unpredictable query plan.

- **LOCKS.** This option is used to specify the number of locks that Microsoft SQL Server allocates for use throughout the server. Locks are used to manage access to database resources such as tables and rows. Set this option to 0 to allow Microsoft SQL Server to dynamically manage lock allocation based on system requirements.
- **AUTO CREATE STATISTICS.** This option allows SQL Server to create new statistics for database columns as needed to improve query optimization. Make sure this option is enabled.
- **AUTO UPDATE STATISTICS.** This allows Microsoft SQL Server to automatically manage database statistics and update them as necessary to achieve proper query optimization. Make sure this option is enabled.

Oracle Database and Siebel EIM

This topic is part of [“Database Guidelines for Optimizing Siebel EIM” on page 143.](#)

The information that follows provides Siebel EIM tuning tips for Oracle Database.

Avoiding Excessive Table Fragmentation

Before running Siebel EIM, consult with an experienced DBA in order to evaluate the amount of space necessary to store the data to be inserted in the EIM tables and the Siebel base tables. Also, for example with Oracle Database, you can make sure that the extent sizes of those tables and indexes are defined accordingly.

Avoiding excessive extensions and keeping a small number of extents for tables and indexes is important because extent allocation and disallocation activities (such as truncate or drop commands) can be demanding on CPU resources.

To check if segment extension is occurring in Oracle Database

- Use the SQL statement that follows to identify objects with greater than 10 extents.

NOTE: Ten extents is not a target number for segment extensions.

```
SELECT segment_name, segment_type, tablespace_name, extents
FROM dba_segments
WHERE owner = (Siebel table_owner)
and extents > 10;
```

To reduce fragmentation, the objects can be rebuilt with appropriate storage parameters. Always be careful when rebuilding objects because of issues such as defaults or triggers on the objects.

Purging an EIM Table

When purging data from an EIM table, use the TRUNCATE command as opposed to the DELETE command. The TRUNCATE command releases the data blocks and resets the high water mark while the DELETE command does not, which causes additional blocks to be read during processing. Also, be sure to drop and recreate the indexes on the EIM table to release the empty blocks.

Disabling Archive Logging

It is recommended that Archive Logging be disabled during initial data loads. You can enable this feature to provide for point-in-time recovery after completing the data loads.

FREELIST Parameter

Multiple Siebel EIM processes can be executed against an EIM table provided they all use different batches or batch ranges. The concern is that you might experience contention for locks on common objects. To run multiple jobs in parallel against the same EIM table, make sure that the FREELIST parameter is set appropriately for the tables and indexes used in the Siebel EIM processing.

NOTE: If you are using Auto Segment Space Mgmt (ASSM) as part of defining table spaces, then the PCTUSED and FREELIST parameters (and FREELIST groups) are ignored.

Applicable database objects include EIM tables and indexes, as well as base tables and indexes. The value of the FREELIST parameter specifies the number of block IDs that will be stored in memory which are available for record insertion. Generally, you set the value to at least half of the intended number of parallel jobs to be run against the same EIM table (for example, a FREELIST setting of 10 permits up to 20 parallel jobs against the same EIM table).

This parameter is set at the time of object creation and the default for this parameter is 1. To check the value of this parameter for a particular object, the following query can be used:

```
SELECT SEGMENT_NAME, SEGMENT_TYPE, FREELISTS
FROM DBA_SEGMENTS
WHERE SEGMENT_NAME=' OBJECT NAME TO BE CHECKED ';
```

To change this parameter, the object must be rebuilt. Again, be careful when rebuilding objects because of issues such as defaults or triggers on the objects.

To rebuild an object

- 1 Export the data from the table with the grants.
- 2 Drop the table.
- 3 Recreate the table with the desired FREELIST parameter.
- 4 Import the data back into the table.
- 5 Rebuild the indexes with the desired FREELIST parameter.

Caching Tables

Another method to improve performance is to put small tables that are frequently accessed in cache. The value of BUFFER_POOL_KEEP determines the portion of the buffer cache that will not be flushed by the LRU algorithm. This allows you to put certain tables in memory, which improves performance when accessing those tables. This also makes sure that after accessing a table for the first time, it will always be kept in the memory. Otherwise, it is possible that the table will get pushed out of memory and will require disk access the next time used.

Be aware that the amount of memory allocated to the keep area is subtracted from the overall buffer cache memory (defined by DB_BLOCK_BUFFERS). A good candidate for this type of operation is the S_LST_OF_VAL table. The syntax for keeping a table in the cache is as follows:

```
ALTER TABLE S_LST_OF_VAL CACHE;
```

Updating Tables

When there are 255 or more NVL functions in an update statement, Oracle Database updates the wrong data due to hash keys overflow. This issue is specific to Oracle Database. To avoid this problem, use less than 255 NVL functions in the update statement.

IBM DB2 for z/OS and Siebel EIM

This topic is part of [“Database Guidelines for Optimizing Siebel EIM” on page 143](#). It describes Siebel EIM tuning tips for IBM DB2 for z/OS.

Review the following list of tuning tips for Siebel EIM:

- Use the IBM DB2 for z/OS load replace option when loading EIM tables.
NOTE: You can also use the IBM DB2 for z/OS load option to purge EIM tables. To do this, run the load option with an empty (null) input file in LOAD REPLACE mode. This purges the specified EIM tables instantly.
- Use separate tablespaces for EIM tables and the base tables.
- For large Siebel EIM loads or where many Siebel EIM tasks execute in parallel, place individual EIM tables in separate tablespaces.
- Use large page sizes for EIM tables and for the larger base tables. Previous experience has determined that a page size of 16 KB or 32 KB provides good performance. The larger page sizes allow more data to be fitted on a single page and also reduces the number of levels in the index B-tree structures.
- Similarly, use large extent sizes for both EIM tables and the large base tables.
- Make sure that the tablespace containers are equitably distributed across the logical and physical disks and across the input/output (I/O) controllers of the database server.
- Use separate bufferpools for EIM tables and the target base tables. Since initial Siebel EIM loads are quite large and there are usually no online users, it is recommended to allocate a significant amount of memory to the EIM table and base table bufferpools.
- After you load new data, reorganize the tables if the data on a disk is out of cluster. Use the REORGCHK command if the results of executing the RUNSTATS command indicate that clustering has deteriorated (clustering index is less than 80% clustered) and that a reorganization of tables is required. For more information about using the REORGCHK command, see 477378.1 (Article ID) on My Oracle Support. This document was previously published as Siebel FAQ 2072.
NOTE: Allocate time to conversion schedules to allow for the reorganization of tables and the gathering of statistics prior to allowing end users access a system containing new data.

- Use IBM DB2 for z/OS snapshot monitors to make sure performance is optimal and to detect and resolve any performance bottlenecks.
 - You can turn off logretain during the initial load. However, you must turn it back on before moving into a production environment.
- NOTE:** When logretain is enabled, you must make a full cold backup of the database.
- For the EIM tables and the base tables involved, alter the tables to set them to VOLATILE. This makes sure that indexes are preferred over table scans.
 - Executing Siebel EIM processes in parallel will cause deadlock and timeout on IBM DB2 for z/OS databases if multiple Siebel EIM processes attempt to update the same catalog table simultaneously. To avoid this, set the UPDATE STATISTICS parameter to FALSE in the Siebel EIM configuration file (.IFB file).
 - Executing UPDATE STATISTICS in each Siebel EIM process consumes significant database server resources. It is recommended that the database administrator updates statistics outside of the Siebel EIM process using the RUNSTATS command.
 - Consider the settings for IBM DB2 for z/OS registry values in [Table 8 on page 149](#).

Table 8. IBM DB2 for z/OS Registry Settings

Registry Value	Setting
DB2_CORRELATED_PREDICATES =	YES
DB2_HASH_JOIN =	NO
DB2_PARALLEL_IO =	""
DB2_STRIPPED_CONTAINERS =	When using RAID devices for table space containers

- Consider the settings for the IBM DB2 for z/OS database manager configuration parameters in [Table 9 on page 149](#):

Table 9. IBM DB2 for z/OS Database Manager Configuration Parameter Settings

Registry Value	Setting
INTRA_PARALLEL =	NO (can be used during large index creation)
MAX_QUERYDEGREE =	1 (can be increased during large index creation)
SHEAPTHRES =	100,000 (depends upon available memory, SORTHEAP setting, and other factors)

- Consider the settings for the IBM DB2 for z/OS database parameters in [Table 10 on page 150](#):

Table 10. IBM DB2 for z/OS Database Parameter Settings

Registry Value	Setting
CATALOGCACHE_SZ =	6400
DFT_QUERYOPT =	3
LOCKLIST =	5000
LOCKTIMEOUT =	120 (between 30 and 120)
LOGBUFSZ =	512
LOGFILESZ =	8000 or higher
LOGPRIMARY =	20 or higher
LOGRETAIN =	NO (only during initial Siebel EIM loads)
MAXLOCKS =	30
MINCOMMIT =	1
NUM_IOCLEANERS =	Number of CPUs in the database server
NUM_IOSERVERS =	Number of disks containing DB2 containers
SORTHEAP =	10240 (This setting is only for initial Siebel EIM loads. During production, set it to between 64 and 256.) The value you specify for SORTHEAP impacts the result of changing the value for SHEAPTHRES. For example, if SORTHEAP = 10000, then you can execute no more than 9 Siebel EIM batches if you set SHEAPTHRES = 100000. If executing concurrent Siebel EIM batches, then make sure to allocate sufficient physical memory so that memory swapping or memory paging do not occur.
STAT_HEAP_SZ =	8000

IBM DB2 for z/OS and Siebel EIM

This topic is part of [“Database Guidelines for Optimizing Siebel EIM” on page 143](#).

For DB2 for z/OS configuration settings, you can find a listing (from the JCL) of the Database Manager Configuration Parameters (DSNZPARM) in *Implementing Siebel Business Applications on DB2 for z/OS*.

Further IBM DB2 for z/OS information is provided in the following topics:

- [“IBM DB2 for z/OS Loading Process for Siebel EIM” on page 151](#)
- [“General Recommendations for the IBM DB2 for z/OS Loading Process” on page 151](#)

IBM DB2 for z/OS Loading Process for Siebel EIM

This topic is part of [“Database Guidelines for Optimizing Siebel EIM”](#) on page 143.

Figure 3 on page 151 illustrates the load process for IBM DB2 for z/OS. For more information, see *Siebel Enterprise Integration Manager Administration Guide*.

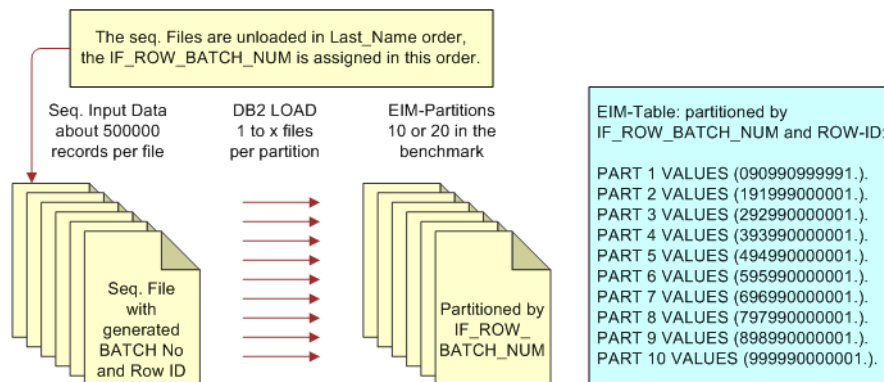


Figure 3. IBM DB2 for z/OS Loading Process for Siebel EIM

General Recommendations for the IBM DB2 for z/OS Loading Process

This topic is part of [“Database Guidelines for Optimizing Siebel EIM”](#) on page 143.

The following general recommendations apply when performing the IBM DB2 for z/OS loading process for Siebel EIM:

- Use the ONLY BASE TABLES and IGNORE BASE TABLES parameters or ONLY BASE COLUMNS and IGNORE BASE COLUMNS parameters in the .IFB files to reduce the amount of processing performed by Siebel EIM. By using the IGNORE BASE COLUMNS option, you allow foreign keys to be excluded, which reduces both processing requirements and error log entries for keys which cannot be resolved. Remember that the key words ONLY and IGNORE are mutually exclusive. For example, the following settings exclude the options IGNORE BASE TABLES and ONLY BASE COLUMNS:

ONLY BASE TABLES = S_CONTACT

IGNORE BASE COLUMNS = S_CONTACT.PR_MKT_SEG_ID

- Import parents and children separately. Wherever possible, load data such as accounts, addresses, and teams at the same time, using the same EIM table.

- Use batch sizes that allow all of the EIM table data in the batch to be stored in the database cache (approximately 2,000 records, 5000 for DB2 for z/OS and OS/390). Siebel EIM can be configured through the use of an extended parameter to use a range of batches, remember to put the variable name into the .IFB file.
- Multiple Siebel EIM processes can be executed against an EIM table, provided they all use different batches or batch ranges. However, the main limit to Siebel EIM performance is not the application server but the database. Contention for locks on common objects can occur if multiple Siebel EIM streams are executed simultaneously for the same base table. Multiple Siebel EIM job streams can run concurrently for different base tables, for example, S_ORG_EXT and S_ASSET.
- Run Siebel EIM during periods of minimum user activity, outside of business hours, if possible. This reduces the load for connected users and makes sure that the maximum processing capacity is available for the Siebel EIM processes.
- Set the system preference (in the Administration - Application screen, and System Preferences view) for Docking: Transaction Logging to FALSE during the initial database load. This reduces transaction activity to the Siebel docking tables, which are used for synchronizing mobile clients.
- Disable the database triggers by removing them through the Server Administration screens. Doing this can also help to improve the throughput rate. Remember to reapply the triggers after the Siebel EIM load has completed, because the lack of triggers will mean that components, such as Workflow Policies and Assignment Manager, will not function for the new or updated data.
- Remember to make sure that the required columns ROW_ID, IF_ROW_STAT, and IF_ROW_BATCH_NUM are correctly populated in the EIM table to be processed. The most efficient time to do this is when populating the EIM table from the data source or staging area, after cleansing the data.
- Unless there are specific processing requirements, make sure the EIM table is empty before loading data into it for Siebel EIM processing. Always make sure that suitable batch numbers are being used to avoid conflicts within the EIM table. If you are using an automated routine, then truncating the EIM table between loads from the data source helps to preserve performance.
- When running Siebel applications on an IBM DB2 for z/OS database, Siebel EIM can sometimes stop responding when updating the S_LST_OF_VAL base table. This is due to a data issue. The BU_ID column in the S_LST_OF_VAL base table must have only one or very few distinct values. That makes the DB2 optimizer perform a table scan through all rows in the S_LST_OF_VAL table when most or all rows have the same BU_ID column value.

To avoid this problem and speed up the query, modify the statistics data by running the following SQL statements:

```
update sysibm.sysindexes set firstkeycard=1000 where name='S_LST_OF_VAL_M2' ;

update sysibm.syscolumns set colcard = 1000 where tbname='S_LST_OF_VAL' and
name='BU_ID' ;
```

NOTE: Depending on the data with which you are working, you might need to run other SQL statements beforehand.

Data Management Guidelines for Optimizing Siebel EIM

The following recommendations apply when performing the Siebel EIM loading process:

- The Siebel EIM mapping chart shows that many of the EIM table columns derive their values not from legacy database fields but from unvarying literal strings. Avoid filling up the EIM tables with this type of information, because it slows down the movement of real legacy data from the EIM tables to the base tables.
- Siebel EIM offers an alternative method for populating base table columns with unvarying literal strings, namely by using the `DEFAULT COLUMN` statement. This approach allows you to specify default literals that must be imported into the base tables without having to retrieve them from the EIM tables.

For example, the Siebel EIM mapping chart shows Default Organization as the constant value for `CON_BU` in `EIM_CONTACT`, which in turn will move into `BU_ID` in `S_CONTACT`.

The same result can be achieved with the setting `DEFAULT COLUMN = CON_BU`, Default Value in the `.IFB` file. There are many other opportunities for moving literal strings from the EIM tables to the `.IFB` file.

Run Parameter Guidelines for Optimizing Siebel EIM

The following recommendations are for setting run parameters when performing the Siebel EIM loading process:

- Do not set `TRIM SPACES` to `FALSE`. Using the `TRIM SPACES` parameter causes trailing spaces to be stored in the Siebel base table. This can lead to inefficient use of disk space because Siebel Business Applications use `VarChar` on virtually all text columns longer than a single character. Setting `TRIM SPACES` to `FALSE` can also waste valuable bufferpool space for the tablespace data.
- Use either the `IGNORE BASE TABLES` parameter or the `ONLY BASE TABLES` parameter to limit the number of tables being inserted into or updated. The `ONLY BASE TABLES` parameter is preferable because the list is usually shorter and it is self-documenting. Using these parameters improves performance because it limits the number of tables Siebel EIM attempts to load and they also save space for tables that will not be used by the user interface.
- Use either the `IGNORE BASE COLUMNS` parameter or the `ONLY BASE COLUMNS` parameter to limit the number of tables being inserted into or updated. The `ONLY BASE COLUMNS` parameter is preferable because the list is usually shorter and it is self-documenting. Using these parameters improves performance because they limit the number of foreign keys Siebel EIM attempts to resolve.
- Set the `USING SYNONYMS` parameter to `FALSE` in the `.IFB` file. This logical operator indicates to Siebel EIM that account synonyms do not require processing during import, which reduces the amount of processing. Do not set the `USING SYNONYMS` parameter to `FALSE` if you plan to use multiple addresses for accounts. Otherwise, Siebel EIM will not attach addresses to the appropriate accounts.

- Suppress inserts when the base table is already fully loaded and the table is the primary table for an EIM table used to load and update other tables. The command format is `INSERT ROWS = table name, FALSE`.
- Suppress updates when the base table is already fully loaded and does not require updates such as foreign key additions, but the table is the primary table for an EIM table used to load and update other tables. The command format is `UPDATE ROWS = table name, FALSE`.

Monitoring the Siebel Server During a Siebel EIM Task

When monitoring the Siebel Server, the assumption is that you have allocated sufficient processor and memory resources for running the Siebel EIM task on the Siebel Servers and Siebel database servers.

If you are using Microsoft Windows Server 2003 as the operating system for the Siebel Server, then you can use the Microsoft Windows Performance Monitor to verify the amount of processor and memory being used by the hardware.

When using a supported UNIX or Linux operating system for the Siebel Server, you can use `vmstat` and `iostat` to verify how much processor and memory is being used by the hardware.

11 Tuning Siebel Remote

This chapter discusses tuning for Siebel Remote that can enhance performance. It contains the following topics:

- [About Siebel Remote on page 155](#)
- [Tuning Siebel Remote Server Components on page 156](#)
- [Tuning the Mobile Web Client in a Siebel Remote Deployment on page 160](#)

Related Book

Siebel Remote and Replication Manager Administration Guide.

About Siebel Remote

Siebel Remote allows Mobile Web Clients (typically operating remotely, in disconnected mode on a laptop) to connect to a Siebel Server and exchange updated data and files, a process known as synchronization. Siebel Remote supports mobile computing by allowing field personnel to share current information with members of virtual teams of other mobile and connected users across the organization.

Siebel Remote uses the following components to manage the exchange of data and files:

- Database Extract (alias DbXtract)
- Generate New Database (alias GenNewDb)
- Parallel Database Extract (alias PDbXtract)
- Synchronization Manager (alias SynchMgr)
- Transaction Merger (alias TxnMerge)
- Transaction Processor (alias TxnProc)
- Transaction Router (alias TxnRoute)

For more information about each of these components, see *Siebel Remote and Replication Manager Administration Guide*.

[“Tuning Siebel Remote Server Components” on page 156](#) discusses how you can configure some of these components to optimize the performance of your Siebel Remote deployment.

Tuning Siebel Remote Server Components

This topic describes how you can improve the performance of certain components on the Siebel Remote Server. It contains the following information:

- [“Increasing Throughput for the Database Extract and Parallel Database Extract Components” on page 156](#)
- [“Tuning the Transaction Router Component” on page 157](#)

Increasing Throughput for the Database Extract and Parallel Database Extract Components

This topic is part of [“Tuning Siebel Remote Server Components” on page 156](#). It provides tips to help you improve the throughput for the Database Extract and Parallel Database Extract components.

Use Parallel Database Extract for extracting a regional database or extracting remote users who have very large visibility presence. Some of the characteristics of Database Extract and Parallel Database Extract follow:

- Parallel Database Extract can support only one task at a time. You cannot run multiple instances of Parallel Database Extract simultaneously, as you can do with Database Extract.
- Parallel Database Extract takes advantage of servers that have multiple CPUs and RAID disc array by using multiple threads.
- Database Extract by default uses a data file type of Bulk Copy format (BCP) to export data. This can lead to a faster time in extracting and initializing than using the data file type DAT.

The following list describes tips to improve the throughput for the Database Extract and Parallel Database Extract components:

- Run multiple concurrent instances of Database Extract

You can increase throughput by running multiple concurrent instances of the Database Extract component. Each instance of the Database Extract component requires a temporary table. This table is called S_DOCK_INITM_N where N equals the value of the parameter TS Table Number (alias TSTableNum). TS Table Number specifies the number of the temporary table that serves the Database Extract component.

For example, if TS Table Number equals 1, then temporary table 1 (S_DOCK_INITM_1) serves the instance of the Database Extract component that is currently executing.

By default, 48 temporary tables are available for use. If you require additional tables, then you create them using Siebel Tools.

The recommended number of temporary tables to use depends on the database platform in use. For example:

- Oracle Database

The number of temporary tables that you use depends on the size of the shared pool that this database server can access. If the size of the shared pool is less than 300 MB, then it is recommended that you use one temporary table and execute one instance of the Database Extract component. If the size of the shared pool is greater than 600 MB, then using one temporary table for each instance of the Database Extract component can increase throughput.

- Microsoft SQL Server and IBM DB2

Use one temporary table for each instance of the Database Extract component that you execute. For example, if you execute 11 instances of the Database Extract component, then use 11 temporary tables.

- Stop Transaction Router component

Stop the Transaction Router component if you are running multiple instances of the Database Extract or Parallel Database Extract component in order to avoid table locking.

- Extract lists of users simultaneously

Separate users to be extracted into groups of 50 to 100 users per list and extract these lists of users simultaneously using the Database Extract and Parallel Database Extract components. These components extract the Enterprise visibility data for all users in the list once.

- Set the Truncate TS table parameter to TRUE

Setting this parameter (alias TruncateTSTable) to TRUE can improve performance, because deletions from S_DOCK_INITM_N tables are usually logged and can take a long time.

- Defragment tables

Defragment tables to which you intend to extract data and defragment S_DOCK_INITM_N tables.

For more information about the components and parameters described here, see *Siebel Remote and Replication Manager Administration Guide*. For more information about performance issues for the Database Extract component, see 477174.1 (Article ID) on My Oracle Support. This document was previously published as Siebel Troubleshooting Steps 15.

Tuning the Transaction Router Component

This topic is part of [“Tuning Siebel Remote Server Components” on page 156](#). It describes how to resolve or avoid performance issues for the Transaction Router component.

Performance issues for Transaction Router can arise from the following sources:

- Visibility-related transactions
- Docking rules and data distribution
- Slow-running queries

- Increasing Transaction Router throughput
- Setting the Id Db Size parameter

For more information about Transaction Router performance issues, consult the following documents:

- 476759.1 (Article ID) on My Oracle Support. This document, previously published as Siebel Troubleshooting Steps 8, describes how to diagnose and resolve Transaction Router performance issues.
- 477816.1 (Article ID) on My Oracle Support. This document, previously published as Siebel Troubleshooting Steps 38, describes how to monitor and manage the transaction backlog for a Siebel Remote implementation.

Visibility-Related Transactions

If you diagnose the root cause of the Transaction Router performance issue to be visibility-related transactions, then consider the following two possible solutions:

- Reextract all mobile users and regional nodes. For more information, see *Siebel Remote and Replication Manager Administration Guide*.
- Allow the Transaction Router component tasks to continue processing until they clear the backlog.

After the Transaction Router has processed all visibility-related transactions, the backlog will be processed more quickly. Starting additional Transaction Router tasks can also improve performance, but do not start more tasks than the Siebel Server or database engine can support.

Docking Rules and Data Distribution

If you diagnose the root cause of the Transaction Router performance issue to be docking rule related transactions, then it is advisable to request assistance from Global Customer Support. Create a service request (SR) on My Oracle Support. Alternatively, you can phone Global Customer Support directly to create a service request or get a status update on your current SR. Support phone numbers are listed on My Oracle Support.

When you log the SR, provide the following pieces of information:

- RDBMS trace of the Transaction Router task
- Transaction Router log files
- .dx files the Transaction Router is processing from the `SIEBEL_ROOT\siebsrvr\Docking\txnproc` directory
- The results from executing the visrule script

For more information about the visrule script, see 476759.1 (Article ID) on My Oracle Support. This document was previously published as Siebel Troubleshooting Steps 8.

Slow-Running Queries

If you diagnose the root cause of the Transaction Router performance issue to be slow-running queries, then consult your database administrator to determine the following:

- All indexes are present and valid on the tables involved in the poor performing queries. To determine if all indexes are valid and present, see *Siebel Data Model Reference* on My Oracle Support.
- Check if the tables and indexes involved in the poor-performing queries require defragmentation.

Increasing Transaction Router Throughput

The following factors can impact throughput for the Transaction Router component:

- Large batch sizes for Siebel Enterprise Manager (Siebel EIM)
It is recommended that, where possible, you reduce the size of the batch that Siebel EIM processes when it imports data. In addition, it is recommended that you log transactions to the Siebel File System rather than the Master Transaction Log (S_DOCK_TXN_LOG). For more information, see *Siebel Enterprise Integration Manager Administration Guide*.
- Large batch size for Siebel Assignment Manager
Where possible, reduce the batch file size for Siebel Assignment Manager, because large batch files can impact the performance of the Transaction Router component. For information, see *Siebel Assignment Manager Administration Guide*.

In both instances, you need to decide if an increase in throughput for the Transaction Router component is more important than a decrease in throughput for the Siebel EIM and Siebel Assignment Manager components before you make changes.

Setting the Id Db Size Parameter

Increasing the size of the visdata.dbf file has been shown to help performance in certain situations. This file is a cached file used by the Transaction Router and Transaction Processor components. The size of the visdata.dbf file is determined by the value of the parameter Id Db Size, which is available for both of these components.

The value for this parameter must be the same in both components. If the value of the Id Db Size parameter is increased in one component, such as to 204800 (200 MB), then it must be changed in the other component.

Tuning the Mobile Web Client in a Siebel Remote Deployment

This topic discusses how you can optimize the performance of a Mobile Web Client in a Siebel Remote deployment. It contains the following information:

- [“Optimizing Application Configuration File Parameters” on page 160](#)
- [“Guidelines for Optimizing Data Synchronization Between Siebel Mobile Web Client and Siebel Remote Client” on page 162](#)
- [“Choosing an Appropriate Routing Model” on page 162](#)

For additional information about performance tuning for Mobile Web Clients, see [Chapter 5, “Tuning Siebel Web Client.”](#)

Optimizing Application Configuration File Parameters

This topic is part of [“Tuning the Mobile Web Client in a Siebel Remote Deployment” on page 160](#). It discusses how you can modify values for the parameters specified in your Siebel application configuration file to optimize the performance of a Mobile Web Client.

DockTxnsPerCommit

The value of this parameter specifies the number of transactions that Siebel Remote applies to the local database before performing a commit. If you have an environment where a large number of transactions are constantly being created, then adjusting the value of DockTxnsPerCommit upwards might decrease the amount of time taken for initialization and synchronization tasks for large quantities of data. In such a scenario, test a variety of values (for example, 1000, 2000, 3000) and determine which value is best for your environment.

The DockTxnsPerCommit parameter appears in the [Local] section of your application configuration file. The default value is 500.

AutoStopDB

Make sure that AutoStopDB is set to FALSE so that the SQL Anywhere database engine continues to execute after the user exits the Siebel application. This reduces the time required to restart the Siebel application at a later time. If AutoStopDB is set to TRUE, then the SQL Anywhere database engine automatically closes down when you exit a Siebel application.

You set the AutoStopDB parameter in the [Local] section of your application configuration file. The default value is FALSE.

Allocating Memory to the SQL Anywhere Database Engine Cache

The amount of memory (especially cache) made available to the SQL Anywhere database engine is one of the major factors that can influence performance. The SQL Anywhere database engine uses memory for many purposes, but one of the major uses is to hold data that is accessed repeatedly, so that it does not have to retrieve data from the database each time it is needed.

You can configure how much memory is available to the cache by setting a value for the `-c` command-line option in the `ConnectionString` parameter of the `siebel.cfg` file. For example, the following entry:

```
-c15m -ch25m
```

allocates a minimum of 15 MB of memory to the cache. The value for the parameter `(ch)` indicates that the amount of memory allocated to the cache can be increased upwards to a maximum of 25 MB.

By default, the values for these parameters are expressed as a percentage of the total memory available. For example, the following entry:

```
-c5p -ch7p
```

specifies that a minimum of 5% of available memory be allocated to the cache memory and a maximum of 7%.

Allocating more memory to the memory cache of the SQL Anywhere database engine reduces the amount of memory that is available to other applications on the local computer.

As a guideline, use the difference between 80% of total computer memory and the amount of memory used by all applications on the computer during regular use. For example, if a local computer that has 512 MB of memory available uses 328 MB during regular use (after all applications including Siebel Business Applications are loaded), then you can allocate 82 MB of memory to the cache of the SQL Anywhere database engine.

Also conduct tests to determine an upper limit for the amount of memory that you can allocate to the SQL Anywhere database engine cache.

CAUTION: Do not increase the amount of memory allocated to the cache to a level that it results in paging. Paging is where memory utilization exceeds the total available memory and can cause reduced performance.

Sort Collation

Set the parameter `SortCollation` to binary to optimize the retrieval of data from the local database. The `SortCollation` parameter is not a default part of the application configuration file. You have to manually add it to the configuration file of your Siebel application. You set the value of `SortCollation` in the `[Local]` section of your application configuration file.

For more information about this parameter, see *Siebel System Administration Guide*. To determine the current status of `SortCollation`, see 477096.1 (Article ID) on My Oracle Support. This document was previously published as Siebel Alert 801.

Guidelines for Optimizing Data Synchronization Between Siebel Mobile Web Client and Siebel Remote Client

This topic is part of [“Tuning the Mobile Web Client in a Siebel Remote Deployment” on page 160](#). It lists some points that can help you optimize data synchronization between your Siebel Mobile Web Client and Siebel Remote Server. Note the following:

- Synchronize frequently

Synchronizing frequently reduces the number of transactions to transmit and commit for each synchronization session. The longer a user waits between synchronization sessions, the more data there is to send.
- Enable TrickleSync on the Siebel Mobile Web Client

Each time a Siebel Mobile Web Client connects to your Siebel Enterprise network, TrickleSync performs database synchronization. For more information, see *Siebel Remote and Replication Manager Administration Guide*.
- Use time-based filters to prevent sending data from server to client that is older than a specific date.

Choosing an Appropriate Routing Model

This topic is part of [“Tuning the Mobile Web Client in a Siebel Remote Deployment” on page 160](#).

One way to increase performance is to reduce the volume of data transmitted to the remote users. This can be best achieved by choosing the appropriate routing model.

Routing model choices provide a useful level of granularity in setting the size of a local database. In general, using appropriate routing models improves overall performance of a remote environment and helps decrease Database Extract times and increase Transaction Router throughput.

If none of the supplied routing models are appropriate, then contact Oracle's Application Expert Services, who will help you to develop a routing model that is appropriate to your environment. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

12 Tuning Customer Configurations

This chapter discusses how you can avoid common performance-related problems in Siebel Business Applications that stem from customer configuration done using Siebel Tools or Siebel scripting languages. It contains the following topics:

- [General Performance Guidelines for Customer Configurations on page 163](#)
- [Analyzing Generated SQL for Performance Issues on page 166](#)
- [Performance Guidelines for Siebel Scripting on page 175](#)
- [Performance Guidelines for Data Objects Layer on page 179](#)
- [Performance Guidelines for Business Objects Layer on page 185](#)
- [Performance Guidelines for User Interface Objects Layer on page 189](#)

Related Books

Siebel Tools Online Help

Configuring Siebel Business Applications

Using Siebel Tools

Siebel Developer's Reference

Siebel Object Types Reference

Siebel Object Interfaces Reference

Siebel eScript Language Reference

Siebel VB Language Reference

General Performance Guidelines for Customer Configurations

This topic provides some general guidelines for customer configuration using Siebel Tools.

Using your hardware resources optimally, and configuring your system appropriately, can help you to achieve your performance goals. Consider your resources and requirements carefully, and test and monitor system performance on a continual basis.

Siebel Business Applications architecture has been designed and tuned for optimal performance, making use of features such as database indexes, data caching, RDBMS cursors, efficient SQL generation, native database APIs, and so on. However, custom configurations can have various potential performance pitfalls, the impact of which might be amplified in environments with large databases and wide data distribution across servers. Follow guidelines presented here and in other documentation to avoid such problems.

The following are some miscellaneous configuration guidelines for maintaining optimal performance:

- **Avoid using sort specifications on nonindexed columns or joined columns.** For more information, see [“Managing Database Indexes in Sorting and Searching” on page 180](#) and other relevant topics.
- **Avoid the use of case insensitivity.** Use of case-insensitive queries can significantly increase the possibility of performance issues due to the additional complexity required at the database level to support case-insensitive database operations.

Prior to enabling case insensitivity, a thorough review of business requirements and performance criteria is highly recommended. In addition, if the feature is enabled, then conduct a performance test with a full copy of the production database. The severity of the performance impact increases with the complexity of the configuration and the size of the production database.

It is also recommended that Oracle's Application Expert Services be engaged to optimize the configuration and review requirements. Case insensitivity is a database platform constraint and must also be reviewed with the database platform vendor. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services. For more information about configuring case insensitivity for an application or for specified fields, see *Siebel Applications Administration Guide*.

- **Limit the use of case insensitivity for queries.** Case-sensitive searches perform better than case-insensitive queries. Siebel Business Applications are case-sensitive by default. You can enable case insensitivity either for the entire application or for specified fields. In general, the larger the database, and the larger the number of records returned by a case-insensitive query, the more that overall database performance is affected. Overall performance is also affected by the number of users who perform case-insensitive queries. End users can also force case-sensitive or case-insensitive queries. For more information about configuring case sensitivity for an application or for specified fields, see *Siebel Applications Administration Guide*.
- **Avoid overly complex user interface configuration.** In general, do not include a large number of applets per view (generally include no more than four applets), or a large number of fields per applet.
- **Limit the number of business components in a view.** An excessive number of different business components used in applets in a view can slow down the display of data upon entry into that view. This is because each of the applets must be populated with data.
- **Limit the number of virtual business components in a view.** Avoid using more than two virtual components in a single view.

- **Limit the number of fields in business components or applets.** There is no set limit on the number of fields in a business component, or number of list columns in a list applet. However, a business component with too many active fields will have degraded performance. Also, in some database systems it is possible to generate a query that is too large to be processed. See also [“Limiting the Number of Active Fields” on page 185](#).

In particular, reduce the number of fields displayed in the master applet on related views. The information is static and might not be necessary. Additional space will be available on the view for supporting data without users needing to scroll. (This will also provide a usability benefit.)

End users can reduce or increase the number of fields displayed in a list applet, by using the Columns Displayed menu option. However, it is best to provide an optimal default number of visible fields for each applet. It is also best to provide the minimum required total number of fields, including those that are hidden by default.

- **Limit the number of required fields.** Required fields are always retrieved from database queries. Consequently, limiting the number of required fields (fields for which the Required user property is TRUE) in your business components can improve performance. See also [“Limiting the Number of Active Fields” on page 185](#).
 - **Limit the number of records returned.** To limit the number of records returned for a business component, you can add a search specification to the business component or to applicable applets or links, or you can define a default predefined query on the view.
 - **Limit the number of joins, extension tables, and primary ID fields in a business component.** Joins degrade performance by causing an extra row retrieval operation in the joined table for each row retrieval in the main table. Extension tables and primary ID fields also use joins, although implied rather than explicitly defined, adding a row retrieval operation for each.
- The more joins, extension tables, and primary ID fields defined in a business component, the higher the number of row retrievals required in tables other than the main table, with a corresponding performance degradation.
- **Limit the use of Link Specification property in fields.** TRUE settings in the Link Specification property in fields might also slow performance. If TRUE, then the field's value is passed as a default value to a field in the detail business component through a link.

This is necessary if the master business component has a link relationship (in the current business object) with one or more detail business components, and these detail business components utilize the “Parent:” expression in the Pre Default Value, Post Default Value, or Calculated Value properties in any fields. The master business component must pass the field value to any detail records displayed.

As with the Force Active property, fields with the Link Specification property set to TRUE will be retrieved every time the business component is queried.

- **Use inner joins rather than outer joins.** Inner joins can be used for joined tables, with a resulting savings in overhead, provided you are guaranteed that all foreign key references are valid.

For example, when the join is from a detail business component to its master, you are guaranteed of the existence of the master. You can configure the join as an inner join by setting the Outer Join Flag property of the Join object definition to FALSE. This improves the performance of queries that use the join. In general, avoid using double outer joins.

- **Configure Cascade Delete appropriately for many-to-many links.** The Cascade Delete property in a Link object definition must be correctly configured for use in a many-to-many link, or the first insertion or deletion in the association applet will be abnormally slow. A link object definition used in a many-to-many relationship is one that contains a non-NULL value for the Inter Table property. The Cascade Delete property in such a link must be set to None.
- **Remove unneeded sort buttons.** Remove sort buttons from list columns in list applets where this capability is not required.
- **Reduce the need to scroll in a view.** Whenever possible, design views that do not require scrolling. (This will also provide a usability benefit.)
- **Provide tuned PDQs.** Provide tuned PDQs (predefined queries) that address most user requirements. Doing so reduces the likelihood of users creating undesirably complex queries. You might also provide guidance to end users on constructing appropriate queries.
- **Cache business services.** Cache business services that must be accessible at all times in a user session. To do this, set the Cache property to TRUE for each applicable Business Service object definition. Caching of business services has an impact on memory, as the services are cached per session. Make sure that only frequently accessed business services in a session are marked cacheable.
- **Avoid calculated fields that do Counts and Sums.** Reduce, where possible, the use of calculated fields that do Counts and Sums. If such fields are active, then they will cause performance degradation.

Related Topics

- [“Analyzing Generated SQL for Performance Issues” on page 166](#)
- [“Performance Guidelines for Siebel Scripting” on page 175](#)
- [“Performance Guidelines for Data Objects Layer” on page 179](#)
- [“Performance Guidelines for Business Objects Layer” on page 185](#)
- [“Performance Guidelines for User Interface Objects Layer” on page 189](#)

Related Book

Configuring Siebel Business Applications

Analyzing Generated SQL for Performance Issues

Performance troubleshooting is an iterative process. You need to consider performance implications during design and development. Note any changes to potentially troublesome areas, such as MVGs, business component sort and search specifications, joins, extension tables, or indexes. You then test the application to determine bottlenecks, using realistic data volumes and distribution in your test environment. Focus your testing effort on the slowest, most important, and most highly configured views.

If a performance problem is detected in testing or production, then your next step is to analyze the SQL statements being generated by Siebel applications. This is one of the most useful diagnostic tools available to you for performance analysis. For more information about analyzing generated SQL for performance issues, see:

- [“About Specifying SQL Logging and SQL Tagging for Siebel Application Object Manager Components” on page 167](#)
- [“Troubleshooting Poor Performing SQL at the Database Level Using Workload Tagging” on page 168](#)
- [“Specifying SQL Spooling in Siebel Developer Web Client” on page 172](#)
- [“Troubleshooting Performance Using SQL Trace Files” on page 172](#)
- [“Troubleshooting Performance Using SQL Query Plans” on page 173](#)

About Specifying SQL Logging and SQL Tagging for Siebel Application Object Manager Components

This topic is part of [“Analyzing Generated SQL for Performance Issues” on page 166](#).

SQL logging and SQL tagging are diagnostic tools that log and tag SQL statements generated and executed by Siebel Business Applications. When used in conjunction with SQL logging, SQL tagging can help administrators trace long-running or slow-performing SQL queries back to the user or action that triggered them. You can also use these tools to review generated SQL in a development or test environment after making configuration changes in Siebel Tools.

- SQL logging controls the level of SQL logging detail in a log file for an object manager-based component. To set SQL logging for a component, set the Object Manager SQL Log (alias ObjMgrSqlLog) event to an appropriate value. For example, specify the event value for a Siebel Application Object Manager component such as Call Center Object Manager (ENU).
- SQL tagging controls whether or not SELECT statements are tagged for an object manager-based component. SQL tagging also controls whether or not to log additional information, such as the name of a business component. SQL tagging is primarily used during the development phase. To set SQL tagging for a component, set the OM SQL Tagging (alias ObjMgrSqlTag) event to an appropriate value. By default, SQL tagging is disabled.

CAUTION: SQL tagging changes the format and content of the SQL SELECT statement by adding a bind, which inhibits the use of Oracle stored outlines, if they exist.

With SQL tagging, data that can be logged includes the server component name, server name, task ID, user ID, flow ID (with Siebel ARM ID), business object name, business component name, and view name. Which elements are logged depends on the level set for the event.

NOTE: SQL tagging requires that you have installed Siebel CRM version 8.1.1.1 or later.

For instructions on how to use SQL tagging and logging to log and tag generated SQL statements at the server component level, see *Siebel System Monitoring and Diagnostics Guide*.

Troubleshooting Poor Performing SQL at the Database Level Using Workload Tagging

This topic is part of [“Analyzing Generated SQL for Performance Issues”](#) on page 166.

Siebel workload tagging is a diagnostic tool that tags SQL statements generated and executed by Siebel Business Applications running on an Oracle database. This tool helps database administrators trace and troubleshoot poor performing SQL or excessive workload at the database level.

NOTE: Siebel workload tagging is available only for Siebel Business Applications running on an Oracle database.

When workload tagging is enabled, specified Siebel Application Object Manager-generated SQL statements (SELECT, INSERT, UPDATE, DELETE) are tagged with workload tagging attributes. The values for these attributes include:

- CLIENT_IDENTIFIER: component_name,servername,taskID,userID,FlowID:SarmID
- ACTION: View Name
- MODULE: Business Component Name

NOTE: The length of the previous attribute values is limited to 64 bytes for CLIENT_IDENTIFIER, 32 bytes for ACTION, and 48 bytes for MODULE. Truncation occurs when the length exceeds these limitations.

When poor performing SQL statements are suspected, the database administrator can use this tagging information to trace details about the generated SQL and the user that initiated it. Database administrators can then use Oracle Enterprise Manager to find the component task and log detail for more in-depth analysis of the performance issue. Alternatively, you can use SQL Plus or similar tools to query the Oracle Database V\$SESSION view to look up the CLIENT_IDENTIFIER, ACTION, and MODULE attributes by using SQL_ID of the SQL statement.

About Enabling and Disabling Workload Tagging

By default, workload tagging is disabled. You choose which SQL statements you want enabled for workload tagging by setting the log level of the OCI SQL Tagging event type for the applicable Siebel Application Object Manager server component as described in [Table 11](#).

Table 11. Log Levels for Enabling and Disabling Siebel Workload Tagging

Log Level	Description
0	Workload tagging is disabled.
1	Workload tagging is disabled. This is the default setting.
2	Workload tagging is enabled only for SELECT statements.
3	Workload tagging is enabled only for INSERT, UPDATE, and DELETE statements.
4	Workload tagging is enabled for SELECT, INSERT, UPDATE, and DELETE statements.

NOTE: The OCI SQL Tagging event type is available to all object manager-based components for supported languages, such as Call Center Object Manager (ENU), Sales Object Manager (DEU), eService Object Manager (FRA), and so on.

You can use either of the following ways to enable and disable workload tagging:

- “Enabling and Disabling Workload Tagging Using the Siebel Application” on page 169
- “Enabling and Disabling Workload Tagging Using the Siebel Server Manager” on page 170

Requirements for Enabling and Disabling Workload Tagging

Before enabling (or disabling) workload tagging, make sure the following requirements are met:

- Siebel Server is up and running.
- Oracle Enterprise Manager is configured for the database server. For information about Oracle Enterprise Manager, see the documentation on Oracle Technology Network at <http://www.oracle.com/technetwork/indexes/documentation/index.html>

Enabling and Disabling Workload Tagging Using the Siebel Application

Use the following procedure to set the log level of the OCI SQL Tagging event type for the applicable Siebel Application Object Manager server component using the Siebel application GUI.

To enable and disable workload tagging using the Siebel application

- 1 Navigate to the Administration - Server Configuration screen, then the Components view.
- 2 Select the appropriate Siebel Application Object Manager for which you want to enable workload tagging. For example, if the application you are using is Siebel Call Center for English, then select Call Center Object Manager (ENU).

3 Click the Events subview, and then select the OCI SQL Tagging event type.

4 Set the Log Level as described in [Table 11 on page 169](#).

After workload tagging is enabled, database administrators can use Oracle Enterprise Manager to diagnose and troubleshoot the problematic generated SQL.

Related Topics

[“Enabling and Disabling Workload Tagging Using the Siebel Application” on page 169](#)

[“Enabling and Disabling Workload Tagging Using the Siebel Server Manager” on page 170](#)

[“Using Workload Tagging to Troubleshoot Poor Performing SQL at the Database Level” on page 171](#)

Enabling and Disabling Workload Tagging Using the Siebel Server Manager

Use the Siebel Server Manager (srvrmgr program) to set the OCI SQL Tagging event type for the applicable Siebel Application Object Manager server component by way of the command-line interface. For more information about Siebel Server Manager, see *Siebel System Administration Guide*.

To enable and disable workload tagging using the Siebel Server Manager

1 Make sure the Siebel Application Object Manager server component for your Siebel application is running.

2 Start the srvrmgr program.

For information on starting the srvrmgr program, see *Siebel System Administration Guide*.

3 Do one of the following:

- To enable workload tagging, run the following command:

```
change evtloglvl OCI Sql Tag = loglevel for comp appobjmgr_lang
```

where:

- *loglevel* determines the types of SQL statements that are tagged (see [Table 11 on page 169](#)).
- *appobjmgr* is the Siebel Application Object Manager server component for your application, and *lang* is the three-letter identifier for the language specific to your environment.

For example, the following command enables workload tagging only for SELECT statements for Siebel Call Center for English:

```
change evtloglvl OCI Sql Tag = 2 for comp sccobjmgr_enu
```

- To disable workload tagging, run the following command:

```
change evtloglvl OCI Sql Tag = 0 for comp appobjmgr_lang
```

where:

- *appobjmgr* is the Siebel Application Object Manager server component for your application, and *lang* is the three-letter identifier for the language specific to your environment.

After workload tagging is enabled, database administrators can then use Oracle Enterprise Manager to diagnose and troubleshoot problematic generated SQL. For information on how to use Oracle Enterprise Manager for workload tagging, see [“Using Workload Tagging to Troubleshoot Poor Performing SQL at the Database Level” on page 171](#).

Using Workload Tagging to Troubleshoot Poor Performing SQL at the Database Level

Use the following procedure to diagnose and troubleshoot poor performing SQL at the database level using Oracle Enterprise Manager. For more information about Oracle Enterprise Manager, see the documentation on Oracle Technology Network.

To use workload tagging to troubleshoot poor performing SQL at the database level

- 1 Make sure workload tagging is enabled.
- 2 Log in to Oracle Enterprise Manager.
- 3 Navigate to the Database Control Top Activity page.
- 4 From the Top Sessions section, select the Top Sessions view, and then drill down on the Session ID you want to troubleshoot.
- 5 Click the General tab.
- 6 In the Client section, note the values in the Current Client ID field.

The format for this field is:

```
component_name, servername, taskID, userID, FlowID: SarmID
```

where:

- *component_name* is the server component alias, for example SCCObjMgr_enu.
- *servername* is the name of the Siebel Server from which the SQL originated.
- *taskID* is the task ID of the user who generated the query.
- *userID* is the login name of the user who generated the query.
- *FlowID* is the flow ID of the task.
- *SarmID* is the SARM ID of the task.

For example:

```
SCCObjMgr_enu, server1, 10485770, SADMI N, 000008814d7f1fe8: 163
```

7 In the Application section, drill down on the following fields for more details:

- Current Module is the business component associated with the SQL.
- Current Action is the view name associated with the SQL.

Specifying SQL Spooling in Siebel Developer Web Client

This topic is part of [“Analyzing Generated SQL for Performance Issues” on page 166](#).

Alternatively, after making configuration changes in Siebel Tools, you can spool the SQL that is generated by the Siebel application during runtime to troubleshoot configuration-related performance issues.

To spool the generated SQL into a trace file, start the Siebel application in the Siebel Developer Web Client (connecting to the Siebel Database) using the command-line option `/s sql_trace_file`. For more information about installing and running the Siebel Developer Web Client, see the *Siebel Installation Guide* for the operating system you are using.

The SQL trace file contains all of the unique SQL statements generated during the current session, and identifies the amount of time spent processing each one. The trace file can be opened in a text editor for examination after the session has ended. The SQL trace file, which is simply a text file holding the spooled SQL from the session, is overwritten during every new session.

You can specify the `/s sql_trace_file` option by modifying properties for the Start menu item or desktop shortcut from which the Siebel application is invoked. The following example shows a command line for spooling generated SQL from Siebel Call Center using the Siebel Developer Web Client:

```
"D:\Siebel\8.2\Client_1\bin\siebel.exe /c
D:\Siebel\8.2\Client_1\bin\enu\agent.cfg /s siebel_sql.txt"
```

If you do not specify a path, then the SQL trace file is created in the Siebel client root bin directory, such as:

```
"D:\Siebel\8.2\Client_1\bin"
```

- For the purpose of spooling SQL, you can create shortcuts for Siebel Developer Web Client to run customer applications such as Siebel eService. For example, the shortcut would point to the application configuration file `eservice.cfg`.
- You can also programmatically start and stop SQL spooling through the Siebel Object Interfaces by using the `TraceOn` and `TraceOff` methods on the `Application` object. For more information about these methods, see *Siebel Object Interfaces Reference*.

Troubleshooting Performance Using SQL Trace Files

This topic is part of [“Analyzing Generated SQL for Performance Issues” on page 166](#).

As described, you can generate SQL trace files related to your configuration changes, such as for a particular view you have configured. Analyze the contents of the SQL trace file to identify any possible performance issues.

As you look through the SQL trace file, be aware of factors such as:

- The number and complexity of SQL statements.
- Execution times for SQL statements. This is the SQL execution time plus the time it takes to return rows. It does not include time for client-side processing.
- Selection criteria in the WHERE clauses, indicating search specifications.
- Sorting criteria in the ORDER BY clauses, indicating sort specifications. (In general, it is better for a query to first filter data using WHERE clauses, in order to reduce the volume of data to be then sorted. Applying sorting criteria that match users' needs reduces the likelihood of users performing their own sort operations, which would require additional system resources.)
- The use of joins.

NOTE: If the same SQL statement is executed repeatedly, then the Siebel application displays the entire statement for the first query. For each subsequent iteration of the same query, only the bind variables are displayed. You can recognize a query that is repeated by the specific set of bind variables it uses.

SQL statements are displayed for all queries, including housekeeping queries. These are queries that are necessary for system operation, such as looking up the user's login to obtain responsibilities, and determining today's alarms in the calendar. You will also see queries to the S_LST_OF_VAL table to populate picklists. Queries that populate views are also present in the SQL trace file, and are easily distinguishable based on the tables they access.

Troubleshooting Performance Using SQL Query Plans

This topic is part of ["Analyzing Generated SQL for Performance Issues" on page 166](#).

If you identify a problematic query in the SQL trace file, then you can obtain more information about it using the database query tool provided with the RDBMS, such as SQL*Plus for Oracle Database.

Copy and paste the SQL statement from the trace file into the database query tool, execute the query against the Siebel Database, then generate a query plan. A query plan is a detailed reporting of various statistics about the query you executed. For an example of generating a query plan against an SQL Anywhere database, see ["Example of Obtaining Query Plan" on page 175](#).

Use query plans to check:

- The use of indexes
- The use of temporary tables
- The use of sequential table scans

Finally, compare your results with a standard application (that is, not custom-configured) in order to identify any potentially slow queries.

You can resolve many performance issues either by modifying search specifications or sort specifications, or by creating new indexes on the base table.

CAUTION: Only specially trained Oracle personnel can modify existing Siebel indexes. This restriction is enforced so that performance in other modules (such as Siebel EIM) is not adversely affected by any index modifications you make to improve query performance through the user interface. For more information, see [“Managing Database Indexes in Sorting and Searching”](#) on page 180.

Consider any potential performance implications before modifying search specification and sort specification properties for a business component. By spooling out the SQL into trace files, you can analyze which indexes are likely to be used when your application queries the business component through each applet.

Run your query plans against datasets that are comparable to the production dataset. You will not obtain useful results analyzing the performance of a query against a 30-record test dataset when the production database has 200,000 records.

You might find it useful to prioritize the views to examine, as follows:

- **First priority.** Views that are known to have the biggest performance bottlenecks.
- **Second priority.** Views that are accessed most frequently.
- **Third priority.** Views that are the most highly configured (as compared to the standard Siebel application).

Comparison with the standard Siebel application provides you with a benchmark for evaluation. It is often very useful to obtain a trace file from the standard Siebel application, following a preselected route through the views. Then you obtain a separate trace file from the custom-configured application, following the same route as closely as possible. The two trace files are compared, noting differences in the bullet items listed previously.

NOTE: When you review a query plan, keep track of the business object to which each query applies. You can tell where each new business object is being opened by searching for the S_APP_QUERY statement. The business object that was accessed is represented using the bind variable statements beneath the query.

Bind variables are the values that determine which records are brought back. The RDBMS substitutes the value of a bind variable into an SQL statement when the same SQL statement is being reused, generally in place of each occurrence of a question mark or series of question marks. For example, a business object bind variable is used in an S_APP_QUERY statement because the purpose of this statement is to open the business object.

Watch for the following indications of potential problems:

- Unnecessary fields are being accessed, especially ones not exposed in the user interface and not needed for calculated fields, or used for passing values to detail records.
- Unnecessary joins are occurring, particularly to tables that are not being accessed.
- Unnecessary multiple joins are being made to the same table. This can indicate duplicate join or Multi Value Link (MVL) object definitions, or joins using the same foreign key.
- Multiple short queries similar to the following:

```
... FROM  
SI EBEL. S_ADDR_PER T1
```

When a short query appears many times, this generally indicates that an MVG without a primary join is being accessed by a list applet. The system is running a secondary query for each master record to obtain its detail records. The secondary queries are the short queries appearing in the log file. This is usually your best diagnostic indicator of the need for a primary join.

When a short query appears only once, it indicates the same situation, but accessed in a form applet. In either case, the cure is a primary join, as explained in [“Using Primary ID Fields to Improve Performance” on page 187](#).

Example of Obtaining Query Plan

The following procedure shows an example of obtaining a query plan when running against a local SQL Anywhere database using the Siebel Mobile Web Client.

To obtain a query plan for an SQL statement in your trace file

- 1 Execute the Interactive SQL (dbisqlc.exe) program, located in the Siebel client installation directory (Siebel Mobile).
- 2 In order to analyze an SQL statement from the SQL trace file, copy the SQL statement and paste it into the Interactive SQL program's Command pane.
- 3 Replace bind variable references with the corresponding bind variable values.
- 4 Click the Execute button.

The query runs against the local SQL Anywhere database. The Statistics pane provides analysis information.

SQL Queries Against Database Data

The database that underlies Siebel applications can be queried to obtain information on a read-only basis.

CAUTION: Update queries must never be directly performed on the Siebel Database. All data manipulation and restructuring must be performed through Siebel Tools or through the Siebel application.

Performance Guidelines for Siebel Scripting

This topic provides performance guidelines for Siebel scripting using Siebel eScript or Siebel VB, or for using declarative alternatives in place of scripts. It contains the following information:

- [“Using Declarative Alternatives to Siebel Scripting” on page 176](#)
- [“Siebel Scripting Guidelines for Optimal Performance” on page 177](#)

Using Declarative Alternatives to Siebel Scripting

This topic is part of [“Performance Guidelines for Siebel Scripting” on page 175](#).

Often, customers use scripts for data validation, responses to data changes, or other purposes that might best be addressed through declarative means: by defining properties or specifying business service method invocation using Siebel Tools.

Scripting is often unnecessary and must be minimized or avoided because it can introduce performance problems, add risk and complexity, require greater maintenance, and duplicate functionality already available in Siebel Business Applications.

For example, the Validation field property, which allows for common VB expressions and comparison operators, can be used to perform field validation or string manipulation of data entered through the user interface or through Siebel Object Interfaces.

Expressions for the Validation property can include methods such as `LoginId()`, `LoginName()`, `LookupValue()`, `ParentFieldValue()`, `PositionId()`, `PositionName()`, `Today()`, and so on.

The Force Case field property can also be useful in a data-validation context, such as to ensure that personal names entered have initial capital letters. For more information on supported expressions and operators, see *Siebel Developer's Reference*.

Setting the Auto Primary property on MVL object definitions can also help you achieve results that you might otherwise use scripting for. For example, if your business requirement is to assign the first record in an MVL as the primary record (for example, primary address or primary owner), then set Auto Primary to the value Default. For more information about using Primary ID fields, see [“Using Primary ID Fields to Improve Performance” on page 187](#) and see *Configuring Siebel Business Applications*.

Scripting can be used in combination with declarative methods, such as to present customized error messages that guide users to enter data appropriately for each field subject to validation rules.

Functionality such as custom responses to data changes, which can often be handled through scripting, might best be addressed through declarative means. Such mechanisms, many of which can be used in combination, include:

- User properties on applets, business components, fields, controls, list columns, and other object definitions (for example: Required, Pre-Default, Post Default, Search Spec, Type Field, or Type Value)
- Siebel Workflow
- State model
- Siebel Personalization
- Run-time events
- Named methods
- Business services
- Visibility configuration

For more scripting guidelines, see *Configuring Siebel Business Applications*.

Siebel Scripting Guidelines for Optimal Performance

This topic is part of [“Performance Guidelines for Siebel Scripting” on page 175](#). It provides guidelines for appropriate use of Siebel scripting using Siebel eScript or Siebel VB. For more information about these and other guidelines, see:

- *Siebel eScript Language Reference*
- *Siebel VB Language Reference*
- *Siebel Object Interfaces Reference*
- *Configuring Siebel Business Applications*

The following are some guidelines for appropriate use of Siebel scripting:

- **Use declarative alternatives.** Generally, try all other possibilities before using scripting to accomplish a functional requirement. See also [“Using Declarative Alternatives to Siebel Scripting” on page 176](#).
- **Use browser scripts for simple client-side functions such as field validation.** Browser scripts are best used to perform simple procedural logic on the client side, such as performing field validation, or displaying blocking messages or alerts to users. Some such uses, particularly field validation, can reduce server round trips. Using more complex browser scripts, however, might reduce performance.

For example, using Set/Get Profile attribute calls, or invoking multiple business service methods, can require more server round trips and lead to performance problems. Adding extra functionality to scripts that display messages can have a similar effect.

NOTE: Setting the Immediate Post Changes field property has a similar effect on server round trips. Use this property only for constrained picklists and calculated fields that must be updated dynamically.

- **Do not return large result sets from server business services to browser scripts.** Browser scripts that invoke server scripts must return simple values or a single record, and must not return large result sets.
- **Minimize scripting on field-level or control-level events.** Field-level or control-level events are fired more often than most other types of events. Consequently, invoking scripts from such events can dramatically impact scalability. Avoid scripting frequent events, or simplify scripts on these events. Examples of such events include BusComp_PreGetFieldValue(), WebApplet_PreCanInvokeMethod(), and WebApplet_ShowControl().
- **Use simple scripts on applet-level and business component-level events.** Scripts written on events for applets or business components (for example, for Change Record events) must be very simple, because such events are fired often. Complex or I/O-intensive operations in such events will adversely affect performance.
- **Caching data in Siebel eScript scripts.** Executing the same SQL statements from various locations in a Siebel eScript script can generate an excessive number of script API calls and a redundant number of business component queries. In order to reduce the performance impact (assuming data does not change between invocations), you can cache a limited set of data within your scripts. (In some cases, you might not want to cache data at the script level, such as if the data that needs to be cached is too complex or too large.)

- **Declare your variables.** Declaring your variables and specifying their data type, as appropriate, can use less memory and improve performance.
- **Destroy any created objects when you no longer need them (Siebel eScript).** Theoretically, the Siebel eScript interpreter takes care of object cleanup, complex code involving many layers of object instantiation can in some cases cause the interpreter not to release objects in a timely manner. Destroying or releasing objects helps to minimize the impact on resources such as server memory.

Explicit destruction of Siebel objects must occur in the procedure in which they are created. To destroy an object in Siebel eScript, set it to NULL, or set the variable that contains it to another value. It is recommended that you destroy objects in reverse order of creation; that is, destroy child objects before you destroy parent objects.
- **Verify your script is defined on the appropriate method.** A script that is not defined on the right method might have a performance impact. For example, if special code needs to be run at the record level when an insert or update is done, then it is better to invoke a script from `BusComp_WriteRecord()` rather than `BusComp_SetFieldValue()`. The reason for this is that `SetFieldValue` events are fired much more often than `WriteRecord` events. Limit your use of specialized invocation methods.
- **Verify your script is implemented in the right view.** A script that is not implemented in the right view might cause significant performance impact. Verify that this script is implemented in the right place in the configuration, based on data manipulations, navigation requirements, and business requirements in general.
- **Avoid redundant repository object settings.** Do not perform unnecessary object validation. Each method invocation you perform has a performance cost. Details on this issue regarding field activation, for example, are provided below.
- **Use the `ActivateField()` method sparingly (Siebel eScript).** Do not activate a field if you will not use it. Use the `ActivateField()` method sparingly. Using this method increases the number of columns retrieved by a query, and can lead to multiple subqueries involving joins. These operations can use a significant amount of memory, and can degrade application performance.

Do not perform any unnecessary field activation (for fields that are already active). Each method invocation you perform has a performance cost.
 - Do not activate system fields, because they are already activated by default. Such fields include `Created`, `Created By`, `Updated`, and so on.
 - Do not activate any other fields that are already active. Check the `Force Active` field property in Siebel Tools to see if you need to activate it.
- **Use the `ExecuteQuery()` method sparingly (Siebel eScript).** Removing calls to execute a business component, using the method `ExecuteQuery()`, can yield significant performance benefit. It is better practice to use shared variables to share values of specific business component records across scripts than to separately invoke `ExecuteQuery()` in each script.
- **Use `SetSearchSpec()` method rather than `NextRecord()` method (Siebel eScript).** You can improve performance by using the `SetSearchSpec()` method to get a specific record, rather than using the `NextRecord()` method to go through a list of retrieved methods until a specific record is found.

- **Use ForwardOnly cursor mode (Siebel eScript).** Use the ForwardOnly cursor mode for ExecuteQuery() unless ForwardBackward is required. Using ForwardBackward uses a significant amount of memory, which can degrade application performance.
- **Use appropriate error handling.** Appropriate error handling can help maintain optimal performance. Although error handling is important, it also has a performance cost. For additional guidelines for using error handling in scripts, see 477766.1 (Article ID) on My Oracle Support. This document was previously published as Siebel Technical Note 514.
- **Avoid nested query loops.** Nested query loops can involve a large number of subqueries and can significantly impact performance. Use this technique very sparingly. Implement a nested query loop in the correct order in order to minimize the number of iterations. Be aware that a nested query loop can be invoked implicitly, depending on how your script is written.
- **Use the *this* object reference (Siebel eScript).** The special object reference *this* is eScript shorthand for "this (the current) object." Use it in place of references to active business objects and components.

For example, in a business component event handler, you must use *this* in place of ActiveBusComp(), usage of which can have a significant performance impact. Refer to the following example:

```
function BusComp_PreQuery()
{
  this.s.ActivateField("Account");
  this.s.ActivateField("Account Location");
  this.s.ClearToQuery();
  this.s.SetSortSpec( "Account(Descending), " +
    "Account Location(Descending)");
  this.s.ExecuteQuery();
  return (ContinueOperation);
}
```

- **Use the Switch construct (Siebel eScript).** The Switch construct directs the program to choose among any number of alternatives you require, based on the value of a single variable. Using this construct offers better performance than using a series of nested If statements, and is easier to maintain.
- **Use the Select Case construct (Siebel VB).** The Select Case construct directs the program to choose among any number of alternatives you require, based on the value of a single variable. Using this construct offers better performance than using a series of nested If statements, and provides other benefits.
- **Test your custom scripts.** Make sure your scripts are fully tested and optimized, and are no more complex than required to meet your business needs.

Performance Guidelines for Data Objects Layer

This topic describes performance guidelines for configuring selected elements in the data objects layer for optimal performance. It contains the following information:

- [“Multilingual LOVs Query and Cache Performance” on page 180](#)
- [“Managing Database Indexes in Sorting and Searching” on page 180](#)
- [“Reusing Standard Columns” on page 182](#)

Multilingual LOVs Query and Cache Performance

This topic is part of [“Performance Guidelines for Data Objects Layer” on page 179](#).

Multilingual List of Values (MLOV) fields are implemented below the business component level. Fields that point to MLOVs with enabled target columns return display values that match the current language setting for the session.

For display, the underlying language-independent code is converted to its corresponding display value using a Siebel application lookup. For searching and sorting, however, a database join to the list of values table (S_LST_OF_VAL) is performed. Make sure that any configuration directly involving the S_LST_OF_VAL table is compatible with your Siebel application MLOV functionality.

When a view with MLOVs is displayed for the first time, a separate query on the S_LST_OF_VAL table is made for each field that has an MLOV. The query obtains all the display values for that MLOV and writes the values to the LOV cache in memory. When the view is subsequently displayed during the same session, the values are obtained from the cache rather than by issuing another query.

NOTE: Displaying multiple records in a list applet that contains one or more MLOV fields will cause memory consumption to increase, and can produce poor performance. The problem manifests particularly when multiple fetches are performed against a given logical result set; that is, you scroll through records. It might also manifest when client-side export is performed to automate this behavior, or anytime the NextRecord method is invoked repeatedly on the business component. It is generally recommended to use MLOV fields sparingly in list applets, or to disable client-side export from list applets containing MLOVs.

For more information on configuring MLOVs, see *Configuring Siebel Business Applications* and *Siebel Global Deployment Guide*.

Managing Database Indexes in Sorting and Searching

This topic is part of [“Performance Guidelines for Data Objects Layer” on page 179](#).

A *database index* is a data structure in the RDBMS that is associated with a table. It provides references to all records in the table for quick lookup and filtering, and is sorted in a particular order for sorting in that order quickly. The Siebel Database Server uses an index to efficiently retrieve and sort the result set of a query.

Indexes provided in the Siebel Data Model are tuned for optimal performance of standard Siebel applications. When you add new business components with custom sorting or filtering requirements, you need to make sure that a database index is present that supports the requirement and delivers the result set efficiently. You might need to add new indexes.

You add indexes using the Index and Index Column object types. The index is added in the database as a result of its being created in Siebel Tools and database extensions being applied.

NOTE: The addition of custom indexes does not always improve performance and can reduce performance in some cases. The incremental value of an index depends in large part on the heterogeneity and distribution of the data.

When data is heterogeneous, all or most of the values are unique (such as with row ID values, which are unique). The less heterogeneous the data (that is, the more homogeneity or repeated instances of values), the less benefit the index offers relative to its costs.

For Boolean fields, indexes generally offer little value. Some performance benefit might be found when querying for the least commonly represented values. Little or no benefit is found when querying on more commonly represented values or values that are evenly distributed. Similar guidelines apply for other homogeneous data, such as fields that are constrained to a list of values.

Indexing generally improves performance of SELECT operations. However, it can significantly reduce performance for batch UPDATE and INSERT operations, such as are performed by Siebel EIM.

Discuss any custom index requirements with Oracle's Application Expert Services. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle's Application Expert Services.

Sort Specification

The Sort Specification property for a business component, picklist, or predefined query orders the records retrieved in a query, and serves as the basis for the ORDER BY clause in the resulting SQL issued. An index needs to be present that supports the order specified in the sort specification. Otherwise, the RDBMS engine physically sorts the entire result set in a temporary table.

The index needs to include the base columns for all of the fields, and to use them in the same order. There can be more columns specified in the index than are used in the sort specification, but the reverse is not true.

For example, the sort specification Last Name, First Name in the Contact business component is supported by at least one index on the S_CONTACT base table. One of these indexes is called S_CONTACT_U1, and it contains the LAST_NAME, FST_NAME, MID_NAME, PR_DEPT_OU_ID, OWNER_PER_ID, and CONFLICT_ID columns, in that order. If you wanted a sort specification that ordered contacts in first-name order, then you would need to create a custom index.

Do not sort on joined columns, because indexes cannot be used.

Search Specification

The Search Specification property for a business component, applet, link, or picklist selectively retrieves rows from the underlying table that meet the criterion specified in the property. The search specification is the basis for the WHERE clause in the resulting SQL issued. An index needs to be present that supports the criterion. Otherwise, the RDBMS might scan through all rows in the table rather than only those to be returned by the query.

The index needs to contain all the columns referenced by fields in the search specification.

In Sales Rep views such as My Accounts or where organization access control is implemented, if the user queries or sorts columns that are denormalized to the intersection table (for example, NAME and LOC in S_ORG_EXT), then performance is likely to be good. The Siebel application uses the intersection to determine visibility to records in the base table, and indexes can be used on the intersection table to improve performance. For related information, see [“Reusing Standard Columns” on page 182](#).

NOTE: If a query or sort includes columns that are not denormalized to the intersection table, then performance is likely to degrade, because indexes are not used.

Reusing Standard Columns

This topic is part of [“Performance Guidelines for Data Objects Layer” on page 179](#).

The architecture and data model of your application has been tuned for best performance. This optimization is achieved by using proper indexes, data caching, and efficient SQL generation, and also by denormalizing columns on certain tables. These denormalized columns are indexed so that the application can improve the performance of complex SQL statements by using these columns for search or sort operations instead of the columns of the original tables.

NOTE: Do not remap existing fields, especially those based on User Key columns, to other columns in the same table.

CAUTION: Do not use custom denormalized columns without the assistance of Oracle’s Application Expert Services. Denormalized columns can improve performance by allowing indexes to be placed directly on an intersection table, rather than on its master or detail table. However, if this is configured improperly, then the data in the denormalized column can become out of sync with its source. This can result in several problems ranging from inconsistent sorting to corrupt data. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle’s Application Expert Services.

Example: Reusing NAME and LOC in S_ORG_EXT Table

The columns NAME and LOC of the S_ORG_EXT table are denormalized into ACCNT_NAME and ACCNT_LOC in the S_ACCNT_POSTN table.

When sorting accounts by name and location in views where the Visibility Applet Type property is set to Sales Rep, the Siebel application uses the denormalized columns ACCNT_NAME and ACCNT_LOC of the S_ACCNT_POSTN table. Doing so allows the use of an index.

If the account name and location were stored in extension columns (for example, X_NAME and X_LOC), then these columns would have to be used for sorting instead of NAME and LOC. Even if these extension columns were indexed, the application could not use an existing index to create the necessary joins and sort the data, because the index is on S_ORG_EXT and not on S_ACCNT_POSTN. Therefore, the result would be a significant decrease in performance.

Query Plan for My Accounts View

The first SQL statement is generated by the standard My Accounts view. The query plan shows that the database uses numerous indexes to execute the statement.

```

SELECT
    T1. LAST_UPD_BY,
    T1. ROW_ID,
    T1. CONFLICT_ID,
    .
    .
    .
    T10. PR_EMP_ID,
    T2. DUNS_NUM,
    T2. HIST_SLS_EXCH_DT,
    T2. ASGN_USR_EXCLD_FLG,
    T2. PTNTL_SLS_CURCY_CD,
    T2. PAR_OU_ID
FROM
    SIEBEL. S_PARTY T1
        INNER JOIN SIEBEL. S_ORG_EXT T2 ON T1. ROW_ID = T2. PAR_ROW_ID
        INNER JOIN SIEBEL. S_ACCNT_POSTN T3 ON (T3. POSITION_ID = ?, 0.05)
    AND T2. ROW_ID = T3. OU_EXT_ID
        INNER JOIN SIEBEL. S_PARTY T4 ON (T4. ROW_ID = T3. POSITION_ID, 0.05)
        LEFT OUTER JOIN SIEBEL. S_PRI_LST T5 ON T2. CURR_PRI_LST_ID = T5. ROW_ID
        LEFT OUTER JOIN SIEBEL. S_INVLOC T6 ON T2. PR_FULFL_INVLOC_ID =
    T6. ROW_ID
        LEFT OUTER JOIN SIEBEL. S_ORG_EXT T7 ON T2. PAR_OU_ID = T7. PAR_ROW_ID
        LEFT OUTER JOIN SIEBEL. S_ORG_EXT_SS T8 ON T1. ROW_ID = T8. PAR_ROW_ID
        LEFT OUTER JOIN SIEBEL. S_INT_INSTANCE T9 ON T8. OWN_INST_ID =
    T9. ROW_ID
        LEFT OUTER JOIN SIEBEL. S_POSTN T10 ON T2. PR_POSTN_ID = T10. PAR_ROW_ID
        LEFT OUTER JOIN SIEBEL. S_USER T11 ON T10. PR_EMP_ID = T11. PAR_ROW_ID
        LEFT OUTER JOIN SIEBEL. S_ADDR_ORG T12 ON T2. PR_ADDR_ID = T12. ROW_ID
        LEFT OUTER JOIN SIEBEL. S_INDUST T13 ON T2. PR_INDUST_ID = T13. ROW_ID

```

```

LEFT OUTER JOIN SIEBEL.S_ASGN_GRP T14 ON T2.PR_TERR_ID = T14.ROW_ID
LEFT OUTER JOIN SIEBEL.S_POSTN T15 ON T3.POSITION_ID = T15.PAR_ROW_ID
LEFT OUTER JOIN SIEBEL.S_USER T16 ON T15.PR_EMP_ID = T16.PAR_ROW_ID
LEFT OUTER JOIN SIEBEL.S_ORG_SYN T17 ON T2.PR_SYN_ID = T17.ROW_ID
LEFT OUTER JOIN SIEBEL.S_ORG_BU T18 ON T2.BU_ID = T18.BU_ID AND
T2.ROW_ID = T18.ORG_ID
LEFT OUTER JOIN SIEBEL.S_PARTY T19 ON T18.BU_ID = T19.ROW_ID
LEFT OUTER JOIN SIEBEL.S_ORG_EXT T20 ON T18.BU_ID = T20.PAR_ROW_ID
WHERE
((T2.INT_ORG_FLG != 'Y' OR T2.PRTNR_FLG != 'N') AND (T3.ACCNT_NAME >= ?))
ORDER BY
T3.POSITION_ID, T3.ACCNT_NAME

```

Query plan :

```

T3(S_ACCNT_POSTN_M1), T2(S_ORG_EXT_P1), T1(S_PARTY_P1), T15(S_POSTN_U2), T10(S_POSTN_U2),
T4(S_PARTY_P1), T12(S_ADDR_ORD_P1), T13(S_INDUST_P1), T7(S_ORG_EXT_U3), T16(S_USER_U2),
T11(S_USER_U2), T17(S_ORG_SYN_P1), T6(S_INVLOC_P1), T5(S_PRI_LST_P1), T14(S_ASGN_GRP_P1),
T18(S_ORG_BU_U1), T19(S_PARTY_P1), T20(S_ORG_EXT_U3), T8(S_ORG_EXT_SS_U1), T9(se)

```

Query Plan for My Accounts View (Different ORDER BY Clause)

The second SQL statement generated in My Accounts, below, has a different ORDER BY clause. Even though the columns NAME and LOC of S_ORG_EXT are indexed, the database cannot use this index. Performance decreases from the use of a temporary table. The same behavior occurs if the ORDER BY clause uses the columns X_NAME and X_LOC instead of NAME and LOC.

The following example shows a different ORDER BY clause than the previous example query plan.

```

WHERE
((T2.INT_ORG_FLG != 'Y' OR T2.PRTNR_FLG != 'N') AND
(T3.ACCNT_NAME >= ?))
ORDER BY
T3.ACCNT_NAME, T3.POSITION_ID

```

Query plan : TEMPORARY TABLE

```

T3(S_ACCNT_POSTN_M1), T2(S_ORG_EXT_P1), T1(S_PARTY_P1), T15(S_POSTN_U2), T10(S_POSTN_U2),
T4(S_PARTY_P1), T12(S_ADDR_ORD_P1), T13(S_INDUST_P1), T7(S_ORG_EXT_U3), T16(S_USER_U2),
T11(S_USER_U2), T17(S_ORG_SYN_P1), T6(S_INVLOC_P1), T5(S_PRI_LST_P1), T14(S_ASGN_GRP_P1),
T18(S_ORG_BU_U1), T19(S_PARTY_P1), T20(S_ORG_EXT_U3), T8(S_ORG_EXT_SS_U1), T9(se)

```


Performance Guidelines for Business Objects Layer

This topic describes performance guidelines for configuring selected elements in the business objects layer for optimal performance. It contains the following information:

- [“Using Cache Data Property to Improve Business Component Performance” on page 185](#)
- [“Limiting the Number of Active Fields” on page 185](#)
- [“Guidelines for Using Calculated Fields” on page 186](#)
- [“Using Properties to Improve Picklist Performance” on page 187](#)
- [“Using Primary ID Fields to Improve Performance” on page 187](#)
- [“How the Check No Match Property Impacts Performance” on page 188](#)

Using Cache Data Property to Improve Business Component Performance

This topic is part of [“Performance Guidelines for Business Objects Layer” on page 185](#).

To cache on the Siebel Application Object Manager the content of a business component for subsequent use in the same user session, the property Cache Data property must be set to TRUE for the business component.

Setting Cache Data to TRUE is appropriate for semi-static data that might be subject to repetitive queries, but that is unlikely to change during the user session.

For some business components, Cache Data is set to TRUE by default. This is done, for example, for the Internal Product business component.

Cache Data must be FALSE for business components representing transactional data that can change within a user session. See also [“Using Properties to Improve Picklist Performance” on page 187](#).

Limiting the Number of Active Fields

This topic is part of [“Performance Guidelines for Business Objects Layer” on page 185](#).

Field object definitions are instantiated for each business component when the business component is instantiated, such as by a user navigating to a view containing an applet based on the business component. All such instantiated fields are included in the SELECT statements in generated SQL that is issued to the Siebel Database, even fields that are not represented in the user interface with a corresponding list column or other field control.

The set of fields that is instantiated includes those for which the Force Active property is set to TRUE. The Force Active setting of TRUE indicates to the system that it must obtain data for the field every time the business component is accessed, even if the field is not displayed in the current applet; this adds the field to the SQL query each time.

When Force Active is set to TRUE, there is an associated performance cost. Force Active affects performance more significantly when fields are based upon MVLs or joins, because the Siebel application has to create the relationships in the SQL query to retrieve data for these columns.

In most cases, the Force Active property is not required. In general, do not set Force Active to TRUE unless strictly necessary.

Use Force Active only when the field must be included in generated queries, but the field does not appear in the user interface.

Guidelines for Using Calculated Fields

This topic is part of [“Performance Guidelines for Business Objects Layer” on page 185](#).

Calculated fields provide a convenient way to access and display data in the user interface that is not directly stored in a table. However, calculated fields have a cost associated with them. Consequently, it is important to use them appropriately to fulfill your requirements, and not to misuse them.

Each calculated field is evaluated whenever the business component is queried to provide a value for the field. Extensive use of calculated fields, or usage in certain contexts, can impact performance. Some guidelines are as follows:

- Use calculated fields sparingly. Be sure there is a valid business case for their usage.
- Minimize the complexity of the expressions defined in your calculated fields.
- Minimize the use of calculated fields that perform Sum, Count, Min, or Max calculations, such as for detail records in an MVG business component. In particular, avoid using such fields in list applets, or in More Info form applets. The cost of using such expressions can be significant depending on the number of detail records.

Whenever data is totaled there are performance implications. It is important to limit the number of records being totaled. For example, totaling the line items in a Quote or Expense report is not resource-consuming. However, summing the expected revenue for all Opportunities is resource-consuming.

The latter occurs when you generate a chart. However, charts tend not to be generated frequently. Accessing the Opportunities list view for routine searches and data entry is done frequently.

CAUTION: Never put a `sum([MVfield])` in a list column. Doing so requires that a separate query be executed for each record in the list, which is a significant performance issue.

- Avoid defining calculated fields using complex expressions that provide different values depending on the current language.
- Avoid using a calculated field to directly copy the value of another field.
- Avoid including calculated fields in search specifications, particularly if the calculated fields use functions that are not supported by the underlying RDBMS.

- If the RDBMS supports the function, then it will have algorithms for performing the calculations efficiently and will return the calculated values with the result set. However, if functions such as EXISTS, Max, or Count are included, then multiple subqueries can be performed, impacting performance.
- If the function is *not* supported in the RDBMS, then the Siebel application might have to rescan the entire result set to perform the desired calculation, considerably increasing the time it takes to obtain the results of the query.

In the first case, the calculations can take place before the results are returned, while, in the second case, they have to be performed in memory (on the Siebel Application Object Manager or client).

NOTE: Even if the calculated field is supported at the RDBMS level, there can be other reasons why a search specification on a calculated field might result in poor performance, such as the lack of an index (for example, when using the LIKE function) supporting the search specification. See [“Managing Database Indexes in Sorting and Searching” on page 180](#).

Using Properties to Improve Picklist Performance

This topic is part of [“Performance Guidelines for Business Objects Layer” on page 185](#).

To cache the content of certain picklists for subsequent use in the same user session, the property Cache Data property must be set to TRUE for the PickList Generic business component. By default, this property is FALSE.

NOTE: Picklists based on PickList Generic display LOV data, which is unlikely to change during the user session and are thus suitable for caching. Picklists based on other business components display data that could change during a user’s session and is thus generally unsuitable for caching.

Also set the Long List property to TRUE for each applicable Pick List object definition. When Long List is TRUE, the focus is not maintained on the current picklist record, thus improving performance for picklists with many records. The default setting of Long List varies for each Pick List object definition.

Using Primary ID Fields to Improve Performance

This topic is part of [“Performance Guidelines for Business Objects Layer” on page 185](#).

MVGs configured without Primary ID fields require separate queries to display each parent record and each set of child records. For example, for a list applet that displays 10 records and two MVGs per record, a total of 21 queries would be required to populate the applet: one query to populate the parent records and 20 additional queries (two per parent record) to populate the MVGs. The number of queries executed is many times the number actually required.

You can avoid unnecessary queries by configuring a Primary ID field on the master business component. The Primary ID field serves as a foreign key from a parent record to one primary child record in the detail business component. This allows the application to perform a single query using a SQL join to display values for the parent record and the primary child record in the applet. In other words, it defers having to perform additional queries for the MVG until the user opens the MVG and displays a list of all child records.

List applets receive the most performance benefit from using Primary ID fields because list applets typically access a large number of records and each record can have one or more MVGs associated with it. The Primary ID field avoids having to submit queries for each MVG for every parent record.

Form applets can also benefit from Primary ID fields, even though in form applets only one parent record is accessed at a time. A Primary ID field allows the application to submit a single query for each new parent record displayed, rather than having to perform multiple queries for every MVG on the form applet. This can improve performance as the user moves from one record to another.

In some circumstances, configuring a Primary ID field is not desirable or feasible:

- When Microsoft SQL Server is being used, and the creation of the primary join would create a double-outer-join situation prohibited by the Microsoft software
- When the only purpose of the multi-value field is to sum detail record values

For information on how to configure Primary ID fields, see *Configuring Siebel Business Applications*.

How the Check No Match Property Impacts Performance

This topic is part of [“Performance Guidelines for Business Objects Layer” on page 185](#).

In most cases, the Check No Match property of a Multi Value Link object definition (used to implement Primary ID fields) must be set to FALSE. Setting the Check No Match property to TRUE could negatively impact performance, especially in situations where most parent records do not have child records defined in an MVG.

The Check No Match property defines whether a separate query must be used to populate an MVG when no child record is found through a primary join.

- When Check No Match is set to FALSE, the application does the following:
 - If a parent record's Primary ID field is invalid or has the value of NULL, then a secondary query is performed to determine if there are child records in the MVG. If there are no child records, then the Primary ID field is set to the value *NoMatchRowId*.
 - If a parent record's Primary ID field has the value *NoMatchRowId*, then the application does not perform a secondary query, because *NoMatchRowId* indicates that there are no child records in the MVG. Avoiding these extra SQL queries improves performance.

NOTE: *NoMatchRowId* is not a permanent setting; the Primary ID field can be updated after it is set to *NoMatchRowId*.

- When Check No Match is set to TRUE, a separate SQL query is executed for each parent record in which the primary join did not find a primary child record. Doing this ensures that the multi-value field does not appear blank unless there are no child records. But executing these extra SQL queries decreases performance.

It is appropriate to set the Check No Match property to TRUE in the following cases:

- When the multi-value group allows records to be added without having to go through the MVG. For example, account addresses might actually be inserted through the Business Address multi-value group on the Contact business component instead of the Account business component.
- When records can be added to a detail business component through Siebel EIM.

For more information about configuring Multi Value Link object definitions, see *Configuring Siebel Business Applications*.

Performance Guidelines for User Interface Objects Layer

This topic describes performance guidelines for configuring selected elements in the user interface objects layer for optimal performance. It contains the following information:

- [“Addressing Performance Issues Related to Grid Layout” on page 189](#)
- [“Maintaining Performance When Using Applet Toggles” on page 189](#)

Addressing Performance Issues Related to Grid Layout

This topic is part of [“Performance Guidelines for User Interface Objects Layer” on page 189](#).

The grid layout feature allows developers to create effective and usable form applets for Siebel views. However, performance can be adversely affected by certain applet design choices.

Typically, such performance problems relate to the alignment of user interface controls such as labels and fields, and stem from the total number of cells in the grid-based form applet, including spacer cells. Performance impact will depend on the number of user interface elements, the applet size, and other factors.

You can optimize user interface performance by:

- Making stacked sets of labels or fields the same width. Doing so can reduce the number of adjacent spacer cells you require.
- Aligning stacked sets of labels consistently.
- Making labels the same height as the adjacent fields.
- Eliminating horizontal or vertical spacer cells you deem unnecessary.

NOTE: Weigh all optional measures against possible usability concerns. Judicious use of spacing in your view layouts is generally appropriate for optimal usability.

For more information about using the grid layout, see *Configuring Siebel Business Applications*.

Maintaining Performance When Using Applet Toggles

This topic is part of [“Performance Guidelines for User Interface Objects Layer” on page 189](#).

Applet toggles are a useful feature where multiple applets based on different business components occupy the same location in a view. Which applet displays at one time depends on a field value in a parent applet (dynamic toggle) or on a user selection (static toggle).

Dynamic toggle applets are based on the same business component, while static toggle applets can be based on different business components.

In general, when configuring applet toggles for your Siebel application, particularly dynamic toggles, you can reduce memory and CPU usage for user application sessions by minimizing the number of applet toggles and fields per applet.

It is important to be aware of potential performance impact of using applet toggles, particularly dynamic toggles:

- When a user selects a record in a parent applet for a dynamic applet toggle, the business component and fields for all of the applet toggles are instantiated and cached in memory, and all of these fields are queried.

This query is used to populate other applet toggles that might be displayed when the user changes the relevant field value in the parent record. However, each time the user selects a different record in the parent applet, all of the fields in the toggle business component are required.

Also note that view layout caching is not performed for views containing dynamic applet toggles.

- When a user navigates to a view containing a static applet toggle, the business component and fields for the default displayed applet is instantiated and cached in memory, and these fields are queried. Other business components are instantiated and cached, and other queries performed, when the user navigates to the other applets in the toggle.

In each case, cached objects remain in memory until the user navigates to a different screen.

13 Tuning the Web Server Computer for All UNIX and Linux Platforms

This chapter describes tuning steps designed to improve the performance and scalability of your Siebel Enterprise Server installation. It contains the following topics:

- [Tuning Microsoft Windows for Enhanced Siebel Server Performance on page 191](#)
- [Tuning the Siebel Server for All UNIX and Linux Platforms on page 192](#)
- [Tuning the Web Server Computer for All Applicable UNIX and Linux Platforms on page 193](#)
- [Tuning the Siebel Web Server Extension for All UNIX and Linux Platforms on page 195](#)
- [Tuning an Apache Web Server for Applicable UNIX and Linux Platforms on page 195](#)
- [Tuning Siebel Business Applications for AIX on page 198](#)
- [Tuning Siebel Business Applications for HP-UX on page 202](#)
- [Tuning Siebel Business Applications for HP-UX on page 202](#)
- [Tuning Siebel Business Applications for Oracle Solaris on page 203](#)

Before doing any of the procedures in this chapter, you must have completed the minimum necessary steps described in the chapters about installing and configuring the Siebel Gateway Name Server and the Siebel Server contained in the *Siebel Installation Guide* for the operating system you are using. For additional information about tuning and monitoring, see *Siebel System Monitoring and Diagnostics Guide* and *Siebel System Administration Guide*.

NOTE: Settings provided in this appendix are based on a controlled lab environment using a standard Siebel application, such as Siebel Call Center for Siebel Industry Applications. The degree of performance gained by using these settings at your site depends on your implementation. Contact your vendor for additional tuning recommendations for your supported operating system platform.

Tuning Microsoft Windows for Enhanced Siebel Server Performance

This topic describes how you can configure settings for your Microsoft Windows operating system to optimize the performance of Siebel Business Applications.

Maximizing Data Throughput

For a server computer on which you install Siebel Enterprise Server software, changing the data throughput setting from *Maximize data throughput for file sharing* (default) to *Maximize data throughput for network applications* can result in the following benefits:

- Better Symmetrical Multi-Processing (SMP) scalability
- Improved networking performance

- Allocation of more physical memory for your Siebel applications

NOTE: Where the Siebel Database is on a server computer running MS SQL Server, the setting *Maximize data throughput for network applications* is set by default, and is also generally recommended for optimal performance. Whether to keep or change this default depends on how you are using the server computer.

For more information on these settings, see Microsoft's documentation.

Turning on the 4GT RAM Tuning Feature

You can expand the per-process address limit from 2 GB to 3 GB. This reduces the amount of physical RAM available to the operating system from 2 GB to 1 GB. The difference (1 GB) is allocated to your applications. This feature is referred to as 4GT RAM tuning. For information on how to configure this setting, see Microsoft's documentation.

NOTE: Each Siebel process (Siebel Application Object Manager) cannot use more than 2 GB of RAM.

Tuning the Siebel Server for All UNIX and Linux Platforms

For all Siebel Server computers running on supported UNIX platforms, setting the environment variables described in this topic can help you manage your server resources appropriately and stay within appropriate CPU-usage limits.

Environment Variable for Siebel Assert Creation

For Siebel Server computers or Web server computers, the environment variable `SIEBEL_ASSERT_MODE` determines whether assert files are created. With the default setting of 0, the creation of assert files is disabled, which conserves disk space and improves performance.

Set this variable to a nonzero value only if you are performing system diagnostics, and only in consultation with Oracle Global Customer Support. For more information about this variable, see *Siebel System Monitoring and Diagnostics Guide*.

Environment Variable for Operating System Resource Limits

Set the environment variable `SIEBEL_OSD_MAXLIMITS` using one of the following methods (define the variable in the applicable profile for the Siebel Server):

- C Shell:

```
setenv SIEBEL_OSD_MAXLIMITS 1
```

- Korn Shell or Bourne Shell:

```
SIEBEL_OSD_MAXLIMITS=1; export SIEBEL_OSD_MAXLIMITS
```

Setting this variable to 1 specifies that operating system maximum values for resources will apply. Such resources might include `coredumpsize`, `cputime`, `filesize`, `descriptors`, `maxmemory`, and others.

Environment Variable for Operating System Latches

Depending on the total number of tasks on the Siebel Server, you might need to set the environment variables described here in order to manage these loads. SIEBEL_OSD_NLATCH controls named latches and SIEBEL_OSD_LATCH controls unnamed latches. Latches, which are similar to mutexes (mutual exclusion objects), are used for communication between processes.

If SIEBEL_OSD_NLATCH and SIEBEL_OSD_LATCH are not defined, then the values are 5000 and 1000, respectively. If these values are sufficient for the total number of tasks on the Siebel Server, then you do not set these variables. Do not set these variables to values lower than their default values.

NOTE: Before changing these variables, stop the Siebel Server, then using the `stop_server` command, run the `cleansync` utility.

Set SIEBEL_OSD_NLATCH and SIEBEL_OSD_LATCH on the Siebel Server computer based on the following formulas (define the variables in the applicable profile for the Siebel Server):

- $\text{SIEBEL_OSD_NLATCH} = (7 \text{ times (cumulative MaxTasks for all components)}) \text{ plus } 1000$
- $\text{SIEBEL_OSD_LATCH} = 1.2 \text{ times (cumulative MaxTasks for all components)}$

Assume, for example, that you have enabled two multithreaded server components on the same Siebel Server: SCCObjMgr_enu and WfProcMgr. For SCCObjMgr_enu, MaxTasks = 700 and, for WfProcMgr, MaxTasks = 150. In this example, parameter values must be as follows:

- $\text{SIEBEL_OSD_NLATCH} = 6950 = (7 \text{ times } (700 \text{ plus } 150)) \text{ plus } 1000$
- $\text{SIEBEL_OSD_LATCH} = 1020 = 1.2 \text{ times } (700 \text{ plus } 150)$

NOTE: Although the formulas presented here are expected to work for most deployments, in certain cases you might need to increase (for example, double) the values of these variables, and then monitor system behavior with the adjusted values. (An error message such as SBL-OSD-00217: Error exceeded maximum number of latches can sometimes be generated even when the variables are set according to the recommended formulas.)

Tuning the Web Server Computer for All Applicable UNIX and Linux Platforms

This topic describes how to configure the thread stack size for a Siebel thread on the Web server computer. The Siebel Web Server Extension or plug-in is responsible for Siebel thread creation in the Web server.

NOTE: The topic describes only fine tuning the stack size of the Siebel thread. It is not applicable to Web server native threads.

The default thread stack size is 512 KB for the Siebel threads in all UNIX and Linux platforms supported for the Web server computer on which you install the Siebel Web Server Extension. (In previous Siebel CRM releases, the default stack size for Siebel threads was 1 MB on the Oracle Solaris and Linux platforms, and 512 KB on the HP-Itanium and AIX platforms.)

In many cases, the default thread stack size for Siebel thread might be larger than necessary. With a larger thread stack size, the httpd process consumes more resources, and thus fewer users can be supported for each httpd process than with smaller sizes.

If you determine that a smaller thread stack size for the Siebel thread in a Web server is more suitable for your Siebel applications, then you can set the SIEBEL_OSD_PTHREAD_STACK_SIZE environment variable at the operating system level to specify the size you require, in bytes.

Configuring the Thread Stack Size for Siebel Thread on the Web Server Computer (for All Web Server Platforms Except Oracle HTTP Server)

Use the following procedure to configure the thread stack size on the Web server computer (for all Web server platforms except Oracle HTTP Server).

To configure the thread stack size on the Web server computer (for all Web server platforms except Oracle HTTP Server)

- 1 Stop the Web server.
- 2 Open a new shell and execute a command similar to the following (this example sets the thread stack size to 64 KB):

C Shell

```
setenv SIEBEL_OSD_PTHREAD_STACK_SIZE 65536
```

Korn or Bourne Shell

```
export SIEBEL_OSD_PTHREAD_STACK_SIZE=65536
```

- 3 Start the Web server.

Configuring the Thread Stack Size on the Web Server Computer for Oracle HTTP Server

Use the following procedure to configure the thread stack size on the Web server computer for Oracle HTTP Server.

To configure the thread stack size on the Web server computer (for Oracle HTTP Server)

- 1 Navigate to the \$ORACLE_HOME/opmn/conf directory.
- 2 Open the opmn.xml file.
- 3 Locate SIEBEL_CODEPAGE.
- 4 Below the line containing SIEBEL_CODEPAGE, insert a line similar to the following (this example sets the thread stack size to 64 KB):

```
<variable id="SIEBEL_OSD_PTHREAD_STACK_SIZE" value="65536"/>
```

- 5 Save the opmn.xml file.

- 6 Restart the Web server.

Tuning the Siebel Web Server Extension for All UNIX and Linux Platforms

You must tune the Siebel Web Server Extension (SWSE) to run Siebel Business Applications on UNIX platforms.

To tune the SWSE for UNIX platforms

- 1 In the SWSE installation directory, navigate to the bin subdirectory.
- 2 Using a text editor such as vi, open the eapps.cfg file for editing.
- 3 Set the appropriate AnonUserName user names and passwords. This will depend on your user authentication strategy. For more information, see *Siebel Security Guide*.
- 4 Set GuestSessionTimeout to 60.

NOTE: This configuration is appropriate for application scenarios where users browse without logging in.

- 5 Restart the Web server for these changes to take effect.

Related Topic

- [“Tuning an Apache Web Server for Applicable UNIX and Linux Platforms” on page 195](#) and applicable Web server topics for particular operating systems.

Tuning an Apache Web Server for Applicable UNIX and Linux Platforms

Several of the Web servers supported for use with Siebel Business Applications are based on the Apache Web Server. These include Oracle HTTP Server, IBM HTTP Server, and HP Apache Web Server, which are supported on certain UNIX and Linux operating systems.

This topic describes recommended parameter settings that are optimized for scalability and performance on an Apache Web Server. You can further adjust these settings at your discretion to optimize the performance of your Web server. For Web server support details, see *Siebel System Requirements and Supported Platforms* on Oracle Technology Network.

For more information about tuning the Web server, see:

- [“Tuning the Siebel Web Server Extension for All UNIX and Linux Platforms” on page 195](#)
- Applicable Web server tuning topics for specific operating systems
- [“Specifying Static File Caching” on page 57](#)

The default setting for ThreadLimit for the Apache Web Server is 64, but this parameter can be set to a much higher number, subject to limitations from kernel settings. Related parameters include ThreadsPerChild, MaxClients, and ServerLimit.

Many parameter settings are subject to available memory and other resources. High values for parameters such as MaxClients or ThreadLimit can increase memory usage. Use lower values to reduce memory usage.

- ThreadLimit = 20000 is the maximum value supported by the Apache Web Server. You can set this parameter to a number your system supports.

NOTE: The ThreadLimit parameter must be executed before other parameters.

- ThreadsPerChild = Number of threads per child. Cannot exceed ThreadLimit.
- MaxClients = Maximum connection.
- ServerLimit = Typically, set this parameter to the value of MaxClients divided by the value of ThreadsPerChild. (If MaxClients and ThreadsPerChild are set the same, then set ServerLimit to 1.)

To configure parameters for an Apache Web Server

- 1 Using a text editor, set values for parameters in the worker.c section of the file `Web_server_install/conf/httpd.conf`, where `Web_server_install` is the root directory in which your Web server is installed. Set the parameter values as shown in the table below:

ThreadLimit	N where N is a value 1 or 1.2 times the maximum number of concurrent users (threads). The value for the applicable parameters can be set equal to or up to 1.2 times the number of concurrent users the Web server must support. The specific values to use for each applicable parameter depend on your organization's requirements.
ServerLimit	MaxClients divided by ThreadsPerChild
StartServers	1
MaxClients	ServerLimit multiplied by ThreadsPerChild
MinSpareThreads	1
MaxSpareThreads	N where N is a value greater than or equal to the sum of MinSpareThreads and ThreadsPerChild.
ThreadsPerChild	N where N is a value less than or equal to ThreadLimit.
MaxRequestsPerChild	0

By default, the Web server provides persistent connections. Consider setting the `KeepAlive` parameter to `Off` to prevent the Web server from providing persistent connections and to free connections for reuse. However, using this setting can incur a cost, because memory will be used to close and set up new TCP/IP connections. Alternatively, you can reduce the value for the `KeepAliveTimeout` parameter so that a thread is reused more quickly.

NOTE: For HP-UX, if you are not using multiplex sessions, then make sure the kernel parameter `max_thread_proc` is set to a number greater than two times N .

- 2 In the file `httpd.conf`, also set the following values:
 - The value for `ServerName` *must* match the Primary Internet Address you used in installing the SWSE.
 - Change the values for `User` and `Group` to a valid computer user and group:
 - Ideally, the user ID has no privileges that allow access to files other than those used by the Siebel application. This user, however, does have full access rights (read, write, execute) to the SWSE installation directory and its subdirectories.
 - It is recommended that the group be created specifically for running this server.
 - The value for `UseCanonicalName` is recommended to be set to `Off`. It *must* be set to `Off` if you are load-balancing your Web servers.

CAUTION: For security reasons, it is recommended not to use root for User or Group.

- If you are not using the CGI functionality of the Web server, then you might want to comment out the line that loads the CGI module. Doing so will make tracking Web server processes simpler, because there will be always one child process. The line is as follows:

```
LoadModule cgi_module modules/mod_cgid.so
```

- Set `MaxKeepAliveRequests` to 0.

Tuning Siebel Business Applications for AIX

This topic provides instructions for configuring and tuning Web servers, OS settings, and Siebel Enterprise Server components so you can run Siebel Business Applications on AIX. It contains the following information:

- [“Tuning the IBM HTTP Server for AIX” on page 198](#)
- [“Tuning the Siebel Server for AIX” on page 198](#)
- [“Tuning Kernel Settings for AIX” on page 200](#)

Tuning the IBM HTTP Server for AIX

This topic is part of [“Tuning Siebel Business Applications for AIX” on page 198](#). It describes recommended values for environment variables that are optimized for scalability and performance on IBM HTTP Server (IHS) Web server. You can further adjust these settings at your discretion to optimize the performance of your Web server. For more information about tuning the Web server, see:

- [“Tuning the Siebel Web Server Extension for All UNIX and Linux Platforms” on page 195](#)
- Applicable Web server tuning topics for particular operating systems
- [“Specifying Static File Caching” on page 57](#)

Tuning the Siebel Server for AIX

This topic is part of [“Tuning Siebel Business Applications for AIX” on page 198](#).

AIX provides several environment variables that can be tuned to optimize Siebel Server performance. These environment variables and their values are used as start parameters when the Siebel Server is started. [Table 12 on page 199](#) and [Table 13 on page 199](#) describe each of these environment variables and their recommended settings.

NOTE: For more information about tuning the Siebel Server, see [“Tuning the Siebel Server for All UNIX and Linux Platforms” on page 192](#).

Table 12. Environment Variables Used for Optimization in `$SIEBEL_ROOT/siebenv`

Environment Variable	Value	Description
AI XTHREAD_SCOPE	S	Controls contention scope. S signifies system-based contention scope (1:1).
AI XTHREAD_MNRATIO	1:1	Controls the M:N ratio of number of kernel threads that must be employed to handle runnable pthreads.
AI XTHREAD_MUTEX_DEBUG	OFF	Maintains a list of active mutexes for use by the debugger.
AI XTHREAD_RWLOCK_DEBUG	OFF	Maintains a list of read-write locks for use by the debugger.
AI XTHREAD_COND_DEBUG	OFF	Maintains a list of condition variables for use by the debugger.

Table 13. Environment Variables Used for Optimization in `$SIEBEL_ROOT/bin/siebmtshw`

Environment Variable	Value	Description
SPI NLOOPTI ME	1000	Controls the number of times to retry a busy lock before yielding to another processor.
YI ELDLOOPTI ME	4	Controls the number of times to yield the processor before blocking on a busy lock (only for libpthreads). Set this variable, at the minimum, to the number of CPUs.
MALLOCTYPE	buckets	<p>Malloc buckets provide an optional buckets-based extension of the default allocator. This feature improves malloc performance for applications that issue large numbers of small allocation requests.</p> <p>When malloc buckets are enabled, allocation requests that fall within a predefined range of block sizes are processed by malloc buckets. All other requests are processed in the usual manner by the default allocator.</p>

Table 13. Environment Variables Used for Optimization in `$SIEBEL_ROOT/bin/siebmtshw`

Environment Variable	Value	Description
MALLOCMULTIHEAP	heaps: <i>n</i>	Controls the number of heaps within the process private segment. <i>n</i> must be equal to the number of processors on the server.
LDR_CNTRL	LOADPUBLIC@IGNOREUNLOAD@MAXDATA=0x50000000	<p>The LOADPUBLIC option directs the system loader to load all modules requested by an application into the global shared library segment. Set LDR_CNTRL in the environment of the user, or, preferably, in the shell script that launches the executable needing the extra memory.</p> <p>NOTE: The IGNOREUNLOAD option is not supported in AIX 5.3 or later.</p>

Tuning Kernel Settings for AIX

This topic is part of [“Tuning Siebel Business Applications for AIX” on page 198](#).

There are several AIX kernel settings you can tune for optimal Siebel Server or Web server performance under AIX. These include the Virtual Memory Management and TCP settings. You must have root privileges to modify these settings. Use the `vmo`, `ioo`, and `no` commands to tune the AIX kernel. For more information about AIX kernel settings, see your operating system vendor's documentation.

To change the kernel settings

- 1 Using a text editor such as `vi`, open the `/etc/rc.net` file for editing.
- 2 Modify settings using `vmo`, as follows:

```
if [ -f /usr/sbi n/vmo ] ; then
    /usr/sbi n/vmo -o minperm%=5 -o maxperm%=30 -o minfree=720 -o maxfree=768
```

In providing values for `minfree` and `maxfree`, use the following formulas:

- $\text{minfree} = \text{number_of_CPUs} \times 120 = 6 \times 120 = 720$
- $\text{maxfree} = \text{number_of_CPUs} \times (120 \text{ plus } \text{maxpgahead}) = 6 \times (120 \text{ plus } 8) = 768$

where:

- *number_of_CPUs* is the number of CPUs on the AIX server you are tuning (for example, 6)
- *maxpgahead* is the value of the *maxpgahead* (`-o maxpgahead=number`) parameter (for example, 8)

NOTE: Although the default value for maxperm is 80, for a Siebel Business Applications deployment it must be set to a value between 30 and 50.

- 3 Modify settings using `ioo`, as follows:

```
if [ -f /usr/sbin/i oo ] ; then
    /usr/sbin/i oo -o numfsbuf=200 -o sync_release_i l ock=1
```

- 4 Modify the network options, as follows:

```
if [ -f /usr/sbin/no ] ; then
    /usr/sbin/no -a rfc1323=1
    /usr/sbin/no -a tcp_sendspace=24576
    /usr/sbin/no -a tcp_recvspace=24576
    /usr/sbin/no -a rfc2414=1
    /usr/sbin/no -a tcp_i ni t_wi ndow=3
    /usr/sbin/no -a use_i sno=0
    /usr/sbin/no -a tcp_nagl e_l i mi t=0
```

- 5 Check the settings for all User Limits (`ulimit`) and make sure that they are set to -1 (unlimited), as follows:

```
ul i mi t -a
```

NOTE: To change the set limits, update the `/etc/security/limits` file by changing all `ulimit` parameter values to -1 (unlimited).

- 6 Save your changes and exit the editor.

- 7 Make sure that the `rpc.statd` and `rpc.lockd` daemons run on the Siebel Server computer and on the server computer where the Siebel File System is located. Then set the number of threads for the `rpc.lockd` daemon on each applicable server computer.

It is recommended to increase the number of `rpc.lockd` daemon threads from the default. If possible, use the maximum number of threads, which is 511. System degradation can occur and logins might be blocked if the `rpc.lockd` daemon is not configured to handle a large number of lock requests. For example, you might execute commands like this:

```
chssys -s rpc.lockd -a 511
stopsrc -s rpc.lockd; startsrc -s rpc.lockd
```

- 8 Restart the server computer to have the new settings take effect.

Tuning Siebel Business Applications for HP-UX

This topic provides instructions for configuring and tuning OS settings and Siebel Enterprise Server components so you can run Siebel Business Applications on HP-UX. It contains the following information:

- [“Tuning Kernel Settings for HP-UX” on page 202](#)

For information about tuning the Web server, see:

- [“Tuning the Siebel Web Server Extension for All UNIX and Linux Platforms” on page 195](#)

NOTE: The functionality described in this topic requires that you install Siebel CRM version 8.2.2.3 or later. For details, see the applicable *Siebel Maintenance Release Guide* on My Oracle Support.

Tuning Kernel Settings for HP-UX

This topic is part of [“Tuning Siebel Business Applications for HP-UX” on page 202](#).

Modify the HP-UX kernel parameters to values like those shown below (suggested guidelines). Use the HP-UX System Administration Manager (SAM) tool to make these changes.

nproc	4096 - 4096
ksi_allloc_max	32768 - (NPROC*8)
max_thread_proc	4096 - 4096
maxdsiz	0x90000000 - 0X90000000
maxdsiz_64bit	2147483648 - 2147483648
maxfiles	4000 - 4000
maxssiz	401604608 - 401604608
maxssiz_64bit	1073741824 - 1073741824
maxtsiz	0x40000000 - 0X40000000
msgmap	4098 - (NPROC+2)
msgmni	4096 - (NPROC)
msgtql	4096 - (NPROC)
ncsize	35840 - (8*NPROC+2048+VX_NCSIZE)
nfile	67584 - (16*NPROC+2048)
ninode	34816 - (8*NPROC+2048)
nkthread	7184 - (((NPROC*7)/4)+16)
nproc	4096 - 4096
nsysmap	8192 - ((NPROC)>800?2*(NPROC): 800)
nsysmap64	8192 - ((NPROC)>800?2*(NPROC): 800)
semnmi	1024 - 1024
semnms	16384 - ((NPROC*2)*2)
semnmu	2048 - 2048
semume	256 - 256
shmmax	0x40000000 Y 0X40000000
shmmni	1024 - 1024
shmseg	1024 Y 1024
vps_ceiling	64 - 64

Setting Permissions for the HP-UX Scheduler

This topic is part of [“Tuning Siebel Business Applications for HP-UX” on page 202](#).

Siebel Business Applications will have better performance on HP-UX if you make the following changes, which allow the Siebel Server to execute the HP-UX scheduler upon startup. You must have root privileges to make these changes.

To set permissions for the HP-UX scheduler

- 1 Add the following line to the `/etc/privgroup` file, creating it if necessary:

```
-g RTSCHED
```

- 2 Save the file and exit.

- 3 Execute the following command:

```
setpri vgrp -f /etc/pri vgroup
```

- 4 Verify that global RTSCHED permissions are set by executing the following command:

```
getpri vgrp
```

If the command is successful, then the system will respond:

```
global pri vi leges: RTSCHED
```

Tuning Siebel Business Applications for Oracle Solaris

This topic provides instructions for configuring and tuning Web servers, OS settings, and Siebel Enterprise Server components so you can run Siebel Business Applications on Oracle Solaris. It contains the following information:

- [“Tuning the Oracle iPlanet Web Server” on page 203](#)
- [“Tuning Kernel Settings for Oracle Solaris” on page 205](#)
- [“Tuning Siebel Application Object Manager Instances for Oracle Solaris” on page 205](#)

Tuning the Oracle iPlanet Web Server

This topic is part of [“Tuning Siebel Business Applications for Oracle Solaris” on page 203](#).

If you have a busy Web server, then some of your users might experience difficulty connecting to your Web server. To address this issue, change the `tcp_conn_req_max_q` and `tcp_conn_req_max_q0` default values, using the `ndd` command. For details on how to use the `ndd` command, see [“Tuning Siebel Application Object Manager Instances for Oracle Solaris” on page 205](#).

Also, tune the Oracle iPlanet Web Server for optimal performance using the following procedure. For more information about tuning the Web server, see [“Specifying Static File Caching” on page 57](#).

To tune the Oracle iPlanet Web Server

- 1 Using a text editor such as vi, open the file *Web_server_Root/config/magnus.conf*, where *Web_server_Root* is the root path of the Oracle iPlanet Web Server.
- 2 Set the parameter RqThrottle to 1024 or a higher value, depending on computer resources such as memory and CPU.

The RqThrottle parameter specifies the maximum number of simultaneous transactions the Web server can handle. The default value is 512. By changing this value to 1024 or a higher value, you can reduce or minimize latencies for the transactions that are performed.

- 3 Add or modify the MaxKeepAliveConnections parameter, setting its value to 1000. The default value is 200.
- 4 Save your modifications to the magnus.conf file.
- 5 Restart the Web server.

After making the changes to the Oracle iPlanet Web Server parameters, change the following parameters on the workstation hosting the Oracle iPlanet Web Server.

- 6 Open the */etc/system* file for editing.
- 7 Set the Oracle Solaris system parameters as shown in the table below:

Parameter	Scope	Default Value	Tuned Value	Comments
rlim_fd_max	/etc/system	1024	65536	Process open file descriptors limit. The tuned value shown is a starting point. A higher value could be used, depending on the expected load (for associated sockets, files, and pipes, if any).
rlim_fd_cur	/etc/system	64	65536	The tuned value shown is a starting point. A higher value could be used, depending on the expected load (for associated sockets, files, and pipes, if any).
kernel_cage_enable	/etc/system	0	1	Enables the kernel cage.

- 8 Restart the workstation hosting the Oracle iPlanet Web Server.

Tuning Kernel Settings for Oracle Solaris

This topic is part of [“Tuning Siebel Business Applications for Oracle Solaris” on page 203](#).

To run Siebel Servers or Web servers in an Oracle Solaris environment, you need to set Oracle Solaris kernel parameters to specific recommended values for particular releases of Oracle Solaris servers. To learn the specific parameter recommendations for Siebel Servers or Web servers running on Oracle Solaris, contact Oracle’s Application Expert Services. Contact your Oracle sales representative for Oracle Advanced Customer Services to request assistance from Oracle’s Application Expert Services.

Several Oracle Solaris kernel parameter settings significantly affect performance of Siebel Business Applications in general, and the Siebel Server in particular.

Oracle Solaris kernel parameters reside in the configuration file `/etc/system`. To change the settings for these parameters, you must manually edit this file, save your changes, and reboot the system.

Normally, the Oracle Solaris kernel memory parameter settings are relatively low. However, for large memory-model applications like the Siebel Server applications, it is recommended that you increase the values assigned to several of these parameters.

CAUTION: If you use the default Oracle Solaris kernel parameters, or lower, to run a Siebel Server in an Oracle Solaris environment, then there is a risk of serious performance problems, resulting in SIGABRT or SIGSEV errors for some Siebel Server components.

To tune the Oracle Solaris kernel settings for Siebel Server

- 1 Using an editor such as vi, open the `/etc/system` file for editing.
- 2 Add or modify the following lines, which are general settings:

```
set rlim_fd_cur = 65536
set rlim_fd_max = 65536
set kernel_cage_enable = 1
```
- 3 Save your changes and exit the editor.
- 4 Restart the server computer to have the new settings take effect.

Tuning Siebel Application Object Manager Instances for Oracle Solaris

This topic is part of [“Tuning Siebel Business Applications for Oracle Solaris” on page 203](#).

Oracle Solaris computers running more than 50 Siebel Application Object Manager instances (multithreaded processes for Siebel Application Object Manager) might experience a situation where one or more of the processes do not start correctly, while the rest start and function normally. The log files for the processes that do not start will indicate that they have not started correctly. If you experience these symptoms, change the `tcp_conn_req_max_q` and `tcp_conn_req_max_q0` default values, then using the `ndd` command.

To change TCP values

1 Log in as root.

2 Issue the `ndd` command:

NOTE: The responses are noted in bold.

```
ndd /dev/tcp
name to get/set ? tcp_conn_req_max_q
val ue ? 1024
name to get/set ? tcp_conn_req_max_q0
val ue? 4096
```

3 Add the following lines to the `/etc/system` file, using any text editor such as `vi`:

```
set tcp:tcp_conn_req_max_q = 1024
set tcp:tcp_conn_req_max_q0 = 4096
```

4 Save your changes and exit the editor.

14 Monitoring Siebel Application Performance with Siebel ARM

This chapter provides an overview of performance monitoring using the Siebel Application Response Measurement (Siebel ARM) feature. It contains the following topics:

- [About Siebel Application Response Measurement on page 207](#)
- [About Siebel ARM Parameters and Variables on page 208](#)
- [Enabling and Configuring Siebel ARM on page 211](#)
- [Guidelines for Converting Siebel ARM Files on page 212](#)

NOTE: Additional diagnostic and monitoring tools are available using the Siebel Management Framework. Deploying the Siebel Management Framework in your Siebel environment is optional. For detailed information about the Siebel Management Framework, see *Siebel System Monitoring and Diagnostics Guide* and the *Siebel Installation Guide* for the operating system you are using.

Related Topic

[Chapter 15, "Analyzing Siebel ARM Data"](#)

About Siebel Application Response Measurement

Siebel ARM is a framework for capturing critical performance data in Siebel Business Applications. This data is saved in binary file format. For more information about Siebel ARM binary files and about analyzing the data in these files, see [Chapter 15, "Analyzing Siebel ARM Data."](#)

Siebel ARM captures response times at key monitoring points within the Siebel Server infrastructure. These Siebel ARM monitoring points are classified in the following distinct areas within the Siebel infrastructure:

- **Web Server Time.** Time duration a request has spent on the Web server.
- **Infra-Network Time.** Time duration between a request from the Web server and the Siebel Server (including the network time).
- **Siebel Server Time.** Time duration for the request to be processed by the Siebel Server and the Siebel Database server (time between Server Thread (SMI) and any database-layer calls).
- **Database Time.** Time duration for any Siebel Database-layer calls.
- **Application-Specific Time.** Time duration spent in application-specific areas of the infrastructure.

The Siebel ARM feature monitors system performance in the infrastructure and application-specific areas in the following list. The areas in [Table 14 on page 208](#) are listed as they appear in Siebel ARM output. The name in parentheses after the area name represents the area symbol, which also appears in Siebel ARM output.

NOTE: MWC in [Table 14](#) represents the Mobile Web Client.

Table 14. Siebel Infrastructure and Application Areas Monitored by Siebel ARM

Area Monitored by Siebel ARM	Area Monitored by Siebel ARM
■ SARM Framework (SARM)	■ Assignment Manager (AM)
■ Web Engine (SWE)	■ Fulfillment Engine (FSFULFILL)
■ Build Web Page (SWEPAGE)	■ Preventative Maintenance Engine (FSPREVMNT)
■ Web Server Plugin (SWSE)	■ Siebel Loyalty (LOY)
■ Database Connector (DBC)	■ Handheld Sync (HHSYNC)
■ Application Server (INFRA)	■ SmartScript (SMARTSCRIPT)
■ Workflow (WORKFLOW)	■ Siebel Anywhere (SIEBANYWHERE)
■ eScripts (SCRIPT)	■ Communications Channel Manager (CSMM)
■ Request Broker (SRB)	■ Communications Server Service (CSS)
■ File System Manager (FSM)	■ Customer/Order Management - Configurator (COMCFG)
■ Business Service (BUSSRVC)	■ EAI Transports (EAITRANSP)
■ Email Response (EMR)	■ MWC Profiler (MWC)
■ Security / Authentication (SEC)	■ Communications Outbound Manager (COM)
■ Object Manager (OBJMGR)	■ Universal Inbox (UINBOX)
■ Siebel Repository (SRF)	

Each area listed contains one or more subareas, which further define the timing and performance of their respective area. The number of areas and subareas present in Siebel ARM files depends on the granularity level, which is configured by the parameter SARM Granularity Level. For more information on this parameter, see [“About Siebel ARM Parameters and Variables” on page 208](#). For more information about the format of Siebel ARM files, see [“About Siebel ARM Files” on page 215](#).

About Siebel ARM Parameters and Variables

The following parameters on the Siebel Server and equivalent environment variables on the Web server enable and configure the Siebel ARM feature. The Siebel ARM parameters and environment variables are equivalent in function and similar in naming convention.

See [Table 15](#) for a listing of each Siebel ARM parameter and its equivalent environment variable. Descriptions of each parameter and environment variable follow the table. For details on enabling Siebel ARM using these parameters and environment variables, see [“Enabling and Configuring Siebel ARM” on page 211](#).

Table 15. Siebel ARM Parameters and Environment Variables

Parameter Display Name	Parameter Alias	Environment Variable Name
SARM Granularity Level	SARMLevel	SIEBEL_SARMLevel
SARM Buffer Size	SARMBufferSize	SIEBEL_SARMBufferSize
SARM Period	SARMPeriod	SIEBEL_SARMPeriod
SARM Max Number of Files	SARMMaxFiles	SIEBEL_SARMMaxFiles
SARM Data File Size	SARMFileSize	SIEBEL_SARMFileSize
Not applicable for Siebel Server	Not applicable for Siebel Server	SIEBEL_SARMThreshold

SARM Granularity Level

This parameter or equivalent environment variable specifies the amount of response measurement detail logged to Siebel ARM files, and effectively enables or disables the Siebel ARM feature. SARM Granularity Level has the following settings:

- **0 (OFF)**. This setting is the default value and disables Siebel ARM.
- **1 (ARM)**. This setting captures general application performance and is based on the application response measurement (ARM) standard. At this level, Siebel ARM collects information such as process and component boundaries, third-party software calls, database measurements, workflow execution, and script performance. Use this level for general performance monitoring.
- **2 (Detail)**. This setting captures the information at level 1 as well as detailed information such as steps of workflow execution, construction of large objects, reading of large files, and crossing significant architectural areas. Use this level for problem diagnostics.

SARM Buffer Size

The Siebel ARM framework uses a buffered data generation mechanism. Siebel ARM collects data and stores it in memory. After the in-memory data size reaches a threshold defined by the SARM Buffer Size parameter or equivalent environment variable, Siebel ARM outputs the stored data to file on a physical disk. SARM Buffer Size is specified in bytes. The default value is 5000000 (5,000,000 bytes, approximately 5 MB). Valid settings correspond to values ranging from 100,000 bytes to 50,000,000 bytes.

NOTE: Siebel ARM also outputs stored data to a file based on elapsed time, which is defined by the SARM Period parameter or equivalent environment variable. The SARM Period setting can determine the size of the data saved to the file, rather than the threshold value defined by SARM Buffer Size.

For example, if the setting of SARM Buffer Size is equivalent to 5 MB and there are five instances (processes) of the component, then the total memory used is 25 MB.

SARM Period

Siebel ARM collects data and stores it in memory. The time period specified by the SARM Period parameter or equivalent environment variable determines when Siebel ARM outputs the stored data to file on a physical disk, regardless of the value set for SARM Buffer Size. The value of SARM Period is specified in minutes, and has a default value of 3 minutes. The valid settings for this parameter range from 1 minute to 60 minutes.

NOTE: Only use SARM Period to output Siebel Server performance data based on elapsed time. Siebel ARM outputs Web server performance data based only on the SARM Buffer Size value. For more information about the SARM Buffer Size, see [“SARM Buffer Size” on page 209](#).

SARM Max Number of Files

This parameter or equivalent environment variable specifies the maximum number of Siebel ARM files created per component instance. The default value is 4, and there is no predefined upper limit to the number of files Siebel ARM creates. (SARM Data File Size configures how large a file becomes before a new file is stored on the physical disk.)

The number of active Siebel ARM files per component process is the value of SARM Max Number of Files plus 1. Siebel ARM removes the oldest file for that process only after the file representing SARM Max Number of Files plus 1 reaches the value of SARM Data File Size. For an example on how to calculate memory usage using these parameters or environment variables, see [“SARM Data File Size.”](#)

SARM Data File Size

This parameter or equivalent environment variable specifies how large a file becomes before Siebel ARM stores data in a new file on the physical disk. The value is specified in bytes. The default value is 15000000 (15,000,000 bytes, approximately 15 MB), and there is no predefined upper limit to file size.

Until the specified size is reached, Siebel ARM continues to append file segments to the current file. When the file limit is reached, Siebel ARM creates a new file. (SARM Max Number of Files configures the number of files maintained by Siebel ARM.)

When Siebel ARM reaches the file number specified by SARM Max Number of Files (that is, the number of files of size SARM Data File Size has reached the value of SARM Max Number of Files), Siebel ARM removes the first file (that is, the oldest file) when the next file reaches the SARM Data File Size limit.

Therefore, the maximum amount of disk space used is approximately SARM Max Number of Files plus 1 times the number of bytes represented by the value of SARM Data File Size. This amount of memory is per process (per component instance).

For example, if SARM Data File Size is 15 MB, SARM Max Number of Files is 4, and there are 5 instances (processes) of the component, then the maximum amount of disk space consumed is approximately 375 MB; that is, 15 MB per file, times 5 files per process, times 5 processes (instances of the component).

Siebel_SARMThreshold Environment Variable

If you are using the Siebel ARM Analyzer Tool, then set the Siebel_SARMThreshold environment variable to 0 on the Web server computer. Otherwise, user session data will not be available to the Siebel ARM Analyzer Tool, and the XML output file will be blank.

Enabling and Configuring Siebel ARM

Enabling and configuring Siebel Application Response Measurement (Siebel ARM) involves two tasks:

- Setting Siebel ARM parameters on the Siebel Server
- Setting Siebel ARM environment variables on the Web server

Setting Siebel ARM Parameters on the Siebel Server

Perform the following procedure to enable and configure Siebel ARM on the Siebel Server. By default, Siebel ARM is disabled.

NOTE: If any of the Siebel ARM parameters you want to set are not visible, then make sure the parameter Show Advanced Objects (alias ShowAdvancedObjects) is set to TRUE for the server component Server Manager (alias ServerMgr), or click Advanced to view the parameters in a server configuration view.

To enable and configure Siebel ARM on the Siebel Server

- 1 Set the parameter SARM Granularity Level (alias SARMLLevel) to a value of 1 or 2 to enable Siebel ARM on the Siebel Server.

You can enable Siebel ARM at either the enterprise, Siebel Server, or server component level. For more information on this parameter and its settings, see [“About Siebel ARM Parameters and Variables” on page 208](#).

- 2 Set the other Siebel ARM-related parameters to configure the Siebel ARM file characteristics on the Siebel Server.

You can configure Siebel ARM at the Siebel Server or server component level. For more information on these parameters, see [“About Siebel ARM Parameters and Variables” on page 208](#).

For more information on setting Siebel Server parameters using the Server Manager GUI or command-line interface, and for background information on parameter administration, see *Siebel System Administration Guide*.

Setting Siebel ARM Environment Variables on the Web Server

Perform the following procedure to enable and configure Siebel ARM on the Web server computer.

To enable and configure Siebel ARM on the Web server

- 1 Set the environment variable SIEBEL_SARMLevel to a value of 1 or 2 to enable Siebel ARM on the computer hosting the Web server. For more information, see the description for SARM Granularity Level in [“About Siebel ARM Parameters and Variables” on page 208](#).
- 2 Set the environment variable SIEBEL_SARMThreshold to a value of 0. For more information, see the description for SIEBEL_SARMThreshold in [“About Siebel ARM Parameters and Variables” on page 208](#).
- 3 Set the other Siebel ARM-related environment variables to configure the Siebel ARM file characteristics on the computer hosting the Web server. For more information on these environment variables, see [“About Siebel ARM Parameters and Variables” on page 208](#).

For more information on setting environment variables on both Windows and UNIX, see *Siebel System Administration Guide*.

Guidelines for Converting Siebel ARM Files

Review the following information as recommendations when converting Siebel ARM files.

- Set the Siebel ARM granularity level to level 1 for monitoring production deployments; set the Siebel ARM granularity to level 2 for diagnostic purposes.
- Set the SARM Max Number of Files parameter to 0 in order to disable Siebel ARM file creation. This scenario can be useful when enabling Siebel ARM for use with third-party application response measurement tools.
- Make sure the Siebel ARM feature has flushed data to the Siebel ARM file before converting the file. The Siebel ARM feature creates an empty Siebel ARM file before data is flushed to the file. For details on this process, see the descriptions for SARM Data File Size and SARM Period in [“About Siebel ARM Parameters and Variables” on page 208](#).
- Change the value of SARM Memory Size Limit (alias SARMMaxMemory) or SARM Period (alias SARMPERIOD) to a lower setting if the Siebel ARM files remain empty on a consistent basis. For details on this process, see the descriptions for SARM Data File Size and SARM Period in [“About Siebel ARM Parameters and Variables” on page 208](#).
- Make sure the Siebel ARM file name and path name, as necessary, are correct when referencing the Siebel ARM files in the commands.

(Siebel ARM Analyzer Tool only) If the Siebel ARM Analyzer Tool cannot convert large Siebel ARM files or the output file is too large, then split the Siebel ARM file by using the -p flag with the Siebel ARM Analyzer Tool. For more information on the -p flag, see [Table 24 on page 237](#).

- Concatenate Siebel ARM files to increase the amount of performance data for a given process. For example, as the Siebel ARM feature can save numerous Siebel ARM binary files for each process, concatenate these files to view performance data for multiple requests for this process. (For details on the number of files saved, see the description for SARM Max Number of Files in ["About Siebel ARM Parameters and Variables"](#) on page 208.)

TIP: Use a third-party utility to concatenate Siebel ARM files on Windows. Use the command `cat list_of_files > filename.sarm` to concatenate Siebel ARM files on UNIX.

NOTE: Only concatenate Siebel ARM files of the same process.

- Gather performance-analysis data on your Siebel application before customizing the application. These baseline measurements provide a good reference when monitoring the performance of your Siebel application after any customizations.
- Run a user session trace analysis if there are performance problems affecting an individual user during a particular session. The user trace session trace data identifies each request the user made and identifies which request required the longest time when compared to a base line.
- Use the performance aggregation data to diagnose performance at a given point in time or for a certain process. Reviewing the data by group can diagnose the area that is performing poorly. After reviewing a high-level view of the performance data, extrapolate a more detailed review by running the comma-separated value analysis. For details on running this analysis, see ["Running Siebel ARM Data CSV Conversion"](#) on page 240.
- Compile performance aggregation data over a period of time to determine a trend analysis.

15 Analyzing Siebel ARM Data

This chapter describes how to analyze Siebel Application Response Measurement (Siebel ARM) data. It contains the following topics:

- [About Siebel ARM Files on page 215](#)
- [Analyzing Siebel ARM Files Using the Siebel ARM Query Tool on page 216](#)
- [Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool on page 236](#)

NOTE: Additional diagnostic and monitoring tools are available using the Siebel Management Framework. Deploying the Siebel Management Framework in your Siebel environment is optional. For detailed information about the Siebel Management Framework, see *Siebel System Monitoring and Diagnostics Guide* and the *Siebel Installation Guide* for the operating system you are using.

Related Topic

[Chapter 14, “Monitoring Siebel Application Performance with Siebel ARM”](#)

About Siebel ARM Files

When enabled, the Siebel ARM feature saves binary Siebel ARM files in the following locations:

- Siebel Server log subdirectory on Windows: `SIEBSRVR_ROOT\log`
- Siebel Server log subdirectory on UNIX: `SIEBSRVR_ROOT/enterprises/EnterpriseServerName/SiebelServerName/log`
- Siebel Web Server Extension log subdirectory: `SWSE_ROOT\log`.

For information about Siebel ARM, see [“About Siebel Application Response Measurement” on page 207](#).

The Siebel ARM feature names the binary data files as in the following example:

`T200401081744_P001768_N0006.sarm`

where:

- T is a constant value, indicating that timing convention information follows.
- 200401081744 indicates the date and time of the Siebel ARM file. This example indicates this file was saved on January 8th, 2004 at 17:44.
- P is a constant value, indicating that process ID information follows.
- 001768 indicates the process ID on which Siebel ARM collects data.
- N is a constant value, indicating that Siebel ARM ID information follows.
- 0006 indicates the Siebel ARM log ID number for the listed process ID. Starts at 0000 and increments until it reaches 9999, at which point it wraps around to 0000.

- .sarm is the Siebel ARM file extension.

NOTE: The Siebel ARM feature creates an empty Siebel ARM file on the Web server before populating it with data. It begins storing data to these files after the feature reaches the value of the `SIEBEL_SARMFileSize` environment variable. For more information, see [“About Siebel ARM Parameters and Variables” on page 208](#).

To analyze the data contained in the binary Siebel ARM files, use one of the following tools, which present the collected data in a readable format:

- **Siebel ARM Query Tool.** This command-line tool allows you to analyze the performance data collected by Siebel ARM. The Siebel ARM Query Tool is the recommended tool for analyzing Siebel ARM data. For more information about the Siebel ARM Query Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#).
- **Siebel ARM Analyzer Tool.** This command-line tool allows you to analyze the performance data collected by Siebel ARM. However, the Siebel ARM Query Tool is recommended for use over the Siebel ARM Analyzer Tool. For more information about the Siebel ARM Analyzer Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

NOTE: Siebel Diagnostic Tool is no longer supported for Siebel CRM version 8.1 and later. The functionality previously provided by Siebel Diagnostic Tool is now available with Oracle® Application Management Suite for Siebel. For more information about using Oracle Application Management Suite for Siebel, see *Oracle Enterprise Manager Getting Started with Application Management Pack for Siebel* in the documentation library on Oracle Technology Network.

Analyzing Siebel ARM Files Using the Siebel ARM Query Tool

This topic describes how to analyze Siebel ARM files using the Siebel ARM Query Tool. This tool converts binary Siebel ARM files into readable output for analysis.

The following topics provide more information about using the Siebel ARM Query Tool:

- [“About the Siebel ARM Query Tool” on page 217](#)
- [“General Commands for the Siebel ARM Query Tool” on page 218](#)
- [“Configuring the Siebel ARM Query Tool” on page 219](#)
- [“Configuring Input for the Siebel ARM Query Tool” on page 220](#)
- [“Configuring Output from the Siebel ARM Query Tool” on page 221](#)
- [“Using Selection Filters with the Siebel ARM Query Tool” on page 224](#)
- [“Aggregating Siebel ARM Data with the Siebel ARM Query Tool” on page 232](#)
- [“Generating Histograms with the Siebel ARM Query Tool” on page 234](#)
- [“Using Macros with the Siebel ARM Query Tool” on page 235](#)

About the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#).

The Siebel ARM Query Tool is a performance-analysis command-line tool that processes the binary Siebel ARM data produced by the Siebel Server. This tool provides many command-line options that allow you to create complex queries.

Comparison with Siebel ARM Analyzer Tool

The Siebel ARM Query Tool is recommended for use over the Siebel ARM Analyzer Tool, which is described in [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#). The Siebel ARM Query Tool processes data faster than the Siebel ARM Analyzer Tool, allows you to create more complex queries, and offers many other benefits summarized in this topic.

Summary of Features of the Siebel ARM Query Tool

Some of the features of the Siebel ARM Query Tool are described below:

- **Source data specification.** You can specify individual Siebel ARM files or directories as input for the Siebel ARM Query Tool. For more information about the parameters that allow you to specify input data, see [“Configuring Input for the Siebel ARM Query Tool” on page 220](#).
- **Multiple output formats.** You can specify that multiple types of output be produced simultaneously. Supported output formats include .TXT, .XML, and .CSV. For more information about how to specify the output format, see [“Configuring Output from the Siebel ARM Query Tool” on page 221](#).
- **Data filtering capability.** You can specify selection filters that include or exclude data from analysis by the Siebel ARM Query Tool. Selection filters provide many ways to specify the types of records you are interested in. Filters are generally not mutually exclusive. That is, you do not need to choose just one filter; you can use multiple filters, in any combinations you require. For more information about using selection filters, see [“Using Selection Filters with the Siebel ARM Query Tool” on page 224](#).
- **Aggregating data.** You can specify the order of aggregation for your output. This includes rollup calculations. Aggregation is the grouping of Siebel ARM records that share common attributes and the computing of statistics on the group. For more information about aggregation, see [“Aggregating Siebel ARM Data with the Siebel ARM Query Tool” on page 232](#).
- **Histograms.** Histograms are a special type of aggregation. You can use a histogram to aggregate the results that the Siebel ARM Query Tool retrieves when you submit a query on Siebel ARM data that can return too many values if you do not use the histogram. For more information about using histograms, see [“Generating Histograms with the Siebel ARM Query Tool” on page 234](#).
- **Macro language.** The Siebel ARM Query Tool supports the use of macros. For more information about using macros, see [“Using Macros with the Siebel ARM Query Tool” on page 235](#).

- **Configurable memory for statistical accuracy.** By default, approximately 20 MB of memory is used for the Siebel ARM Query Tool's internal buffer used for statistical calculation. You can increase this amount of memory in order to increase the statistical accuracy of results.

For example, increasing the amount of memory for the tool to 500 MB reduces the statistical error level from 1% to 0.2%. For more information about configuring memory allocation for the Siebel ARM Query Tool, see [“Configuring the Siebel ARM Query Tool” on page 219](#).

- For more command-line options for the Siebel ARM Query Tool, see [“General Commands for the Siebel ARM Query Tool” on page 218](#) and [“Configuring the Siebel ARM Query Tool” on page 219](#).

General Commands for the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#). It describes the general options for use with the Siebel ARM Query Tool.

Siebel ARM Query Tool options include commands to display online help and to display progress information to the command window. [Table 16 on page 218](#) describes these commands.

Table 16. General Flags for the Siebel ARM Query Tool

Flag	Description
-hel p	Prints help for the Siebel ARM Query Tool (the sarmquery command).
-copyri ght	Prints copyright information for the Siebel ARM Query Tool.
-ti ps	Prints the command-line syntax to accomplish common aggregations, reports, and conversions.
-macrosyntax	The Siebel ARM Query Tool supports a macro language. This flag prints the syntax of the macro language. For more information about using macros, see “Using Macros with the Siebel ARM Query Tool” on page 235 .
-pl anonl y	Prints an execution plan for a query, without executing the query.
-qui et	Prints results only to the output console. No progress information appears when you specify this flag. You can also specify to save progress information to a file, as noted in “Configuring Output from the Siebel ARM Query Tool” on page 221 . The following example saves progress information to the file verbose.txt: <pre>> sarmquery -output verbose=verbose.txt -input data.sarm -aggregate time=1</pre>

Configuring the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#). It describes how you can configure the Siebel ARM Query Tool. An example of configuring the Siebel ARM Query Tool is modifying the amount of memory that the Siebel ARM query uses, to improve statistical accuracy.

[Table 17 on page 219](#) describes the flags that you can use to configure the Siebel ARM Query Tool. All the listed options are preceded by the option `-config`.

Table 17. Configuration Flags for the Siebel ARM Query Tool

Flag	Description
<code>file=<i>macro_file_name</i></code>	Specifies the file that contains a macro. For more information about using macros, see “Using Macros with the Siebel ARM Query Tool” on page 235 .
<code>macro=<i>macro_name(string)</i></code>	<p>Executes the specified macro and passes the specified string as the first argument to the macro.</p> <p>Before you can use this flag, you must specify the file that contains the macro, by using the file flag:</p> <pre>> sarmquery -config file=<i>macro_file_name</i></pre>
<code>gmt=0</code>	Parses all timestamps (on the command line) and displays all time in Greenwich Mean Time (GMT) to the command window. If you do not specify this flag, then the Siebel ARM Query Tool uses local time.
<code>gmt=[+-]<i>HHMM</i></code>	<p>Parses and reports all times as an offset from GMT. Offsets are specified in the HHMM notation, where <i>HH</i> is the hours (0023) and <i>MM</i> is the minutes (00 to 59).</p> <p>For example, to report Pacific Standard Time (PST), use the following flag:</p> <pre>> sarmquery -config gmt=-0800</pre>
<code>datlimit=<i>limit</i></code>	Specifies a maximum number of records to return, where <i>limit</i> is the maximum number. When the maximum number is exceeded, an error message (data limit exceeded) appears, and the Siebel ARM Query Tool terminates.

Table 17. Configuration Flags for the Siebel ARM Query Tool

Flag	Description
timelimit=e	<p>Specifies the maximum number of seconds that the Siebel ARM Query Tool can execute during the data processing phase. When the maximum number of seconds elapses, the Siebel ARM Query Tool exits.</p> <p>The Siebel ARM Query Tool processes in two phases:</p> <ul style="list-style-type: none"> ■ File collection phase. The tool identifies all required files. The value that you specify for the timelimit flag does not affect this phase; the time that the Siebel ARM Query Tool takes for this phase cannot be constrained. ■ Data processing phase. The value that you specify for the timelimit flag affects the length of this phase.
memlimit= <i>memory</i>	<p>Specifies the amount of memory, in megabytes, that the Siebel ARM Query Tool can allocate to an internal buffer that the Siebel ARM Query Tool uses to compute aggregated statistics.</p> <p>You can increase the amount of memory for the buffer in order to increase statistical accuracy for the results returned by the tool.</p> <p>By default, the Siebel ARM Query Tool uses approximately 20 MB of memory. You can allocate a value between 5 MB and 500 MB.</p> <p>With the default memory allocation, a statistical error level of approximately 1% can be expected in the output results. Increasing the amount of memory to 500 MB reduces the statistical error level to 0.2%.</p> <p>NOTE: It is impossible to completely eliminate statistical error.</p> <p>The amount of memory that you allocate for the Siebel ARM Query Tool does not affect the performance of the tool. Any additional memory that you allocate is applied to the buffer that the Siebel ARM Query Tool uses for statistical calculation.</p>

Configuring Input for the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#). It describes how you can specify the Siebel ARM file(s) that the Siebel ARM Query Tool uses. The Siebel ARM Query Tool converts binary Siebel ARM files into readable output for analysis.

Table 18 on page 221 describes the available input options. All the listed options are preceded by the option `-input`.

Table 18. Input Flags for the Siebel ARM Query Tool

Flag	Description
<i>sarmfile</i>	Specifies a binary Siebel ARM file (.sarm).
<i>directory</i>	Specifies a directory that contains Siebel ARM files. The Siebel ARM Query Tool processes all Siebel ARM files that it finds in the specified directory.
<i>stdin</i>	<p>A literal keyword that tells Siebel ARM query to read a list of Siebel ARM file names from standard input. You specify one Siebel ARM file or a directory name per line.</p> <p>The following are examples of valid input using <code>stdin</code> or other input flags:</p> <pre>> sarmquery . > sarmquery -input d:\sarmdata > sarmquery d:\sarmdata\srvr1 d:\sarmdata\srvr2 sarmfile1.data > dir /s /b *.sarm sarmquery -input stdin</pre>

Configuring Output from the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#).

This topic discusses the output options that are available to you when you use the Siebel ARM Query Tool. [Table 19 on page 222](#) describes the flags that you can use to configure output.

Siebel ARM Query Tool can simultaneously write output to files in the .TXT, .XML, and .CSV formats. The following example shows using the Siebel ARM Query Tool to write different types of information to different files:

```
> sarmquery -input d:\sarmdata -select subarea=infra_entry -select resp=1000 -output
header=hdr.csv -output sarm=sarm.csv -aggregate area -output agg=agg.xml
```

The preceding command writes:

- Information about the header metadata from the Siebel ARM files to the file `hdr.csv`
- Siebel ARM data to the file `sarm.csv`
- Aggregate information about the area to the file `agg.xml`

You can also specify a maximum number of lines to write to a file. When the file contains the maximum number of lines specified, the Siebel ARM Query Tool creates a new file with the same file name plus *N*, where *N* equals a number. The following example illustrates the use of this option:

```
> sarmquery -output sarm=sarmdata.csv#20000
```

This command writes Siebel ARM data in comma-separated values (CSV) format to the file named `sarmdata.csv`. The optional value `#20000` writes a maximum of 20000 lines to `sarmdata.csv`. Once `sarmdata.csv` contains 20000 lines, the Siebel ARM Query Tool writes to a new file with the name `sarmdata_0002.csv`, and so on, until all data is output.

All flags in [Table 19 on page 222](#) are preceded by the command option `-output`.

Table 19. Output Flags for the Siebel ARM Query Tool

Flag	Description
<code>fdr=filename</code>	<p>Converts all FDR files (specified using <code>-input</code>) to CSV format and writes them to the specified file.</p> <p>NOTE: For more information about FDR (flight data recorder) files, see <i>Siebel System Monitoring and Diagnostics Guide</i>.</p>
<code>fdrhdr=filename</code>	Converts all the FDR headers to CSV format and writes them to the specified file.
<code>error=filename</code>	By default, the Siebel ARM Query Tool writes output to <code>stderr.txt</code> . This command redirects error messages to the specified file.
<code>debug=filename</code>	By default, Siebel ARM query does not write debug messages. This command enables the Siebel ARM Query Tool to write debug messages to the file specified file.
<code>dbglines=filename</code>	This command enables the Siebel ARM Query Tool to write more debug information.
<code>verbose=filename</code>	By default, the Siebel ARM Query Tool writes output to <code>stderr.txt</code> . This command redirects verbose output to the specified file.
<code>map=filename</code>	Creates a file which lists all areas and subareas for Siebel ARM.
<code>header=filename</code> <code>appheader=filename</code> <code>webheader=filename</code> <code>cliheader=filename</code>	<p>The available flags are:</p> <ul style="list-style-type: none"> ■ header. Writes all header information to the specified file. ■ appheader. Writes header information from Siebel ARM files generated by the Siebel Server. ■ webheader. Writes header information from Siebel ARM files generated by the Siebel Web Server Extension (SWSE). ■ cliheader. Writes header information from the Siebel ARM files generated by clients. <p>The following example writes information about all headers to <code>all.csv</code> and then writes information for each specific part to a different file:</p> <pre>> sarmquery -input dir -output header=all.csv -output appheader=app.csv -output webheader=web.csv -output cliheader=cli.csv</pre>

Table 19. Output Flags for the Siebel ARM Query Tool

Flag	Description
sarm= <i>filename</i> appsarm= <i>filename</i> websarm= <i>filename</i> clisarm= <i>filename</i>	<p>The available flags are:</p> <ul style="list-style-type: none"> ■ sarm. Writes all Siebel ARM data to the specified file. ■ appsarm. Writes all Siebel ARM data generated by the Siebel Server to the specified file. ■ websarm. Writes all Siebel ARM data generated by the Siebel Web Server Extension (SWSE) to the specified file. ■ clisarm. Writes all Siebel ARM data generated by clients. <p>The following example writes Siebel ARM data generated by the Siebel Server and the SWSE to different files:</p> <pre>> sarmquery -output appsarm=app.csv -output websarm=web.csv</pre>
agg= <i>filename</i> iagg= <i>filename</i> sagg= <i>filename</i>	<p>The available flags are:</p> <ul style="list-style-type: none"> ■ agg. Writes the aggregation report to the specified file. ■ sagg. Writes a subset of the aggregation report to the specified file. This flag specifies exclusive metrics. ■ iagg. Writes a subset of the aggregation report to the specified file. This flag specifies inclusive metrics. <p>To use the options, you must have aggregated Siebel ARM data by using the flag -aggregate.</p>
avgincresp= <i>filename</i> pctcount= <i>filename</i> pctincresp= <i>filename</i> pctselftime= <i>filename</i>	<p>The available flags are:</p> <ul style="list-style-type: none"> ■ avgincresp. Writes the average response time to the specified file. ■ pctcount. Writes the percentage of server requests that complete within the specified time to the specified file. <p>The following example writes the percentages of server requests that complete within 100 milliseconds, between 100 milliseconds to 500 milliseconds, and so on:</p> <pre>> sarmquery -histogram resp=100, 500, 1000, 2000, 5000 -output pctcount=stdout.txt</pre> <ul style="list-style-type: none"> ■ pctselftime. Writes the percentage of time spent in internal parts of the Siebel Server. This is also known as SLA fingerprint.

Using Selection Filters with the Siebel ARM Query Tool

This topic is part of “[Analyzing Siebel ARM Files Using the Siebel ARM Query Tool](#)” on page 216. It describes how to use selection filters with the Siebel ARM Query Tool to constrain the data that the tool retrieves.

Siebel ARM files contain one or more segments. Each segment has a header and a body section. The header section contains metadata describing the data contained in the body section. For example, the metadata describes the time range when the data in the body section was collected and the amount of data in the body section.

You can use selection filters to filter the metadata and the actual performance data.

■ [Table 20 on page 224](#) describes the filters that you can use for the metadata in Siebel ARM files.

■ [Table 21 on page 227](#) describes the filters that you can use for data in Siebel ARM files.

Note the following points when you formulate queries:

■ String filter searches are case-insensitive and accept wild cards. For example, the following query retrieves Siebel ARM records associated with the users *johndoe*, *JOHNDOE*, and *JoHnDoE*:

```
> sarmquery -select user=JohnDoe
```

■ String filters also accept leading and trailing wild cards. For example, each of the following queries retrieves records associated with JohnDoe, and possibly others:

```
> sarmquery -select user="*doe"
```

```
> sarmquery -select user="j ohn*"
```

```
> sarmquery -select user="*hndo*"
```

NOTE: Wild cards can only be leading or trailing. Wild cards in the middle of a pattern do not retrieve results. For example, the pattern *jo*ndoe* does not retrieve *johndoe*.

■ You can combine selection filters to retrieve records that match multiple conditions. For example, the following selection filter retrieves all script execution records that executed for at least 5 seconds:

```
> sarmquery -select area=script -select resp=5000
```

All the options listed in [Table 20 on page 224](#) are preceded by the option -select.

Table 20. Header Filters for the Siebel ARM Query Tool

Flag	Description
<code>component=<i>componentname</i></code>	Selects headers from Siebel ARM files that were generated by the named component.

Table 20. Header Filters for the Siebel ARM Query Tool

Flag	Description
fillfactor= <i>percent</i> Or: fillfactor= <i>minpct,maxpct</i>	<p>Use this flag to specify a percentage value; the Siebel ARM Query Tool retrieves headers that are <i>percent</i> full.</p> <p>Alternatively, you can specify two arguments, a minimum percentage <i>minpct</i> and a maximum percentage <i>maxpct</i>, to retrieve a range of values.</p> <p>Generally, headers have a capacity to hold a certain number of Siebel ARM records. The percentage value is the ratio of actual number of records to that capacity.</p>
host= <i>hostname</i>	Selects headers from Siebel ARM files generated on the named host.
procid= <i>integer</i>	Selects headers from Siebel ARM files generated by a process whose ID (PID) is <i>integer</i> .
segcapacity= <i>nrecs</i> Or: segcapacity= <i>min,max</i>	<p>Select headers whose capacity to hold Siebel ARM records matches the number of records specified by <i>nrecs</i>, which is specified in number of records.</p> <p>Alternatively, you can specify two arguments, a minimum number of records and a maximum number of records, to retrieve a range of values.</p>
segduration= <i>nsecs</i> Or: segduration= <i>min,max</i>	<p>Select headers whose duration matches the number of seconds specified. A header duration is the number of seconds the file segment was in memory before it was flushed.</p> <p>Alternatively, you can specify two arguments, a minimum number of seconds and a maximum number of seconds, to retrieve a range of values.</p>
segid= <i>min</i> Or: segid= <i>min,max</i>	<p>Specify a single argument to retrieve the headers with an internal segment ID of at least <i>min</i>.</p> <p>Alternatively, you can specify two arguments, <i>min</i> and <i>max</i>, to retrieve a range of values.</p>
segsize= <i>size</i> Or: segsize= <i>min,max</i>	<p>A segment size is the size of the header and body in bytes. Specify a single argument of <i>size</i> to retrieve all headers whose segment size is <i>size</i> bytes.</p> <p>Alternatively, you can specify two arguments, <i>min</i> and <i>max</i>, to retrieve a range of values.</p>
server= <i>servername</i>	Retrieve headers from Siebel ARM files generated on the specified Siebel Server <i>servername</i> .

Table 20. Header Filters for the Siebel ARM Query Tool

Flag	Description
<code>sourcetype= <i>value</i></code>	<p>Retrieve headers from Siebel ARM files generated by the specified type of server or process.</p> <p>Alternatively, specify one of the following parameters in place of <i>value</i> to retrieve the headers:</p> <ul style="list-style-type: none"> ■ app. Retrieves headers generated by Siebel Servers. ■ web. Retrieves headers generated by the Siebel Web Server Extension (SWSE). ■ cli. Retrieves headers generated by other client programs, such as the Siebel Mobile Web Client.
<code>threshold= <i>min</i></code> Or: <code>threshold= <i>min,max</i></code>	<p>A threshold is a value represented in milliseconds. In the Siebel ARM framework, any performance record whose total response time duration is less than the threshold amount is discarded. The threshold setting at the time the Siebel ARM file was generated is saved in the header.</p> <p>Specify a single argument to retrieve all headers that contain Siebel ARM records whose threshold was at least <i>min</i> milliseconds.</p> <p>Specify two arguments to retrieve a range of values between the minimum threshold value and the maximum threshold value.</p>
<code>starttime= <i>start</i></code>	<p>Specify a start time to retrieve headers that contain Siebel ARM records that ended after the specified start time.</p> <p>NOTE: Time filters are compared against the generation start time for Siebel ARM records.</p> <p>You can specify the start time in the following ways:</p> <ul style="list-style-type: none"> ■ A string in the format <i>YYYY-MM-DD hh:mm:ss</i>. For example: <pre>> sarmquery -select starttime="2006-02-13 17:05:00"</pre> ■ A number interpreted in Universal Time Coordinated (UTC). For example: <pre>> sarmquery -select starttime=1108083900</pre> ■ A negative number to indicate the number of seconds from current time. For example, this command retrieves headers that contain data generated no more than 300 seconds ago: <pre>> sarmquery -select starttime=-300</pre>

Table 20. Header Filters for the Siebel ARM Query Tool

Flag	Description
<code>endtime=end</code>	<p>Specify an end time to retrieve headers that contain Siebel ARM records that were generated before the specified time <i>end</i>.</p> <p>NOTE: Time filters are compared against the generation end time for Siebel ARM records.</p> <p>You can specify the end time in the following ways:</p> <ul style="list-style-type: none"> ■ A string in the format <i>YYYY-MM-DD hh:mm:ss</i>. For example: <code>> sarmquery -select endtime="2006-02-13 17:05:00"</code> ■ A number interpreted in Universal Time Coordinated (UTC). For example: <code>> sarmquery -select endtime=1108083900</code> ■ A negative number to indicate the number of seconds from current time. For example, this command retrieves headers that contain data generated no more than 300 seconds ago: <code>> sarmquery -select endtime=-300</code> ■ A positive number to indicate the number of seconds after the start time. For example, this command retrieves headers that contain data generated more than 600 seconds after the start time: <code>> sarmquery -select starttime=+600</code>

Table 21 on page 227 describes the filters that you can use when you formulate a query to retrieve Siebel ARM data. All the listed options are preceded by the option `-select`.

Table 21. Siebel ARM Query Tool Flags for Record Selection

Flag	Description
<code>clickid=ID</code>	The Web server Siebel ARM files (<code>-select sourcetype=web</code>) contain a user click ID (also known as operation ID). This filter retrieves all records where <code>clickid</code> equals <i>ID</i> .
<code>foreign</code>	Use this flag to retrieve records whose parent records are from a different process. This is frequently the case for Siebel ARM data generated from batch mode components.
<code>orphan</code>	Selects records that are not root records and whose parent record was not found in the input Siebel ARM file(s).

Table 21. Siebel ARM Query Tool Flags for Record Selection

Flag	Description
<code>level=level</code> Or: <code>level=min,max</code>	Specify one argument to retrieve records whose Application Program Interface (API) level is at least equal to <i>level</i> . Specify two arguments to retrieve a range of values between the minimum and maximum level specified. The API level indicates the importance of a Siebel ARM record, as follows: <ul style="list-style-type: none"> ■ <code>level=1</code> is equivalent to <code>SARMLLevel=1</code> and indicates high level information. ■ <code>level=2</code> is equivalent to <code>SARMLLevel=2</code> and indicates detailed performance information. ■ <code>level=3</code> is equivalent to <code>SARMLLevel=3</code> and indicates debug or internal performance information.
<code>parentid=ID</code>	Select Siebel ARM records whose parent ID is <i>ID</i> .
<code>rootid=ID</code>	Selects Siebel ARM records whose root ID (also known as <i>request ID</i> in interactive mode components) is <i>ID</i> .
<code>sarmid=sarmid</code> Or: <code>sarmid=procid.sarmid</code> Or: <code>sarmid=segid.procid.sarmid</code>	Retrieves Siebel ARM records primary ID, as follows: <ul style="list-style-type: none"> ■ <i>sarmid</i> is a number uniquely identifying this performance record within a process or component, ■ <i>procid</i> is the process ID (same as <code>-select procid=procid</code>), and ■ <i>segid</i> is the segment ID (same as <code>-select segid=segid,segid</code>)
<code>sessionid=sessid</code>	All session-based component performance data contain an attribute known as the session ID. All performance data from all processes that have the same session ID are assumed to belong to a single session of a single user. The Siebel ARM Query Tool does a case-insensitive search using <i>sessid</i> . Wild cards are acceptable.
<code>taskid=taskid</code>	Select all records that associated with task <i>taskid</i> .
<code>threadid=threadid</code>	Select all records that were created by the thread whose operating system thread ID is <i>threadid</i> .
<code>user=username</code>	Select all records that were created by the user named <i>username</i> , which is typically the login name. The search is case-insensitive and wild cards are acceptable.

Table 21. Siebel ARM Query Tool Flags for Record Selection

Flag	Description
<code>area=area</code> <code>area=areacode</code> <code>subarea=sub</code> <code>subarea=subcode</code> <code>pararea=parea</code> <code>pararea=pareacode</code> <code>parsubarea=psarea</code> <code>parsubarea=pscode</code>	<p>Two filters, <code>area</code> and <code>subarea</code>, identify each Siebel ARM record.</p> <p><code>Area</code> and <code>subarea</code> are logical sections of the Siebel Server.</p> <p>For example, the Siebel Web Engine (SWE) <code>area</code> creates Web pages. Siebel ARM records associated with the SWE describe Web page creation performance. Similarly, the database connector (DBC) <code>area</code> represents the connection to the enterprise database. Siebel ARM records associated with DBC indicate the performance of database queries.</p> <p>For example, the following command retrieves all Siebel ARM records associated with database queries:</p> <pre>> sarmquery -select area=DBC</pre> <p>You can retrieve a complete list of <code>areas</code>, <code>subareas</code> and descriptions from the Siebel ARM Query Tool. For example, the following command saves the complete list to the file <code>map.csv</code>:</p> <pre>> sarmquery -output map=map.csv</pre> <p>If you know the numeric <code>area</code> or <code>subarea</code> codes, then you can use them directly in this command.</p> <p>Otherwise, you can use the string form of the symbol. When using the string form, you do not need to use the entire text. You can use a partial text string, provided that it uniquely identifies the <code>area</code> or <code>subarea</code>.</p> <p>The filters <code>pararea</code> and <code>parsubarea</code> are similar to <code>area</code> and <code>subarea</code>, except that they select Siebel ARM records whose parent <code>area</code> and <code>subarea</code> (respectively) are <code>parea</code> and <code>psarea</code>. For example, the following command retrieves all Siebel ARM records whose parent <code>area</code> is SWE:</p> <pre>> sarmquery -select pararea=swe</pre>
<code>children=0</code>	Use this flag to retrieve Siebel ARM records that have no child records.
<code>children=count</code> Or: <code>children=min,max</code>	Specify one argument to retrieve Siebel ARM records that have child records equal to at least <i>count</i> . Specify two arguments to retrieve the records that have child records between <i>min</i> and <i>max</i> .
<code>cputime=ms</code> Or: <code>cputime=min,max</code>	Specify one argument to retrieve Siebel ARM records that spent at least <i>ms</i> milliseconds on the CPU. Specify two arguments to retrieve the Siebel ARM records that spent between <i>min</i> and <i>max</i> milliseconds on the CPU.

Table 21. Siebel ARM Query Tool Flags for Record Selection

Flag	Description
depth= <i>depth</i> Or: depth= <i>min,max</i>	Specify one integer argument to retrieve Siebel ARM records that are no more than <i>depth</i> from the root node. Specify two integer arguments to retrieve a range of records.
descendants= <i>count</i> Or: descendants= <i>min,max</i>	Specify one integer argument to retrieve Siebel ARM records that have at least <i>count</i> descendants. Specify two arguments to retrieve a range of values.
instance= <i>string</i> detail= <i>string</i> int1= <i>int</i>	<p>Use these filters to retrieve Siebel ARM records where the filter is as follows:</p> <ul style="list-style-type: none"> ■ instance. This filter retrieves Siebel ARM records where instance metadata equals <i>string</i>. Instance metadata typically contains the names of things, such as the view names, screen names, user name, workflow name and script name. ■ detail. This filter retrieves Siebel ARM records where the detail metadata equals <i>string</i>. Detail metadata typically contains identification information. For example, it might contain a database row ID or the name of a business component method. ■ int1. This filter retrieves Siebel ARM records where the int1 filter equals <i>int</i>. The <i>int1</i> metadata typically contains counter values, task IDs, or other unspecified information. <p>NOTE: The values that the preceding metadata fields reference depend on the associated area or subarea.</p>
memusage= <i>excl</i> Or: memusage= <i>min,max</i>	<p>Specify one argument to retrieve all Siebel ARM records where memory allocated or deallocated was larger than <i>excl</i> bytes.</p> <p>Specify two arguments to retrieve a range of Siebel ARM records that record the memory allocation or deallocation events within the specified range.</p> <p>For example, the following command retrieves all Siebel ARM records where the associated event recorded a memory allocation of 1 MB (or larger) or a deallocation of 1 MB or larger:</p> <pre>> sarmquery -select memusage=1000000</pre>
pctcpu= <i>percent</i> Or: pctcpu= <i>min,max</i>	Specify a single argument to retrieve Siebel ARM records that indicate a CPU consumption of the percentage <i>percent</i> . Specify two arguments to retrieve a range between <i>min</i> and <i>max</i> percent.

Table 21. Siebel ARM Query Tool Flags for Record Selection

Flag	Description
resptime= <i>ms</i> Or: resptime= <i>min,max</i>	<p>Specify one argument to retrieve Siebel ARM records where the inclusive wall clock time consumed is <i>ms</i> milliseconds. Specify two arguments to retrieve a range between <i>min</i> and <i>max</i>.</p> <p>An inclusive wall clock time is the time spent in a specific part of the architecture and includes the time spent in all of its descendants. For example, the response time of a SWE area would be the time spent in constructing a Web page including the time to evaluate business component events and database access.</p>
selftime= <i>ms</i> Or: selftime= <i>min,max</i>	<p>With one argument, select Siebel ARM records where the exclusive time consumed is <i>ms</i> milliseconds. With two arguments, <i>min</i> and <i>max</i> specify a range.</p> <p>An exclusive time is time spent in an area, and only in that area, excluding the time spent in its descendants. Hence, the self time of a SCRIPT area would be the time spent in the scripting engine, and not including time spent in the database, object manager, workflow, and so on.</p>
starttime= <i>start</i> Or: endtime= <i>end</i>	<p>The semantics and syntax are similar to the time filters described for headers, except the filters select Siebel ARM records based on the timestamps.</p>
tree=all tree=parent tree=ancestor tree=children tree=descendant	<p>These are a specialized set of selection filters that do not just operate on a single record basis but on record sets. In most cases, they serve to replace a performance record with an entire set.</p> <p>In order to understand these operators, it is important to realize that Siebel ARM records form a tree of an execution thread with parent and children nodes:</p> <ul style="list-style-type: none"> ■ tree=all replaces a selected record by all records in its tree. ■ tree=parent replaces a selected record by its parent. The record itself is discarded. ■ tree=ancestor replaces a selected record by all Siebel ARM records leading from the selected record to the root of the tree. The selected record is also included in the set. ■ tree=children replaces a selected record by its immediate children. The selected record is not included. ■ tree=descendant replaces a selected record by all of its descendants, including the selected record.

Aggregating Siebel ARM Data with the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#). It describes the aggregation options available to you when you use the Siebel ARM Query Tool. Aggregation is the grouping of Siebel ARM records that share common attributes and the computing of statistics on the group.

Multiple -aggregate options can be specified, which are interpreted as subaggregations.

The Siebel ARM Query Tool can compute the total, maximum, minimum, average, and the count of contributing records for the following items:

- Inclusive Response Time
- Exclusive Response Time (self time)
- Inclusive CPU Time
- Exclusive CPU Time
- Inclusive Memory Usage
- Exclusive Memory Usage

NOTE: Aggregation is computed on Siebel ARM records retrieved by the selection filters.

All the options listed in [Table 22 on page 232](#) are preceded by the option -aggregate.

Table 22. Siebel ARM Query Tool Flags for Aggregation

Flag	Description
area	Siebel ARM records that have the same value for area are grouped together.
subarea	Siebel ARM records that have the same value for subarea are grouped together. Note that it is usually more common to also group by area when grouping by subarea, or filter on area, as with the following examples: <ul style="list-style-type: none"> > sarmquery -aggregate area -aggregate subarea > sarmquery -select area=DBC -aggregate subarea

Table 22. Siebel ARM Query Tool Flags for Aggregation

Flag	Description
instance	<p>Siebel ARM records that have the same value of instance metadata are grouped together.</p> <p>Instance metadata typically contains names of things such as scripts, workflows and views. This means that the value for the instance metadata depends on the area and subarea. As a result, aggregate instance must either be preceded by an aggregation of area and subarea or a filter on area and subarea, as with the following example:</p> <ul style="list-style-type: none"> > sarmquery -aggregate area -aggregate instance > sarmquery -select area=script -aggregate instance
server component host procid	<p>Aggregates Siebel ARM records that have the same value for:</p> <ul style="list-style-type: none"> ■ Server ■ Component ■ Host ■ Procid
user	<p>Aggregates Siebel ARM records that have the same value for user name. User name is case sensitive.</p>
sessionid	<p>Aggregates Siebel ARM records that have the same value for session ID.</p> <p>TIP: This flag can be useful when used in conjunction with the flag for user, as in the following examples:</p> <ul style="list-style-type: none"> > sarmquery -aggregate user -aggregate sessionid > sarmquery -select user=andy -aggregate sessionid
clickid	<p>Aggregates Siebel ARM records that have the same value for clickid.</p> <p>Typically in an interactive component, a user has multiple sessions and a session can have multiple requests. Sometimes a single user action on the client (browser) can result in multiple requests. In such cases, each of those requests is associated with the same unique ID known as the clickid.</p> <p>This ID is generated and associated with requests even if a user action results in a single request.</p> <p>TIP: Use this flag in conjunction with user and sessionid, as in the following examples.</p> <ul style="list-style-type: none"> > sarmquery -aggregate user -aggregate sessionid - aggregate clickid > sarmquery -select user=andy -aggregate sessionid - aggregate clickid

Table 22. Siebel ARM Query Tool Flags for Aggregation

Flag	Description
<code>time=<i>interval</i></code>	<p>Aggregates Siebel ARM records by their timestamps over the interval specified by <i>interval</i>, where <i>interval</i> is a value in minutes.</p> <p>For example, if you specify 5, then Siebel ARM records with timestamps of 12:00 to 12:05 form one aggregation, those with timestamps of 12:05 to 12:10 form another aggregation, and so on.</p> <p>The following example command displays a report to the command window that shows the average response time of the application server over intervals of 15 minutes:</p> <pre>> sarmquery -select source=app -select subarea=infra_entry -aggregate time=15</pre>

Generating Histograms with the Siebel ARM Query Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#).

Histograms are a special type of aggregation. You can use a histogram to aggregate the results that the Siebel ARM Query Tool retrieves when you submit a query on Siebel ARM data that can return too many values if you do not use the histogram.

For example, the following query:

```
> sarmquery -aggregate resptime
```

This example might retrieve millions of rows: the first row returns the number of requests that completed in one millisecond, and the second row returns the number of requests that completed in two milliseconds.

A more effective query generates a histogram, as in the following example:

```
> sarmquery -histogram resptime=100, 200, 300, 400, 500
```

This example returns six rows. The first row aggregates the number of requests that complete in less than 100 milliseconds. The second row aggregates the number of requests that completed between 100 and 200 milliseconds, and so on.

In this example, the values 100, 200, 300, and so on, are known as *bin endpoints*.

Table 23 on page 235 describes the flags that you can use with the histogram parameter. All the listed options are preceded by the option `-histogram`.

Table 23. Siebel ARM Query Tool Flags for Histograms

Flag	Description
<code>resptime= val1, val2 ... valN</code>	Creates a histogram of response (inclusive) times, where the arguments <i>val1</i> and <i>val2</i> specify the bin endpoints. <i>val1</i> and <i>val2</i> are expressed in milliseconds.
<code>selftime= val1, val2... valN</code>	Creates a histogram of self time (exclusive) times, where <i>val</i> and <i>valN</i> specify the endpoints. The values are expressed in milliseconds.
<code>pctcpu= val1, val2... valN</code>	Creates a histogram of the percentage of CPU time consumed. The arguments are expressed as integers between 0 and 100.
<code>children= val1, val2... valN</code>	Creates a histogram of width of the execution trees.
<code>depth= val1, val2... valN</code>	Creates a histogram of the depth of the execution trees.

Using Macros with the Siebel ARM Query Tool

This topic is part of “[Analyzing Siebel ARM Files Using the Siebel ARM Query Tool](#)” on page 216.

The Siebel ARM Query Tool supports the use of macros. A macro allows you to save and reuse complex queries to a file which you can subsequently use as input for the Siebel ARM Query Tool.

To create a macro, you need to:

- Write the syntax of your query in a .CFG file (for example, `macro.cfg`).
Each part of the query must appear on a separate line in the .CFG file.
- Define a name for the macro which you insert before the first line of your query in the .CFG file.

For example, the following query that retrieves data on the average login time for users can be input as one entry from the command line:

```
> sarmquery -input d:\sarmdata -select source=web -select subarea=swse_login -select tree=all -select depth=1 -aggregate user -output avgincl resp=stdout.txt
```

However, in the .CFG file the preceding query is assigned the macro name `[Logi n]` and appears as follows:

```
[Logi n]
-input d:\sarmdata
-select source=app
-select subarea=obj_mgr_sess_rel ogi n
```

```
-select tree=all  
-select depth=1  
-aggregate user  
-output avginclresp=stdout.txt
```

For the Siebel ARM Query Tool to use a macro, you need to specify two parameters:

- The location of the .CFG file that contains the macro

NOTE: A .CFG file can contain multiple macros.

- The name of the macro to execute

The following example requests that the Siebel ARM Query Tool execute the Login macro which is saved in the macro.cfg file.

```
> sarmquery -config file=macro.cfg -config macro=Login
```

The Siebel ARM Query Tool reads the input arguments for the macro as if you had specified the individual parts of the macro on the command line. This means that you can specify additional arguments that might not be in your macro. The following example executes the Login macro and in addition specifies a time slice:

```
> sarmquery -config file=macro.cfg -config macro=Login -select starttime="2004-06-10 10:00:00" -select endtime="2004-06-11 09:59:59"
```

For more information about configuration commands, see [“Configuring the Siebel ARM Query Tool” on page 219](#).

For a complete description of the macro language syntax and other options, execute the following command from the command line:

```
> sarmquery -macrosyntax
```

Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool

This topic describes how to analyze Siebel ARM files using the Siebel ARM Analyzer Tool. This tool converts binary Siebel ARM files into readable output for analysis.

NOTE: The Siebel ARM Query Tool, which is described in [“Analyzing Siebel ARM Files Using the Siebel ARM Query Tool” on page 216](#), is recommended for use over the Siebel ARM Analyzer Tool.

For more information about the Siebel ARM Analyzer Tool, see [“About the Siebel ARM Analyzer Tool” on page 237](#).

The following topics describe how to generate different types of analysis output using the Siebel ARM Analyzer Tool:

- [“Running Performance Aggregation Analysis” on page 238](#)
- [“Running Call Graph Generation” on page 239](#)

- [“Running User Session Trace” on page 239](#)
- [“Running Siebel ARM Data CSV Conversion” on page 240](#)

The following topics provide information about the different types of analysis output that the Siebel ARM Analyzer Tool generates. The analyzer tool produces output in either XML or CSV formats based on the type of conversion analysis. For more information, see:

- [“About Call Graph Generation Analysis and Data” on page 249](#)
- [“About User Session Trace Analysis and Data” on page 251](#)
- [“About Siebel ARM to CSV Conversion Data” on page 253](#)

About the Siebel ARM Analyzer Tool

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

The Siebel ARM Analyzer Tool parses the files created by the Siebel ARM feature and generates extensible markup language (XML) analytic results or comma-separated value (CSV) results. Run the Siebel ARM Analyzer Tool manually at the command-line. For an overview of the types of output this tool can generate, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

This command-line utility resides in the bin subdirectory (BIN) of the Siebel Server root directory as the executable program sarmanalyzer.exe on Microsoft Windows or sarmanalyzer on UNIX. [Table 24 on page 237](#) describes the parameters to use with this tool.

CAUTION: Monitoring the Siebel application can result in large Siebel ARM files. In some cases, the Siebel ARM Analyzer Tool cannot allocate enough memory to convert extremely large binary Siebel ARM files. In this situation, use the -p flag with the Siebel ARM Analyzer Tool to split the Siebel ARM file into smaller files. For information on this flag, see [Table 24 on page 237](#).

Table 24. Siebel ARM Analyzer Tool Flags

Flag	Description
-hel p	Use this flag with the Siebel ARM Analyzer Tool to list and describe the available flags.
-f	Use this flag with a Siebel ARM file argument to run a performance aggregation analysis. For details, see “Running Performance Aggregation Analysis” on page 238 .
-o	Use this flag to name the output path and file resulting from the analysis of the Siebel ARM binary file. Make sure to include the correct file extension based on the selected analysis, that is, either XML or CSV.
-d	Use this flag and the argument XML or CSV to indicate the type of output file format: extensible markup language (XML) or a comma-delimited list (CSV).
-a	Use this flag with the arguments AREA or DETAILS when running a performance aggregation analysis. For more information on this analysis, see “Running Performance Aggregation Analysis” on page 238 .

Table 24. Siebel ARM Analyzer Tool Flags

Flag	Description
-i	Use this flag with a directory argument when running a user session trace analysis. For more information on this analysis, see “Running User Session Trace” on page 239 .
-s	Use this optional flag to denote a start time for a user session trace. The format of the time argument is as follows: <i>yyyy-mm-dd hh:mm:ss</i> . Use this flag with the -e flag to create a time range.
-e	Use this optional flag to denote an end time for a user session trace. The format of the time argument is as follows: <i>yyyy-mm-dd hh:mm:ss</i> . Use this flag with the -s flag to create a time range.
-p	Use this optional flag to split large Siebel ARM files into smaller sizes. Use a value of 0 to 50 as the flag argument, which denotes the size, in megabytes, of the reduced files. The default value is 14 MB. The Siebel ARM Analyzer Tool uses the default value if the flag argument is 0. The split files are suffixed with <i>_Snnnnn</i> , where <i>nnnn</i> is the split sequence number.

Running Performance Aggregation Analysis

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

Use the following procedure to obtain performance aggregation analysis output. For a description of the performance aggregation analysis and output, see [“About Performance Aggregation Analysis and Data” on page 241](#). For more information on running the Siebel ARM Analyzer Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#) and [“About the Siebel ARM Analyzer Tool” on page 237](#).

To run a performance aggregation analysis

- 1 Navigate to the bin subdirectory within the Siebel Server root directory.
- 2 Run the Siebel ARM Analyzer Tool using the following command:


```

      sarmanalyzer -o output_file_name.xml -a aggregate_argument -f sarm_file_name.sarm
      
```

 where:
 - *output_file_name.xml* is the name and path of the XML output file.
 - *aggregate_argument* is either AREA or DETAILS, depending on which area you want the Siebel ARM post-process tools to aggregate data from. For more information, see [“About Performance Aggregation Analysis and Data” on page 241](#).
 - *sarm_file_name.sarm* is the name and path of the binary Siebel ARM file. Use a comma-delimited list to aggregate data from more than one Siebel ARM file.
- 3 Review the XML output in the file named *output_file_name.xml*. For more information on analyzing the performance aggregation analysis XML output, see [“About Performance Aggregation Analysis and Data” on page 241](#).

Running Call Graph Generation

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

Use the following procedure to obtain call graph generation analysis output. For a description of the call graph generation analysis and output, see [“About Call Graph Generation Analysis and Data” on page 249](#). For more information on running the Siebel ARM Analyzer Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

To run a call graph generation analysis

- 1 Navigate to the bin subdirectory within the Siebel Server root directory.
- 2 Run the Siebel ARM Analyzer Tool using the following command:

```
sarmanalyzer -o output_file_name.xml -d xml -f sarm_file_name.sarm
```

where:
 - *output_file_name.xml* is the name and path of the XML output file.
 - *-d xml* identifies the call graph generation analysis.
 - *sarm_file_name.sarm* is the name and path of the binary Siebel ARM file.
- 3 Review the XML output in the file named *output_file_name.xml*. For more information on analyzing the call graph analysis XML output, see [“About Call Graph Generation Analysis and Data” on page 249](#).

Running User Session Trace

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

Use the following procedure to obtain user session trace analysis output. Before running this analysis, manually collect Siebel Server and Web server Siebel ARM files and store in a common directory. Use this directory as an argument with the Siebel ARM Analyzer Tool. For a description of the user session trace analysis and output, see [“About User Session Trace Analysis and Data” on page 251](#). For more information on running the Siebel ARM Analyzer Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

TIP: To reduce the amount of data logged, use the time frame parameters (*-s* start time and *-e* end time).

To run a user session trace analysis

- 1 Navigate to the bin subdirectory within the Siebel Server root directory.

- 2 Run the Siebel ARM Analyzer Tool using the following command:

```
sarmanalyzer -o output_file_name.xml -u user_name -i SARM_File_Directory -s  
start_time -e end_time
```

where:

- *output_file_name.xml* is the name and path of the XML output file.
 - *user_name* is the User ID of the session you want to trace.
 - *SARM_File_Directory* is the directory containing the Siebel ARM files of the Web Server and the Siebel Server.
 - *start_time* is an optional variable defining a start time of a time range for the user session trace. The argument format is as follows: *yyyy-mm-dd hh:mm:ss*.
 - *end_time* is an optional variable defining the end time of a time range for the user session trace. The argument format is as follows: *yyyy-mm-dd hh:mm:ss*.
- 3 Review the XML output in the file named *output_file_name.xml*. For more information on analyzing user session trace XML output, see [“About User Session Trace Analysis and Data” on page 251](#).

Running Siebel ARM Data CSV Conversion

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

Use the following procedure to obtain a comma-separated value (CSV) analysis output. For a description of the CSV conversion analysis and output, see [“About Siebel ARM to CSV Conversion Data” on page 253](#). For more information on running the Siebel ARM Analyzer Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

To run a Siebel ARM data to CSV conversion analysis

- 1 Navigate to the bin subdirectory within the Siebel Server root directory.
- 2 Run the Siebel ARM Analyzer Tool using one of the following commands:

```
sarmanalyzer -o output_file_name.csv -d csv -f sarm_file_name.sarm
```

where:

- *output_file_name.csv* is the name and path of the CSV output file.
 - *-d csv* identifies the Siebel ARM data CSV conversion analysis.
 - *sarm_file_name.sarm* is the name and path of the binary Siebel ARM file or files.
- 3 Review the CSV output in the file named *output_file_name.csv*. For more information on analyzing CSV data, see [“About Siebel ARM to CSV Conversion Data” on page 253](#).
- NOTE:** Running a CSV conversion can create large output files that, in some cases, cannot be read by third-party software. Use the *-p* flag to split large Siebel ARM files. For more information on this flag, see [Table 24 on page 237](#).

About Siebel ARM Analyzer Output Files

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236.](#)

Running the Siebel ARM Analyzer Tool produces output files of either extensible markup language (XML) or comma-separated value (CSV) format depending on the type of Siebel ARM file conversion. For details on converting Siebel ARM files and running the Siebel ARM Analyzer Tool, see [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236.](#)

Use an XML editor or Web browser to view the XML output files, which result from several types of analyses. Values of timing measurements are included among the XML tags.

Use third-party software (for example, a spreadsheet program) to view the output files that result from the conversion of Siebel ARM files to CSV files. Tags and values of timing measurements are included.

Siebel ARM records all timings included in both the XML and CSV output in milliseconds.

About Performance Aggregation Analysis and Data

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236.](#)

Performance aggregation analysis is a compilation of the data contained in a Siebel ARM binary file. Siebel ARM files group performance data based on the instrumented areas. For information and a listing of instrumented areas, see [“About Siebel Application Response Measurement” on page 207.](#) For details on creating this format of Siebel ARM output, see [“Running Performance Aggregation Analysis” on page 238.](#)

Running a performance aggregation analysis of a Siebel ARM file results in an extensible markup language (XML) output file. This file contains timing data for the instrumented areas.

The amount of information contained in the performance aggregation analysis XML output is dependent on the argument used for the -a flag when performing the analysis (either AREA or DETAILS) and the setting for the SARM Granularity Level parameter. For information on this parameter, see [“About Siebel ARM Parameters and Variables” on page 208.](#)

The performance aggregation XML output file contains the following tag schema when the -a flag argument is set to DETAILS. If the -a flag argument is set to AREA when running the analysis, then the tag schema is the same minus the <NumberOfSubAreas> and <SubArea> information.

```
<Area>
  <Name>
  <Symbol >
  <NumberOfSubAreas>
  <Invocations>
    <Recursive>
    <NonRecursive>
  <ResponseTime>
    <Total >
    <Average>
    <StandardDeviation>
    +<Maximum>
    +<Minimum>
```

- <ExecutionTime>
 - <Total >
 - <Calls>
 - <Average>
 - <Maximum>
 - <Minimum>
 - <PercentOfResponse>
- <RecursiveTime>
 - <Total >
 - <Calls>
 - <Average>
 - <Maximum>
 - <Minimum>
 - <PercentOfResponse>
- <InclusiveMemory>
 - <Total >
 - <Average>
 - <StandardDeviation>
 - +<MaxAllocated>
 - +<MaxDeallocated>
- <ExclusiveMemory>
 - <Total >
 - <Average>
 - <StandardDeviation>
 - +<MaxAllocated>
 - +<MaxDeallocated>
- <SubArea>
 - <Name>
 - <Symbol >
 - <NumberOfInstances>
 - +<Invocations>
 - +<ResponseTime>
 - +<ExecutionTime>
 - +<Memory>
 - +<Instance>
 - +<Parents>
 - +<Children>
- <Parents>
 - <NumberOfParents>
 - <ParentArea>
 - <Name>
 - <Symbol >
 - +<InvocationsFromParents>
 - +<ResponseTime>
 - +<Memory>
- <Children>
 - <NumberOfChildren>
 - <ChildArea>
 - <Name>
 - <Symbol >
 - +<InvocationsOfChild>
 - +<ResponseTime>
 - +<Memory>

For descriptions of each of the tags, see [Table 25 on page 243](#).

Table 25. Performance Aggregation Analysis Tags

Tag	Description
Area	Specifies performance data captured for a specific area of the Siebel ARM architecture. There can be one or more areas captured with performance data. For more information on Siebel ARM areas, see “About Siebel Application Response Measurement” on page 207 .
Name	Name of the area containing performance data. For a listing of area names, see “About Siebel Application Response Measurement” on page 207 .
Symbol	Symbol of the area containing performance data. For a listing of symbol names, see “About Siebel Application Response Measurement” on page 207 .
NumberOfSubAreas	A count of subareas within the area that contain data. This figure also indicates the number of <SubArea> tags appearing under the particular <Area> tag.
Invocations	<p>Number of times this area was called during the monitoring period.</p> <ul style="list-style-type: none"> ■ Recursive. One of the key features of Siebel ARM is the capability to handle recursion. An example of a recursive call is if a workflow step calls a Siebel Application Object Manager function, which also invokes another workflow step. When accounting for the number of times the workflow layer is called, Siebel ARM uses two metrics: Recursive and NonRecursive. In the previous example, Recursive is 1 and NonRecursive is also 1. When calculating the response time, only the root-level call is accounted for, that is, the first workflow call to the Siebel Application Object Manager function. When calculating execution time, both calls are accounted for. ■ Nonrecursive. Number of times an instrumentation area is called. This tag helps identify how fast it takes a layer to respond to a request.

Table 25. Performance Aggregation Analysis Tags

Tag	Description
ResponseTime	<p>Specifies the time spent for a request to enter and exit an instrumentation area (layer) including calls to other child areas. Also called inclusive time in other commercial profiling tools. Other tags in this area include:</p> <ul style="list-style-type: none"> ■ Total. Total time spent by requests through this instrumentation area (layer). ■ Average. Average response time for a request. ■ StandardDeviation. The standard deviation value of request times through this area. ■ +<Maximum>. The maximum time spent by a request in this area. Expand this tag to review further details on the specific Siebel ARM node where this time was spent. For more information on Siebel ARM node tags, see “About Call Graph Generation Analysis and Data” on page 249. ■ +<Minimum>. The minimum time spent by a request in this area. Expand this tag to review further details on the specific Siebel ARM node where this time was spent. For more information on Siebel ARM node tags, see “About Call Graph Generation Analysis and Data” on page 249.
ExecutionTime	<p>Specifies the total time spent in a particular instrumentation area, not including the time spent in the descendant layers. It is also called exclusive time in other commercial profiling tools. Other tags in this area include:</p> <ul style="list-style-type: none"> ■ Total. Total time spent for a request to enter and exit an instrumentation area (layer). ■ Calls. Total number of calls including both recursive and nonrecursive calls. ■ Average. Average time spent for a request to enter and exit an instrumentation area (layer). ■ Maximum. Maximum time for a request to enter and exit an instrumentation area (layer). ■ Minimum. Minimum time for a request to enter and exit an instrumentation area (layer). ■ PercentageofResponse. Percentage of the total response time spent in the area.

Table 25. Performance Aggregation Analysis Tags

Tag	Description
RecursiveTime	<p>Specifies the total time spent in recursive calls within this area. That is, the time spent in this area when it calls itself.</p> <p>Other tags in this area include:</p> <ul style="list-style-type: none"> ■ Total. Total time spent for recursive requests. ■ Calls. Number of recursive calls. ■ Average. Average time spent for a recursive request. ■ Maximum. Maximum time spent by a recursive request. ■ Minimum. Minimum time spent by a recursive request. ■ PercentageofResponse. Percentage of the total response time spent recursively in the area.
InclusiveMemory	<p>Specifies amount of memory used by requests that enter this area and any child or descendent areas. The memory value is recorded in bytes.</p> <p>Other tags in this area include:</p> <ul style="list-style-type: none"> ■ Total. Total memory usage by requests in this area. ■ Average. Average memory usage by requests in this area. ■ StandardDeviation. The standard deviation value of memory usage in this area. ■ +<MaxAllocated>. Expand this tag to reveal further data on Siebel ARM node where maximum memory was allocated. ■ +<MaxDeallocated>. Expand this tag to reveal further data on Siebel ARM node where memory was deallocated. <p>For more information on Siebel ARM node tags, see “About Call Graph Generation Analysis and Data” on page 249.</p>

Table 25. Performance Aggregation Analysis Tags

Tag	Description
Excl usi veMemory	<p>Specifies amount of memory used by requests that enter only this area. The memory value is recorded in bytes.</p> <p>Other tags in this area include:</p> <ul style="list-style-type: none"> ■ Total. Total memory usage by a request in this area. ■ Average. Average memory usage by a request in this area. ■ StandardDeviation. The standard deviation value of memory usage in this area. ■ +<MaxAllocated>. Expand this tag to reveal further data on Siebel ARM node where maximum memory was allocated. ■ +<MaxDeallocated>. Expand this tag to reveal further data on Siebel ARM node where memory was deallocated. <p>For more information on Siebel ARM node tags, see “About Call Graph Generation Analysis and Data” on page 249.</p>

Table 25. Performance Aggregation Analysis Tags

Tag	Description
SubArea	<p>Specifies performance data captured for a specific subarea of the given area. There can be one or more subareas captured with performance data under a given area.</p> <ul style="list-style-type: none"> ■ Name. Name of the subarea containing performance data. ■ Symbol. Symbol of the subarea containing performance data. ■ NumberOfInstances. A count of instances within the subarea that contain data. This figure also indicates the number of <Instance> tags appearing under the particular <SubArea> tag. An instance is a further level of detail defining the subarea. ■ Invocations. Number of times this subarea was called during the monitoring period. ■ +<ResponseTime>. Specifies the time spent for requests to enter and exit the subarea. Expand this tag to review further timing details. These tags are the same as those defined for the area ResponseTime tag. ■ +<ExecutionTime>. Specifies the time spent in the subarea. Expand this tag to review further timing details. These tags are the same as those defined for the area ExecutionTime tag. ■ +<InclusiveMemory>. Specifies amount of memory used by requests that enter this subarea and any child or descendent areas. The memory value is recorded in bytes. Expand this tag to review further memory details. The expanded tags are the same as those defined for the area InclusiveMemory tag. ■ +<ExclusiveMemory>. Specifies amount of memory used by requests that enter only this subarea. The memory value is recorded in bytes. Expand this tag to review further memory details. The expanded tags are the same as those defined for the area ExclusiveMemory tag. ■ +<Instance>. An instance is another level of detail defining the subarea. Expand this tag to review further the instance's details. These tags are the same as those defined for the area tag. ■ +<Parents>. Specifies the parents of the subarea; that is, those areas that called the subarea. Expand this tag to review further parent subarea details. These tags are the same as those defined for the area Parents tag. ■ +<Children>. Specifies the children of the subarea; that is, those areas called by the subarea. Expand this tag to review further parent subarea details. These tags are the same as those defined for the area Children tag.

Table 25. Performance Aggregation Analysis Tags

Tag	Description
Parents	<p>Specifies the parents of the subarea; that is those areas that called the given area. This information helps identify the caller or callers of an area and the total time and number of calls the area contributed to its parent's response time. Other tags in this area include:</p> <ul style="list-style-type: none"> ■ NumberOfParents. A count of parent areas calling the given area. ■ ParentArea. Specifies performance data captured for a specific parent area of the Siebel ARM architecture. There can be one or more parent areas captured with performance data. ■ Name. Name of the parent area calling the given area. ■ Symbol. Symbol of the parent area calling the given area. ■ +<InvocationsFromParents>. Number of times the given area was called by the parent area. Expand this tag for further timing details. ■ +<ResponseTime>. Specifies the time spent for a request to enter and exit the parent area. Expand this tag for further parent area response time details. ■ +<Memory>. Specifies the amount of memory used by parent area. Expand this tag to review further parent subarea details.
Children	<p>Specifies the areas called by a parent area; that is, those areas called by the given area. Expanding an area's children information determines response time break downs within each of the children. Other tags in this area include:</p> <ul style="list-style-type: none"> ■ NumberOfChildren. A count of child areas called by the given area. ■ ChildArea. Specifies performance data captured for a specific child area of the Siebel ARM architecture. There can be one or more child areas captured with performance data. ■ Name. Name of the child area called by the given area. ■ Symbol. Symbol of the child area called by the given area. ■ +<InvocationsOfChild>. Number of times the child area was called by the given area. Expand this tag for further timing details. ■ +<ResponseTime>. Specifies the time spent for a request to enter and exit the child area. Expand this tag for further child area response time details. ■ +<Memory>. Specifies the amount of memory used by child area. Expand this tag to review further child subarea details.

About Call Graph Generation Analysis and Data

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

A call graph generation analysis constructs a map of call references. Each node in the call map represents an instrumentation instance, that is, response times for an individual request through an instrumented area. For information on instrumented areas, see [“About Siebel Application Response Measurement” on page 207](#). For details on creating this format of Siebel ARM output, see [“Running Call Graph Generation” on page 239](#).

Running a call graph generation analysis of a Siebel ARM file results in an extensible markup language (XML) output file. For a given Siebel ARM file, the Siebel ARM Analyzer Tool constructs a map with call references. Each node in the call map represents an instrumentation instance. Use this option to generate an XML file containing all the calls made by each component (if that component captures response time data).

The XML output file contains the following tag schema, which records the details of the calls.

```
<SarmNode>
  <SarmID>
  <TypeLevel>
  <RootID>
  <ParentSarmID>
  <ParentTimeID>
  <ParentProcID>
  <AreaCodeSymbol>
  <AreaDescription>
  <SubAreaCodeSymbol>
  <SubAreaDescription>
  <Count>
  <Duration>
  <Pool edMemoryUsage>
  <Pool edMemoryCalls>
  <SystemMemoryUsage>
  <SystemMemoryCalls>
  <Appl nt1>
  <Appl nt2>
  <AppStri ng1>
  <AppStri ng2>
  +<ChildNode>
</SarmNode>
```

For descriptions on each of the tags, see [Table 26 on page 249](#).

Table 26. Call Graph Generation Analysis Tags

Tag	Description
SarmNode	Data contained within this tag represents an instance of a Siebel ARM node, which is an instrumented area of the Siebel ARM architecture. Each Siebel ARM node can have zero to many nodes as its descendants.
SarmID	A unique number representing the Siebel ARM node.

Table 26. Call Graph Generation Analysis Tags

Tag	Description
TypeLevel	The granularity level at which Siebel ARM records the Siebel ARM node information. For more information on granularity level, “About Siebel ARM Parameters and Variables” on page 208 .
RootID	The SarmID of the root Siebel ARM node.
ParentSARMID	The parent SarmNode from which the request traveled.
ParentTimeID	A unique ID number that generates from the starting time of the corresponding parent Siebel ARM node.
ParentProcID	The parent process ID, that is, the OS (operating system) process ID for the Siebel component.
AreaCodeSymbol	Symbol of the instrumentation area within the Siebel architecture. For information on Siebel architecture areas, see “About Siebel Application Response Measurement” on page 207 .
AreaDescription	Name of the instrumentation area within the Siebel architecture. For information on Siebel architecture areas, see “About Siebel Application Response Measurement” on page 207 .
SubAreaCodeSymbol	Symbol of the subarea within an area of the Siebel architecture. For information on Siebel architecture areas, see “About Siebel Application Response Measurement” on page 207 .
SubAreaDescription	Name of the subarea within an area of the Siebel architecture. For information on Siebel architecture areas, see “About Siebel Application Response Measurement” on page 207 .
Count	Number of times Siebel ARM accesses this Siebel ARM Node.
SubArea	Detailed instrumentation within an area of the Siebel architecture. For example, Siebel ARM captures response time for invoking a method (Invoke Method) or executing a step (Step Execution) within a Workflow execution.
StartTime	Internal representation of the Siebel ARM record timestamp.
Duration	Total time to execute the instrumented area.
PooledMemoryUsage	Amount of memory consumed from or released to the Siebel High Performance memory allocator.
PooledMemoryCalls	The number of calls made to the High performance memory allocator.
SystemMemoryUsage	Amount of memory consumed from or released to the operating system.
SystemMemoryCalls	The number of calls made to the operating system.
UserInt1	Context information captured at the point of instrumentation. The value depends on the instrumented area.
UserInt2	Context information captured at the point of instrumentation. The value depends on the instrumented area.

Table 26. Call Graph Generation Analysis Tags

Tag	Description
UserString	Context information captured at the point of instrumentation. The value depends on the instrumented area. For example, name of the method invoked or workflow process initialized.
AppInt1 and AppInt2	Context integer value captured at the point of instrumentation. The value depends on the instrumented area.
AppString1 and AppString2	Context string value captured at the point of instrumentation. The value depends on the instrumented area. For example, name of the method invoked or workflow process initialized.
+ <ChildNode>	Expand this tag to reveal performance details on descendent nodes of the given node. The descendent nodes are defined the same as the parent node, that is, the tag definitions are the same as previously described.

About User Session Trace Analysis and Data

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236](#).

Running a user session trace analysis using Siebel ARM files from the Web server and the Siebel Server results in an extensible markup language (XML) output file. The XML output file contains detailed information on each of the SWE requests made by the user identified when running the Siebel ARM file conversion.

If the user logs onto the system multiple times, then the output shows that there are multiple sessions. The SWE requests are grouped into specific login sessions and sorted by the time the requests were made. For more details on the Siebel ARM architecture, see [“About Siebel Application Response Measurement” on page 207](#). For details on creating this format of Siebel ARM output, see [“Running User Session Trace” on page 239](#).

The XML output file contains the following tag schema, which records the details of user session trace. The user session trace data also contains the tag schema of the performance aggregation analysis. For details on those tags, see [“About Performance Aggregation Analysis and Data” on page 241](#).

```

<UserID>
  <Session>
    <SessionID>
      <UserID>
        <SWERequest>
          <ReqID>
            <TotalServerTime>
            <WebServerTime>
            <NetworkTime>
            <SiebelServerTime>
            <DatabaseTime>
            <DatabaseCalls>
            +<SiebelServerDetail>

```

For descriptions on each of the tags specific to the user session trace analysis, see [Table 27 on page 252](#). For descriptions of the tags that are also a part of the performance aggregation analysis, see [Table 25 on page 243](#).

Table 27. User Session Trace Tag Descriptions

Tag	Description
UserID	User login name. For example, SADMIN.
Session	Specifies performance data captured for a specific user session contained within this tag.
SessionID	Refers to a unique user session ID in hexadecimal format. The first component of the session ID refers to the server ID, the second refers to the process ID, and the last section to the task ID. For example: !1.2b40.182b Server ID = !1 Process ID = 2b40 (2b40 is 11072 in decimal format and represents the operating system process ID number.) Task ID = 182b (182b is 6187 in decimal format and represents the task ID number.)
UserActionID	Data contained within this tag represents a specific individual action or request of the user.
ID	Number that identifies the specific user action or request in sequence for that particular user session.
SWERequest	Specifies performance timing data for the specific user action or request.
ReqID	An incremental numeric ID number corresponding to a Siebel Web Server Extension (SWSE) plug-in request.
TotalServerTime	Total request time on the servers (includes Web server, Siebel Server, and network time).
WebServerTime	Total time spent on the Web server for a given request.
NetworkTime	Total time spent between the Web server and the Siebel Server. This time can also include some Siebel infrastructure time routing the request to the handling Siebel Server task.
SiebelServerTime	Time spent on the Siebel Server.
DatabaseTime	The time spent on the network when communicating to the database.
DatabaseCalls	Number of calls to the Siebel Server database connector layer.
+<SiebsrvrDetail>	Response time and execution time for each of the architectural areas of instrumentation for a given session. For more information on these tags, see "About Performance Aggregation Analysis and Data" on page 241 and "About Call Graph Generation Analysis and Data" on page 249 .

About Siebel ARM to CSV Conversion Data

This topic is part of [“Analyzing Siebel ARM Files Using the Siebel ARM Analyzer Tool” on page 236.](#)

CSV format is a comma-separated file without any interpretation or aggregation. The CSV file contains data organized under column headers. Use third-party software tools to view this output, for example, a spread sheet. For details on creating this format of Oracle’s Siebel ARM output, see [“Running Siebel ARM Data CSV Conversion” on page 240.](#) For a listing and description of these column headers, see the definitions of the tags for the call graph analysis in [“About Call Graph Generation Analysis and Data” on page 249.](#) Information can be reviewed and organized by these columns. See [Figure 4 on page 253](#) for an example of CSV data.

	B	C	D	E	F	G	H	I	J	K	L
1	ThreadID	IsRoot	Type(level)	RootID	ParentSarmID	ParentTimeID	ParentProcID	AreaCodeSymbol	AreaDesc	SubAreaCodeSymbol	SubAreaDesc
2	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
3	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
4	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
5	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
6	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
7	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
8	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
9	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
10	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
11	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
12	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
13	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
14	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
15	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
16	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
17	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
18	5300	N	Detail(2)	1	1	1074205585	1848	Area_SARM	SARM Framework	Sub_SARM_IO	Flush SARM Buffer To Di
19	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se
20	5300	N	Sarm(1)	1	1	1074205585	1848	Area_SWSE	Web Server Plugin	Sub_SWSE_SENDMSG	Send message to app se

Figure 4. Example of CSV Data

Index

Numerics

4GT RAM tuning 192

A

action interval, setting for Workflow Action Agent 88

action interval, setting for Workflow Monitor Agent 88

activity records

session communications and performance 77

Siebel Email Response and performance 80

agents, concurrent communications users 69

AIX

tuning IBM HTTP Server 198

tuning kernel settings 200

tuning Siebel Server 198

All Mode user property, setting for performance 120

AOM

See Siebel Application Object Manager (AOM)

applets

applet toggles, and performance 189

as memory consumer 35

grid layout, and performance 189

application configuration, network capacity 54

architecture

general flow, steps 18

generic, graphic 14

tuning architecture and infrastructure 15

user request flow, processing flow 18

architecture planning

database layout 123

database sizing guidelines 122

requirements 122

archive logging, disabling 147

asynchronous Workflow mode, pros and cons 90

B

base tables

loading data directly 121

batches

controlling size 140

bind variables, and potential problems 174

browser caching

behavior 57

managing 56

view layout caching 59

business components

Cache Data Property, and performance 185

Cache Data property, and performance 185

monitoring conditions 89

business objects layer

business objects layer, guidelines for

Cache Data Property, and performance 185

Cache Data property, and performance 185

calculated fields 186

Check No Match property, and performance 188

Primary ID fields, and performance 187

properties, using to improve picklist performance 187

setting Force Active property to FALSE 185

business objects, caching by EAI Siebel

Adapter 118

business services, invoking through

Workflow Process Manager 92

C

Cache Data Property, and performance 185

Cache Data property, and performance 185

caching

caching persistent view layout 61

disabling view layout caching 62

managing browser cache 56

memory preloading cached views 62

retrieving current view layout 63

setting view layout cache size 61

Siebel Configurator default behavior 102

Siebel Configurator, types supported 101

SQL cursor cache 32

SQL data caches 32

through Workflow Process Manager 92

tuning Siebel AOM caches 32

view layout caching 59

view layout caching in memory 60

views, caching 63

WebSphere MQ Transport, improving outbound performance 115

calculated fields, guidelines 186

- Call Center, parameter example settings** 32
- call references**
 - call graph generation analysis tags 249
 - call graph generation data, about 249
 - call graph generation, running 239
- ChannelCleanupTimer parameter** 76
- Check No Match property, and performance** 188
- child business component, monitoring conditions** 89
- Client** 162
- client, providing hardware resources** 55
- columns**
 - reusing standard columns 182
 - reusing standard columns example 182
- CommConfigCache parameter (AOM)** 73
- CommConfigManager parameter** 73
- CommConfigMgr**
 - server component 69, 73
- CommInboundProcessor**
 - server component 69, 78
- CommInboundRcvr**
 - configuring threads 80
 - server component 69, 78
- CommLogDebug parameter (AOM)** 75
- CommLogFile parameter (AOM)** 75
- CommMaxLogKB parameter (AOM)** 75
- CommMaxMsgQ parameter (AOM)** 76
- CommOutboundMgr**
 - server component 69, 78
- CommReleaseLogHandle parameter (AOM)** 75
- CommReqTimeout parameter (AOM)** 76
- CommSessionMgr**
 - component tuning, guidelines for 73
 - logging parameters 75
 - running on AOM computer 71
 - server component 68
- communications**
 - See Siebel Communications Server
- communications configuration**
 - performance 74
 - session communications, configuring logging 75
- Communications Configuration Manager**
 - server component 69
- Communications Inbound Processor**
 - server component 69, 78
- Communications Inbound Receiver**
 - server component 69, 78
- Communications Outbound Manager**
 - server component 69, 78
- Communications Server**
 - See Siebel Communications Server
- Communications Session Manager**
 - server component 68
- components** 155
- concurrent communications users, and performance** 69
- concurrent users, defined** 20
- configuration**
 - Siebel Web Client guidelines 56
- configuration parameters, MS SQL Server** 145
- Configurator**
 - See Siebel Configurator
- Configurator cache**
 - refreshing cache with attribute definition changes 111
 - refreshing cache with product changes 109
 - refreshing cache with product class changes 110
 - updating cache with attribute definition changes 110
 - updating cache with product class changes 109
- Configurator memory caching**
 - determining caching parameters rough size 107
- connection pooling**
 - assigning 41
 - assigning specialized database connection pooling 43
 - configuring SISNAPI connection pooling 47
 - configuring specialized database connection pooling 42
 - for SQL statements 41
 - multiplexing 40
 - shared connection pooling, configuration example 41
 - specialized connection pooling scenario 43
 - specialized database connection pooling 42
 - specialized database connection pooling example 43
 - specialized database connections 37
- converting**
 - ARM files, guidelines for 212
 - Siebel ARM files 236
- CPU**
 - hardware resource defined 26
 - tuning, guidelines for Siebel AOM components 27
- CSV conversion**
 - data, about 253
 - running 240
- customer configurations, guidelines for**
 - accessible views, limiting 163
 - analyzing SQL for performance, about 166

- business components in a view 164
 - business components or applets fields, limiting 164
 - cache business services 166
 - Cascade Delete, configuring 165
 - extension tables, limiting 165
 - inner joins, using 165
 - joins, limiting 165
 - Link Specification property in fields, limiting 165
 - number of records returned, limiting 165
 - number of required fields, limiting 165
 - primary ID, limiting 165
 - providing tuned PDQs 166
 - reducing scrolling 166
 - removing unneeded buttons 166
 - screen tabs, limiting 163
 - Siebel scripting performance 177
 - Siebel scripting, declarative alternatives 176
 - specifying SQL spooling in Siebel Developer Web Client 172
 - SQL queries against the database 175
 - SQL query plan example 175
 - SQL query plans, using to troubleshoot 173
 - SQL trace files, using to troubleshoot 172
 - user interface configuration 164
 - customer data, volume and performance** 69
- ## D
- data management recommendations** 153
 - data objects layer, guidelines for**
 - database indexes, sorting and searching 180
 - reusing standard columns 182
 - reusing standard columns example 182
 - data objects layer, performance guidelines**
 - multilingual LOVs query and cache performance 179
 - data, customer data volume and performance** 69
 - database authentication, and database connections** 41
 - database client libraries, as memory consumers** 35
 - database connections**
 - assigning shared connections 41
 - assigning specialized database connection pooling 43
 - configuring shared database connection pooling 40
 - configuring SISNAPI connection pooling 47
 - configuring specialized database connection pooling 42
 - nonpooled database connections 36
 - pooled database connections 37
 - shared connection pooling, configuration example 41
 - Siebel Application Object Manager
 - assumptions 36
 - specialized connection pooling scenario 43
 - specialized database connection pooling 42
 - specialized database connection pooling example 43
 - Database Extract**
 - increasing throughput 156
 - database indexes**
 - Search specification property 181
 - Sort specification property 181
 - sorting and searching 180
 - databases**
 - layout 123
 - planning guidelines 122
 - DB2 for z/OS** 150
 - DbXtract**
 - See Database Extract
 - dedicated server**
 - AOM, configuring for Siebel Configurator deployments 96
 - running CommSessionMgr 71
 - running Siebel Configurator 96
 - deleting**
 - note, using EIM 121
 - DSMaxCursorSize parameter (AOM)** 34
 - DSMaxFetchArraySize parameter (AOM)** 34
 - DSPreFetchSize parameter (AOM)** 34
- ## E
- EAI Object Manager**
 - caution, running two sessions in parallel 118
 - disabling logging 117
 - EAI Siebel Adapter, running in parallel 118
 - EAI Siebel Adapter performance**
 - analyzing SQL produced by EAI 117
 - caching business objects 118
 - caution, running two sessions in parallel 118
 - disabling logging 117
 - minimizing integration object size 117
 - reviewing scripting 117
 - running in parallel 118
 - EIM processes**
 - implementing sequence 128
 - separating by operation 130
 - testing 126
 - EIM tables**
 - about 121
 - caching tables 147
 - controlling records 141

- creating proper statistics 138
- disabling archive logging 147
- fixing fragmentation 143
- fixing Oracle Database tables 146
- indexes 137
- purging MS SQL Server tables 144
- purging Oracle Database table 146
- rebuilding an object 147
- setting FREELIST parameter 147
- updating tables 148
- using parallel data load 144
- EIM usage planning**
 - mapping into Siebel Business Applications 125
 - term definition 124
 - testing EIM processes 126
- email processing directories, managing** 80
- Email Response**
 - See Siebel Email Response
- employee applications, and message bar** 64
- EnableCDA parameter (AOM)** 34
- EnableSIBusyCursor** 64
- EnableSIBusyCursor parameter** 64
- EnableViewCache parameter** 62
- EnableViewCache parameter, disabling view layout caching** 62
- Enterprise Application Integration**
 - See Siebel EAI, tuning for performance
- eProdCfgObjMgr server component** 38, 93
- escalation action request table** 85
- escalation request table** 85
- escalation state table** 85
- eService, parameter example settings** 32
- EXEC, disabling triggers** 142
- exporting data**
 - note, using EIM 121

F

- File System Manager**
 - server component 78
- files**
 - Siebel ARM files, converting 236
- first login, and network consideration** 54
- Force Active property**
 - setting to FALSE 185
- fragmentation**
 - fixing MS SQL Server 143
 - fixing Oracle Database tables 146
- FREELIST parameter, setting** 147
- FSMSrvr**
 - server component 78

G

- graphical format, viewing log data** 84
- grid layout, and performance** 189

H

- hardware resources, and performance** 26
- heartbeat messages, using Push Keep Alive communications driver** 76
- high interactivity applications**
 - settings 56
 - view layout caching 59
- HP Apache Web Server**
 - specifying static file caching 58
- HP-UX**
 - tuning Apache Web Server 202
 - tuning kernel settings 202
 - tuning scheduler 203
 - tuning SWSE 195
- HTTP Inbound Transport, improving performance** 116

I

- IBM AIX. See AIX**
- IBM DB2 for z/OS**
 - loading process 151
 - optimization tips 148
 - performance tuning 151
- IBM HTTP Server**
 - specifying static file caching 58
 - tuning for AIX 198
- IBM WebSphere MQ Transport, improving performance**
 - outbound messages and caching 115
 - queue, testing and options 114
 - running inbound WebSphere MQ messages 115
 - setting performance tracing 116
- Id Db Size parameter (Transaction Router and Transaction Processor)** 159
- IFB file**
 - checking optimization 129
 - optimizing 129
 - using to test performance 132
 - checking optimization 129
 - optimizing 129
 - using to test performance 132
- importing**
 - note, using EIM 121
- inbound calls processed per hour** 69
- inbound communications, about** 67
- inbound email messages processed per hour** 78
- indexes**

- caching tables 147
- disabling archive logging 147
- dropping for performance 139
- on EIM tables 137
- rebuilding an object 147
- setting FREELIST parameter 147
- updating tables 148

Internet SMTP/POP3 Server communications driver 80

K

kernel settings

- tuning for AIX 200
- tuning for HP-UX 202
- tuning for Oracle Solaris 205

L

latency, defined 20

layout caching

- See view layout caching

Linux, performance tuning for an Apache Web Server 195

LogDebug parameter (CommSessionMgr) 75

LogFile parameter (CommSessionMgr) 75

logging parameters

- AOM 75
- CommSessionMgr 75

logging, configuring for Siebel Email Response 81

logical database layout 123

logs

- Workflow Agent trace files, using 84

M

mapping guidelines 125

Maximize data throughput for network applications 191

MaxLogKB parameter (CommSessionMgr) 75

MaxMTServers parameter

- calculation formula 30
- configuring 28
- configuring guidelines 30
- example settings 31
- formula variables 31
- settings, effects 29

MaxTasks parameter

- calculation formula 30
- configuring 28
- configuring guidelines 30
- example settings 31

- for CommSessionMgr component 73
- formula variables 31
- settings, effects 29

memory

- AOM memory consumers 35
- Configurator memory caching, parameters for configuring 104
- determining Configurator Cache memory rough size 107
- guidelines for tuning Siebel Application Object Manager components 27
- hardware resources defined 26
- preloading cached views 62
- running workflows in Workflow Process Manager 90
- running workflows locally 90
- using view layout caching 60

MemProtection parameter (AOM) 34

merging data

- note, using EIM 121

message bar, managing performance 64

Microsoft

- IIS, specifying static file caching 58
- Internet Explorer, recommended settings 56

Microsoft SQL Server

- See MS SQL Server

MinMTServers parameter

- calculation formula 30
- configuring 28
- configuring guidelines 30
- example settings 31
- formula variables 31
- settings, effects 29

MinTrxDBConns parameter, setting 42

MLOV query and cache performance 179

Mobile Web Client

- sizing for caching parameters 107
- tuning in a Siebel Remote deployment 160

MS SQL Server

- configuration parameters 145
- fixing table fragmentation 143
- purging tables 144
- using parallel data load 144
- using TempDB 145

MT server

- See multithreaded process

multilingual LOVs query and cache performance 179

multithreaded process

- defined 20
- and threads 22

multithreaded server

- See multithreaded process

N**network capacity**

- application configuration 54
- first login 54
- view layout caching 54

nonpooled database connections 36**NUM_IPTABLE_LOAD_CUTOFF**

- extended parameter 141

O**object, rebuilding 147****optimizing**

- EIM implement sequence 128
- general guidelines 127
- IBM DB2 for z/OS 148
- MS SQL Server 143
- Oracle Database 146

Oracle Database

- avoiding excessive table fragmentation 146
- disabling archive logging 147
- Oracle optimizer mode 146
- purging tables 146
- rebuilding an object 147
- setting FREELIST parameter 147
- updating tables 148

Oracle Enterprise Manager

- enabling and disabling for Siebel workload tagging 169
- troubleshooting poor performing SQL using Siebel workload tagging 169

Oracle HTTP Server

- specifying static file caching 58

Oracle iPlanet Web Server

- specifying static file caching 59
- tuning for Oracle Solaris 203

Oracle Solaris

- tuning kernel settings 205
- tuning Object Manager instances 205
- tuning Oracle iPlanet Web Server 203
- tuning Siebel Business applications 203
- tuning SWSE 195

outbound calls processed per hour 69**outbound communications, about 67****P****parallel data load, using for tables 144****Parallel Database Extract**

- increasing throughput 156

parallel, running tasks 142**parameters**

- Siebel ARM 208
- Siebel Server Siebel ARM parameters 209

parent business component, monitoring**conditions 89****performance**

- about and example 13
- about DB2 for z/OS 150
- architecture planning requirements 122
- batches, controlling size 140
- caching tables 147
- communications configurations, improving 74
- controlling records in tables 141
- creating proper statistics 138
- data management recommendations 153
- database layout 123
- database sizing guidelines 122
- disabling archive logging 147
- disabling triggers 142
- dropping indexes 139
- EIM implementing sequence 128
- EIM tables indexes 137
- EIM usage planning 124
- IBM DB2 for z/OS optimization 148
- IBM DB2 for z/OS, loading process 151
- IBM DB2 for z/OS, performance tuning 151
- monitoring the Siebel Server 154
- MS SQL Server 143
- MS SQL Server configuration parameters 145
- NUM_IPTABLE_LOAD_CUTOFF 141
- optimizing guidelines 127
- optimizing SQL 132
- Oracle Database 146
- purging MS SQL Server tables 144
- purging Oracle Database tables 146
- rebuilding an object 147
- recommended run parameters 153
- running tasks in parallel 142
- screen pop performance, improving 77
- session communications 69
- session communications tuning, guidelines for 72
- setting FREELIST parameter 147
- Siebel Application Object Manager (AOM), context 15
- Siebel Communications Server, about tuning 16
- Siebel Configurator, about tuning 16
- Siebel EAI, about tuning 16
- Siebel Tools, about tuning 17
- Siebel Web Client, about tuning 16
- Siebel Workflow, about tuning 16
- SQLPROFILE, using 136
- terminology 19
- testing Siebel Web Client 54
- third-party product 70

- updating tables 148
- USE ESSENTIAL INDEX HINTS 132
- USE INDEX HINTS 132
- using parallel data load 144
- USING SYNONYMS Parameter 141
- using TempDB 145
- Workflow Action Agent, setting action interval 88
- Workflow Agents, running on multiple servers 87
- Workflow Monitor Agent and Workflow Action Agent 86
- Workflow Monitor Agent, setting action interval 88
- workflow policy groups, managing Siebel Server load 86
- workflow policy groups, setting sleep interval 87
- performance aggregation analysis**
 - data, about 241
 - running 238
 - tags 243
- performance drivers, and deploying Siebel Application Object Manager** 24
- performance guidelines** 185
- performance tracing**
 - WebSphere MQ Transport 116
 - workflows 119
- persistent view layout caching and preloading**
 - logic 61
 - note, disabling 57
 - note, settings 56
- physical database layout** 123
- picklists, improving performance** 187
- policies, viewing all logs executed** 84
- Policy Frequency Analysis view** 84
- pooled database connections** 37
- predefined queries (PDQ), as memory consumer** 35
- Primary ID fields, and performance** 187
- process, defined and example** 20
- properties, improving picklist performance** 187
- purging EIM tables**
 - MS SQL Server 144
 - Oracle Database 146
- Push Keep Alive communications driver** 76

R

RAID

- performance tuning 123

RDBMS

- See database connections

records

- controlling in tables 141

ReleaseLogHandle parameter (CommSessionMgr) 75

REMOVE, disabling triggers 142

repositories

- improving loading 141

resources

- local computer resources 53
- server resources, conserving 73

response time, defined 20

rows

- recommended for single batch 140

run parameters, recommended 153

S

S_ESCL_ACTN_REQ table 85

S_ESCL_REQ table 85

S_ESCL_STATE table 85

S_ORG_EXT, improving performance 138

scalability, about and example 13

scheduler, tuning HP-UX 203

screen pop performance

- improving 77

scripting

- declarative alternatives, using 176
- performance guidelines 177

scripts, as memory consumers 35

Search Specification parameter

- minimizing usage 89
- recommended indexing fields 89

Search Specification property, managing database indexes 181

SearchSpec parameter

- minimizing usage 89
- recommended indexing fields 89

Server Request Broker

- server component and example 68
- tuning 49

session communications

- about 67
- activity creation, performance impact 77
- Communications Configuration Manager 69
- communications configuration, improving performance 74
- Communications Inbound Processor 69
- Communications Inbound Receiver 69
- Communications Outbound Manager 69
- Communications Session Manager 68
- configuring logging 75
- guidelines for tuning 72
- performance factors 69

- screen pop performance, improving 77
- Siebel Email Response 77
- third-party product modules 69
- topology 71
- session connections, improving availability** 76
- session timeouts and memory** 35
- SessPerSisnConn parameter, about** 47
- shared database connections**
 - configuring pooling 40
 - pooled database connection 37
- Siebel AOM, tuning instances for Oracle Solaris** 205
- Siebel Application Object Manager**
 - about and example 21
 - cache tuning 32
 - component tuning, guidelines for 73
 - configuring SISNAPI connection pooling 47
 - configuring thread pooling 44
 - database connections, assumptions 36
 - deployments, topology considerations 27
 - logging parameters 75
 - memory consumers 35
 - modules, communicating 22
 - parameter settings, effects 29
 - parameter values, calculation formulas 30
 - parameter, example settings 31
 - parameters, configuring guidelines 30
 - performance drivers 24
 - running Siebel Configurator in 95
 - Siebel Configurator, configuring for dedicated deployments 96
 - tuning activities 23
 - tuning for CPU and memory utilization 27
 - using thread pooling 44
- Siebel Application Object Manager (AOM)**
 - assigning shared connections 41
 - assigning specialized database connection pooling 43
 - concurrent users and performance 24
 - configuring shared connection pooling 40
 - configuring specialized database connection pooling 42
 - conserving server resources 73
 - context 15
 - hardware resources and performance 26
 - infrastructure 22
 - nonpooled database connections 36
 - parameter values, formula variables 31
 - parameters, configuring 28
 - parameters, relationship to each other 28
 - pooled database connections 37
 - running CommSessionMgr on same computer 71
 - server component defined 68
 - shared connection pooling, configuration example 41
 - Siebel application deployment 26
 - Siebel Configurator elements 93
 - specialized connection pooling example 43
 - specialized connection pooling scenario 43
 - specialized database connection pooling 42
 - think time and performance 25
 - tuning Server Request Broker 49
- Siebel Application Object Manager component, tuning** 73
- Siebel ARM**
 - ARM data CSV conversion, running 240
 - ARM to CSV conversion data, about 253
 - call graph generation analysis tags 249
 - call graph generation data, about 249
 - converting ARM files, guidelines for 212
 - data, about 241
 - enabling and configuring process 211
 - files, converting 236
 - parameters and variables 208
 - performance aggregation analysis tags 243
 - performance aggregation data, about data 241
 - Siebel Server Siebel ARM parameters 209
 - user session trace data tag 251
 - user session trace data, about 251
 - using to monitor transactions 55
- Siebel ARM analyzer tool**
 - about 237
 - ARM data CSV conversion, running 240
 - ARM to CSV conversion data, about 253
 - call graph generation analysis tags 249
 - call graph generation data, about 249
 - call graph generation, running 239
 - data, about 241
 - performance aggregation data tags 243
 - performance aggregation data, about data 241
 - running 236
 - running performance aggregation analysis 238
 - user session trace data, about 251
 - user session trace tags 251
 - user session trace, running 239
- Siebel ARM Query Tool**
 - about 217
 - aggregating Siebel ARM data 232
 - command-line parameters 218
 - configuring input 220
 - configuring output 221
 - configuring the tool 219
 - generating histograms 234

- using filters 224
- using macros 235
- Siebel Assignment Manager** 78
- Siebel Business applications**
 - configuring for Workflow performance 90
 - tuning for Oracle Solaris 203
 - tuning HP Apache Web Server for HP-UX 202
 - tuning HP-UX scheduler 203
 - tuning kernel setting for AIX 200
 - tuning kernel setting for HP-UX 202
 - tuning Siebel AOM instances for Oracle Solaris 205
 - tuning Siebel Server for AIX 198
 - tuning Siebel Web Server Extension for Oracle Solaris 195
- Siebel Business Applications, mapping into guidelines** 125
- Siebel Call Center, parameter example settings** 32
- Siebel Client, defined** 51
- Siebel Communications Server**
 - communications supported activities 67
 - component topology 71
 - tuning architecture and infrastructure 16
- Siebel Configurator**
 - caching, default behavior 102
 - caching, types of 101
 - components 93
 - customizable products and classes 100
 - defining topology, considerations for 95
 - determining rough size of Configurator cache 107
 - memory caching parameters 104
 - performance factors 94
 - running dedicated servers 96
 - running Siebel Application Object Manager component in 95
 - Siebel Application Object Manager, configuring for dedicated deployments 96
 - tuning 99
 - tuning, about 16
 - tuning, guidelines for 98
- Siebel Database, communicating with** 22
- Siebel Developer Web Client, specifying SQL spooling** 172
- Siebel EAI, tuning**
 - about 113
 - All Mode Sort user property, setting 120
 - checking disks 120
 - creating business components 120
 - EAI Siebel Adapter performance 117
 - guidelines for 113
 - HTTP Inbound Transport performance 116
 - IBM WebSphere MQ Transport performance 114
 - improving Workflow Process Manager performance 119
 - optimizing database queries 120
 - optimizing messages 120
 - tuning, about 16
 - turn off logging 120
 - virtual business component performance 118
- Siebel Email Response**
 - activity creation and performance 80
 - CommInboundRcvr threads, configuring 80
 - email processing directories, managing 80
 - inbound email processed per hour 78
 - infrastructure 77
 - key server components 78
 - logging, configuring 81
 - Siebel Assignment Manager 78
 - third-party email server 78
 - topology 79
 - tuning, guidelines for 79
 - volume of customer data 78
- Siebel Enterprise Application Integration**
 - See Siebel EAI, tuning
- Siebel eService, parameter example settings** 32
- Siebel File System**
 - Siebel Configurator component 93
- Siebel Internet Session application programming interface**
 - See SISNAPI connection pooling
- Siebel modules, supporting multiple modules** 52
- Siebel Product Configuration Object Manager** 38, 93
- Siebel Remote**
 - about 155
 - data synchronization, guidelines for
 - optimizing between Siebel Mobile Client and Siebel Remote 162
 - increasing throughput for Database Extract and Parallel Database Extract 156
 - routing model 162
 - server components 155
 - server components, tuning 156
 - tuning mobile web client 160
- Siebel scripting**
 - declarative alternatives, using 176
 - performance guidelines 175, 177
- Siebel Server**
 - communications components 68
 - Communications Configuration Manager 69
 - Communications Inbound Processor 69

- Communications Inbound Receiver 69
- Communications Outbound Manager 69
- Communications Session Manager 68
- Siebel ARM parameters 209
- third-party product modules 69
- tuning for AIX 198
- tuning kernel setting for Solaris 205
- Workflow Agents, running on multiple servers 87
- workflow policy groups, creating to manage 86
- Siebel Server, monitoring** 154
- Siebel Tools, about tuning** 17
- Siebel user request flow**
 - processing flow 18
 - steps 18
- Siebel Web Client, tuning**
 - about 16
 - client hardware resources 55
 - configuration guidelines 56
 - disabling view layout caching 62
 - guidelines for 53
 - HP Apache Web Server, static file caching 58
 - IBM HTTP Server, static file caching 58
 - local computer resources 53
 - managing browser cache 56
 - memory, preloading cached views 62
 - message bar, managing performance 64
 - Microsoft IIS, static file caching 58
 - Oracle HTTP Server, static file caching 58
 - Oracle iPlanet Web Server, static file caching 59
 - persistent view layout caching 61
 - retrieving current view layout 63
 - setting view layout cache size 61
 - Siebel Client, defined 51
 - static file caching 57
 - supporting multiple Siebel modules 52
 - testing performance 54
 - tuning system components 55
 - view layout caching 59
 - views, and layout caching 63
 - Web server and network capacity 54
- Siebel Web Engine (SWE)**
 - user session trace, running 239
- Siebel Web Server Extension**
 - communicating with 22
 - tuning for Oracle Solaris 195
- Siebel Workflow, tuning**
 - about 83
 - architecture and infrastructure 16
 - components 83
 - configuring 90
 - escalation action request table 85
 - escalation request table 85
 - escalation state table 85
 - monitoring memory overhead 90
 - parent and child business components, monitoring 89
 - performance tracing 119
 - Policy Frequency Analysis view, using 84
 - Search Specification parameter, minimizing usage 89
 - Siebel Server load, creating workflow policy groups 86
 - tuning Workflow Process Manager 91
 - work policy groups, setting sleep interval 87
 - Workflow Action Agent, setting action interval 88
 - Workflow Agents, running on multiple servers 87
 - Workflow Monitor Agent and Workflow Action Agent 86
 - Workflow Monitor Agent, setting action interval 88
 - workflow policies, monitoring 84
 - workflow policies, using logs and files 84
- Siebel Configurator**
 - and database connection pooling 38
- SISNAPI connection pooling, configuring sleep interval, setting for workflow policy groups** 87
- Solaris. See Oracle Solaris**
- Sort Specification property, managing database indexes** 181
- specialized database connections** 37
- SQL**
 - connection pooling for SQL statements 41
 - note, support 121
 - time-intensive statements 136
- SQL cursor cache** 32
- SQL data caches** 32
- SQL logging and SQL tagging for Application Object Manager components, about** 167
- SQL, analyzing for performance**
 - about 166
 - Siebel Developer Web Client, specifying SQL pooling 172
 - SQL queries against the database 175
 - SQL query plan example 175
 - SQL query plans, using to troubleshoot 173
 - SQL trace files, using to troubleshoot 172
- SQL, poorly performing query** 89
- SQL, produced by EAI** 117
- SQLPROFILE, using** 136
- SRBroker (Server Request Broker)**
 - server component and example 68

- tuning 49
- standard column**
 - reusing 182
 - reusing example 182
- standard interactivity client, and best performance** 55
- static file caching**
 - about specifying 57
 - specifying on HP Apache Web Server 58
 - specifying on IBM HTTP Server 58
 - specifying on Microsoft IIS 58
 - specifying on Oracle HTTP Server 58
 - specifying on Oracle iPlanet Web Server 59
- synonyms**
 - USING SYNONYMS Parameter 141
- T**
- tables**
 - caching 147
 - updating 148
- task**
 - defined 20
 - used interchangeably with thread 22
- TempDB, using** 145
- term definition, guideline** 124
- terminology** 19
- testing EIM processes** 126
- testing performance**
 - Siebel Web Client 54
- think time and performance** 25
- think time, defined and example** 20
- third-party products**
 - email server 78
 - performance 70
 - product modules 69
- thread**
 - defined 20
 - used interchangeably with task 22
- thread pooling for AOM**
 - note, handling overhead 44
 - note, recommendation 44
- thread pooling for Siebel Application Object Manager**
 - configuring 44
 - using 44
- throughput, defined** 20
- tools**
 - See Siebel Tools
- Transaction Router**
 - increasing throughput 159
 - tuning 157
- transactions per second (TPS), throughput defined** 20

- TrickleSync** 162
- triggers, disabling** 142
- troubleshooting**
 - about DB2 for z/OS 150
 - batches, controlling size 140
 - caching tables 147
 - controlling in tables 141
 - creating proper statistics 138
 - data management recommendations 153
 - disabling archive logging 147
 - disabling triggers 142
 - dropping indexes 139
 - EIM tables indexes 137
 - IBM DB2 for z/OS optimization 148
 - IBM DB2, loading process 151
 - monitoring the Siebel Server 154
 - MS SQL Server 143
 - MS SQL Server configuration
 - parameters 145
 - NUM_IPTABLE_LOAD_CUTOFF 141
 - optimizing SQL 132
 - Oracle Database 146
 - purging MS SQL Server tables 144
 - purging Oracle Database tables 146
 - rebuilding an object 147
 - recommended run parameters 153
 - running tasks in parallel 142
 - setting FREELIST parameter 147
 - SQL query plan example 175
 - SQL query plans, using to troubleshoot 173
 - SQL trace files, using to troubleshoot 172
 - SQLPROFILE, using 136
 - updating tables 148
 - USE ESSENTIAL INDEX HINTS 132
 - USE INDEX HINTS 132
 - using parallel data load 144
 - USING SYNONYMS Parameter 141
 - using TempDB 145
 - workload tagging, using to troubleshoot 168
- tuning system components** 55

U

- UNIX, performance tuning**
 - Siebel Web Server Extension for Oracle
 - Solaris 195
 - tuning Business applications for Oracle
 - Solaris 203
 - tuning HP-UX scheduler 203
 - tuning IBM HTTP Server for AIX 198
 - tuning kernel setting for HP-UX 202
 - tuning kernel settings for AIX 200
 - tuning Siebel AOM instances for Oracle
 - Solaris 205

- tuning Siebel Server Extension for AIX 198
- tuning Web server for HP-UX 202
- UNIX, performance tuning for an Apache Web Server** 195
- UPDATE STATISTICS parameter**
 - running tasks in parallel 142
- USE ESSENTIAL INDEX HINTS** 132
- USE INDEX HINTS parameter**
 - using 132
- user communications actions per second** 69
- user interface objects layer, guidelines for**
 - applet toggles, and performance 189
- user interface objects layer, performance guidelines**
 - grid layout, and performance 189
- user per AOM, as memory consumer** 35
- user request flow**
 - processing flow 18
 - steps 18
- user session trace analysis**
 - data, about 251
 - running 239
 - tag descriptions 251
- USING SYNONYMS parameter**
 - using 141

V

- variables**
 - Siebel ARM 208
- view caching**
 - See view layout caching
- view layout caching**
 - about 59
 - caching persistent view layout 61
 - disabling view layout caching 62
 - memory, preloading cached views 62
 - network capacity 54
 - retrieving current view layout 63
 - setting cache size 61
 - view layout caching in memory 60
 - views, caching 63
- ViewPreloadSize parameter** 62
- ViewPreloadSize parameter, setting** 62
- views, and layout caching** 63
- virtual business component performance, improving** 118
- virtual memory management** 200
- vmtune values, changing** 200
- volume of customer data** 78

W

- Web client**
 - See Siebel Web Client

- Web server**
 - tuning IBM HTTP Server for AIX 198
- WebSphere MQ Transport, improving performance**
 - outbound messages and caching 115
 - queue, testing and options 114
 - running inbound WebSphere MQ messages 115
 - setting performance tracing 116
- WebTemplateVersion parameter** 61
- WebTemplateVersion parameter, and persistent layout caching** 61
- Windows**
 - IIS, specifying static file caching 58
 - Internet Explorer, recommended settings 56
- workflow**
 - See Siebel Workflow
- Workflow Action Agent**
 - defining 86
 - setting action interval 88
- Workflow Agent**
 - process, monitoring 84
 - running on multiple servers 87
 - trace files, using 84
- Workflow Monitor Agent**
 - defining 86
 - policies, viewing all policies executed 84
 - setting action interval 88
- workflow policies**
 - about 83
 - log and files 84
 - monitoring 84
- workflow policy groups**
 - optimal sleep interval, setting 87
 - using to manage Siebel Server load 86
 - Workflow Monitor Agent and Workflow Action Agent 86
- Workflow Process Manager**
 - caching business services 92
 - performance issues 119
 - tuning 91
 - using to monitor memory overhead 90
 - workflow, performance tracing 119
- workflow processes**
 - about 83
- workflow processes, tuning**
 - configuring Siebel Business applications 90
 - monitoring memory overhead 90
 - parent and child business components, monitoring 89
 - Search Specification parameter, minimizing usage 89
 - Workflow Process Manager 91
- workload tagging**

- enabling and disabling using the Siebel application 169
- enabling and disabling using the Siebel Server Manager 170
- enabling and disabling, about 169
- enabling and disabling, requirements for 169
- using to troubleshoot poor performing SQL at the database level 168, 171

